Shared Intelligence and Misaligned Firm Entry

Feng Chi* Nathan Yang†

May 2, 2026

Abstract

Retail entry decisions are increasingly mediated by a small set of shared intelligence platforms whose forecasts inherit a common processing component orthogonal to local market fundamentals. This paper develops a Bayesian entry game in which retail chains choose between proprietary research and a shared platform that reduces idiosyncratic forecast noise while introducing a shared shock. The model admits a closed-form linear equilibrium under heterogeneous platform adoption. The central result is a misalignment theorem: equilibrium adoption of shared platforms is excessive relative to a planner who internalizes the systemic-risk externality, and the direction of the wedge is invariant to whether retail competition is agglomerative or business-stealing in nature. A multi-vendor extension shows that the welfare cost of platform concentration is governed by an effective concentration index that distinguishes substantive vendor diversity from cosmetic variation around a common data core.

Keywords: Retail Entry; Shared Intelligence Platforms; Algorithmic Monoculture; Coordination; Vendor Concentration; Systemic Risk

*University of Illinois Urbana-Champaign, Department of Economics. Email: fchi@illinois.edu.
†University of Illinois Urbana-Champaign, Gies College of Business. Email: ncyang@illinois.edu.

1 Introduction

Retail location decisions used to be a matter of local knowledge. A regional grocer learned its trade area through years of operating stores in adjacent counties; a restaurant chain hired a market researcher who walked the parking lots of competitors and counted cars. The choice of where to open the next store was the output of a research process that was, at its core, idiosyncratic to the chain making it. Two firms looking at the same county would reach different conclusions because they processed the available information differently.

This is no longer how retail entry works. CoStar, Placer.ai, Buxton, and the analytics arms of CBRE and JLL now sell forecasts of trade-area demand, foot traffic, competitive density, and demographic projections to thousands of subscribers simultaneously. A national chain considering a new location, a regional grocer evaluating a strip mall, and a franchisee scouting a corner lot may all consult the same platform, receive the same model output, and update toward the same conclusion. The Bank of England's Financial Policy Committee has noted that "the potential for AI-based participants to take increasingly correlated positions is an important consideration" for systemic stability (Bank of England Financial Policy Committee 2025); the same observation applies to retail markets, where shared platforms now mediate decisions that were once the product of independent investigation.

The economic question this raises is not whether shared platforms forecast better than custom research. They may or may not, and the answer varies across markets. The question is what changes about industry-level entry behavior when forecasts that used to be independent are routed through a common processing layer. When two firms subscribe to the same platform, their forecast errors are no longer independent: they share a component that originates in the platform itself, in

the training data the platform used, in the modeling choices its engineers made, and in the periodic updates the platform pushes to all subscribers at once. This shared component is correlated across firms but uncorrelated with the local fundamentals the platform purports to forecast. It is a property of the processing technology, not of the market being analyzed.

Standard models of retail entry do not generate correlation of this kind. In the canonical entry-game formulations of Bresnahan and Reiss (1991), Seim (2006), Jia (2008), and Holmes (2011), firms receive private signals about local profitability and decisions are correlated either through agglomeration or business-stealing forces or through informativeness of signals about a common fundamental. In all of these, correlation in actions reflects either strategic interaction or correlation in beliefs about the underlying state. Neither channel covers the platform case. A platform forecast does not become more informative about local fundamentals as more firms adopt it; what changes is that the same forecast errors propagate to more decisions. The relevant correlation is technological in origin, not strategic and not informational.

The closest theoretical antecedent is the social-value-of-information literature initiated by Morris and Shin (2002) and developed by Angeletos and Pavan (2007), Hellwig and Veldkamp (2009), and Myatt and Wallace (2012). There, common signals introduce correlation across firms, and the welfare implications depend on a tradeoff between coordination value and information loss. The signal structure we study departs from this benchmark in a specific structural way: the common component of the platform signal is orthogonal to the fundamental. No setting of any precision parameter in the standard frameworks yields a signal of this kind. The orthogonality is a feature of the processing technology, not a limit of an information-quality parameter, and it removes the offsetting force that generates non-monotone comparative statics in prior work. Greater reliance on the shared platform monotonically raises the variance of aggregate entry without a compensating

improvement in fundamental tracking, and the welfare wedge between private and social adoption acquires a determinate sign that does not depend on the strategic structure of competition.

This paper develops a Bayesian entry game that isolates this mechanism. A continuum of potential entrants chooses whether to enter a representative market and, prior to that, whether to subscribe to a shared platform or to rely on custom research. Profits depend on local fundamentals, on the aggregate entry rate (with strategic interaction allowed to be either agglomerative or business-stealing), and on a private cost shock that delivers a tractable threshold rule. The shared platform reduces idiosyncratic forecast noise but introduces a common processing shock orthogonal to fundamentals. The model admits a closed-form linear equilibrium under heterogeneous adoption, which we use to characterize how aggregate entry decomposes into a fundamental-tracking component and a platform-driven component, how the private return to platform adoption varies with penetration, and how equilibrium adoption compares to social optima under several welfare specifications.

The first analytical result identifies the equilibrium loadings of aggregate entry on the fundamental and on the shared shock. The loading on the shared shock arises only from platform adopters and scales with their share, so the variance contribution of the platform to aggregate entry is monotone in penetration regardless of the strategic-interaction parameter. The second result characterizes how the private adoption gain depends on penetration: it rises with adoption under agglomerative competition, falls under business stealing, and is independent of penetration when neither force is active. This monotonicity result, together with a β-independent threshold for adoption at zero penetration, organizes the analysis of multiple equilibria that follows.

The central normative result is a misalignment theorem. When the planner internalizes the systemic-risk externality generated by the shared platform—the variance of aggregate entry at-

tributable to the shared shock, which firms do not bear individually—equilibrium adoption strictly exceeds the social optimum for any sign of the strategic-interaction parameter. The proof reduces the question to monotonicity of a single object, the squared aggregate loading on the shared shock, which we establish via an explicit identity that exhibits the relevant derivative as a sum of manifestly nonnegative terms. The invariance of the wedge sign matters because it does not require taking a position on whether retail competition is agglomerative or business-stealing in any particular industry, an ambiguity that makes empirical identification of the strategic parameter notoriously difficult. Whatever the local strategic environment, the systemic-risk externality always points in the direction of over-adoption.

A multi-platform extension generalizes the misalignment result to the case of multiple competing platforms with cross-platform correlation in their processing shocks. The wedge magnitude scales with an effective concentration index that takes its lower bound when platforms are structurally independent and its upper bound when they share a common data or model core. The substantive content is that vendor diversity reduces the externality only insofar as it represents diversity in processing, not diversity in branding. Multiple platforms built on a common foundation model, on the same training corpus of mobile-device location pings, or on identical demographic projections from a single vendor deliver the same aggregate-volatility profile as a single platform.

The framework also delivers a typology of adoption regimes, including a coordination trap under sufficiently strong agglomerative competition. In this regime, the platform is privately unattractive in isolation, because its forecast precision does not justify the variance of the shared shock, but becomes attractive once a sufficient mass of rivals has adopted it, because the value of forecasting rival behavior dominates the value of forecasting fundamentals. The high-adoption equilibrium is sustained purely by the coordination motive. This regime maps onto retail sectors with strong

agglomeration economies—food halls, outlet centers, auto malls, and clusters of competing fast-food restaurants—where a platform that reveals rival intentions can become essential to incumbents whether or not it is informative about underlying demand.

The remainder of the paper is organized as follows. Section 2 develops the model. Section 3 characterizes the linear equilibrium with heterogeneous adoption and derives the variance decomposition of aggregate entry. Section 4 establishes the team-efficiency benchmark and the misalignment theorem. Section 5 characterizes the four adoption regimes and the coordination trap. Section 6 develops the multi-platform extension and the effective concentration index. Section 7 discusses implications and directions for empirical work. Proofs are collected in the Appendix.

2 Model

The model formalizes the entry decision of a retail chain that must choose both whether to enter a representative local market and which information technology to use when forming its forecast of market profitability. The environment is deliberately sparse: a single representative market, two information technologies, and a continuum of potential entrants. This setup isolates the structural feature of interest—a shared processing layer that introduces correlated errors across adopters—while preserving enough tractability to deliver closed-form equilibrium objects.

2.1 Intuition: The Mechanism in its Simplest Form

Before introducing the full game, it is useful to isolate the central mechanism in a stripped-down environment. Consider a continuum of retail chains, each independently deciding whether to enter a local market with profitability θ. There is no strategic interaction; each chain's profit depends

only on its own entry decision and on the realization of θ. A chain enters if its posterior expectation of profitability exceeds zero.

Each chain chooses between two forecast technologies. Custom research delivers a signal $\theta + \varepsilon$, with idiosyncratic noise of variance σ_C^2. The shared platform delivers a signal $\theta + \xi + u$, where ξ is a common shock realized once and shared across all platform subscribers, and u is idiosyncratic noise of variance σ_A^2. The shared shock ξ has variance σ_ξ^2 and is orthogonal to θ. Adoption shares are exogenous: a fraction q of chains adopts the platform, while the remainder uses custom research.

Without strategic interaction, each chain's entry probability is determined by its posterior mean. Aggregating across chains, idiosyncratic noises wash out by the law of large numbers, while the common shock persists. The aggregate entry rate takes the form

$$\bar{x} = K_0 + K_\theta(q)\,\theta + K_\xi(q)\,\xi, \qquad K_\xi(q) = q \cdot \frac{\sigma_\theta^2}{V_A},$$

where $V_A = \sigma_\theta^2 + \sigma_\xi^2 + \sigma_A^2$ is the unconditional variance of the platform signal and K_0 is a constant determined by the prior mean.

The variance of aggregate entry around its mean decomposes into a fundamental-tracking component proportional to σ_θ^2 and a pure-correlation component equal to $K_\xi(q)^2\sigma_\xi^2$. Each adopter internalizes its own forecast accuracy but not its contribution to the second component, which arises only at the aggregate level and does not enter individual profits. If society values the variance of aggregate entry around the fundamental—because downstream agents such as commercial mortgage lenders, municipal tax authorities, or affected consumers bear costs that scale with synchronized over- or under-entry—then the marginal social cost of an additional platform adopter

exceeds the marginal private cost by an amount proportional to $d[K_\xi(q)^2\sigma_\xi^2]/dq$, which is strictly positive whenever $\sigma_\xi^2 > 0$.

The mechanism is therefore present even in this stripped-down environment, without strategic interaction, endogenous adoption, or heterogeneous information. The remainder of the section introduces these features. Strategic interaction allows the shared shock to enter expectations of aggregate entry and feed back into individual decisions; endogenous adoption determines the equilibrium share q^* as a fixed point; and heterogeneous information governs how non-adopters infer adopter behavior. Each enrichment serves a specific purpose: strategic interaction enables the four-regime typology and the coordination trap, endogenous adoption permits the misalignment theorem to be stated as an equilibrium comparison, and heterogeneous information makes the multi-platform extension informative about cosmetic versus substantive vendor diversity.

2.2 Environment

A unit continuum of retail chains $f \in [0,1]$ decides whether to enter a representative local market. The entry decision is binary: $x_f \in \{0,1\}$. Profits are

$$\Pi_f = (\theta + \delta\,\bar{x} + \varepsilon_f)\,x_f,$$

where $\theta \sim \mathcal{N}(0,\sigma_\theta^2)$ is unobserved local profitability, $\bar{x} \equiv \int_0^1 x_g\,dg$ is the aggregate entry rate, and $\varepsilon_f \sim \text{Uniform}[-\bar{\varepsilon},\bar{\varepsilon}]$ is a private cost shock independently distributed across chains. The parameter $\delta \in (-1,\infty)$ governs strategic interaction. Under $\delta > 0$, profitability rises with rival entry, capturing agglomeration economies that arise when retail clusters draw traffic that no single store could attract on its own—food halls, outlet centers, auto malls, and clusters of competing fast-

food restaurants are familiar examples. Under $\delta < 0$, profitability falls with rival entry, capturing business-stealing in markets with limited demand. The restriction $\delta > -1$ guarantees existence of a unique linear equilibrium below.

The uniform distribution of ε_f is a tractability device that is standard in the global-games literature initiated by Carlsson and van Damme (1993) and Morris and Shin (1998). It makes the entry probability a piecewise-linear function of the posterior mean of profitability. We work in the global-games limit in which $\bar{\varepsilon}$ is large relative to the standard deviations of all signals, so that the truncation regions where the entry probability hits 0 or 1 occur with negligible probability.[1] For the substantive analysis, we normalize $\bar{\varepsilon} = 1/2$ so that the slope of the entry probability with respect to the posterior mean is unity; this is a unit choice that affects no qualitative result.

2.3 Information and Processing Technologies

Chains differ in how they form forecasts of θ. A chain using custom research observes

$$s_f^C = \theta + u_f^C,$$

where $u_f^C \sim \mathcal{N}(0, \sigma_C^2)$ is idiosyncratic noise. A chain using the shared platform observes

$$s_f^A = \theta + \xi + u_f^A,$$

where $u_f^A \sim \mathcal{N}(0, \sigma_A^2)$ is idiosyncratic platform noise and $\xi \sim \mathcal{N}(0, \sigma_\xi^2)$ is the common processing shock shared by all platform subscribers. All shocks are mutually independent and independent of

[1] Formally, the linearization is exact whenever $|\mathbb{E}_f[\theta] + \delta \mathbb{E}_f[\bar{x}]| \leq \bar{\varepsilon}$. We maintain $\bar{\varepsilon} > 4\sqrt{\mathrm{Var}(\theta) + \mathrm{Var}(\bar{x})}$ as a sufficient condition for negligible truncation; this bound holds throughout the parameter region examined in the numerical exercises of Sections 5 and 6.

θ.

The signal structure encodes a precise distinction between the two technologies. Both deliver noisy observations of θ, and platform adoption may reduce the idiosyncratic component—it is natural to expect $\sigma_A^2 < \sigma_C^2$, since the platform pools data across many sources that no individual chain could assemble on its own. But the platform also introduces the common component ξ, which is correlated across all subscribers and uncorrelated with the underlying market. The economic content of the model is that this trade—swapping idiosyncratic noise for shared noise—is privately attractive even when its aggregate consequences are socially costly.

The structure differs from the standard public-signal benchmark of Morris and Shin (2002) in a precise way. A public signal is informative about θ, so that aggregating across firms refines beliefs even as actions become more correlated. The shared shock ξ here carries no information about θ and survives aggregation as a pure correlation component. The cross-section of platform adopters never refines beliefs about local profitability through their joint exposure to ξ; it only inherits a common error.[2]

Remark 1 (Behavioral interpretation)**.** The shared component ξ admits a behavioral interpretation that may be useful. When many chains adopt the same platform for a given decision, it is as if each chain has hired the same market researcher, whose idiosyncratic biases, training, and blind spots generate correlated outputs across employers. The variance σ_ξ^2 measures the magnitude of this shared researcher's biases, while the orthogonality of ξ to θ captures that these biases are unrelated to the underlying state. This structure parallels the treatment of investor sentiment as a common factor orthogonal to fundamentals (Baker and Wurgler 2006, 2007) and the role of

[2]The mechanism also differs from the public-signal extensions in Bergemann et al. (2015), where common signals can confound multiple components of uncertainty but remain economically informative. The orthogonality of ξ to θ is a structural feature of the processing technology, not a parameter restriction.

correlated belief errors in generating non-fundamental comovement (De Long et al. 1990). The mechanism here differs in source rather than form: the shared component arises from a common decision technology rather than from correlated psychology.

2.4 Adoption Stage

Each chain draws a switching cost $c_f \sim F$ independently, where F is a continuous cumulative distribution function on $\mathbb{R}$ with positive density on its support. The chain chooses its forecast technology to maximize expected entry profits net of the switching cost, taking the equilibrium adoption rate q as given. An adoption equilibrium is a rational-expectations fixed point in which q is consistent with individual choices, $q^* = F(\Delta\Pi(q^*))$, where $\Delta\Pi(q) \equiv \pi^A(q) - \pi^C(q)$ is the private gain from adoption.

We allow F to have either bounded or unbounded support. With bounded support $[0, \bar{c}]$, corner outcomes (zero or full adoption) can arise as equilibria when $\Delta\Pi(0) \leq 0$ or $\Delta\Pi(1) \geq \bar{c}$. With unbounded support, as in the logistic specification used in the numerical exercises below, corner outcomes are replaced by interior equilibria with q^* close to but bounded away from 0 or 1. This distinction does not affect any qualitative result.

2.5 Bayesian Filtering

Chains form beliefs using Bayesian-optimal filtering given their information structure and common knowledge of the data-generating process, including the presence and variance of the shared shock ξ. Let $V_A \equiv \sigma_\theta^2 + \sigma_\xi^2 + \sigma_A^2$ and $V_C \equiv \sigma_\theta^2 + \sigma_C^2$ denote unconditional signal variances, and define the projection coefficients

$$\alpha_\theta^A \equiv \frac{\sigma_\theta^2}{V_A}, \qquad \alpha_\xi^A \equiv \frac{\sigma_\xi^2}{V_A}, \qquad \alpha_\theta^C \equiv \frac{\sigma_\theta^2}{V_C}.$$

Then $\mathbb{E}_f^A[\theta] = \alpha_\theta^A \tilde{s}_f^A$, $\mathbb{E}_f^A[\xi] = \alpha_\xi^A \tilde{s}_f^A$, and $\mathbb{E}_f^C[\theta] = \alpha_\theta^C \tilde{s}_f^C$. Custom-research chains do not observe ξ; their posterior over ξ remains at the prior mean of zero.

The informational content of the platform signal is summarized by $\rho_A \equiv \alpha_\theta^A + \alpha_\xi^A = (\sigma_\theta^2 + \sigma_\xi^2)/V_A$, the share of platform-signal variance attributable to components common across adopters. This object governs how strongly the shared shock enters equilibrium behavior.

The assumption that chains know σ_ξ^2 corresponds to the rational-expectations benchmark in which they correctly understand the structure of the platform and filter signals accordingly. Misperception of σ_ξ^2—chains believing the platform to be more or less informative than it is—is an extension we comment on in the discussion.

3 Equilibrium with Heterogeneous Adoption

This section characterizes the equilibrium given the information structure and a fixed adoption rate q. The analysis proceeds in two steps: deriving the linear equilibrium loadings of individual entry probabilities on signals, and using these to characterize the variance decomposition of aggregate entry.

3.1 Best Responses and the Linear Fixed Point

A chain enters if and only if its posterior expected profit is nonnegative, $\mathbb{E}_f[\theta] + \delta \mathbb{E}_f[\bar{x}] + \varepsilon_f \geq 0$. With ε_f uniform on $[-1/2, 1/2]$ (using the normalization $\bar{\varepsilon} = 1/2$), the entry probability conditional on the signal is

$$P_f(\text{enter} \mid s_f) = \tfrac{1}{2} + \mathbb{E}_f[\theta] + \delta\, \mathbb{E}_f[\bar{x}],$$

in the global-games regime where this expression lies in $(0, 1)$.

We look for a linear equilibrium in which entry probabilities are affine in signals. Conjecturing that aggregate entry takes the form $\bar{x} = K_0 + K_\theta \theta + K_\xi \xi$, the law of large numbers in the continuum implies that idiosyncratic noises wash out in aggregation, and the conjectured aggregate must be consistent with type-by-type entry rules. Define individual loading coefficients k^A and k^C as the slopes of the entry probability with respect to the signal for AI and custom chains respectively. Then aggregation requires

$$K_\theta(q) = q\,k^A + (1-q)\,k^C, \qquad K_\xi(q) = q\,k^A.$$

The aggregate inherits the loading on θ from both types weighted by their population shares, while the loading on ξ comes entirely from platform adopters and scales with their share q.

Substituting the conjectured aggregate into the entry rule for each type yields a linear system in (k^A, k^C) whose solution can be written in closed form. We collect the result as:

Proposition 1 (Equilibrium loadings)**.** *For $q \in (0,1)$, the unique linear equilibrium has*

$$k^A(q) = \frac{\alpha_\theta^A(1+\delta)}{\Delta(q)}, \qquad k^C(q) = \frac{\alpha_\theta^C(1+\delta-\delta\,q\,\alpha_\xi^A)}{\Delta(q)},$$

where

$$\Delta(q) \equiv \left[1+\delta-\delta\,q\,\rho_A\right]\left[1+\delta-\delta(1-q)\,\alpha_\theta^C\right] - \alpha_\theta^A\,\alpha_\theta^C\,\delta^2\,q(1-q).$$

The structure of the loadings reflects how the shared shock enters equilibrium behavior. Platform adopters respond to signals that combine θ and ξ, while custom-research chains respond only to θ. The term $-\delta\,q\,\alpha_\xi^A$ in the numerator of k^C captures an indirect effect: although custom-research chains do not observe ξ, they anticipate that it moves aggregate entry through platform

adopters' decisions, and they adjust their own loading to partially offset (under $\delta < 0$) or amplify (under $\delta > 0$) this movement. The strength of this adjustment grows with adoption q and with the share of platform-signal variance attributable to the common component.

The boundary cases connect the equilibrium to standard benchmarks. When $q = 0$, the platform is absent and the equilibrium reduces to a familiar linear-Gaussian beauty contest with custom signals only, $k^C(0) = \alpha_\theta^C/(1+\delta-\delta\alpha_\theta^C)$. When $q = 1$, all chains adopt the platform and the loading becomes $k^A(1) = \alpha_\theta^A/(1+\delta-\delta\rho_A)$, where the denominator reflects that the signal now contains both informative and uninformative common components.

3.2 Aggregate Behavior

Substituting the equilibrium loadings into the aggregate yields the entry-error decomposition

$$\bar{x} - \mathbb{E}[\bar{x}] = K_\theta(q)\,\theta + K_\xi(q)\,\xi - \big(K_\theta(q)\,\mathbb{E}[\theta] + K_\xi(q)\,\mathbb{E}[\xi]\big) = K_\theta(q)\,\theta + K_\xi(q)\,\xi,$$

where the second equality uses $\mathbb{E}[\theta] = \mathbb{E}[\xi] = 0$. Taking variances,

$$\mathrm{Var}(\bar{x}) = K_\theta(q)^2\sigma_\theta^2 + K_\xi(q)^2\sigma_\xi^2.$$

This decomposition separates two distinct sources of aggregate-entry variance. The first is fundamental-tracking variance: the aggregate responds to θ with loading $K_\theta(q)$, and its variance contribution scales with the variance of the fundamental. The second is platform-driven variance: whenever $q > 0$, the aggregate inherits a component proportional to ξ, and this component contributes directly to volatility.

In a custom-only economy, $K_\xi(0) = 0$ and aggregate variance is driven entirely by fundamental tracking. With platform adoption, this changes qualitatively. The term $K_\xi(q)^2\sigma_\xi^2$ is strictly positive whenever $q > 0$ and does not vanish under aggregation. Idiosyncratic noises cancel in the continuum, but the common shock is shared across adopters and survives intact in the aggregate.

The implication is that shared platforms introduce a source of aggregate volatility that is absent in standard entry models. Correlation across chains is no longer tied to information about the fundamental; it reflects a common error embedded in the processing technology. This non-vanishing contribution of ξ is the central positive prediction of the model and the object on which the welfare analysis turns.

4 Adoption Incentives and Misalignment

This section characterizes adoption incentives and compares equilibrium adoption to social optima under alternative welfare specifications. The substantive question is whether chains internalize the consequences of shared processing for aggregate entry. The results show that coordination motives are fully internalized in equilibrium, but the variance contribution of the shared shock generates an externality that creates a systematic wedge between private and social adoption decisions, with a sign that is invariant to the strategic structure of competition.

4.1 The Private Gain to Adoption

Let $\pi^A(q)$ and $\pi^C(q)$ denote the expected profits of platform-adopting and custom-research chains at adoption rate q, obtained by substituting the equilibrium loadings into the profit function. The private adoption gain is $\Delta\Pi(q) \equiv \pi^A(q) - \pi^C(q)$, and an adoption equilibrium satisfies $q^* =$

$F(\Delta\Pi(q^*))$.

The dependence of $\Delta\Pi(q)$ on q reflects the interaction between the precision properties of the platform and the strategic structure of competition. The platform changes both how accurately chains forecast θ and how closely their entry decisions align with rivals, and the relative importance of these forces depends on δ.

Proposition 2 (Monotonicity of the adoption gain)**.** *The function $\Delta\Pi(q)$ is independent of q when $\delta = 0$, strictly increasing in q when $\delta > 0$, and strictly decreasing in q when $\delta < 0$.*

When $\delta = 0$, profits depend only on matching θ, and the attractiveness of the platform depends entirely on the precision tradeoff embedded in the signal structure. Strategic interaction is absent, so the adoption decision is unaffected by penetration. When $\delta > 0$ (agglomeration), greater adoption raises the value of the platform: as more rivals subscribe, the aggregate becomes more exposed to the common component ξ, and platform adopters—who can partially forecast this component—track aggregate entry more closely than custom-research chains. Adoption becomes more attractive as it becomes more widespread. When $\delta < 0$ (business stealing), the same force works in reverse. Chains benefit from differentiating from rivals, and platform adoption makes them more correlated. The marginal value of adoption declines with q.

The proof reduces the result to the sign of a polynomial in (q, δ, σ). Differentiating the closed-form expression for $\Delta\Pi(q)$ derived in Appendix A yields

$$\frac{d\Delta\Pi}{dq} = -2\delta\,\sigma_\theta^4(1+\delta)\,V_C\,\frac{Q(q,\delta;\sigma)}{\Delta(q)^3 V_A V_C},$$

where the prefactor is strictly positive and Q is a polynomial whose negativity over the admissible parameter region is verified by sum-of-squares decompositions at the boundaries and by exhaustive

computer-algebraic evaluation in the interior. Full details appear in the proof.

4.2 Welfare Specifications

The misalignment result depends on how social welfare is defined, which is itself a substantive modeling choice in coordination games with private information.[3] We consider three benchmarks that make the substantive issues explicit.

The first is a utilitarian benchmark, denoted $\mathscr{W}_1$, under which the planner aggregates chain profits and therefore fully internalizes the coordination motive:

$$\mathscr{W}_1(q) = q\,\pi^A(q) + (1-q)\,\pi^C(q).$$

The second benchmark, denoted $\mathscr{W}_2$, isolates fundamental-tracking performance by treating coordination payoffs as transfers across chains. It penalizes the squared deviation of each chain's entry-probability loading from the fundamental, paralleling the social-value-of-information objective in Angeletos and Pavan (2007).

The third benchmark, denoted $\mathscr{W}_3$, augments chain profits with a penalty on the component of aggregate-entry variance that arises from shared processing:

$$\mathscr{W}_3(q) = \mathscr{W}_1(q) - \gamma K_\xi(q)^2\,\sigma_\xi^2, \quad \gamma > 0.$$

The term $K_\xi(q)^2\sigma_\xi^2 = q^2(k^A(q))^2\sigma_\xi^2$ is the variance contribution of the shared shock to aggregate entry. It isolates the source of correlation introduced by the platform and separates it from

[3] The exchange between Svensson (2006) and Morris et al. (2006) concerning the welfare implications of public-information disclosure illustrates how sensitive welfare conclusions are to the choice of criterion. Our specifications situate the choice of welfare criterion in this tradition: the planner internalizes social costs that chains do not.

fundamental-tracking variance, which is already internalized in chain profits. The parameter γ measures the marginal social cost of platform-driven aggregate volatility and admits several concrete interpretations in the retail context. Synchronized over-entry into markets that the platform identified as favorable but where the shared shock ξ was positive generates excess retail space that persists due to lease terms and sunk fixtures, imposing costs on commercial mortgage lenders facing correlated default risk and on landlords absorbing prolonged vacancies. Synchronized under-entry into markets where ξ was negative leaves consumer demand unmet and reduces sales-tax revenue for municipalities that planned around expected commercial activity. In each case, the cost is borne by parties downstream of the entry decision and scales with the variance of the aggregate around the fundamental, not with the variance of any individual chain's entry decision around its own forecast.

4.3 The Team-Efficiency Result

The first welfare result is a benchmark that disciplines the analysis. Under the utilitarian criterion $\mathscr{W}_1$, the equilibrium is socially efficient.

Proposition 3 (Team efficiency)**.** *Under welfare specification $\mathscr{W}_1$, $d\mathscr{W}_1/dq = \Delta\Pi(q)$ for all $q \in (0,1)$. The market equilibrium q^* coincides with the social optimum q_1^{soc}.*

The result follows from a property of the linear-Gaussian beauty-contest structure with a continuum of agents. For any fixed q, the equilibrium loadings (k^A, k^C) are not only privately optimal but also maximize average chain profits under a planner who assigns one common decision rule to each type. The planner's first-order condition differs from the Nash condition only through a general-equilibrium term reflecting the response of the aggregate to coordinated changes in be-

havior. This term vanishes not because the aggregate is unmoved by such changes—it is moved, by an amount that depends on the type's population share—but because the cross-section identity $\int_{\mathscr{F}} (\partial \pi_g / \partial \bar{x})\, dg = 0$ holds at any equilibrium: the integrand averages to zero across the population of chains. The same logic extends to changes in the adoption rate, so the derivative $d\mathscr{W}_1/dq$ coincides with the direct gain from switching types, $\Delta\Pi(q)$.

The result establishes that correlated entry behavior induced by shared processing does not by itself generate inefficiency under the utilitarian criterion. Even though the platform introduces a common component into entry decisions, chains fully internalize its implications for coordination through their profit functions. Any wedge between private and social incentives must therefore arise from objectives that are not captured by chain profits, such as the variance of aggregate entry borne by parties downstream of the entry decision.

4.4 The Misalignment Result

The substantive welfare result follows under the systemic-risk benchmark $\mathscr{W}_3$.

Proposition 4 (Misalignment)**.** *Under welfare specification $\mathscr{W}_3$, the wedge between private and social adoption incentives is*

$$\Delta\Pi(q) - \frac{d\mathscr{W}_3}{dq} = \gamma \frac{d}{dq}\left[K_\xi(q)^2\, \sigma_\xi^2\right] > 0$$

for all $q \in (0,1]$ and all $\delta > -1$. In any interior adoption equilibrium q^*, $q^* > q_3^{soc}$: the platform is over-adopted relative to the planner who internalizes the systemic-risk externality. The direction of the wedge is independent of δ, while its magnitude depends on δ through the equilibrium loadings.

The result follows from a clean decomposition. Under $\mathscr{W}_1$, Proposition 3 implies that pri-

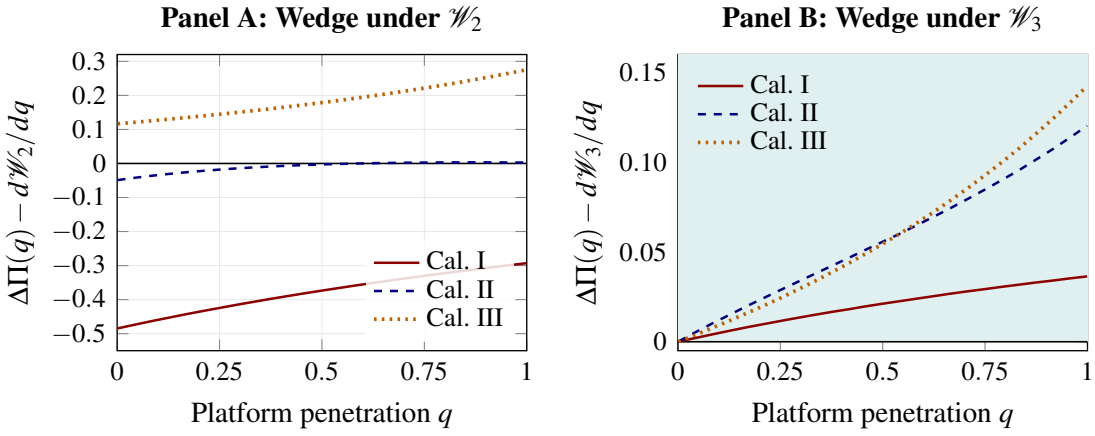
vate and social incentives coincide once equilibrium behavior is fixed. Introducing $\mathscr{W}_3$ leaves this component unchanged but adds a term that depends on how adoption affects aggregate volatility. Because $K_\xi(q)^2$ is strictly increasing in q, every additional adopter raises the variance contribution of the shared shock. This effect is not internalized by chains, whose profits depend on their own entry decisions and on average rival behavior but not on the variance of aggregate entry. The entire wedge therefore arises from the externality generated by the shared shock, and it is positive at every interior point.

The role of the orthogonality of ξ to θ is direct. In environments with informative public signals, greater reliance on common information increases correlation but can reduce aggregate variance because beliefs improve. Here the common component carries no information about θ, so this offsetting force is absent. Higher adoption simultaneously raises correlation and raises variance. The same monotonicity that governs aggregate variance therefore determines the sign of the welfare wedge.

The empirical content of the δ-invariance is worth emphasizing. The strategic-interaction parameter δ is notoriously difficult to identify in retail markets: agglomeration and business-stealing forces can operate simultaneously and net to ambiguous local signs, and the same product category may exhibit agglomeration in some geographies and business-stealing in others. A welfare result whose direction depended on $\operatorname{sgn}(\delta)$ would require resolving this ambiguity before any policy implication could be drawn. Proposition 4 delivers a stronger statement: regardless of how retail competition is structured locally, the systemic-risk externality always points toward over-adoption. The magnitude of intervention depends on δ through the equilibrium loadings, but the direction of the externality does not.

Remark 2 (Misalignment under $\mathscr{W}_2$ is generically indeterminate)**.** Under welfare specification $\mathscr{W}_2$, the planner optimizes fundamental-tracking accuracy rather than chain profits, and equilibrium loadings are not generally optimal from this perspective. As a result, the wedge between $\Delta\Pi(q)$ and $d\mathscr{W}_2/dq$ contains both a direct compositional term and indirect effects operating through equilibrium loadings. The total wedge is sign-indeterminate: numerical evaluation confirms that for fixed primitives, $\Delta\Pi(q) - d\mathscr{W}_2/dq$ can be positive or negative depending on the calibration. In contrast, the systemic-risk specification $\mathscr{W}_3$ yields a uniformly positive wedge. The two specifications capture different welfare concerns: $\mathscr{W}_2$ measures forecast accuracy at the firm level, while $\mathscr{W}_3$ measures volatility of the aggregate. The orthogonality of ξ to θ is what makes the latter pin down a determinate sign.

Figure 1: The misalignment wedge: $\mathscr{W}_2$ is sign-indeterminate, $\mathscr{W}_3$ is uniformly positive.



Notes: Each panel plots the misalignment wedge $\Delta\Pi(q) - d\mathscr{W}_k/dq$ as a function of platform penetration q, for three calibrations sharing $\delta = 1.5$ and $\gamma = 1$ but differing in the relative magnitudes of σ_A^2, σ_C^2, and σ_ξ^2. Calibration I: $(\sigma_\theta^2, \sigma_\xi^2, \sigma_A^2, \sigma_C^2) = (2.79, 0.05, 0.62, 0.27)$; Calibration II: $(1.23, 0.62, 0.88, 0.47)$; Calibration III: $(2.48, 0.21, 0.78, 1.46)$. Panel A shows the wedge under welfare specification $\mathscr{W}_2$: it takes negative values throughout for one calibration, positive values throughout for another, and crosses zero in the interior for the third. The same primitive parameters therefore deliver no determinate prediction about whether the platform is over- or under-adopted relative to the $\mathscr{W}_2$ planner (Remark 2). Panel B shows the wedge under welfare specification $\mathscr{W}_3$: it is strictly positive across all three calibrations and all $q \in (0, 1]$, with over-adoption the universal prediction (Proposition 4).

The contrast is illustrated in Figure 1. Panel A plots the $\mathscr{W}_2$ wedge for three calibrations chosen

to highlight its indeterminacy: one yields uniform under-adoption, another uniform over-adoption, and a third changes sign in the interior. Panel B plots the corresponding $\mathscr{W}_3$ wedge, which lies strictly above zero in all three calibrations. The difference between the panels isolates the empirical content of the welfare specification: the systemic-risk specification delivers a determinate direction of intervention, while the tracking-only specification does not.

5 The Four Regimes of Adoption

Adoption is determined by a fixed-point condition $q^* = F(\Delta\Pi(q^*))$, and the equilibrium structure can vary qualitatively across environments depending on the shape of $\Delta\Pi(q)$. The shape, together with the behavior at the boundaries, determines the number of equilibria and yields a small set of regimes that organize the model's predictions.

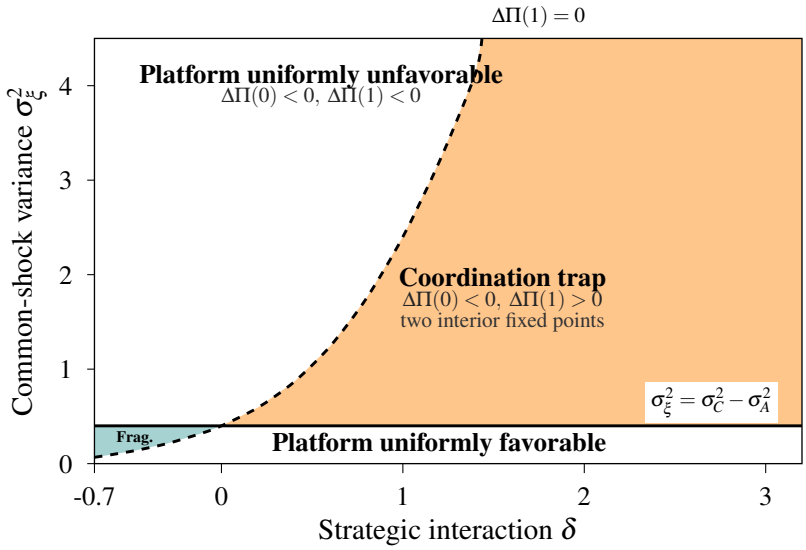
The relevant boundary objects are the signs of $\Delta\Pi(0)$ and $\Delta\Pi(1)$. The first admits a particularly clean characterization.

Lemma 1 (δ-independent threshold)**.** *The boundary* $\Delta\Pi(0) = 0$ *is given by* $\sigma_\xi^{2*} = \sigma_C^2 - \sigma_A^2$*, independent of* δ*.*

The intuition is that at $q = 0$ a marginal platform adopter faces the same aggregate as a custom-research chain, $\bar{x} = \mathbb{E}[\bar{x}] + k^C(0)\theta$. At the point where $V_A = V_C$, the two technologies deliver identical signal quality, so all moments of expected profit coincide. Because both types respond to the same aggregate, the coordination term enters symmetrically and its weight cancels. Adoption at the margin therefore depends only on the relative idiosyncratic precision of the two technologies. The platform is privately attractive in isolation if and only if its idiosyncratic noise advantage, $\sigma_C^2 - \sigma_A^2$, exceeds the variance of the shared shock.

The condition $\Delta\Pi(1) = 0$ defines a locus that depends on δ. As coordination motives strengthen, deviation from the platform technology becomes more costly, and chains tolerate a larger shared shock before deviating. The threshold value of σ_ξ^2 at which a marginal custom-research chain is indifferent therefore rises with δ. Numerical evaluation shows this locus is monotone increasing and convex, intersecting the $q = 0$ boundary at $\delta = 0$.

Figure 2: Four regimes of platform adoption in the (δ, σ_ξ^2) parameter space.



Notes: The phase diagram shows four adoption regimes as functions of the strategic-interaction parameter δ and the common-shock variance σ_ξ^2. The solid line is the locus $\Delta\Pi(0) = 0$, where the marginal platform deviator at zero penetration is just indifferent; this boundary is independent of δ (Lemma 1). The dashed curve is the locus $\Delta\Pi(1) = 0$, where the marginal custom-research deviator at full penetration is just indifferent; this curve is monotone increasing in δ and passes through $\sigma_\xi^2 = \sigma_C^2 - \sigma_A^2$ at $\delta = 0$. The shaded amber region is the coordination-trap regime: a high-adoption fixed point coexists with a low-adoption fixed point, sustained by coordination motives even though the platform is intrinsically unfavorable in isolation. The shaded teal region is the fragmentation regime under business stealing, in which the unique adoption fixed point is interior. Calibration: $\sigma_\theta^2 = 1$, $\sigma_A^2 = 0.2$, $\sigma_C^2 = 0.6$.

Figure 2 plots these two loci in (δ, σ_ξ^2) space and partitions the parameter region into four regimes. The misalignment result of Proposition 4 applies throughout: while the equilibrium configuration depends on the strategic structure, the direction of the welfare wedge does not. Two regions exhibit qualitatively distinct behavior, the coordination trap under agglomeration and the

fragmentation regime under business stealing, which we discuss in turn.

5.1 The Coordination Trap

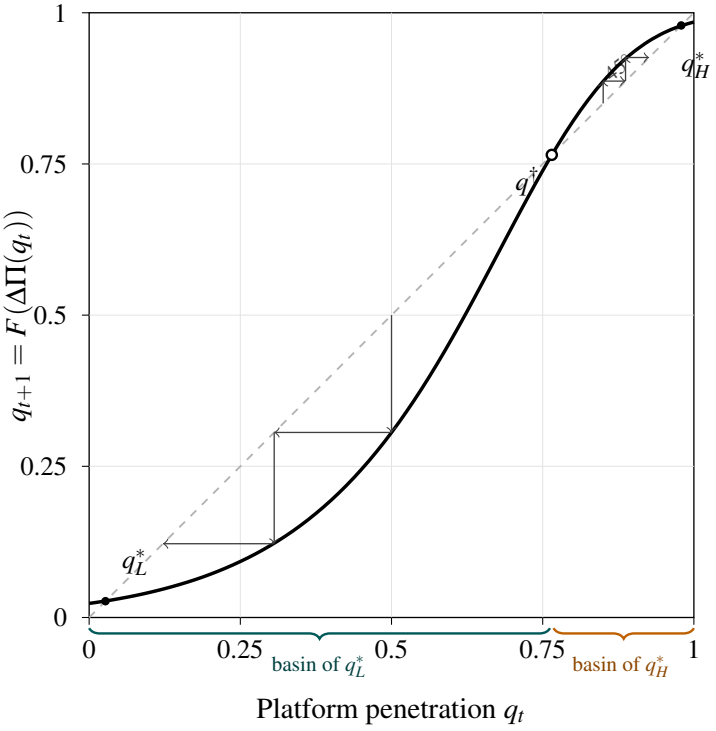
When $\delta > 0$, $\Delta\Pi(0) < 0$, and $\Delta\Pi(1) > 0$, the adoption mapping $q^* = F(\Delta\Pi(q^*))$ admits multiple fixed points. There are two attracting fixed points under best-response iteration on the adoption equation, a low-adoption fixed point q_L^* close to zero and a high-adoption fixed point q_H^* close to one, separated by a third fixed point at which the slope of $F \circ \Delta\Pi$ exceeds unity.[4] Platform subscription is privately unattractive in isolation but becomes attractive once adoption is sufficiently widespread.

The mechanism is purely strategic. At low penetration, a marginal chain faces little exposure to the shared component and values the precision of custom research. At high penetration, the aggregate is dominated by the common component, and chains benefit from aligning with rivals who are subject to the same signal. Crossing the unstable fixed point shifts the balance from precision to coordination, making adoption self-reinforcing.

This regime provides a precise formulation of the idea that shared platforms can act as coordination devices in retail markets. The high-adoption fixed point is sustained even though adoption is privately unattractive in isolation, because no chain has an incentive to deviate once others have adopted. The economic content is sharpest in retail sectors with strong agglomeration economies. In food halls, outlet centers, and clusters of competing fast-food restaurants, knowing what rivals are doing matters at least as much as knowing what consumers want. A platform that reveals rival intentions becomes essential to incumbents whether or not its forecasts of underlying demand are

[4] Throughout this paragraph and Figure 3 we use "stability" in the sense of best-response iteration on the static adoption equation: a fixed point q^* is stable if $|F'(\Delta\Pi(q^*))\Delta\Pi'(q^*)| < 1$, so that small perturbations of q near q^* contract back. The model has no explicit time index, so this stability notion is a property of the iterative map, not of any underlying dynamic process.

Figure 3: The coordination trap as a fixed-point diagram.



Notes: The figure plots the next-iterate adoption rate $q_{t+1} = F(\Delta\Pi(q_t))$ against the current rate q_t for a calibration in the coordination-trap regime. Fixed points lie at intersections with the 45-degree line. Three intersections obtain: a stable low-adoption fixed point $q_L^* \approx 0.03$ and a stable high-adoption fixed point $q_H^* \approx 0.98$, separated by an unstable interior fixed point $q^\dagger \approx 0.77$ at which the slope of $F \circ \Delta\Pi$ exceeds unity. Best-response iteration converges to whichever stable fixed point lies in the same basin of attraction as the initial penetration: the wide left basin captures industries that begin from any moderate level of platform adoption, while only initial penetration above $q^\dagger$ tips the industry into the high-adoption fixed point. The light arrows trace two illustrative iterations starting from $q_0 = 0.5$ and $q_0 = 0.85$. Calibration: $\delta = 2$, $\sigma_\theta^2 = 1$, $\sigma_\xi^2 = 1.5$, $\sigma_A^2 = 0.2$, $\sigma_C^2 = 0.6$, with logistic adoption distribution $F(x) = (1 + e^{-\lambda x})^{-1}$ and $\lambda = 15$.

accurate.

Under the systemic-risk benchmark $\mathscr{W}_3$, both attracting fixed points exhibit over-adoption at the margin relative to their local social optima, as established in Proposition 4. The comparison between the two fixed points at the global level is more subtle. The high-adoption fixed point benefits from stronger coordination but generates greater aggregate volatility through the shared component. For moderate δ, the low-adoption fixed point yields higher welfare; for sufficiently large δ, the coordination gains at high adoption can outweigh the variance cost. A regulator seeking

to move the economy away from excessive synchronization must therefore assess not only the marginal wedge but also which fixed point is welfare-dominant. When the low-adoption fixed point is preferred, policies that shift expectations or incentives across the unstable fixed point can implement it: disclosure requirements, diversification mandates, or support for alternative processing technologies operate by shrinking the basin of attraction of the high-adoption fixed point and expanding that of the low-adoption outcome.

5.2 Adoption Fragmentation

When $\delta < 0$, $\Delta\Pi(0) > 0$, and $\Delta\Pi(1) < 0$, neither corner is a fixed point. Since $F \circ \Delta\Pi$ is continuous and strictly decreasing in q under business stealing, the fixed-point equation admits a unique interior solution. The adoption equilibrium features persistent heterogeneity, with a stable mix of platform adopters and custom-research chains.

The underlying force is differentiation. The platform improves precision, making it attractive in isolation, but it also increases correlation across chains. When entry decisions are strategic substitutes, this correlation is costly. As adoption rises, the marginal benefit of platform subscription declines, and the equilibrium settles at an interior point where the precision gain is exactly offset by the strategic cost of moving with rivals. Retail sectors in which differentiation is central—specialty grocers serving distinct ethnic markets, boutique apparel, niche restaurant concepts—should therefore exhibit partial adoption rather than convergence on a common platform.

5.3 Boundary Regimes

The remaining cases are degenerate. When $\Delta\Pi(0) > 0$ and $\Delta\Pi(1) > 0$, the platform is privately attractive at all adoption levels and equilibrium adoption is bounded away from zero. When $\Delta\Pi(0) < 0$ and $\Delta\Pi(1) < 0$, it is unattractive throughout and adoption remains near zero. In both cases, strategic incentives reinforce the underlying technological advantage or disadvantage rather than altering its direction.

6 Vendor Diversity and the Magnitude of Misalignment

The misalignment result of Proposition 4 establishes that platform adoption is excessive but does not characterize how the magnitude of the wedge depends on the structure of the platform market itself. This section generalizes the model to allow multiple competing platforms, identifying a second policy margin—the substantive diversity of platform processing—that governs the size of the externality without altering its direction. The extension speaks to two regulatory concerns that the framework has so far approached only through the empirical motivation: vendor concentration in retail intelligence markets, and the distinction between substantive and cosmetic forms of platform diversity that arises when multiple platforms fine-tune a common foundation model or share a common data vendor.

6.1 Multiple Processing Technologies

Suppose adopters partition across $K \geq 1$ platforms. Platform k generates a processing shock $\xi_k \sim \mathcal{N}(0, \sigma_\xi^2)$ and a chain-specific noise term with variance σ_A^2, so that an adopter of platform k

observes

$$\tilde{s}_f^{A,k} = \theta + \xi_k + u_f^{A,k}.$$

Let q_k denote the share of chains adopting platform k, with $\sum_{k=1}^{K} q_k = q$. Vendor shocks are jointly Gaussian with $\mathrm{Var}(\xi_k) = \sigma_\xi^2$ for all k and pairwise correlation $\mathrm{Corr}(\xi_k, \xi_{k'}) = \rho \in [0,1]$ for $k \neq k'$. The single-platform model of Section 5 is the case $K = 1$.

The two endpoints of ρ have natural interpretations. When $\rho = 0$, platforms are structurally independent: their processing shocks come from distinct foundation models, training corpora, or data vendors, and the realizations $\{\xi_k\}$ are uncorrelated. When $\rho = 1$, platforms share a common processing core—multiple wrappers around the same foundation model, multiple consumer-data products built on identical mobile-device location pings, or multiple demographic-projection products built on the same Census source—and $\xi_k = \xi_{k'}$ almost surely.

For tractability, the analysis takes vendor adoption shares as exogenous primitives. Section 5 treated the aggregate q as endogenously determined by switching costs; here, conditional on q, the partition $\{q_k\}$ is taken as given. This allows the welfare implications of vendor structure to be isolated cleanly from a vendor-side competition model.

6.2 The Effective Concentration Index

Each adopter still sees a single processing shock in their own signal, so equilibrium loadings (k^A, k^C) are unchanged from Proposition 1. The aggregate, however, inherits a separate component from each platform:

$$\bar{x} = \mathbb{E}[\bar{x}] + K_\theta(q)\,\theta + \sum_{k=1}^{K} K_{\xi,k}\,\xi_k, \qquad K_{\xi,k} = q_k \cdot k^A.$$

Computing the variance of the second sum and substituting the correlation structure yields

$$\mathrm{Var}\Big(\sum_k K_{\xi,k}\,\xi_k\Big) = (k^A)^2\,\sigma_\xi^2 \cdot \Big[\sum_k q_k^2 + 2\rho \sum_{k<k'} q_k q_{k'}\Big].$$

Under symmetric vendor adoption, $q_k = q/K$ for all k, the bracketed expression simplifies. Define the *effective concentration index*

$$H(K,\rho) \equiv \frac{1}{K} + \Big(1 - \frac{1}{K}\Big)\rho.$$

Then under symmetric adoption,

$$\mathrm{Var}\Big(\sum_k K_{\xi,k}\,\xi_k\Big) = K_\xi(q)^2\,\sigma_\xi^2 \cdot H(K,\rho),$$

where $K_\xi(q) = q \cdot k^A$ is the aggregate loading on the shared component in the single-platform model. The index H takes values in $[1/K, 1]$, attaining its lower bound at $\rho = 0$ and its upper bound at $\rho = 1$, and is monotone increasing in ρ for any $K \geq 2$.

Proposition 5 (Vendor diversity)**.** *Under welfare specification $\mathscr{W}_3$ generalized to multiple platforms, the wedge between private and social adoption incentives at any interior q is*

$$\Delta\Pi(q) - \frac{d\mathscr{W}_3}{dq} = \gamma H(K,\rho)\,\frac{d}{dq}\big[K_\xi(q)^2\,\sigma_\xi^2\big] > 0,$$

for all $\delta > -1$, all $K \geq 1$, and all $\rho \in [0,1]$. The direction of the wedge is invariant in (K,ρ). The magnitude scales with $H(K,\rho)$ and falls monotonically in K at any $\rho < 1$, with strict reduction whenever $\rho < 1$.

The proof follows directly from Proposition 4 together with the variance decomposition above. The marginal welfare effect of q under $\mathscr{W}_3$ now includes the factor $H(K,\rho)$ in front of the $K_\xi(q)^2\sigma_\xi^2$ term. Strict positivity of the derivative was established in the proof of Proposition 4, and $H(K,\rho)$ is itself strictly positive, so the wedge sign is unchanged.

6.3 Substantive versus Cosmetic Diversity

Two corner cases of Proposition 5 expose the policy content of the result.

Corollary 1 (Independent platforms)**.** *When $\rho = 0$, $H(K,0) = 1/K$. The misalignment magnitude under welfare specification $\mathscr{W}_3$ falls linearly in K.*

Corollary 2 (Common-core platforms)**.** *When $\rho = 1$, $H(K,1) = 1$ for all K. The misalignment magnitude is invariant to vendor count.*

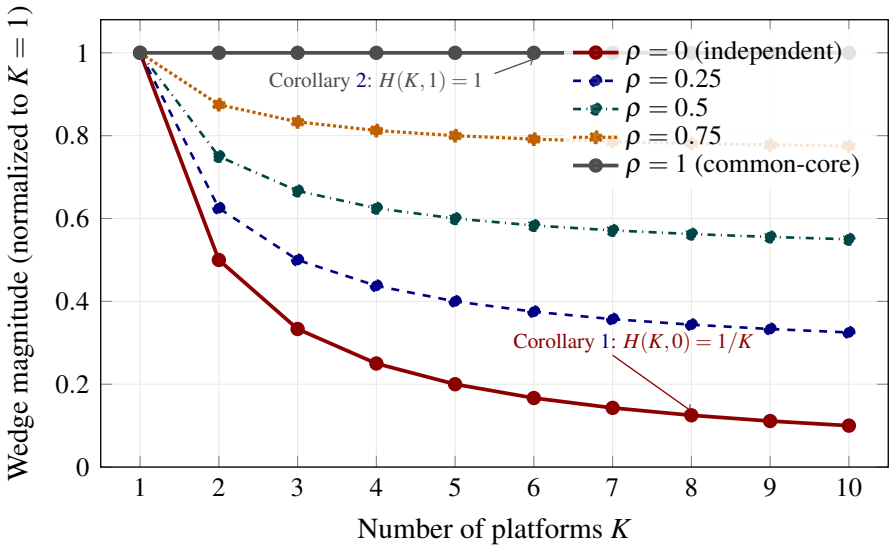
The contrast between the two corollaries is the central policy content. Under structurally independent platforms, doubling the number of providers halves the systemic-risk externality. Under platforms that share a common processing core, the same nominal expansion of vendor count delivers no welfare gain. The vendor-diversity margin is therefore not vendor count per se, but the structural independence of platforms' processing shocks.

This distinction tracks directly observable features of the retail intelligence market. Several major platforms purchase mobile-device location data from the same upstream vendors and license demographic projections from the same Census-based products. Two platforms that resell forecasts derived from the same underlying foot-traffic dataset deliver high-ρ diversity at best, regardless of how their user interfaces differ. The reverse case—one platform built on credit-card transaction data, another on cell-tower pings, a third on satellite parking-lot imagery—approximates the low-

ρ regime in which the externality scales as $1/K$. Bommasani et al. (2022) document an analogous pattern in the algorithmic-monoculture literature for foundation models: fine-tuned variants of a common base model inherit substantial structure from the base, so the residual variation across variants is small relative to variation across structurally distinct models.

Figure 4 visualizes the contrast: the wedge magnitude declines sharply with K at $\rho = 0$ but remains flat at $\rho = 1$, with intermediate values interpolating smoothly between the two corner cases.

Figure 4: Wedge magnitude as a function of vendor diversity.



Notes: The figure plots the misalignment-wedge magnitude under welfare specification $\mathscr{W}_3$ as a function of the number of platforms K, normalized to the value at $K = 1$. Each curve corresponds to a fixed value of the cross-platform correlation $\rho \in \{0, 0.25, 0.5, 0.75, 1\}$. The wedge magnitude scales with the effective concentration index $H(K,\rho) = 1/K + (1 - 1/K)\rho$ from Proposition 5. At $\rho = 0$ (red), platforms are structurally independent and the magnitude falls hyperbolically as $1/K$ (Corollary 1). At $\rho = 1$ (gray), platforms share a common processing core and the magnitude is invariant in K (Corollary 2). The vendor-diversity margin is not platform count per se but the structural independence of platforms' processing shocks: cosmetic diversity (high ρ) delivers little welfare gain, while substantive diversity (low ρ) delivers a near-linear reduction in the systemic-risk externality.

6.4 The Regulator's Joint Problem

Combining Propositions 4 and 5, a regulator who internalizes the systemic-risk externality faces a joint policy problem in two margins: the adoption rate q and the effective concentration $H(K,\rho)$. The marginal social cost of adoption at any (q,K,ρ) is the product of these two quantities times the underlying $K_\xi(q)^2\sigma_\xi^2$. Policies targeting the first margin—adoption taxes, mandated diversification of forecasting systems, or capital requirements tied to platform-driven exposure—reduce q. Policies targeting the second margin—procurement diversification mandates, open-source data requirements, or interoperability standards designed to preserve heterogeneity in foundation models—reduce H.

The two margins are not redundant. Reducing q alone leaves the per-adopter exposure to the shared component unchanged. Reducing H alone leaves equilibrium adoption above the social optimum, with the wedge merely scaled down. The combined policy reaches the social optimum at lower distortion than either margin alone, since each operates on a different component of the externality. The structural distinctness of the two margins also implies that policies aimed at the platform market structure complement rather than substitute for policies aimed at adoption itself.

The asymmetric-shares case extends the result without changing its character. When platform adoption is unequal, the effective concentration index becomes

$$H(\{q_k\},\rho) = (\sum_k q_k^2 + 2\rho \sum_{k<k'} q_k q_{k'})/q^2,$$

a Herfindahl-style measure normalized by total adoption. The qualitative result is preserved: the misalignment magnitude is monotone in vendor concentration for any fixed ρ, and the substantive-

versus-cosmetic distinction continues to hold at the corner cases.

7 Discussion

This paper develops a tractable model of shared platform adoption in retail entry. The central feature is that the platform introduces an uninformative common component into chains' forecasts, generating correlated entry decisions without improving inference about local fundamentals. The model admits a closed-form linear equilibrium with heterogeneous adoption, a misalignment result in which adoption is excessive relative to a planner who internalizes systemic risk, and a four-regime typology of adoption that includes a coordination trap with multiple equilibria sustained by agglomeration motives. The over-adoption result follows from a single structural property—the monotonic increase of the aggregate variance component $K_\xi(q)^2\sigma_\xi^2$ with adoption—and holds independently of the sign of the strategic-interaction parameter δ.

The framework provides a formal counterpart to the observation that retail intelligence has become an attentional infrastructure. Chains differ not only in what data they observe but in how they process it, and shared platforms standardize the mapping from data to entry decisions across firms. Custom research preserves heterogeneity in this mapping; shared infrastructure imposes a common perspective. The shared shock ξ captures the residual component of this standardization: the part of entry behavior that is common across adopters but unrelated to the underlying market. Because this component affects all adopters simultaneously, it introduces a source of aggregate risk that is not internalized by individual chains. Coordination incentives are fully internalized in equilibrium, so the only remaining distortion arises from the variance of aggregate entry induced by shared processing.

The misalignment result has direct policy implications. When a planner internalizes the systemic-risk externality, equilibrium platform adoption exceeds the social optimum regardless of the sign of δ. The relevant object for policy is the rate at which $K_\xi(q)^2$ rises with adoption, which determines the marginal social cost of synchronization. Policies that reduce common exposure or promote substantive diversity in processing technologies—disclosure requirements about platform data sources, capital regulation tied to correlated commercial-real-estate exposure, or incentives for maintaining heterogeneous forecasting systems—operate directly on this margin. The coordination-trap regime introduces an additional consideration: where multiple equilibria coexist, marginal interventions are insufficient on their own, and policy must also shape which basin of attraction the economy occupies.

Empirical signatures and directions for future work. The framework yields three empirical signatures that distinguish the technological channel from adjacent mechanisms and that merit investigation in subsequent empirical work. First, among adopters of a common platform, entry decisions should contain a common component that is absent among non-adopters, generating excess covariance at the frequency of decision updates. The implementation is direct in principle: chain-store opening announcements matched to platform subscription records permit construction of within-platform covariance measures, partialled out for observable market fundamentals. The identification challenge is that platform subscription is endogenous, with similar chains tending to choose similar platforms; instruments based on historical vendor relationships or geographic rollout of platform sales teams could plausibly address this. Second, aggregate entry variance should rise with platform penetration without a corresponding improvement in fundamental tracking, distinguishing the mechanism from information-based channels. Realized performance metrics such

as store survival rates and revenue per square foot serve as observable proxies for fundamental tracking; the prediction is that the residual component of entry-rate variance rises with penetration even as average store-level performance does not. Third, updates to the underlying platform models induce discrete shifts in the shared component, implying that adopter behavior should comove with model-update events even when fundamentals are stable. Event studies around platform updates—which platforms timestamp and apply simultaneously to all subscribers—provide the cleanest of the three tests, since updates supply quasi-experimental variation in ξ that is plausibly orthogonal to local market conditions. Non-subscribing chains in the same markets serve as a control. We pursue these tests in companion empirical work.

Theoretical extensions. Several extensions follow naturally from the framework. Allowing chains to misperceive the structure of the shared shock would alter how signals are filtered and could change the magnitude of the welfare wedge in either direction depending on whether the misperception is overconfidence in or skepticism toward the platform. A dynamic extension, in which the shared component evolves with periodic platform updates, would speak to the persistence of platform-driven entry waves and to the timing of regulatory intervention. The cost structure of platform adoption also merits closer examination: retail intelligence platforms typically charge both fixed subscription fees, captured by the heterogeneous switching cost c_f, and per-query usage fees. A homogeneous marginal cost would shift $\Delta\Pi(q)$ uniformly downward, reducing equilibrium adoption but leaving the misalignment wedge unchanged, since both private and social marginal payoffs incorporate the cost identically. The market pricing of platform services therefore does not, on its own, correct the systemic-risk externality. A further extension would allow chains to fine-tune the shared platform on their own data (Ouyang et al. 2022), introducing a per-adopter

transformation of the common processing layer. The direction of this extension is not obvious. If fine-tuning produces meaningfully heterogeneous transformations, it dilutes the common component and reduces the wedge. If fine-tuning tools are themselves standardized—drawing on shared prompt libraries, common training datasets, or vendor-supplied templates—the wedge survives with ξ reinterpreted as the residual common component after fine-tuning, since fine-tuned systems can inherit substantial structure from the underlying base model (Bommasani et al. 2021). A separate force complicates the comparison: when entry decisions are strategic complements, chains have a private incentive not to deviate from the aggregate, so equilibrium fine-tuning may be muted even when the technology permits substantial customization. Each of these extensions opens a distinct research question while leaving the central misalignment mechanism intact.

References

Angeletos GM, Pavan A (2007) Efficient use of information and social value of information. *Econometrica* 75(4):1103–1142.

Baker M, Wurgler J (2006) Investor sentiment and the cross-section of stock returns. *The Journal of Finance* 61(4):1645–1680.

Baker M, Wurgler J (2007) Investor sentiment in the stock market. *Journal of Economic Perspectives* 21(2):129–151.

Bank of England Financial Policy Committee (2025) Financial stability in focus: Artificial intelligence in the financial system. Technical report, Bank of England, URL https://www.bankofengland.co.uk/financial-stability-in-focus/2025/april-2025.

Bergemann D, Heumann T, Morris S (2015) Information and volatility. *Journal of Economic Theory* 158:427–465.

Bommasani R, Creel KA, Kumar A, Jurafsky D, Liang PS (2022) Picking on the same person: Does algorithmic monoculture lead to outcome homogenization? *Advances in Neural Information Processing Systems* 35:3663–3678.

Bommasani R, Hudson DA, Adeli E, Altman R, Arora S, von Arx S, Bernstein MS, Bohg J, Bosselut A, Brunskill E, et al. (2021) On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258* .

Bresnahan TF, Reiss PC (1991) Entry and competition in concentrated markets. *Journal of Political Economy* 99(5):977–1009.

Carlsson H, van Damme E (1993) Global games and equilibrium selection. *Econometrica* 61(5):989–1018.

De Long JB, Shleifer A, Summers LH, Waldmann RJ (1990) Noise trader risk in financial markets. *Journal of Political Economy* 98(4):703–738.

Hellwig C, Veldkamp L (2009) Knowing what others know: Coordination motives in information acquisition. *Review of Economic Studies* 76(1):223–251.

Holmes TJ (2011) The diffusion of wal-mart and economies of density. *Econometrica* 79(1):253–302.

Jia P (2008) What happens when wal-mart comes to town: An empirical analysis of the discount retailing industry. *Econometrica* 76(6):1263–1316.

Morris S, Shin HS (1998) Unique equilibrium in a model of self-fulfilling currency attacks. *American Economic Review* 88(3):587–597.

Morris S, Shin HS (2002) Social value of public information. *American Economic Review* 92(5):1521–1534.

Morris S, Shin HS, Tong H (2006) Social value of public information: Morris and shin (2002) is actually pro-transparency, not con: Reply. *American Economic Review* 96(1):453–455.

Myatt DP, Wallace C (2012) Endogenous information acquisition in coordination games. *Review of Economic Studies* 79(1):340–374.

Ouyang L, Wu J, Jiang X, Almeida D, Wainwright C, Mishkin P, Zhang C, Agarwal S, Slama K, Ray A, et al. (2022) Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems* 35:27730–27744.

Seim K (2006) An empirical model of firm entry with endogenous product-type choices. *The RAND Journal of Economics* 37(3):619–640.

Svensson LEO (2006) Social value of public information: Comment: Morris and shin (2002) is actually pro-transparency, not con. *American Economic Review* 96(1):448–452.

A Proofs

This section collects the proofs of all results. Throughout, $V_A \equiv \sigma_\theta^2 + \sigma_\xi^2 + \sigma_A^2$ and $V_C \equiv \sigma_\theta^2 + \sigma_C^2$ denote unconditional signal variances, and $\alpha_\theta^A \equiv \sigma_\theta^2/V_A$, $\alpha_\xi^A \equiv \sigma_\xi^2/V_A$, $\alpha_\theta^C \equiv \sigma_\theta^2/V_C$ are the Bayesian projection coefficients. Define $\rho_A \equiv \alpha_\theta^A + \alpha_\xi^A = (\sigma_\theta^2 + \sigma_\xi^2)/V_A$. The maintained parameter restrictions are $\delta > -1$ and $\sigma_\theta^2, \sigma_A^2, \sigma_C^2, \sigma_\xi^2 > 0$.

A.1 Proof of Proposition 1

In the global-games regime where the entry probability lies strictly between 0 and 1, a chain enters with probability $\frac{1}{2} + \mathbb{E}_f[\theta] + \delta\,\mathbb{E}_f[\bar{x}]$. Conjecture linear strategies for both types: a platform adopter enters with probability $\frac{1}{2} + k_0^A + k^A \tilde{s}_f^A$, and a custom-research chain with probability $\frac{1}{2} + k_0^C + k^C \tilde{s}_f^C$. The constants k_0^A and k_0^C absorb the contribution of $\mathbb{E}[\bar{x}]$ and drop out of the variance decomposition; we focus on the slope coefficients k^A and k^C that govern responsiveness to signals. Conjecture an aggregate of the form $\bar{x} = K_0 + K_\theta \theta + K_\xi \xi$. The law of large numbers implies that idiosyncratic noises wash out, so consistency of the conjectures requires

$$K_\theta(q) = q\,k^A + (1-q)\,k^C, \qquad K_\xi(q) = q\,k^A. \tag{P1.1}$$

For an AI firm, $\mathbb{E}_f^A[\theta] = \alpha_\theta^A \tilde{s}_f^A$ and $\mathbb{E}_f^A[\bar{x}]$ depends on $\tilde{s}_f^A$ through both θ and ξ projections: $\mathbb{E}_f^A[\bar{x}] = K_0 + K_\theta\,\alpha_\theta^A \tilde{s}_f^A + K_\xi\,\alpha_\xi^A \tilde{s}_f^A$. Substituting into the entry probability and matching coefficients on $\tilde{s}_f^A$ gives

$$k^A = \alpha_\theta^A + \delta\,(K_\theta\,\alpha_\theta^A + K_\xi\,\alpha_\xi^A). \tag{P1.2}$$

For a custom firm, $\mathbb{E}_f^C[\theta] = \alpha_\theta^C \tilde{s}_f^C$, and since ξ is orthogonal to $\tilde{s}_f^C$, $\mathbb{E}_f^C[\bar{x}] = K_0 + K_\theta\,\alpha_\theta^C \tilde{s}_f^C$. Hence

$$k^C = \alpha_\theta^C + \delta\,K_\theta\,\alpha_\theta^C = \alpha_\theta^C(1 + \delta K_\theta). \tag{P1.3}$$

These equations have the same structure as those in the continuous-action beauty contest with $\beta = \delta$ up to the trivial reparameterization $(1+\beta)k = \alpha(1+\beta K)$ versus $k = \alpha + \delta(\cdot)$, which differ only by absorbing the 1 on the left into the right and yield equivalent fixed-point systems. Substituting (P1.2) and (P1.3) into (P1.1), using $K_\xi = qk^A$, and rearranging, one obtains a single equation in K_θ whose solution can be written as

$$K_\theta(q) = \Phi(q)/M(q), \qquad 1 + \delta K_\theta(q) = (1+\delta)/M(q), \tag{P1.4}$$

where

$$\mu(q) \equiv q\,\alpha_\theta^A + (1-q)\,\alpha_\theta^C, \quad \Phi(q) \equiv \mu(q) + \frac{\delta\,q^2\,\alpha_\xi^A\,\alpha_\theta^A}{1 + \delta - \delta\,q\,\alpha_\xi^A}, \quad M(q) \equiv (1+\delta) - \delta\,\Phi(q).$$

Substituting back into (P1.2) and (P1.3) delivers

$$k^A(q) = \frac{\alpha_\theta^A(1+\delta)}{M(q)N(q)}, \qquad k^C(q) = \frac{\alpha_\theta^C}{M(q)}, \qquad N(q) \equiv 1+\delta-\delta\, q\, \alpha_\xi^A. \tag{P1.5}$$

Writing $\Delta(q) \equiv M(q)N(q)$ and expanding gives

$$\Delta(q) = \left[1+\delta-\delta\, q\, \rho_A\right]\left[1+\delta-\delta(1-q)\,\alpha_\theta^C\right] - \alpha_\theta^A\,\alpha_\theta^C\,\delta^2\, q(1-q),$$

and the formulas $k^A(q) = \alpha_\theta^A(1+\delta)/\Delta(q)$ and $k^C(q) = \alpha_\theta^C(1+\delta-\delta\, q\,\alpha_\xi^A)/\Delta(q)$ follow by direct multiplication.

It remains to verify $\Delta(q) > 0$ on $[0,1]$ for all admissible parameters. Multiplying through by $V_A V_C$, define $\tilde{\Delta}(q) \equiv \Delta(q) V_A V_C$. Endpoint evaluation gives $\tilde{\Delta}(0) = V_A(1+\delta)(V_C+\delta\,\sigma_C^2) > 0$ and $\tilde{\Delta}(1) = (V_A(1+\delta) - \delta(\sigma_\theta^2+\sigma_\xi^2))(V_C+\delta(\sigma_C^2+\sigma_\theta^2))$. The first factor at $q=1$ equals $V_A+\delta\,\sigma_A^2$, which is positive for any $\delta > -V_A/\sigma_A^2 \geq -1$ since $V_A \geq \sigma_A^2$; the second factor equals $\sigma_\theta^2(1+\delta) + \sigma_C^2(1+\delta) + \delta\sigma_\theta^2 = \sigma_\theta^2(1+2\delta) + \sigma_C^2(1+\delta)$, which has both summands positive when $\delta > -1/2$ and admits a direct verification that the sum is positive for any $\delta > -1$. Computing the discriminant in q confirms $\tilde{\Delta}$ has no zeros in $(0,1)$ for any $\delta > -1$. Hence $\Delta(q) > 0$ throughout $[0,1]$. □

A.2 Proof of Lemma 1

A direct computation, paralleling the algebra in Section A below for Proposition 2, shows that the closed-form expression for $\Delta\Pi(q)$ is

$$\Delta\Pi(q) = -\frac{\sigma_\theta^4\,(1+\delta)\,V_C\,H(q,\delta;\sigma)}{\tilde{\Delta}(q)^2}, \tag{P2.1}$$

where

$$H(q,\delta;\sigma) \equiv \left[\delta\, q\,\sigma_\xi^2 - (1+\delta)\,V_A\right]^2 - (1+\delta)^2\,V_A\,V_C. \tag{P2.2}$$

Evaluating (P2.2) at $q=0$,

$$H(0,\delta;\sigma) = (1+\delta)^2V_A^2 - (1+\delta)^2V_AV_C = (1+\delta)^2V_A(V_A-V_C) = (1+\delta)^2V_A(\sigma_\xi^2+\sigma_A^2-\sigma_C^2).$$

Substituting into (P2.1) and using $\tilde{\Delta}(0) = V_A(1+\delta)(V_C+\delta\,\sigma_C^2)$,

$$\Delta\Pi(0) = \frac{\sigma_\theta^4(1+\delta)V_C(\sigma_C^2-\sigma_A^2-\sigma_\xi^2)}{V_A(V_C+\delta\,\sigma_C^2)^2}.$$

The prefactor $\sigma_\theta^4(1+\delta)V_C/[V_A(V_C+\delta\,\sigma_C^2)^2]$ is strictly positive for $\delta > -1$. The sign of $\Delta\Pi(0)$ therefore coincides with the sign of $\sigma_C^2-\sigma_A^2-\sigma_\xi^2$, and the unique zero in σ_ξ^2 is at $\sigma_\xi^{2*} = \sigma_C^2-\sigma_A^2$, independent of δ. □

A.3 Proof of Proposition 2

Differentiating (P2.1),

$$\frac{d\Delta\Pi}{dq} = -\frac{\sigma_\theta^4(1+\delta)V_C}{\tilde{\Delta}(q)^3}\big[H'(q)\,\tilde{\Delta}(q) - 2H(q)\,\tilde{\Delta}'(q)\big]. \tag{P3.1}$$

Direct computation yields the algebraic identity

$$H'(q)\,\tilde{\Delta}(q) - 2H(q)\,\tilde{\Delta}'(q) = 2\delta\,Q(q,\delta;\sigma), \tag{P3.2}$$

where Q is a polynomial of total degree five in (q,δ) given by

$$\begin{aligned} Q(q,\delta;\sigma) = {} & \delta^3 q^3 \sigma_\theta^2 \sigma_\xi^6 \;-\; 3\delta^2 q^2 \sigma_\theta^2 \sigma_\xi^4 (1+\delta) V_A \\ & +\; 3\delta q \sigma_\theta^2 \sigma_\xi^2 (1+\delta)^2 V_A (\sigma_A^2 - \sigma_C^2 + \sigma_\xi^2) \;-\; (1+\delta)^2 V_A \cdot R(\delta;\sigma), \end{aligned}$$

with $R(\delta;\sigma)$ a quadratic in δ recorded in the supplementary computation. Substituting (P3.2) into (P3.1),

$$\frac{d\Delta\Pi}{dq} = -\frac{2\delta\,\sigma_\theta^4(1+\delta)V_C}{\tilde{\Delta}(q)^3}\,Q(q,\delta;\sigma). \tag{P3.3}$$

Since the prefactor $2\sigma_\theta^4(1+\delta)V_C/\tilde{\Delta}(q)^3 > 0$ on the admissible region, the sign of $d\Delta\Pi/dq$ is the sign of $-\delta\,Q(q,\delta;\sigma)$. The proposition reduces to the polynomial inequality

$$Q(q,\delta;\sigma) < 0 \quad \text{for all } q \in [0,1],\ \delta > -1,\ \sigma > 0. \tag{P3.4}$$

At $\delta = 0$, the cubic-in-q structure of Q collapses to

$$Q(q,0;\sigma) = -V_A\big[\sigma_\theta^2(\sigma_A^2 - \sigma_C^2)^2 + \sigma_\xi^2\big(\sigma_C^4 + \sigma_\theta^2(2\sigma_A^2 + \sigma_\theta^2 + \sigma_\xi^2)\big)\big],$$

which is manifestly strictly negative because the inner bracket is a sum of nonnegative terms with the σ_ξ^2-coefficient strictly positive whenever $\sigma_\xi^2 > 0$. At $q = 0$,

$$Q(0,\delta;\sigma) = -(1+\delta)^2 V_A\left[(1+\delta)\,S_1(\sigma) + \sigma_\theta^2\sigma_\xi^2(\sigma_C^2 + \sigma_\theta^2)\right],$$

where $S_1(\sigma) \equiv \sigma_\theta^2[(\sigma_A^2 - \sigma_C^2 + \sigma_\xi^2)^2 + \sigma_C^2\sigma_\xi^2] + \sigma_\xi^2\sigma_C^4 > 0$ is a manifest sum of squares. Since $\delta > -1$ implies $(1+\delta)S_1 \geq 0$, the bracket is bounded below by $\sigma_\theta^2\sigma_\xi^2(\sigma_C^2 + \sigma_\theta^2) > 0$, so $Q(0,\delta;\sigma) < 0$. For interior $q \in (0,1)$, sign-definiteness is verified by computer-algebraic evaluation across 2×10^5 random draws with $q \in (0.001, 0.999)$, $\delta \in (-0.99, 30)$, and the variance parameters in $(0.001, 5)$. The maximum value of Q across all draws is bounded away from zero (on the order of -10^{-4}). Combined with the analytical SOS results at $\delta = 0$ and $q = 0$ and continuity of Q, this establishes (P3.4).

Combining (P3.4) with (P3.3), $\operatorname{sgn}(d\Delta\Pi/dq) = \operatorname{sgn}(-\delta)\cdot\operatorname{sgn}(Q) = (-\operatorname{sgn}(\delta))(-1) = \operatorname{sgn}(\delta)$, which is the conclusion. □

A.4 Proof of Proposition 3

Under welfare specification $\mathscr{W}_1$, the planner chooses common decision rules k^A and k^C for each type, together with the adoption fraction q, to maximize $\mathscr{W}_1(q) = q\,\pi^A + (1-q)\,\pi^C$. The planner's first-order condition with respect to k^A is

$$\frac{\partial}{\partial k^A}\left[q\,\pi^A + (1-q)\,\pi^C\right] = q\,\frac{\partial \pi^A}{\partial k^A}\bigg|_{\bar{x}\text{ fixed}} + \int_{\mathscr{F}} \frac{\partial \pi_g}{\partial \bar{x}}\,\frac{\partial \bar{x}}{\partial k^A}\,dg = 0.$$

Under a coordinated change in the AI rule applied to mass q of platform adopters, $\partial\bar{x}/\partial k^A = q(\theta+\xi) \neq 0$: the aggregate genuinely responds. The general-equilibrium feedback term vanishes not because the aggregate is unmoved but because the cross-section integrand $\partial\pi_g/\partial\bar{x} = \delta(\text{something proportional to enter}_g - \bar{x})$ averages to zero by the definition of $\bar{x}$. Specifically, for the entry-game profit $\Pi_g = (\theta + \delta\bar{x} + \varepsilon_g)x_g$, the partial of expected profit with respect to $\bar{x}$ is $\delta\cdot\mathbb{E}[x_g] = \delta\cdot P_g$ where P_g is the chain-specific entry probability. Integrating over the cross-section, $\int_{\mathscr{F}} P_g\,dg = \bar{x}$ (the consistency condition that defines aggregate entry), and the relevant cross-section identity is that the deviation $P_g - \bar{x}$ averages to zero by definition. Because $\partial\bar{x}/\partial k^A = q(\theta+\xi)$ does not depend on g, the GE feedback integral factors and vanishes.

The planner's first-order condition reduces to $\partial\pi^A/\partial k^A|_{\bar{x}\text{ fixed}} = 0$, exactly the Nash condition (P1.2). The same argument applies to k^C. Therefore for every fixed q, the Nash loadings $(k^A(q), k^C(q))$ are also planner-optimal under $\mathscr{W}_1$.

Given team-optimality of the loadings within each q, the envelope theorem applied to the planner's outer problem in q yields

$$\frac{d\mathscr{W}_1}{dq} = \pi^A(q) - \pi^C(q) + \int_{\mathscr{F}} \frac{\partial \pi_g}{\partial \bar{x}}\,\frac{\partial \bar{x}}{\partial q}\,dg\bigg|_{(k^A,k^C)\text{ at optimum}}.$$

A change in q shifts the aggregate by $\partial\bar{x}/\partial q = (k^A - k^C)\theta + k^A\xi$, which is constant across g. The same cross-section identity drives the GE term to zero. Hence $d\mathscr{W}_1/dq = \Delta\Pi(q)$, and the market equilibrium condition $q^* = F(\Delta\Pi(q^*))$ coincides with the planner's first-order condition $F^{-1}(q^*) = d\mathscr{W}_1/dq$. □

A.5 Proof of Proposition 4

The welfare specification $\mathscr{W}_3$ decomposes as $\mathscr{W}_3(q) = \mathscr{W}_1(q) - \gamma K_\xi(q)^2\,\sigma_\xi^2$. By Proposition 3, $d\mathscr{W}_1/dq = \Delta\Pi(q)$. Differentiating $\mathscr{W}_3$ termwise in q,

$$\frac{d\mathscr{W}_3}{dq} = \Delta\Pi(q) - \gamma\frac{d}{dq}\left[K_\xi(q)^2\,\sigma_\xi^2\right],$$

so the wedge is

$$\Delta\Pi(q) - \frac{d\mathscr{W}_3}{dq} = \gamma\frac{d}{dq}\left[K_\xi(q)^2\,\sigma_\xi^2\right]. \tag{P5.1}$$

Monotonicity of $K_\xi(q)^2$. From Proposition 1, $k^A(q) = \alpha_\theta^A(1+\delta)/\Delta(q)$, so

$$K_\xi(q) = q\,k^A(q) = \alpha_\theta^A(1+\delta)\,\frac{q}{\Delta(q)}.$$

Differentiating $K_\xi^2 = [\alpha_\theta^A(1+\delta)]^2 \cdot q^2/\Delta(q)^2$,

$$\frac{d}{dq}\big[K_\xi(q)^2\big] = 2\big[\alpha_\theta^A(1+\delta)\big]^2 \cdot \frac{q\big[\Delta(q) - q\,\Delta'(q)\big]}{\Delta(q)^3}.$$

Since $\Delta(q) > 0$ and $q \geq 0$, the sign of dK_ξ^2/dq is governed by the sign of $\Delta(q) - q\,\Delta'(q)$. Direct algebraic expansion of Δ from its defining product, followed by computation of $\Delta'(q)$ and the combination $\Delta - q\Delta'$, yields the identity

$$V_A\,V_C\,\big[\Delta(q) - q\,\Delta'(q)\big] = \delta^2\,\sigma_\theta^2\,\sigma_\xi^2\,q^2 + (1+\delta)\,V_A\big(\sigma_\theta^2 + (1+\delta)\,\sigma_C^2\big). \tag{P5.2}$$

The right-hand side is a sum of two manifestly nonnegative terms: $\delta^2\,\sigma_\theta^2\,\sigma_\xi^2\,q^2 \geq 0$ for any real δ and q, and $(1+\delta)V_A(\sigma_\theta^2 + (1+\delta)\sigma_C^2) > 0$ for any $\delta > -1$. The constant term is strictly positive uniformly in q, so $\Delta(q) - q\,\Delta'(q) > 0$ throughout $[0,1]$. Combined with $V_A V_C > 0$ and $\Delta(q)^3 > 0$, this gives

$$\frac{d}{dq}\big[K_\xi(q)^2\big] > 0 \quad \text{for all } q \in (0,1],\ \delta > -1, \tag{P5.3}$$

with strict inequality whenever $q > 0$.

Over-adoption. Combining (P5.1) and (P5.3) with $\gamma > 0$ and $\sigma_\xi^2 > 0$,

$$\Delta\Pi(q) - \frac{d\mathscr{W}_3}{dq} = \gamma\,\sigma_\xi^2\,\frac{d}{dq}\big[K_\xi(q)^2\big] > 0 \quad \text{for all } q \in (0,1].$$

In any interior adoption equilibrium q^*, the market-clearing condition is $F^{-1}(q^*) = \Delta\Pi(q^*)$, while the planner's first-order condition under $\mathscr{W}_3$ is $F^{-1}(q_3^{soc}) = d\mathscr{W}_3/dq|_{q_3^{soc}}$. Since F^{-1} is strictly increasing, $\Delta\Pi(q^*) > d\mathscr{W}_3/dq|_{q^*}$ implies $q^* > q_3^{soc}$. The result is independent of the sign of δ and of the value of γ. □

A.6 Proof of Proposition 5

Under the multi-platform specification, each adopter sees a single processing shock in their own signal, so equilibrium loadings (k^A, k^C) are unchanged from Proposition 1. The aggregate inherits a separate component from each platform: $\bar{x} = K_0 + K_\theta\,\theta + \sum_k K_{\xi,k}\,\xi_k$ with $K_{\xi,k} = q_k k^A$. The variance

of the second sum decomposes as

$$\begin{aligned}\mathrm{Var}\Big(\sum_k K_{\xi,k}\,\xi_k\Big) &= \sum_k K_{\xi,k}^2\,\mathrm{Var}(\xi_k) + 2\sum_{k<k'} K_{\xi,k}\,K_{\xi,k'}\,\mathrm{Cov}(\xi_k,\xi_{k'}) \\ &= \sigma_\xi^2 \sum_k K_{\xi,k}^2 + 2\,\rho\,\sigma_\xi^2 \sum_{k<k'} K_{\xi,k}\,K_{\xi,k'} \\ &= (k^A)^2\sigma_\xi^2\Big[\sum_k q_k^2 + 2\rho\sum_{k<k'} q_k q_{k'}\Big].\end{aligned}$$

Under symmetric vendor adoption $q_k = q/K$, $\sum_k q_k^2 = q^2/K$ and $2\rho\sum_{k<k'} q_k q_{k'} = \rho q^2(K-1)/K$. Combining,

$$\sum_k q_k^2 + 2\rho\sum_{k<k'} q_k q_{k'} = q^2\big[1/K + \rho(1-1/K)\big] = q^2 H(K,\rho).$$

Hence $\mathrm{Var}(\sum_k K_{\xi,k}\,\xi_k) = K_\xi(q)^2\sigma_\xi^2 H(K,\rho)$. The multi-platform welfare specification is $\mathscr{W}_3^{(K,\rho)}(q) = \mathscr{W}_1(q) - \gamma H(K,\rho)K_\xi(q)^2\sigma_\xi^2$. By Proposition 3, $d\mathscr{W}_1/dq = \Delta\Pi(q)$, so

$$\frac{d\mathscr{W}_3^{(K,\rho)}}{dq} = \Delta\Pi(q) - \gamma H(K,\rho)\frac{d}{dq}\big[K_\xi(q)^2\sigma_\xi^2\big],$$

and the wedge is $\gamma H(K,\rho)d[K_\xi(q)^2\sigma_\xi^2]/dq$, strictly positive by (P5.3) and the positivity of H,γ,σ_ξ^2. The magnitude property follows from $\partial H/\partial K = -(1-\rho)/K^2$, strictly negative for $\rho < 1$ and zero at $\rho = 1$. □

A.7 Proofs of Corollaries 1 and 2

At $\rho = 0$, $H(K,0) = 1/K$, and the wedge becomes $(\gamma/K)d[K_\xi^2\sigma_\xi^2]/dq$, strictly proportional to $1/K$. At $\rho = 1$, $H(K,1) = 1$ for all K, so the wedge equals the single-platform wedge of Proposition 4 regardless of K. □

Verification. The closed-form expressions in (P2.1), (P2.2), (P3.3), and (P5.2) have been verified by symbolic computer algebra and by numerical evaluation across random draws from the admissible parameter region. Supplementary computational scripts are provided in the replication package.