Shared Intelligence and Misaligned Firm Entry

Feng Chi* Nathan Yang†

May 2, 2026

Abstract

Retail entry decisions are increasingly mediated by shared intelligence platforms whose forecasts embed a common component orthogonal to local fundamentals. We develop a Bayesian entry game in which chains choose between proprietary research and a shared platform that reduces idiosyncratic noise but introduces a common shock. The model admits a closed-form equilibrium with heterogeneous adoption. Equilibrium adoption is excessive under two externalities: aggregate-volatility costs and correlated-evaluation costs on locations. The wedge is invariant to whether competition is agglomerative or business-stealing, and under the second externality its magnitude is independent of the strategic parameter. A multi-vendor extension yields an effective concentration index governing both externalities.

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 were historically grounded in local knowledge. A regional grocer learned its trade area through years of operating stores in adjacent counties; a restaurant chain relied on market researchers who walked competitor parking lots and counted cars. Entry decisions emerged from processes that were, at their core, idiosyncratic to each firm. Two chains evaluating the same county could reach different conclusions because they processed the available information differently.

This is no longer how retail entry operates. CoStar, Placer.ai, Buxton, and the analytics divisions of CBRE and JLL now sell forecasts of trade-area demand, foot traffic, competitive density, and demographics to large numbers of subscribers simultaneously. A national chain, a regional grocer, and a franchisee may all consult the same platform, receive the same model output, and update in the same direction.[1] The Bank of England's Financial Policy Committee notes 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 concern applies in retail markets, where shared platforms increasingly mediate decisions that were once the product of independent investigation.

The central mechanism can be summarized simply. When firms rely on a shared system, they do not just learn the same things; they make the same mistakes. The common component introduced by the platform survives aggregation and induces correlated behavior that is *unrelated* to underlying market conditions.

Recent empirical work provides suggestive evidence that shared information environments can

[1] Industry evidence suggests that a small number of providers are widely used across retail and commercial real estate. For example, Esri reports that seven of the ten largest U.S. retailers rely on its geospatial analytics platform, while firms such as Placer.ai provide location intelligence services to clients spanning retail, real estate, and hospitality sectors. These platforms are also frequently deployed through intermediaries such as commercial real estate advisory firms, implying that many decision-makers rely on overlapping data sources and processing systems. See https://www.esri.com/en-us/industries/retail/overview and https://www.placer.ai/company/customers.

shape firm behavior in systematic ways. Firms that invest more heavily in AI exhibit increases in systematic risk, becoming more correlated with aggregate outcomes in a manner not observed for other technologies (Babina et al. 2024). At the same time, recent work highlights the economic importance of common data sources in firm decision-making. Using the decennial U.S. Census as a source of variation in information quality, Chi (2025) shows that firms relying on shared demographic data are more likely to make misinformed investment decisions when that data becomes outdated, leading to systematically worse outcomes across cohorts of entrants. Because many firms access such data directly or through common intermediaries, declines in information quality propagate across decisions, generating correlated mistakes rather than independent errors. Taken together, these findings point to a tension. Shared intelligence can improve individual decision-making while inducing correlated behavior across firms. What remains unclear is whether this correlation reflects firms reacting to the same underlying information, or instead reacting to the same errors.

The question is not whether shared platforms produce more accurate forecasts than custom research. Accuracy may vary across contexts. The relevant question is how industry-level entry behavior changes when forecasts that were previously independent are routed through a common processing layer. When firms rely on the same platform, their forecast errors acquire a shared component originating in the platform itself—its data, modeling choices, and updates. This component is correlated across firms but unrelated to local fundamentals. It reflects the structure of the processing technology rather than the environment being studied. As a result, firms may not only act similarly, but may be wrong in similar ways.

Standard models of retail entry do not generate correlation of this form. In canonical frameworks (Bresnahan and Reiss 1991, Seim 2006, Jia 2008, Holmes 2011), firms receive private information about profitability, and correlation arises either from strategic interaction or from shared information about a common underlying state. In each case, correlation is tied to economically relevant information. Even when firms coordinate or follow similar strategies, their errors remain

largely independent. Platform-based forecasts differ in a fundamental way. They do not become more informative as adoption rises. Instead, the same errors propagate across more decisions. The resulting correlation is technological in origin, not strategic and not informational.

A useful analogy arises in financial markets. Large asset managers increasingly rely on shared risk-management systems such as BlackRock's Aladdin platform, which processes portfolio data and produces risk assessments used by many institutions simultaneously. When a large fraction of the market evaluates risk through the same system, the concern is not only that participants respond to similar information, but that they may inherit the same model-driven distortions. In this case, correlation arises from a shared processing layer rather than from common fundamentals or explicit coordination. This phenomenon, often described as model monoculture, highlights the same mechanism studied here in a different setting.

Two strands of theory provide useful benchmarks. The first is the social-value-of-information literature (Morris and Shin 2002, Angeletos and Pavan 2007, Hellwig and Veldkamp 2009, Myatt and Wallace 2012), where common signals introduce correlation that trades off coordination benefits against information loss. The signal structure we study departs from this benchmark in a specific way: the common component is unrelated to the underlying state. Because it carries no information, the usual tradeoff is absent. The second is the algorithmic-monoculture literature (Kleinberg and Raghavan 2021, Peng and Garg 2024), which emphasizes a different cost of correlated decision-making: when agents rely on the same system, independent evaluations disappear, reducing the chance that a favorable case is recognized by at least one decision-maker. We show that both types of cost arise in the environment studied here.

This paper develops a tractable model that isolates this mechanism. A continuum of potential entrants chooses whether to enter a representative market and, before doing so, whether to rely on a shared platform or on custom research. The platform improves precision by reducing idiosyncratic noise but introduces a common component shared across adopters. The model admits a closed-form equilibrium with heterogeneous adoption, which allows us to characterize how aggregate

entry reflects both underlying fundamentals and platform-driven variation, how adoption incentives change as platforms become more widespread, and how equilibrium adoption compares to socially optimal levels.

The equilibrium has two central features. First, the influence of the shared component on aggregate entry increases with platform adoption. As more firms rely on the platform, the aggregate becomes more exposed to the same underlying error. Second, the private value of adoption depends on the strategic environment. When firms benefit from aligning with one another, adoption becomes more attractive as it spreads; when firms benefit from differentiation, the opposite holds. These forces together generate a range of adoption patterns, including environments in which widespread adoption is sustained by coordination motives alone.

The main normative result is a robust misalignment. The shared component generates two distinct external costs that are not internalized by firms. One operates at the aggregate level: synchronized entry decisions create volatility that affects lenders, landlords, and local governments. The other operates at the level of individual locations: when many firms rely on the same system, a location may be rejected by all of them if the shared component is unfavorable, even if it is attractive on fundamentals. In both cases, equilibrium platform adoption is excessive. Importantly, this conclusion does not rely on distinguishing among different sources of coordination or strategic interaction. Instead, it follows from the presence of a common error component that firms do not internalize.

This distinction clarifies the contribution of the model. Many mechanisms can generate correlated behavior or even excessive adoption. What is distinctive here is that correlation arises in errors rather than only in actions, and this generates welfare losses that persist even when firms optimally respond to one another. The empirical implication is that the key object is not simply whether firms move together, but whether they move together in ways that are unrelated to fundamentals.

Extending the model to multiple platforms shows that the magnitude of the externality depends

on how independent these platforms are. If platforms rely on distinct data and methods, their errors are less correlated, and the externality is smaller. If they share a common data or modeling core, additional platforms do not reduce correlation. What matters is not the number of providers but the diversity of the underlying processing technologies.

The framework also highlights a coordination problem in adoption. When firms benefit from aligning with one another, a platform can become widely adopted even if it is not attractive in isolation. In such cases, the industry can settle in a high-adoption outcome sustained purely by coordination motives. This pattern is most relevant in sectors with strong agglomeration economies, where anticipating competitors' behavior is central to entry decisions.

2 Model

The model formalizes the entry decision of a retail chain that chooses both whether to enter a representative market and which information technology to use when forecasting profitability. The environment is intentionally sparse: a single market, two information technologies, and a continuum of potential entrants. This isolates the feature of interest (i.e, a shared processing layer that introduces correlated errors across adopters) while retaining tractability sufficient for closed-form equilibrium analysis.

A key distinction in the analysis is between primitives and outcomes. The variance of the shared component is a primitive feature of the information structure. By contrast, platform adoption is an endogenous equilibrium outcome that reflects firms' responses to this primitive. The normative analysis therefore operates on the underlying variance of shared processing rather than on the realized level of adoption itself.

2.1 Environment

A unit continuum of retail chains $f \in [0,1]$ decides whether to enter a representative local market. The decision is binary, $x_f \in \{0,1\}$, and 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 an idiosyncratic cost shock.

The parameter $\delta \in (-1, \infty)$ governs strategic interaction. When $\delta > 0$, profitability rises with rival entry, capturing agglomeration economies: clusters such as food halls, outlet centers, or fast-food strips draw demand that no single store could attract alone. When $\delta < 0$, profitability falls with rival entry, capturing business-stealing in markets with limited demand. The restriction $\delta > -1$ ensures existence of a unique linear equilibrium.

The uniform distribution of ε_f is a standard global-games device (Carlsson and van Damme 1993, Morris and Shin 1998) that renders entry probabilities linear in posterior beliefs. We work in the regime where $\bar{\varepsilon}$ is large relative to signal noise, so truncation at 0 or 1 occurs with negligible probability.[2]

For convenience, we normalize $\bar{\varepsilon} = 1/2$, so that entry probability responds one-for-one to the posterior mean. This normalization is without loss of generality and has no effect on the qualitative results.

[2] Formally, linearity holds whenever $|\mathbb{E}_f[\theta] + \delta \mathbb{E}_f[\bar{x}]| \leq \bar{\varepsilon}$. We maintain $\bar{\varepsilon} > 4\sqrt{\text{Var}(\theta) + \text{Var}(\bar{x})}$ as a sufficient condition; this is satisfied throughout the parameter region studied below.

2.2 Information and Processing Technologies

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

$$\tilde{s}_f^C = \theta + u_f^C,$$

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

$$\tilde{s}_f^A = \theta + \xi + u_f^A,$$

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

The distinction between the two technologies is not only a matter of precision but of structure. Both signals are noisy measures of θ, and platform adoption may reduce idiosyncratic variance, so it is natural to expect $\sigma_A^2 < \sigma_C^2$ when data are pooled across users. At the same time, the platform introduces a common component ξ that is perfectly correlated across adopters and orthogonal to the underlying market. Adoption therefore replaces idiosyncratic noise with shared noise. This trade can be privately attractive while generating aggregate effects that no individual chain internalizes.

This structure differs fundamentally from the public-signal benchmark of Morris and Shin (2002). A public signal induces correlation because it is informative about θ, so aggregation improves inference even as actions become more aligned. Here the common component carries no information about θ and survives aggregation as a pure error term. The cross-section of platform adopters does not refine beliefs through joint exposure to ξ; it inherits a common mistake.[3]

It is useful to interpret the shared component as a latent processing factor. This factor summarizes the common distortions induced by a shared decision technology. Unlike fundamentals,

[3]This differs also from environments such as Bergemann et al. (2015), where common signals may confound multiple sources of uncertainty but remain informative. The orthogonality of ξ to θ is a property of the processing technology itself.

which average out in the cross section, the shared processing factor persists in aggregate behavior and directly shapes equilibrium outcomes.

Remark 1 (Behavioral interpretation)**.** The shared component ξ admits a useful interpretation. When many chains rely on the same platform, it is as if they employ the same analyst, whose biases, training, and blind spots affect all users simultaneously. The variance σ_ξ^2 measures the magnitude of these shared distortions, while orthogonality to θ captures that they are unrelated to the underlying market. This parallels models of investor sentiment as a common factor orthogonal to fundamentals (Baker and Wurgler 2006, 2007) and correlated belief errors (De Long et al. 1990). The difference is in the source: the correlation arises from a shared decision technology rather than from shared psychology.

2.3 Adoption Stage

Each chain draws a switching cost $c_f \sim F$ independently, where F is a continuous distribution on $\mathbb{R}$ with positive density on its support. The chain chooses its forecasting technology to maximize expected entry profits net of this cost, taking the adoption rate q as given. An adoption equilibrium is a rational-expectations fixed point in which individual choices reproduce the aggregate, so that $q^* = F(\Delta\Pi(q^*))$, 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 can arise: zero adoption when $\Delta\Pi(0) \leq 0$ and full adoption when $\Delta\Pi(1) \geq \bar{c}$. With unbounded support, such as the logistic specification used below, these are replaced by interior equilibria with q^* close to but bounded away from 0 or 1. This distinction does not affect any qualitative result.

2.4 Bayesian Filtering

Chains form beliefs using Bayesian filtering given their information 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 projection coefficients

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

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 and therefore assign it its prior mean.

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 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: agents correctly understand the structure of the platform and filter signals accordingly. Allowing for misperception of σ_ξ^2 would alter how signals are weighted but not the underlying source of correlation; we return to this extension in the discussion.

3 Equilibrium with Heterogeneous Adoption

This section characterizes equilibrium behavior for a fixed adoption rate q. We first derive the linear equilibrium in which entry probabilities are affine in signals, and then use this structure to characterize aggregate behavior.

3.1 Best Responses and the Linear Fixed Point

A chain enters if its posterior expected profit is nonnegative, $\mathbb{E}_f[\theta] + \delta \mathbb{E}_f[\bar{x}] + \varepsilon_f \geq 0$. With $\varepsilon_f \sim$ Uniform$[-1/2, 1/2]$, the implied entry probability is

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

in the region where this expression lies strictly between 0 and 1.

We look for an equilibrium in which entry probabilities are linear in signals and the aggregate is linear in the underlying shocks. Conjecture

$$\bar{x} = K_0 + K_\theta \theta + K_\xi \xi.$$

In a continuum, idiosyncratic noise averages out, so consistency requires that aggregate loadings match the population-weighted contributions of each type. Let k^A and k^C denote the slopes of entry probability with respect to signals for platform adopters and custom chains. Then

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

Substituting the conjectured aggregate into each type's decision rule yields a linear system in (k^A, k^C) with a unique solution.

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 distinction between the two loadings reflects how the shared shock propagates through expectations. Platform adopters respond directly to both θ and ξ, while custom chains respond only to θ. The term $-\delta\,q\,\alpha_\xi^A$ in k^C captures an indirect effect: although custom chains do not observe ξ, they anticipate its impact on aggregate entry through adopters. They therefore adjust their responses to offset this effect under $\delta < 0$ and to reinforce it under $\delta > 0$. This adjustment becomes stronger as adoption rises and as the shared component accounts for a larger share of signal variance.

The boundary cases recover familiar benchmarks. When $q = 0$, the platform is absent and the equilibrium reduces to the standard linear-Gaussian entry game with private signals, $k^C(0) = \alpha_\theta^C/(1+\delta-\delta\alpha_\theta^C)$. 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 contains both informative and uninformative common components.

3.2 Aggregate Behavior

Substituting the equilibrium loadings into the aggregate yields

$$\bar{x} = K_0 + K_\theta(q)\,\theta + K_\xi(q)\,\xi.$$

Using $\mathbb{E}[\theta] = \mathbb{E}[\xi] = 0$, the deviation from the mean is

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

so aggregate variance is

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

This expression separates two distinct components. The first reflects fundamental tracking: the aggregate loads on θ, and its contribution scales with σ_θ^2. The second reflects the shared processing layer: the aggregate loads on ξ, and its contribution scales with σ_ξ^2.

In a custom-only economy, $K_\xi(0) = 0$, so all aggregate variation reflects fundamentals. Once platform adoption is positive, this changes qualitatively. The term $K_\xi(q)^2\sigma_\xi^2$ is strictly positive and does not vanish under aggregation. Idiosyncratic noise washes out in the continuum, but the common shock survives intact.

The key implication is that shared platforms introduce aggregate volatility that is orthogonal to fundamentals. Correlation across chains is no longer tied to information; it reflects a common

error embedded in the processing technology. The term $K_\xi(q)^2\sigma_\xi^2$ is therefore the central object of the model and the source of the welfare distortions analyzed below.

4 Adoption Incentives and Misalignment

This section characterizes adoption incentives and compares equilibrium adoption to social optima under alternative welfare criteria. The central question is whether chains internalize the consequences of shared processing for parties outside the chain population. The analysis shows that coordination motives are fully internalized in equilibrium, but the shared shock generates two distinct externalities that are not. Each externality alone implies over-adoption, so the result does not depend on a particular welfare specification.

4.1 The Private Gain to Adoption

Let $\pi^A(q)$ and $\pi^C(q)$ denote expected profits for platform adopters and custom-research chains at adoption rate q, obtained by substituting equilibrium behavior into the profit function. The private gain from adoption is $\Delta\Pi(q) \equiv \pi^A(q) - \pi^C(q)$, and an adoption equilibrium satisfies the fixed-point condition $q^* = F(\Delta\Pi(q^*))$.

The dependence of $\Delta\Pi(q)$ on q reflects two forces. The platform changes forecast precision and also changes how closely a chain's entry decision aligns with rivals. The relative importance of these forces is governed by the strategic-interaction parameter 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 $\delta = 0$, chains care only about matching θ, so adoption depends solely on the signal-precision tradeoff and is unaffected by penetration. When $\delta > 0$, adoption becomes more attractive as it spreads: the aggregate becomes more exposed to the shared component ξ, and platform

adopters (who partially forecast this component) track aggregate entry more closely than custom chains. When $\delta < 0$, the same mechanism reverses. Chains benefit from differentiating from rivals, and platform adoption increases correlation, reducing its marginal value as q rises.

The result follows from the sign of the derivative of $\Delta\Pi(q)$. Differentiating the closed-form expression 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 that is strictly negative over the admissible parameter region. The sign of $d\Delta\Pi/dq$ is therefore determined by δ. Details are provided in the proof.

4.2 Two Externalities, Two Affected Parties

Identifying the welfare costs of shared processing requires identifying who bears them. Chain profits, captured by the utilitarian benchmark below, do not exhaust the relevant population. The shared shock ξ enters equilibrium behavior in two distinct ways, and each operates on a different group outside the chain population. We treat these as separate externalities rather than as alternative parameterizations of a single welfare object.

The first externality operates at the aggregate level. It affects agents whose payoffs depend on the variance of aggregate entry around fundamentals. In the retail context, this includes commercial mortgage lenders exposed to correlated default risk, landlords facing synchronized vacancies, and municipal authorities planning around expected commercial activity. The relevant object is the platform-driven component of aggregate variance, $K_\xi(q)^2\sigma_\xi^2$. As established in the previous section, this term is increasing in adoption and does not vanish under aggregation. The cost borne by this group therefore scales with the variance contribution of the shared shock.

The second externality operates at the level of individual locations. Locations (e.g., census

tracts, lease parcels, and micro-markets) are evaluated by multiple chains as potential entry targets. When chains use the shared platform, their signals contain the same realization of ξ, so their decisions are correlated. A location that is favorable on fundamentals but receives an unfavorable realization of ξ is rejected by all platform adopters, with no independent evaluation within that group. Custom-research chains evaluate independently, so such a location retains a positive probability of entry when some chains do not adopt the platform. As adoption rises, the share of correlated evaluations increases and the probability that a fundamentally attractive location receives at least one entry attempt declines. The cost borne by this group is the foregone activity in locations passed over due to correlated decisions.

This second mechanism corresponds to the externality emphasized by Kleinberg and Raghavan (2021) and Peng and Garg (2024). In those settings, the decided-about agents are individuals, and the cost is denial of service. Here the agent is a location and the cost is unrealized economic activity. The mapping differs in interpretation but not in structure: in both cases, the externality arises from correlated evaluation errors generated by a shared decision technology.

The two externalities differ in the objects they act on. The first depends on aggregate behavior and therefore on equilibrium loadings; the second depends only on the correlation structure of signals and the probability of shared exposure. Both originate in the same primitive feature of the model (i.e., the orthogonality of ξ to θ) but operate through distinct channels. We formalize each in turn.

4.3 Welfare Specifications

We consider three welfare specifications. The first is a utilitarian benchmark that aggregates chain profits. The second and third correspond to the two externalities identified above. A fourth specification, isolating fundamental-tracking performance, is introduced for comparison in Remark 2.

The utilitarian benchmark, denoted $\mathscr{W}_1$, aggregates chain profits:

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

This specification captures the welfare of chains alone. Any wedge between private and social incentives under $\mathscr{W}_1$ would reflect a failure to internalize the effect of one chain's behavior on others, which Proposition 3 shows does not arise.

The aggregate-volatility specification, denoted $\mathscr{W}_3$, augments $\mathscr{W}_1$ with a penalty on the platform-driven component of aggregate-entry variance:

$$\mathscr{W}_3(q) = \mathscr{W}_1(q) - \gamma_3\,K_\xi(q)^2\,\sigma_\xi^2, \qquad \gamma_3 > 0.$$

The term $K_\xi(q)^2\sigma_\xi^2 = q^2(k^A(q))^2\sigma_\xi^2$ is the contribution of the shared shock to aggregate variance. It depends on equilibrium loadings and therefore on the strategic-interaction parameter δ. The parameter γ_3 measures the marginal social cost of platform-induced volatility.

The decided-about-agent specification, denoted $\mathscr{W}_4$, instead penalizes correlated evaluation at the level of individual locations:

$$\mathscr{W}_4(q) = \mathscr{W}_1(q) - \gamma_4\,q^2\,\frac{\sigma_\xi^2}{V_A}, \qquad \gamma_4 > 0.$$

This term decomposes into two components. The factor q^2 is the probability that two chains evaluating the same location are both platform adopters. The factor σ_ξ^2/V_A is the conditional correlation of platform-signal errors across adopters. Their product, $q^2\sigma_\xi^2/V_A$, is the expected correlation in evaluation errors at penetration q. Unlike the aggregate-volatility term, it depends only on signal-structure primitives and not on equilibrium responses.

The distinction between $\mathscr{W}_3$ and $\mathscr{W}_4$ is structural. Both penalize consequences of the shared

shock ξ, but they operate on different objects.[4] Aggregate volatility is a property of the realized entry rate and therefore depends on equilibrium behavior. Encounter-weighted correlation is a property of the signal structure and adoption rate alone. The first externality is borne by agents exposed to aggregate outcomes; the second by locations subject to correlated evaluations. These are distinct channels, and results derived under one do not mechanically carry over to the other.

4.4 The Team-Efficiency Result

The first welfare result provides a benchmark. 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 structure with a continuum of agents. For any fixed q, the equilibrium loadings (k^A, k^C) maximize average chain profits subject to the information structure. A planner who assigns a common decision rule to each type faces the same optimization problem as individual chains, except for a general-equilibrium term capturing how coordinated changes affect the aggregate.

This term vanishes because of a cross-sectional identity. Although the aggregate responds to coordinated changes in behavior, the marginal payoff with respect to $\bar{x}$ integrates to zero across chains, $\int_{\mathscr{F}} (\partial \pi_g / \partial \bar{x})\, dg = 0$, so the planner's first-order condition coincides with the Nash condition. The same argument applies to changes in the adoption rate, implying $d\mathscr{W}_1/dq = \Delta\Pi(q)$.

The implication is that shared processing does not create inefficiency within the chain population. Even though the platform induces correlated behavior, chains fully internalize its effects on profits through equilibrium responses. Any wedge between private and social incentives must

[4] The two penalties are not nested. Their ratio is $(k^A(q))^2 V_A$, which depends on δ through equilibrium loadings, so no choice of (γ_3, γ_4) equates them across environments.

therefore arise from costs borne outside the chain population. The externalities identified in Section 4.2 operate through this channel and drive the misalignment results that follow.

4.5 The Misalignment Theorem

The central welfare result is that platform adoption is excessive under either externality. We state both cases together to highlight their common structure and their different invariance properties.

Theorem 1 (Robust misalignment)**.** *Fix* $\delta > -1$ *and* $\sigma_\theta^2, \sigma_A^2, \sigma_C^2, \sigma_\xi^2 > 0$.

(i) ***Under aggregate-volatility welfare*** $\mathscr{W}_3$*:*

$$\Delta\Pi(q) - \frac{d\mathscr{W}_3}{dq} = \gamma_3 \frac{d}{dq}\left[K_\xi(q)^2\,\sigma_\xi^2\right] > 0 \quad \textit{for all } q \in (0,1].$$

The wedge sign is invariant in δ*, while its magnitude varies with* δ *through equilibrium loadings.*

(ii) ***Under decided-about-agent welfare*** $\mathscr{W}_4$*:*

$$\Delta\Pi(q) - \frac{d\mathscr{W}_4}{dq} = \gamma_4 \frac{d}{dq}\left[q^2 \frac{\sigma_\xi^2}{V_A}\right] = \frac{2\,\gamma_4\, q\,\sigma_\xi^2}{V_A} > 0 \quad \textit{for all } q \in (0,1].$$

Both the wedge sign and its magnitude are invariant in δ.

In any interior equilibrium q^*, $q^* > q_3^{soc}$ *and* $q^* > q_4^{soc}$.

The two parts have the same structure: the wedge is the derivative of the externality with respect to adoption. In each case, the relevant object is increasing in q, so marginal adoption raises a cost that chains do not internalize.

The difference lies in what that object depends on. Under $\mathscr{W}_3$, the externality is $K_\xi(q)^2\sigma_\xi^2$, the platform-driven component of aggregate variance. This object combines a primitive term, σ_ξ^2,

with an amplification term, $K_\xi(q)^2 = q^2(k^A(q))^2$, which depends on equilibrium responses and therefore on δ. The wedge sign is invariant, but its magnitude inherits this dependence.

Under $\mathscr{W}_4$, the externality is $q^2\sigma_\xi^2/V_A$, which separates these elements. The encounter probability q^2 scales mechanically with adoption, and the within-adopter correlation σ_ξ^2/V_A depends only on signal primitives. Equilibrium loadings do not enter. The wedge is therefore independent of δ in both sign and magnitude.

This distinction clarifies the source of the result. The over-adoption conclusion does not depend on how chains respond to the shared shock; it follows from the existence of a common component that is orthogonal to fundamentals and from the fact that adoption increases exposure to it. The two welfare criteria capture this through different objects, but both imply the same direction.

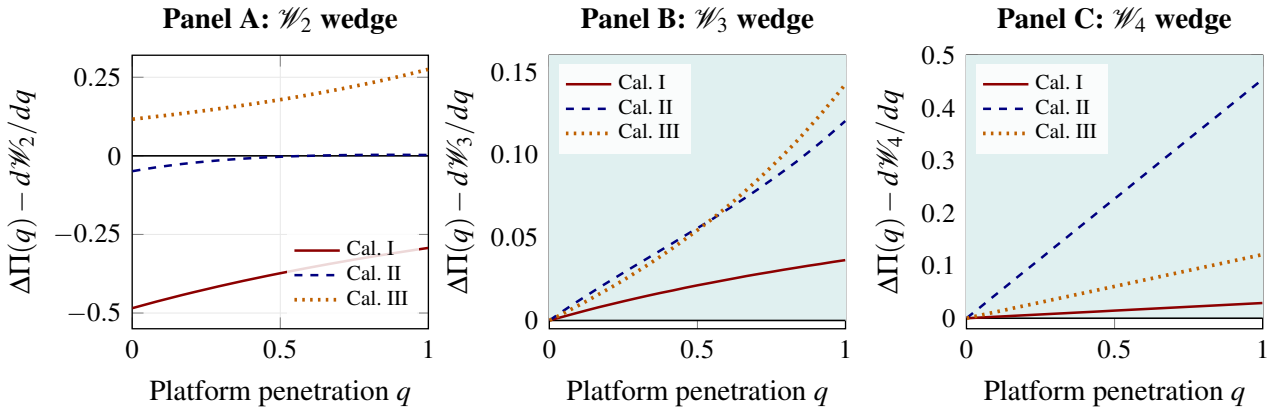
That the result holds under both specifications is substantive. Each corresponds to a different externality acting on a different population. A result that depended on one specification alone would leave open whether it was driven by a particular welfare choice. The coincidence across the two shows that over-adoption is a property of the model's primitives rather than of its welfare construction.

The invariance under $\mathscr{W}_4$ has a further implication. The strategic parameter δ is difficult to identify in retail markets, where agglomeration and business-stealing effects coexist and vary across settings. A welfare result that depends on δ would require resolving this ambiguity. Under $\mathscr{W}_4$, no such identification is needed: the externality depends only on signal structure and adoption, not on equilibrium responses.

Remark 2 (Misalignment under fundamental-tracking welfare $\mathscr{W}_2$ is generically indeterminate)**.** A fourth specification isolates forecast accuracy at the chain level, treating coordination payoffs as transfers and penalizing deviations from the fundamental in the spirit of Angeletos and Pavan (2007). Under this criterion, equilibrium loadings are not generally optimal, and the wedge $\Delta\Pi(q) - d\mathscr{W}_2/dq$ contains both direct and indirect components. Its sign is indeterminate and varies with the calibration. The contrast with Theorem 1 highlights the role of externalities: over-adoption

arises only when the planner internalizes costs borne outside the chain population.

Figure 1: The misalignment wedge under three welfare specifications.



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 $\gamma_k = 1$ 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 fundamental-tracking welfare $\mathscr{W}_2$: 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 aggregate-volatility welfare $\mathscr{W}_3$: it is strictly positive across all three calibrations and all $q \in (0, 1]$, with magnitude depending on δ through equilibrium loadings (Theorem 1, part i). Panel C shows the wedge under decided-about-agent welfare $\mathscr{W}_4$: it is strictly positive, linear in q, and invariant in δ (Theorem 1, part ii). Panels B and C jointly establish that the over-adoption result is robust to the choice of welfare specification across two structurally distinct externalities.

The contrast across the three panels of Figure 1 isolates what the welfare specification contributes. Panel A shows that not every reasonable criterion delivers a determinate sign: $\mathscr{W}_2$ measures an object internal to the chain population (forecast accuracy), and the wedge is calibration-dependent. Panels B and C show that criteria targeting external costs deliver a uniformly positive wedge. The over-adoption result therefore reflects a structural feature of the model rather than a particular welfare choice: whenever the planner internalizes costs imposed on parties outside the chain population, platform adoption is excessive.

4.6 Strategic Interaction versus Shared Error

Strategic interaction and shared processing play distinct roles in the model. The parameter δ governs how firms respond to one another and thereby shapes equilibrium adoption. The variance of the shared component, σ_ξ^2, governs whether correlated behavior imposes costs on parties outside the chain population. These two forces operate through analytically distinct channels.

Remark 3 (Slope–level decomposition)**.** Fix all parameters except (δ, σ_ξ^2). Proposition 2 implies that the sign of $d\Delta\Pi/dq$ is determined by δ alone, so δ determines whether the adoption mapping is increasing, decreasing, or constant in q, and therefore whether multiple equilibria are geometrically possible. The variance σ_ξ^2 enters private incentives through the signal variance V_A and the Bayesian weight ρ_A, but does not affect the slope direction of $\Delta\Pi(q)$. By contrast, Theorem 1 implies that the welfare wedge is strictly positive if and only if $\sigma_\xi^2 > 0$ and $q > 0$, with sign invariant to δ, and under $\mathscr{W}_4$ also magnitude invariant to δ. Hence strategic interaction determines equilibrium selection, while shared processing determines whether each equilibrium is socially distorted.

The decomposition can be expressed directly in terms of two objects. The slope of the adoption mapping, $d\Delta\Pi(q)/dq$, governs equilibrium selection by determining whether adoption is self-reinforcing or self-limiting. The welfare wedge, $\Delta\Pi(q) - d\mathscr{W}_k/dq$, governs distortion by measuring the marginal cost of adoption not internalized by firms. Strategic interaction operates through the slope, while shared processing operates through the wedge.

The distinction can be seen by considering two limiting cases. First, suppose there is no strategic interaction. When $\delta = 0$, firms care only about matching underlying demand and have no incentive to coordinate their actions. In this case, Proposition 2 implies that the private return to adoption is independent of penetration. Nevertheless, if $\sigma_\xi^2 > 0$, Theorem 1 implies that platform adoption is still excessive. The externality arises entirely from the shared component: even without coordination motives, routing decisions through a common system induces correlated errors that affect outsiders.

Second, suppose there is no shared component. When $\sigma_\xi^2 = 0$, the platform reduces idiosyncratic noise without introducing correlation unrelated to fundamentals. In this case, both $\mathscr{W}_3$ and $\mathscr{W}_4$ coincide with $\mathscr{W}_1$, since the penalty terms proportional to σ_ξ^2 vanish. Even if $\delta > 0$ and multiple equilibria are possible, Proposition 3 implies that all equilibria are socially efficient in the sense that they satisfy the planner's first-order condition and impose no costs on external parties. The coordination trap that arises under agglomeration is therefore a feature of the positive equilibrium structure, not a welfare distortion.

Taken together, these observations imply that coordination and shared processing are separable mechanisms. Strategic interaction determines whether firms have an incentive to align their decisions; the shared component determines whether the resulting comovement is costly. The presence of multiple equilibria does not by itself indicate a market failure, and the absence of multiple equilibria does not rule one out. A distortion arises only when the common component is present and marginal adoption increases exposure to it.

This decomposition clarifies the interpretation of Figure 2. The horizontal axis, governed by δ, determines the slope direction of $\Delta\Pi(q)$ and therefore whether multiplicity is possible. The vertical axis, governed by σ_ξ^2, determines both the level of $\Delta\Pi(q)$ at the boundaries—and hence which regime in fact obtains—and whether any given level of adoption generates costs for external parties. The coordination-trap region therefore describes a positive prediction about adoption dynamics; it does not, by itself, identify a welfare problem.

A similar reinterpretation applies to Figure 3. The high-adoption equilibrium arises when $\delta > 0$ and σ_ξ^2 is large enough that $\Delta\Pi(0) < 0$, so that the platform is unattractive in isolation and coordination motives sustain adoption. If $\sigma_\xi^2 = 0$, this outcome is privately optimal and socially efficient, and the planner has no externality-based reason to prefer one equilibrium over another. The trap becomes a policy concern only when $\sigma_\xi^2 > 0$, because high adoption amplifies the exposure of external parties to the shared component.

The decomposition also sharpens the empirical implications of the model. Strategic interaction

alone generates correlation in actions, but this correlation reflects responses to underlying conditions and weakens once fundamentals are controlled for. In contrast, the shared component generates correlation in errors: firms move together in ways that are unrelated to fundamentals. Standard coordination or network models therefore predict comovement that diminishes once fundamentals are accounted for, whereas the shared-processing mechanism predicts residual comovement that persists after conditioning on observable conditions. In particular, synchronized shifts in behavior following updates to a common platform, absent coincident or anticipated changes in fundamentals, are indicative of the shared-error channel rather than strategic interaction alone. Moreover, because the wedge under $\mathscr{W}_4$ is independent of δ in both sign and magnitude, the externality can be quantified without resolving whether the industry exhibits agglomeration or business stealing, a parameter that is typically difficult to identify in retail data.

Finally, the decomposition clarifies the policy margins. Policies that alter the strategic environment, such as agglomeration policies, density restrictions, or zoning rules, operate through δ and may change whether coordination arises. Policies that affect the structure of shared processing, such as vendor diversification or data-source requirements, operate through σ_ξ^2 and reduce the externality directly. Holding q fixed, an increase in σ_ξ^2 strictly increases the welfare wedge under either criterion, so policies that reduce the variance of the shared component operate on a well-defined technological margin. Section 6 formalizes this margin by showing that the wedge scales with an effective concentration index $H(K,\rho)$ that captures the probability that adopters share the same source of error. The two policy levers are therefore complementary: altering δ affects equilibrium selection, while altering σ_ξ^2 affects the welfare consequences of any given equilibrium. To summarize, δ determines how firms respond to one another, while σ_ξ^2 determines whether responding in the same way creates a problem.

5 The Four Regimes of Adoption

Adoption is determined by the fixed-point condition $q^* = F(\Delta\Pi(q^*))$. The equilibrium structure is therefore governed by the shape of $\Delta\Pi(q)$ together with its values at the boundaries. These objects determine the number and stability of fixed points and yield a small set of regimes that organize the model's predictions.

The key boundary conditions are the signs of $\Delta\Pi(0)$ and $\Delta\Pi(1)$. The first admits a particularly simple characterization.

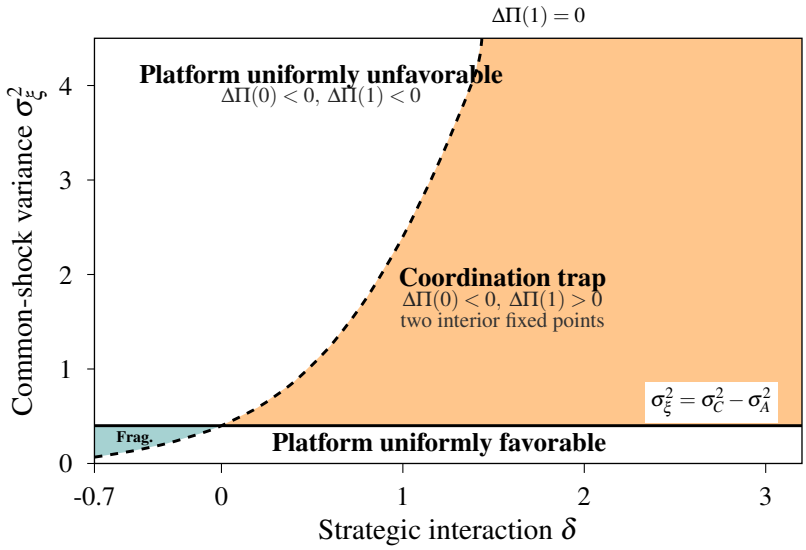
Lemma 1 (δ-independent threshold)**.** *The boundary* $\Delta\Pi(0) = 0$ *is given by* $\sigma_\xi^{2*} = \sigma_C^2 - \sigma_A^2$*, independent of* δ*.*

The intuition is that at $q = 0$ a marginal adopter faces the same aggregate as a non-adopter, $\bar{x} = \mathbb{E}[\bar{x}] + k^C(0)\theta$. When $V_A = V_C$, the two technologies deliver identical signal quality, so expected profits coincide. Because both types respond to the same aggregate, the coordination term cancels, and the adoption decision depends only on relative signal precision. The platform is therefore attractive in isolation if and only if its idiosyncratic advantage, $\sigma_C^2 - \sigma_A^2$, exceeds the variance of the shared shock.

The condition $\Delta\Pi(1) = 0$ defines a second boundary that depends on δ. As coordination motives strengthen, deviation from the platform becomes more costly, so chains tolerate a larger shared shock before switching away. The corresponding threshold in σ_ξ^2 is increasing in δ. Numerical evaluation shows that this locus is monotone and convex, intersecting the $q = 0$ boundary at $\delta = 0$.

Figure 2 plots these two loci in (δ, σ_ξ^2) space and partitions the parameter region into four regimes. The equilibrium configuration varies across regions, but the misalignment result of Theorem 1 applies throughout: while strategic structure determines the shape of adoption, it does not affect the direction of the welfare wedge. Two regions exhibit distinct behavior, such as the

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: $\sigma_\theta^2 = 1$, $\sigma_A^2 = 0.2$, $\sigma_C^2 = 0.6$.

coordination trap under agglomeration and the fragmentation regime under business stealing.

5.1 The Coordination Trap

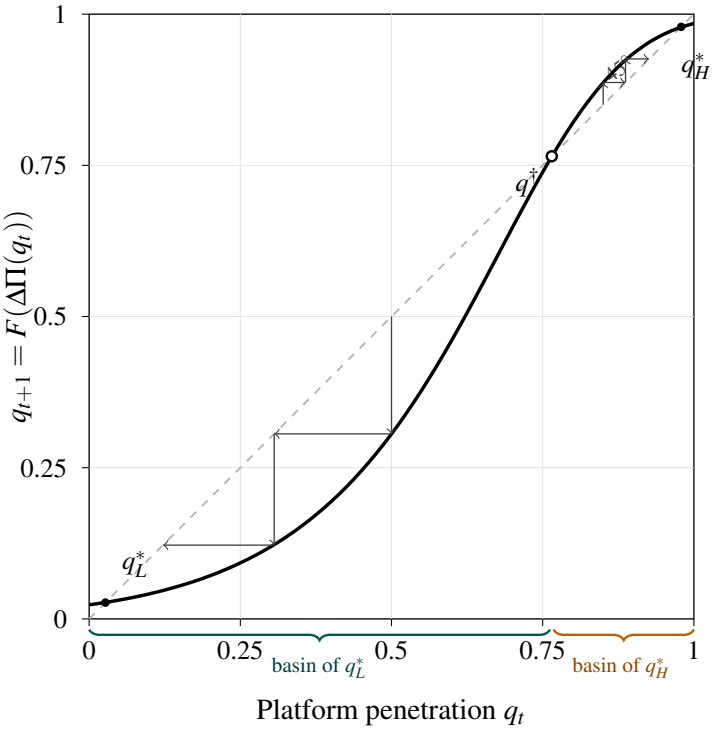
When $\delta > 0$, $\Delta\Pi(0) < 0$, and $\Delta\Pi(1) > 0$, the fixed-point condition $q^* = F(\Delta\Pi(q^*))$ admits multiple solutions. There are two stable fixed points under best-response iteration, a low-adoption equilibrium q_L^* near zero and a high-adoption equilibrium q_H^* near one, separated by an unstable interior threshold at which the slope of $F \circ \Delta\Pi$ exceeds unity.[5]

In this region, platform adoption is privately unattractive in isolation but becomes attractive once adoption is sufficiently widespread. The mechanism is strategic. At low penetration, exposure

[5]Throughout this paragraph and Figure 3, stability refers to best-response iteration on the adoption mapping: a fixed point q^* is stable if $|F'(\Delta\Pi(q^*))\Delta\Pi'(q^*)| < 1$.

to the shared component is limited, and chains prefer the precision of custom research. At high penetration, aggregate entry is driven largely by the shared component, and chains benefit from aligning with rivals who respond to the same signal. Crossing the unstable threshold shifts the dominant force from precision to coordination, making adoption self-reinforcing.

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 $\lambda = 15$.

This regime provides a precise formulation of shared platforms as coordination devices. 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 adopt. The mechanism is strongest in sectors with agglomeration economies. In food halls, outlet centers, and clusters of competing fast-

food restaurants, forecasting rival behavior can be as valuable as forecasting demand, so a platform that reveals common signals becomes essential even if it is not informative about fundamentals.

Under either external-party welfare benchmark where $\mathscr{W}_3$ or $\mathscr{W}_4$, both attracting fixed points exhibit over-adoption at the margin relative to their local optima, as established in Theorem 1. The global comparison between them depends on the welfare object. Under $\mathscr{W}_3$, the high-adoption fixed point improves coordination but raises aggregate volatility, so welfare depends on the balance between these effects; for moderate δ, the low-adoption fixed point is preferred, while for sufficiently large δ, coordination gains can dominate. Under $\mathscr{W}_4$, the comparison is sharper. The externality scales with q^2 independently of δ, while coordination gains depend on equilibrium responses, so the high-adoption fixed point is more likely to be welfare-dominated.

A regulator seeking to reduce synchronization must therefore assess both the marginal wedge and the welfare ranking of fixed points. When the low-adoption fixed point is preferred, policy must shift the economy across the unstable threshold. Measures such as disclosure requirements, diversification mandates, or support for alternative technologies operate by shrinking the basin of attraction of the high-adoption outcome and expanding that of the low-adoption equilibrium.

5.2 Adoption Fragmentation

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

The mechanism is the mirror image of the coordination case. The platform improves precision, making it attractive in isolation, but also increases correlation across chains. When entry decisions are strategic substitutes, this correlation is costly. As adoption rises, the marginal benefit declines, and the equilibrium settles where precision gains are exactly offset by the cost of moving with rivals. Sectors where differentiation is central (e.g., specialty grocers, boutique retail, or niche

restaurant concepts) should therefore exhibit partial adoption rather than convergence to 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 penetration 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 forces reinforce the underlying technological advantage or disadvantage rather than altering its direction.

6 Vendor Diversity and the Magnitude of Misalignment

The misalignment result of Theorem 1 establishes that platform adoption is excessive but leaves open how large the wedge is. This section studies how the magnitude of misalignment depends on the structure of the platform market. We extend the model to allow multiple competing platforms and identify a second policy margin: the diversity of processing technologies. This margin affects the size of the externality without changing its direction.

The extension speaks to two related concerns. The first is vendor concentration in retail intelligence markets. The second is the distinction between substantive and cosmetic diversity: multiple platforms may differ in interface or branding while relying on the same underlying data or models. The model makes this distinction precise.

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 idiosyncratic noise $u_f^{A,k}$ 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 using platform k, with $\sum_{k=1}^K q_k = q$. The shocks $\{\xi_k\}$ are jointly Gaussian with $\text{Var}(\xi_k) = \sigma_\xi^2$ and pairwise correlation $\text{Corr}(\xi_k, \xi_{k'}) = \rho \in [0,1]$ for $k \neq k'$. The single-platform model is the case $K = 1$.

The parameter ρ captures the extent of shared processing across platforms. When $\rho = 0$, platforms are structurally independent: their shocks arise from distinct models, datasets, or data vendors. When $\rho = 1$, platforms share a common core, so that $\xi_k = \xi_{k'}$ almost surely. Intermediate values correspond to partial overlap in underlying inputs or methods.

For tractability, vendor shares $\{q_k\}$ are treated as exogenous. Section 5 determines total adoption q endogenously; here we hold q fixed and study how its allocation across platforms affects welfare. This isolates the role of vendor structure from a full model of platform competition.

6.2 The Effective Concentration Index

Each adopter continues to observe a single processing shock, so the equilibrium loadings (k^A, k^C) are unchanged from Proposition 1. The difference arises in aggregation. With multiple platforms, the aggregate inherits a separate component from each:

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

Taking variances and using the correlation structure,

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

Under symmetric adoption, $q_k = q/K$, this expression collapses to a single index. Define

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

Then

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

where $K_\xi(q) = q\,k^A$. The index $H(K,\rho)$ lies in $[1/K,1]$, attaining $1/K$ when platforms are independent ($\rho = 0$) and 1 when they share a common core ($\rho = 1$). It is increasing in ρ for any $K \geq 2$.

Proposition 4 (Vendor diversity)**.** *Under either welfare specification, the wedge between private and social adoption incentives scales with the effective concentration index. Specifically:*

(i) Under $\mathscr{W}_3$:

$$\Delta\Pi(q) - \frac{d\mathscr{W}_3}{dq} = \gamma_3\,H(K,\rho)\,\frac{d}{dq}\big[K_\xi(q)^2\,\sigma_\xi^2\big] > 0.$$

(ii) Under $\mathscr{W}_4$:

$$\Delta\Pi(q) - \frac{d\mathscr{W}_4}{dq} = \gamma_4\,H(K,\rho)\,\frac{2q\,\sigma_\xi^2}{V_A} > 0.$$

For all $\delta > -1$, $K \geq 1$, and $\rho \in [0,1]$, the wedge remains positive. Its magnitude is proportional to $H(K,\rho)$ and therefore decreases in K whenever $\rho < 1$.

The result follows from two parallel decompositions. Under $\mathscr{W}_3$, the aggregate-variance term factors as $K_\xi(q)^2\sigma_\xi^2$ times $H(K,\rho)$. Under $\mathscr{W}_4$, the same index governs the expected correlation across adopters: with probability $1/K$ two adopters share a platform and correlation is σ_ξ^2/V_A,

while with probability $1 - 1/K$ they use different platforms and correlation is $\rho\,\sigma_\xi^2/V_A$. In both cases, $H(K,\rho)$ captures the probability-weighted exposure to a common processing shock.

The implication is that vendor structure enters only through this index. The two externalities differ in the objects they penalize, but both depend on the same sufficient statistic: the likelihood that adopters share the same source of error. Vendor diversity reduces the externality only insofar as it reduces this probability.

6.3 Substantive versus Cosmetic Diversity

Two corner cases of Proposition 4 isolate the policy content.

Corollary 1 (Independent platforms)**.** *When $\rho = 0$, $H(K,0) = 1/K$. The misalignment magnitude under either welfare specification 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.*

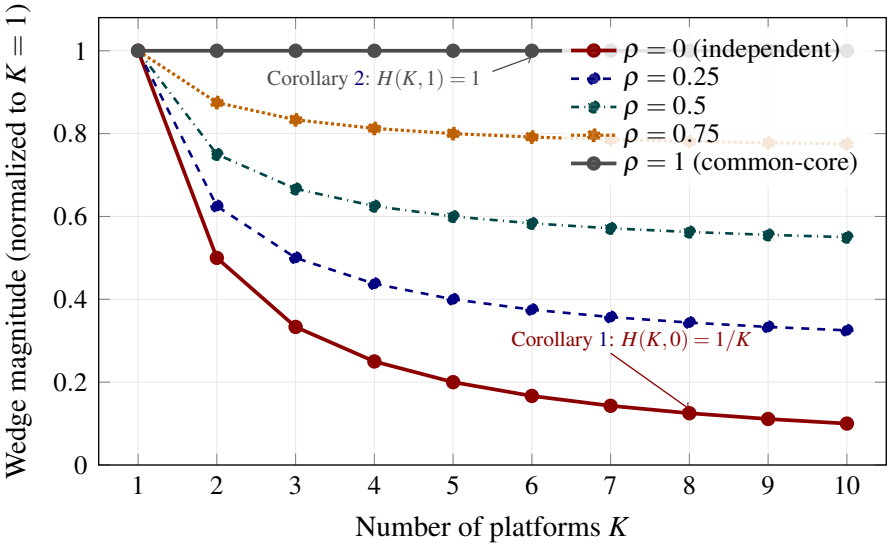
The contrast is stark. When platforms are independent, increasing K reduces the externality proportionally: each additional platform lowers the probability that adopters share the same error. When platforms share a common core, increasing K has no effect: adopters remain exposed to the same underlying shock regardless of how many intermediaries exist. Vendor diversity therefore matters only through the independence of processing technologies, not through the number of providers.

This distinction is directly observable in retail intelligence markets. Many platforms rely on the same upstream data (e.g., mobile-device location pings or standardized demographic projections) so nominal variety masks high ρ. By contrast, platforms built on distinct data sources, such as transaction records, mobility traces, or satellite imagery, approximate the low-ρ regime. Bommasani et al. (2022) document an analogous pattern for foundation models: fine-tuned variants of a common base model exhibit limited independent variation relative to structurally distinct models.

The relevant margin is therefore not the number of platforms but the diversity of underlying inputs and methods.

Figure 4 illustrates this result. The wedge declines sharply with K when $\rho = 0$ and remains flat when $\rho = 1$, with intermediate values interpolating between these cases. Substantive diversity reduces the externality; cosmetic diversity does not.

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



Notes: The figure plots the misalignment-wedge magnitude as a function of the number of platforms K, normalized to the value at $K = 1$. By Proposition 4, this normalized magnitude equals the effective concentration index $H(K,\rho) = 1/K + (1 - 1/K)\rho$ under either welfare specification ($\mathscr{W}_3$ or $\mathscr{W}_4$); the figure therefore applies to both. Each curve corresponds to a fixed value of the cross-platform correlation $\rho \in \{0, 0.25, 0.5, 0.75, 1\}$. 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 either externality.

6.4 The Regulator's Joint Problem

Combining Theorem 1 and Proposition 4, a regulator who internalizes either externality faces a problem in two margins: the adoption rate q and the effective concentration $H(K,\rho)$. The marginal social cost of adoption is the product of these two objects and the underlying penalty

term ($K_\xi(q)^2\sigma_\xi^2$ under $\mathscr{W}_3$, $q^2\sigma_\xi^2/V_A$ under $\mathscr{W}_4$). Policies that reduce q (i.e., adoption taxes, diversification mandates, or capital requirements tied to platform exposure) operate on the first margin. Policies that reduce H, such as procurement diversification, data-source requirements, or interoperability standards that preserve heterogeneity in underlying models, operate on the second.

The two margins are distinct. Reducing q lowers exposure but leaves the correlation structure unchanged. Reducing H lowers correlation but leaves adoption above the social optimum. The joint policy reaches the optimum with lower distortion than either margin alone. Policies targeting platform structure therefore complement, rather than substitute for, policies targeting adoption.

The same logic extends to asymmetric vendor shares. When adoption is uneven, the effective concentration index becomes

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

a normalized Herfindahl measure. The qualitative result is unchanged: the wedge increases with concentration for any fixed ρ, and the distinction between substantive and cosmetic diversity continues to hold.

7 Discussion

This paper develops a tractable model of shared platform adoption in retail entry. The central feature is a common component in forecasts that is orthogonal to fundamentals, generating correlated decisions without improving inference. The model yields a closed-form equilibrium with heterogeneous adoption, a robust misalignment result under two distinct externalities, and a four-regime typology that includes a coordination trap. The over-adoption result holds for all δ, and under the decided-about-agent criterion its magnitude is independent of δ.

The framework provides a formal account of shared intelligence as an attentional infrastructure. Chains differ not only in the data they observe but in how they process it. Shared platforms

standardize this mapping, while custom research preserves heterogeneity. The shared shock ξ captures the residual component of this standardization. Because it is common across adopters and unrelated to fundamentals, it generates two external costs that chains do not internalize. Aggregate-volatility costs fall on agents exposed to synchronized entry. Correlated-evaluation costs fall on locations that are jointly screened by adopters. Coordination incentives are internalized in equilibrium, so these externalities are the sole source of inefficiency.

References

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

Babina T, Fedyk A, He AX, Hodson J (2024) Artificial intelligence and firms' systematic risk. *SSRN Working Paper No. 4868770* .

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.

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.

Chi F (2025) Information waves and firm investment. *Available at SSRN 4269750* .

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.

Kleinberg J, Raghavan M (2021) Algorithmic monoculture and social welfare. *Proceedings of the National Academy of Sciences* 118(22):e2018340118.

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.

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

Peng K, Garg N (2024) Monoculture in matching markets. *Advances in Neural Information Processing Systems* 37:81959–81991.

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

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 Theorem 1

We prove parts (i) and (ii) in turn. Both rely on the team-efficiency identity $d\mathscr{W}_1/dq = \Delta\Pi(q)$ from Proposition 3.

Part (i): aggregate-volatility welfare $\mathscr{W}_3$. The welfare specification $\mathscr{W}_3$ decomposes as $\mathscr{W}_3(q) = \mathscr{W}_1(q) - \gamma_3 K_\xi(q)^2\,\sigma_\xi^2$. Differentiating termwise and substituting the team-efficiency identity,

$$\Delta\Pi(q) - \frac{d\mathscr{W}_3}{dq} = \gamma_3\,\frac{d}{dq}\left[K_\xi(q)^2\,\sigma_\xi^2\right]. \tag{P5.1}$$

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 $\Delta(q)^3 > 0$,

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

with strict inequality whenever $q > 0$. Combining (P5.1) with (P5.3) and the positivity of γ_3 and σ_ξ^2 yields strict positivity of the wedge for all $q \in (0,1]$.

Part (ii): decided-about-agent welfare $\mathscr{W}_4$. The welfare specification $\mathscr{W}_4$ decomposes as $\mathscr{W}_4(q) = \mathscr{W}_1(q) - \gamma_4\,q^2\,\sigma_\xi^2/V_A$. Note that σ_ξ^2/V_A depends only on signal-structure primitives and not on q, so it factors out of the differentiation. Differentiating termwise and substituting the team-efficiency identity,

$$\Delta\Pi(q) - \frac{d\mathscr{W}_4}{dq} = \gamma_4\,\frac{\sigma_\xi^2}{V_A}\,\frac{d}{dq}\big[q^2\big] = \frac{2\,\gamma_4\,q\,\sigma_\xi^2}{V_A}. \tag{P5.4}$$

The right-hand side is strictly positive for all $q \in (0,1]$ and is independent of δ in both sign and magnitude. The independence is structural: equilibrium loadings do not enter the welfare object $q^2\sigma_\xi^2/V_A$, only the encounter probability q^2 and the primitive within-adopter conditional correlation σ_ξ^2/V_A do. The wedge therefore inherits no dependence on δ from the welfare-side derivative.

Over-adoption. For each part, the market-clearing condition for an interior adoption equilibrium q^* is $F^{-1}(q^*) = \Delta\Pi(q^*)$, while the planner's first-order condition is $F^{-1}(q_k^{soc}) = d\mathscr{W}_k/dq|_{q_k^{soc}}$ for $k \in \{3,4\}$. Since F^{-1} is strictly increasing, $\Delta\Pi(q^*) > d\mathscr{W}_k/dq|_{q^*}$ implies $q^* > q_k^{soc}$. □

A.6 Proof of Proposition 4

The proposition has two parts; both rely on factoring the multi-platform aggregate variance into a single-platform component times the effective concentration index $H(K,\rho)$.

Variance decomposition under multi-platform. 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,

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

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

Encounter-correlation decomposition under multi-platform. For a randomly selected pair of distinct platform-adopting chains f,g with f adopting platform k and g adopting platform k', the conditional correlation of their signal errors is

$$\mathrm{Corr}(\tilde{s}_f^{A,k} - \theta,\ \tilde{s}_g^{A,k'} - \theta) = \begin{cases} \sigma_\xi^2/V_A & \text{if } k = k', \\ \rho\cdot\sigma_\xi^2/V_A & \text{if } k \neq k'. \end{cases}$$

Under symmetric vendor adoption with $q_k = q/K$, the probability that two randomly drawn adopters use the same vendor is $\sum_k (q_k/q)^2 = 1/K$, and the probability they use different vendors is $1-1/K$. The expected within-encounter conditional correlation is therefore

$$\frac{1}{K}\cdot\frac{\sigma_\xi^2}{V_A} + \left(1-\frac{1}{K}\right)\cdot\rho\cdot\frac{\sigma_\xi^2}{V_A} = H(K,\rho)\cdot\frac{\sigma_\xi^2}{V_A}.$$

Multi-platform wedges. The multi-platform welfare specifications are

$$\mathscr{W}_3^{(K,\rho)}(q) = \mathscr{W}_1(q) - \gamma_3\,H(K,\rho)\,K_\xi(q)^2\,\sigma_\xi^2, \qquad \mathscr{W}_4^{(K,\rho)}(q) = \mathscr{W}_1(q) - \gamma_4\,H(K,\rho)\,q^2\,\frac{\sigma_\xi^2}{V_A}.$$

By Proposition 3, $d\mathscr{W}_1/dq = \Delta\Pi(q)$, so the wedges are

$$\Delta\Pi(q) - \frac{d\mathscr{W}_3^{(K,\rho)}}{dq} = \gamma_3 H(K,\rho)\frac{d}{dq}\left[K_\xi(q)^2\sigma_\xi^2\right]$$

and

$$\Delta\Pi(q) - \frac{d\mathscr{W}_4^{(K,\rho)}}{dq} = \gamma_4 H(K,\rho)\cdot\frac{2q\sigma_\xi^2}{V_A}.$$

Both are strictly positive by (P5.3), (P5.4), and the positivity of $H, \gamma_3, \gamma_4, \sigma_\xi^2$. 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 under either welfare specification becomes proportional to $1/K$. At $\rho = 1$, $H(K,1) = 1$ for all K, so the wedge equals the single-platform wedge of Theorem 1 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.