Shared Intelligence and Misaligned Firm Entry

Feng Chi* Nathan Yang†

May 5, 2026

Abstract

Firm entry decisions are increasingly mediated by shared intelligence platforms whose forecasts embed a common component unrelated to local fundamentals. We develop a Bayesian entry game in which firms choose between proprietary research and a shared platform that reduces idiosyncratic noise but introduces a common shock. The model yields a closed-form equilibrium with heterogeneous adoption. Adoption is excessive under two externalities (i.e., aggregate volatility and correlated evaluation across locations) both of which persist even when firms optimally respond to one another. Under correlated evaluation, the wedge depends only on adoption and signal structure. A multi-vendor extension shows that only substantive, not cosmetic, diversity mitigates these externalities.

Keywords: Firm 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

Firm 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. As a result, firms may not only act similarly, but may be wrong in similar ways. These correlated errors need not be isolated: when the underlying system evolves slowly, they can persist across periods, generating clusters of similar decisions over time.

Recent empirical work provides suggestive evidence that shared information environments 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

[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.

for other technologies (Babina et al. 2024). At the same time, Chi (2025) shows that firms relying on common data sources can make systematically misinformed decisions when those inputs become outdated, leading to correlated mistakes across cohorts of entrants. Because such data are often accessed through common intermediaries, declines in information quality propagate across decisions rather than remaining idiosyncratic. 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.

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, 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. Correlation in this case 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.

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, allowing a decomposition of aggregate entry into fundamental and platform-driven components and a characterization of how adoption incentives

vary with penetration.

The main normative result is a robust misalignment. The shared component generates two distinct external costs that firms do not internalize. 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, evaluations of any given location become correlated, reducing the effective independence of the screening process. In both cases, equilibrium platform adoption is excessive. This conclusion does not depend on the nature of competition; it follows from the presence of a common error component that firms do not internalize.

The framework delivers three additional insights. First, the strength of the externality depends on the structure of the platform market: vendor diversity reduces correlated errors only when platforms are substantively independent. Second, shared processing alters the temporal structure of outcomes, generating persistence in aggregate behavior even when fundamentals are uncorrelated over time. Third, switching costs introduce a dynamic policy margin: when switching becomes easier, the economy converges more rapidly toward high-adoption regimes, accelerating the accumulation of external costs.

Finally, the model identifies a coordination trap: a high-adoption equilibrium can be sustained by coordination motives even when the platform is privately unattractive in isolation. This regime is most relevant in sectors with strong agglomeration economies, where anticipating competitors' behavior is central to entry decisions. Because the externality is invariant to the strategic environment, both equilibria exhibit over-adoption at the margin, with the welfare ranking depending on the planner's objective.

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 increases with rival entry, capturing agglomeration economies: clusters such as food halls, outlet centers, or fast-food strips attract demand that no single store could generate in isolation. When $\delta < 0$, profitability decreases 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 serves as a standard global-games device (Carlsson and van Damme 1993, Morris and Shin 1998), ensuring that entry probabilities are linear in posterior beliefs. We focus on the regime in which the linear approximation is informative.[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 does not affect any qualitative results.

[2]Formally, linearity holds whenever $|\mathbb{E}_f[\theta] + \delta \mathbb{E}_f[\bar{x}]| \leq \bar{\varepsilon}$ pointwise; a sufficient condition is $\bar{\varepsilon} > 4\sqrt{\text{Var}(\theta) + \text{Var}(\bar{x})}$. Some calibrations used in the figures below pick variance values that lie outside this strict bound to display the qualitative regimes more clearly. The closed-form expressions of Proposition 1 are well-defined analytical functions of the primitives wherever $\Delta(q) > 0$, and the qualitative comparative statics established in Theorem 1 and Propositions 2, 5, 7 extend to a truncated-equilibrium model in which entry probabilities are clipped to $[0,1]$, with the wedge sign and monotonicity-in-adoption properties preserved over the calibrated parameter region.

2.2 Information and Processing Technologies

The two information technologies in the model map directly to a distinction that has become central in the empirical literature on firm AI: whether firms build forecasting capability internally or source it from external vendors. Babina (2026) identifies this make-versus-buy margin as a key organizing dimension of firm AI activity, emphasizing that vendor-provided solutions tend to reduce differentiation across adopters because competing firms rely on similar tools. Empirical evidence is consistent with this characterization. Firms that invest more heavily in AI exhibit higher systematic risk (Babina et al. 2024), and firms that rely on external AI providers behave differently from those that build internally (Calvino and Fontanelli 2026). What this literature does not address is the welfare implication of shared vendor solutions when they are adopted at scale. The signal structure below provides a formal counterpart to this empirical margin. Custom research generates idiosyncratic forecast errors that differ across firms, while platform-based forecasting replaces a portion of this idiosyncratic variation with a component shared across adopters. This shared component captures, in reduced form, the loss of differentiation emphasized in the empirical literature and is the source of the correlated behavior studied in the model.

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 unrelated 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, 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 that is unrelated 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

[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.

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*

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

intercept $K_0 = 1/(2(1-\delta))$, and equilibrium denominator

$$\Delta(q) \equiv (1 - \delta\, q\, \rho_A)\,(1 - \delta(1-q)\,\alpha_\theta^C) - \alpha_\theta^A\,\alpha_\theta^C\,\delta^2\, q(1-q), \tag{2}$$

where $\rho_A = \alpha_\theta^A + \alpha_\xi^A = (\sigma_\theta^2 + \sigma_\xi^2)/V_A$. The equilibrium exists and is unique whenever $\Delta(q) > 0$, which holds for all $q \in [0,1]$ when $\delta \in (-1, 1 + \sigma_C^2/\sigma_\theta^2)$.

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, $k^C(0) = \alpha_\theta^C/(1-\delta\alpha_\theta^C)$. When $q = 1$, all chains adopt the platform and the loading becomes $k^A(1) = \alpha_\theta^A/(1-\delta\rho_A)$, 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 unrelated 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

$$\frac{d\Delta\Pi}{dq} = \delta \cdot \frac{2\,\sigma_\theta^4\,\sigma_\xi^2}{V_A^3 V_C^2} \cdot R(q;\sigma,\delta),$$

where R is a polynomial that is strictly positive over the admissible parameter region $\delta \in (-1,1)$. 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 evaluations of any given location are positively correlated. As platform adoption rises, the share of pairwise evaluations of a location that share this correlated component grows. The cost borne by this group is the loss of independent evaluation that would otherwise distribute entry attempts more evenly across locations.

This mechanism is in the spirit of the externality emphasized by Kleinberg and Raghavan (2021) and Peng and Garg (2024). In those settings, the decided-about agents are individuals, and the relevant statistic is the probability that a favorable case is recognized by at least one decision-maker. Here the decided-about agent is a location, and the formal welfare object we adopt is the expected pairwise correlation of evaluation errors across encounters at a location, which scales as $q^2\sigma_\xi^2/V_A$. The two formulations capture related but distinct welfare consequences of correlated decision technologies; we focus on pairwise correlation because it admits a direct closed form in the model and depends only on signal-structure primitives. A missing-victim formulation in the strict Kleinberg-Raghavan sense would require specifying the number of evaluators per location and would yield a functional form proportional to $1-(1-q)^k$, which scales differently in adoption.

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; it is a positive scalar with units of inverse variance, to be calibrated to the empirical context.

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 parameter $\gamma_4 > 0$ is a positive scalar measuring the marginal social cost of correlated evaluation.

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.

[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.

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 given the information structure. A planner who assigns common decision rules to each type faces the same problem as individual chains, up to a general-equilibrium term capturing how coordinated changes affect the aggregate.

This term vanishes because only deviations from the aggregate matter at the margin. The marginal payoff with respect to $\bar{x}$ for chain g is δP_g, whose average equals $\delta\bar{x}$ by the consistency condition. The planner's first-order condition therefore isolates the deviation $\delta(P_g - \bar{x})$, which integrates to zero across chains. Since the induced change in $\bar{x}$ is common across agents, the general-equilibrium feedback cancels, and the planner's 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. Although the platform induces correlated behavior, chains fully internalize its effects on profits through equilibrium responses. Any wedge between private and social incentives must 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*** $\mathscr{W}_4$*:*

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

The wedge function depends only on adoption q and signal-structure primitives, not on equilibrium loadings; in particular, evaluated at any fixed q, it is 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.[5]

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 function itself, so the wedge evaluated at any fixed q is independent of δ. The wedge *at the equilibrium* $q^*(\delta)$ depends on δ only through the equilibrium adoption rate, not through any feedback of strategic interaction on the marginal externality.

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 unrelated 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

[5]The wedge admits a direct Pigouvian interpretation. A tax $\tau_k(q) = d\mathscr{W}_k/dq - \Delta\Pi(q)$ on platform adoption reduces the private return by the marginal externality, aligning the decentralized condition with the planner's first-order condition $F^{-1}(q) = d\mathscr{W}_k/dq$. Although q is endogenous, the wedge therefore corresponds to a well-defined policy instrument: any intervention that lowers the marginal private return to adoption by $\tau_k(q^*)$, such as adoption taxes, capital requirements indexed to platform exposure, or licensing fees, implements the constrained optimum. In this sense, comparative statics in q can be interpreted as comparative statics in the strength of the corrective instrument.

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.

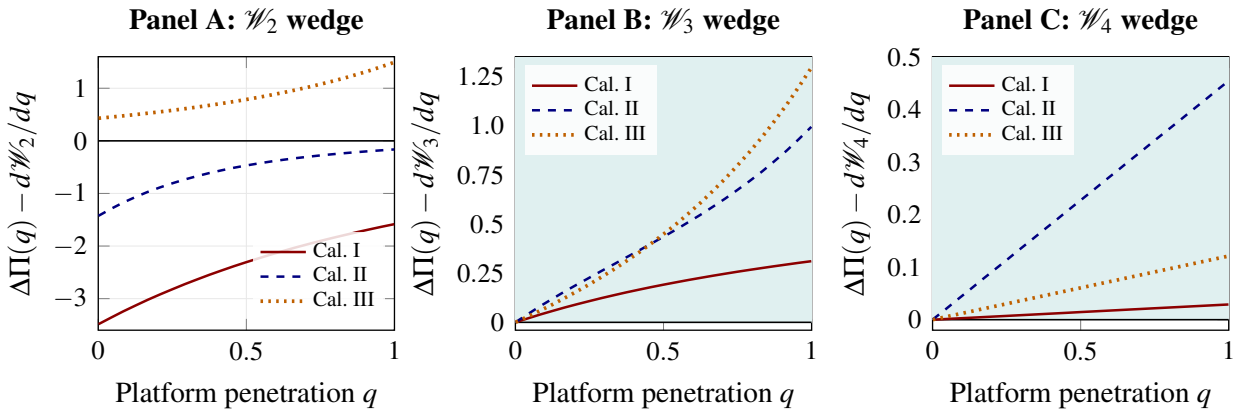
Remark 2 (Misalignment under fundamental-tracking welfare $\mathscr{W}_2$ is generically indeterminate)**.** A fourth specification isolates forecast accuracy at the chain level, treating coordination payoffs as transfers and penalizing deviations from the fundamental in the spirit of Angeletos and Pavan (2007). Specifically,

$$\mathscr{W}_2(q) = -q\,(k^A(q)-1)^2 - (1-q)\,(k^C(q)-1)^2,$$

which penalizes deviations of each type's signal loading from one (perfect tracking). The index follows the canonical numbering of welfare specifications in the social-value-of-information literature: $\mathscr{W}_1$ is the utilitarian benchmark and $\mathscr{W}_2$ is the fundamental-tracking benchmark of Angeletos and Pavan (2007), with the externality-augmented specifications $\mathscr{W}_3$ and $\mathscr{W}_4$ added here. Unlike $\mathscr{W}_3$ and $\mathscr{W}_4$, $\mathscr{W}_2$ is not an additive perturbation of $\mathscr{W}_1$, so equilibrium loadings (k^A, k^C) are not generally optimal under $\mathscr{W}_2$ and the team-efficiency identity $d\mathscr{W}_1/dq = \Delta\Pi(q)$ does not transfer. The wedge $\Delta\Pi(q) - d\mathscr{W}_2/dq$ therefore picks up indirect terms in $(k^A)'(q)$ and $(k^C)'(q)$, and its sign is calibration-dependent. The contrast with Theorem 1 highlights the role of externalities: over-adoption arises only when the planner internalizes costs borne outside the chain population.

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

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 $\delta = 0.7$ 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 two calibrations and positive values throughout for the third, with magnitude varying markedly across calibrations. 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.

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 $\mathscr{W}_4$ the wedge function evaluated at any fixed q depends only on adoption and signal-structure primitives, and not on equilibrium loadings. Hence strategic interaction determines equilibrium selection, while shared processing determines whether each equilibrium is socially distorted.

The decomposition can be summarized by 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 becomes transparent in limiting cases. When $\delta = 0$, firms have no incentive to coordinate and care only about matching underlying demand. Proposition 2 then implies that the private return to adoption is independent of penetration. Nevertheless, if $\sigma_\xi^2 > 0$, Theorem 1 implies that platform adoption remains 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.

When $\delta < 0$, the same point holds under strategic substitution. Proposition 2 implies that $\Delta\Pi(q)$ is strictly decreasing in q, so the adoption mapping is self-limiting and the fixed-point equation admits a unique interior solution (Section 5). Firms have a private incentive to differentiate from rivals, which limits the equilibrium level of adoption. This force works in the same direction as the planner's objective, but it is incomplete. Theorem 1 continues to apply: the welfare wedge is strictly positive at any interior $q^* > 0$. The equilibrium therefore settles at a lower level of adoption than under strategic complements, yet still above the social optimum. Strategic substitution mitigates but does not eliminate the distortion.

Conversely, 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 vanish. Even if $\delta > 0$ and multiple equilibria arise, 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 emerges under agglomeration is therefore a feature of the positive equilibrium structure rather than a welfare distortion.

Taken together, these cases show 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 such alignment 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 Figures 2 and 3. The horizontal axis in Figure 2, 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 whether any given level of adoption generates costs for external parties. The coordination-trap region therefore describes a positive prediction about adoption dynamics rather than a welfare problem.

Similarly, the high-adoption equilibrium in Figure 3 arises when $\delta > 0$ and σ_ξ^2 is large enough that $\Delta\Pi(0) < 0$, so that coordination motives sustain adoption even though the platform is unattractive in isolation. If $\sigma_\xi^2 = 0$, this outcome is privately optimal and socially efficient. It becomes a policy concern only when $\sigma_\xi^2 > 0$, because high adoption amplifies exposure to the shared component.

The decomposition also sharpens the empirical implications of the model. Strategic interaction generates correlation in actions that 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 after conditioning on observable conditions, whereas the shared-processing mechanism predicts residual comovement that persists. 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. Moreover, because the wedge under $\mathscr{W}_4$ is independent of δ, the externality can be quantified without identifying the strategic environment, a parameter that is typically difficult to measure in retail data.

Finally, the decomposition clarifies the policy margins. Policies that alter the strategic environment (e.g., agglomeration policies, density restrictions, or zoning rules) operate through δ and affect equilibrium selection. Policies that affect the structure of shared processing (e.g., 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.

4.7 An Identifying Test

The decomposition in Section 4.6 implies that strategic interaction and shared processing leave distinct empirical fingerprints. This subsection translates that distinction into a falsifiable prediction. The mechanism studied here generates a pattern that canonical coordination or network-effects models cannot produce: a discrete movement in aggregate entry that coincides with a change in the shared processing technology, occurs without a corresponding movement in fundamentals, and scales with platform penetration.

The natural setting for such a movement is a platform-side update event. Revisions to a foundation model, changes in vendor methodology, or updates to the data inputs that feed platform forecasts alter the processing layer without changing local market conditions. In the model, such events correspond to shifts in the realization of the shared component, which propagate through adopters' signals and into their entry decisions. The resulting movement in aggregate entry is therefore pinned down by the equilibrium loadings: it reflects the extent to which adoption exposes firms to the shared component, rather than any change in underlying fundamentals.

Proposition 4 (Model-update signature)**.** *Fix $q \in [0,1]$. Suppose an event E at time τ shifts the realization of the shared component by $\Delta\xi$, holding the realization of the fundamental and all model primitives fixed. The induced movement in aggregate entry is*

$$\Delta\bar{x}_\tau \;=\; K_\xi(q)\cdot\Delta\xi, \qquad K_\xi(q) = q\,k^A(q),$$

with $k^A(q)$ the platform loading from Proposition 1. The magnitude $|\Delta\bar{x}_\tau|$ is strictly increasing in q on $(0,1]$ and equals zero at $q=0$.

In any model without a shared error component, aggregate entry depends only on fundamentals and equilibrium responses to them. Such models therefore predict no discrete movement in aggregate entry at the time of E unless fundamentals themselves change, and no systematic relationship between such movements and platform penetration. The shared-processing mechanism predicts both: a residual shift in aggregate entry after controlling for fundamentals, with magnitude increasing in measured adoption.

Implementing the test requires three ingredients. First, the event must be identified from platform-side records (e.g., release notes, methodology changes, or vendor disclosures) rather than inferred from entry data, to avoid selection bias. Second, fundamentals must be controlled for using observable correlates, so that the residual movement can be attributed to the shared component rather than to a common demand shock. Third, cross-sectional variation in platform penetration provides the scaling dimension: the magnitude of the residual shift should increase with adoption intensity.

Two extensions sharpen the empirical design. A frequency-domain approach exploits the regular cadence of platform updates, predicting concentration of aggregate entry at the update frequency, with peak power increasing in the variance contribution of the shared component. A cross-platform approach uses asynchronous updates across providers: when one platform updates while another does not, adopters of the updating platform should respond more strongly, with differences scaling in vendor shares.

Proposition 4 provides the model's sharpest falsification claim. The over-adoption result of Theorem 1 is empirically non-specific, as many mechanisms can generate excessive adoption. The model-update signature is not. It requires a common component orthogonal to fundamentals, a feature absent from standard models of coordination or herding.

Two qualifications are warranted. Proposition 4 is a *theoretical* identifying restriction. Its empirical implementation requires (i) exogenous platform events identified from platform-side records rather than inferred from entry data, (ii) controls for fundamentals that remove the omitted-variable threat that platforms update in anticipation of unobserved demand shifts, and (iii) a defensible cross-sectional measure of adoption at the relevant unit (firm, sector, or geography). Each of these is a standard event-study identification concern, and the test is only as credible as the IV-validity of platform updates and the completeness of fundamental controls. Recent regulatory

discourse has begun to anchor the platform-event-as-instrument interpretation in financial-stability policy (Bank of England Financial Policy Committee 2025), but the analogous infrastructure for retail intelligence platforms remains underdeveloped. The contribution of the model is to provide a structural restriction that maps cleanly onto observable variation in adoption, not a turnkey procedure: the empirical work required to take the restriction to data is non-trivial and is left to future investigation.

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 boundary values. These jointly determine the number and stability of fixed points, yielding 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.

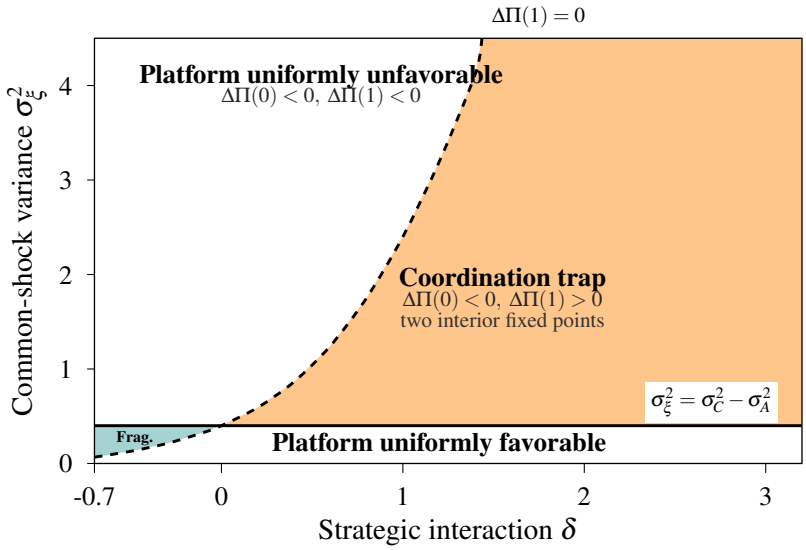
Lemma 1 (δ-independent threshold)**.** *The boundary* $\Delta\Pi(0) = 0$ *is given by* $\sigma_\xi^{2*} = \sigma_C^2 - \sigma_A^2$*, independent of* δ*.*

The intuition is straightforward. At $q = 0$, a marginal adopter faces the same aggregate as a non-adopter, $\bar{x} = \mathbb{E}[\bar{x}] + k^C(0)\theta$, so the coordination channel does not differentiate the two technologies. The adoption decision therefore depends only on signal quality. When $V_A = V_C$, the platform and custom research deliver identical precision, and expected profits coincide. The platform is thus attractive in isolation if and only if its reduction in idiosyncratic noise, $\sigma_C^2 - \sigma_A^2$, exceeds the variance of the shared component.

The second boundary, defined by $\Delta\Pi(1) = 0$, depends on δ. As coordination motives strengthen, deviation from the platform becomes more costly, so chains are willing to tolerate a larger shared component before switching away. The corresponding threshold in σ_ξ^2 is therefore 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. While equilibrium configurations differ across regions, the misalignment result of Theo-

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

rem 1 applies throughout: strategic interaction shapes how adoption is sustained, but not whether it is excessive. The regimes therefore differ in the level and structure of adoption rather than in the direction of the welfare distortion. Two cases are particularly informative: the coordination trap under agglomeration and the fragmentation regime under business stealing.

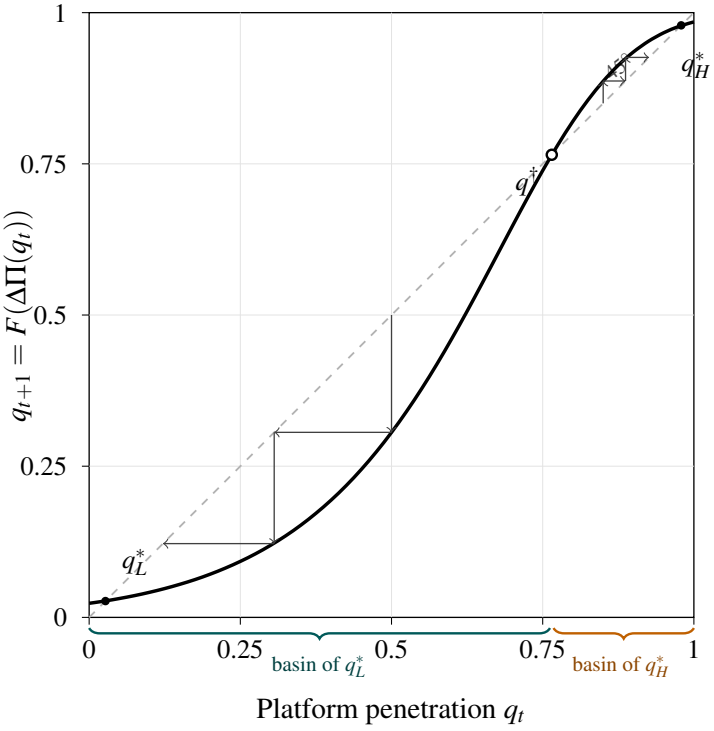
5.1 The Coordination Trap

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.[6]

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

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 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.

The welfare implication of this regime mirrors that of the coordination trap, despite the very different equilibrium structure. By Theorem 1, the unique fragmentation equilibrium exhibits over-

adoption under both $\mathscr{W}_3$ and $\mathscr{W}_4$. Strategic substitution does not insulate the industry from the externality; it only reduces the level of adoption at which it is realized.

The two regimes of Section 5 therefore differ in the equilibrium level of q^*, not in the direction of the welfare distortion. The result of Section 4 holds uniformly for all $\delta > -1$. Business stealing is thus not an exception, but one instance of a broader pattern: strategic forces affect how much adoption occurs, but not whether that adoption is excessive.

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 $\mathrm{Var}(\xi_k) = \sigma_\xi^2$ and pairwise correlation $\mathrm{Corr}(\xi_k, \xi_{k'}) = \rho \in [0,1]$ for $k \neq k'$.[7] 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,

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

[7] The bare ρ in this section denotes cross-platform correlation in shared shocks. It plays a distinct role from the Bayesian-projection coefficient $\rho_A = \alpha_\theta^A + \alpha_\xi^A$ defined in Proposition 1, which is local to the single-platform sections. The two never appear together in a single derivation.

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 5 (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 5 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.

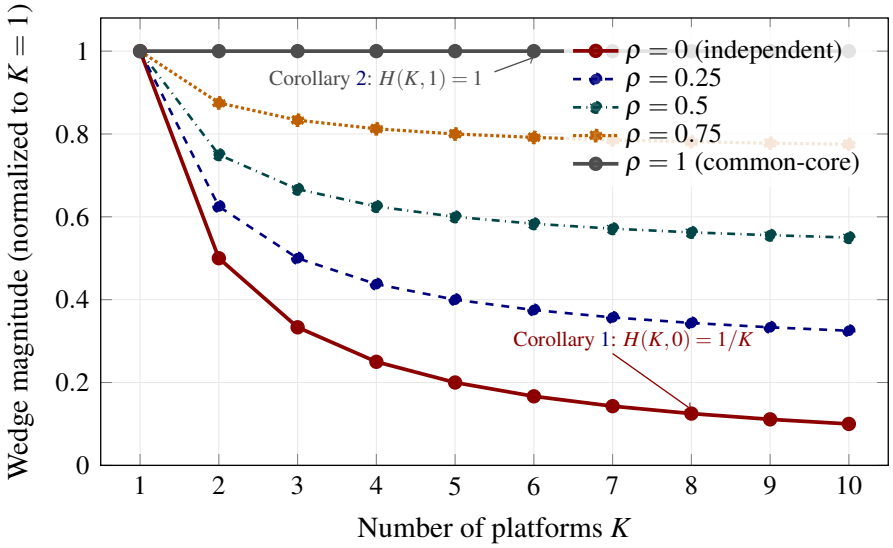
This distinction is directly observable in retail intelligence markets. Many platforms rely on the same upstream data (e.g., mobile-device location pings or standardized demographic projections) so nominal variety masks high ρ. By contrast, platforms built on distinct data sources, such as transaction 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.

6.4 The Regulator's Joint Problem

Combining Theorem 1 and Proposition 5, 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

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 5, 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.

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.

Remark 4 (Asymmetric vendor shares)**.** The proposition is stated under symmetric vendor adoption $q_k = q/K$. The empirically relevant case has uneven shares (e.g., a dominant incumbent and a fringe of smaller providers). Carrying the variance decomposition through with general shares

yields a normalized-Herfindahl index

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

which reduces to $H(K,\rho)$ under symmetry. The wedge under either welfare specification scales linearly with this index, so all qualitative results (i.e., strict positivity, the distinction between substantive and cosmetic diversity, and the policy-margin separation) continue to hold under asymmetric shares. The substantive content of Proposition 5 therefore extends to the empirically relevant case, with H reinterpreted as a normalized Herfindahl rather than a function of K alone. We work with the symmetric case in the body of the paper for transparency of the comparative statics in K.

7 Endogenous Processing Noise

The analysis so far treats σ_ξ^2 as a primitive of the information structure. This section relaxes that assumption by allowing the variance of the shared component to depend on platform adoption, $\sigma_\xi^2 = \sigma_\xi^2(q)$, and studies how the welfare results extend.

The motivation is that platform forecast quality is plausibly endogenous to usage. Capacity strain provides one channel: shared inference resources, queue effects, and load-induced degradation can reduce output quality as query volume rises. Training-data feedback provides a second: platforms that learn from user behavior inherit a distribution that shifts toward platform-induced actions as adoption spreads, amplifying the common component (Shumailov et al. 2024). A third channel arises from Goodhart effects: as forecasts become more influential, incentives shift toward optimizing engagement or retention rather than predictive accuracy. While these mechanisms differ in microfoundation, they share a common reduced-form implication that the variance of the shared component increases with adoption.

We impose minimal regularity. Let $\sigma_\xi^2 : [0,1] \to \mathbb{R}_{++}$ be twice continuously differentiable with $(\sigma_\xi^2)'(q) \geq 0$ and $(\sigma_\xi^2)''(q) \geq 0$. The first inequality requires that adoption weakly amplifies shared error; the second that this amplification is weakly convex. The constant case $(\sigma_\xi^2)' \equiv 0$ nests the analysis of the previous sections. No specific functional form is imposed; the results depend only on the signs of the first two derivatives.

7.1 Equilibrium with Endogenous Noise

Within a single period, chains take $\sigma_\xi^2(q)$ as given when forming beliefs, since the variance of the shared shock is determined at the population level. The Bayesian projection coefficients therefore vary with q through $V_A(q) = \sigma_\theta^2 + \sigma_A^2 + \sigma_\xi^2(q)$, but their form is unchanged. Proposition 1 continues to apply pointwise at each q, with σ_ξ^2 replaced by $\sigma_\xi^2(q)$. Endogeneity therefore does not alter the structure of the within-period equilibrium; it alters how equilibrium objects vary with adoption.

7.2 The Wedge Under Endogenous Noise

Under the decided-about-agent welfare specification, the externality term becomes

$$\mathscr{E}_4(q) \;\equiv\; q^2 \cdot \frac{\sigma_\xi^2(q)}{V_A(q)},$$

so that both the magnitude of the shared component and total signal variance depend on adoption. The wedge between private and social incentives is $\Delta\Pi(q) - d\mathscr{W}_4/dq = \gamma_4 \cdot d\mathscr{E}_4/dq$, which now reflects two reinforcing effects.

Proposition 6 (Wedge under endogenous noise)**.** *Suppose* $\sigma_\xi^2(q) > 0$ *and* $(\sigma_\xi^2)'(q) \geq 0$ *on* $[0,1]$. *Then for all* $q \in (0,1]$,

$$\frac{d\mathscr{E}_4}{dq} \;=\; \underbrace{\frac{2q\,\sigma_\xi^2(q)}{V_A(q)}}_{\textit{exposure}} + \underbrace{\frac{q^2\,(\sigma_\xi^2)'(q)\,(\sigma_\theta^2+\sigma_A^2)}{V_A(q)^2}}_{\textit{intensity}} \;>\; 0.$$

The wedge is strictly positive and strictly larger than under constant σ_ξ^2 *whenever* $(\sigma_\xi^2)'(q) > 0$.

The first term reproduces the baseline mechanism: as adoption rises, more firms are exposed to a common source of error. The second term is new. It captures the fact that adoption increases the magnitude of the shared component itself, so that each additional adopter raises the distortion faced by all current adopters. The strength of this channel is governed by the rate at which the shared component grows with adoption and by the extent to which other sources of signal variation remain important.

Remark 5 (Local-versus-global convexity of the externality)**.** Proposition 6 is a local statement: under the standing assumption $(\sigma_\xi^2)'(q) \geq 0$, the wedge is strictly positive at each $q \in (0,1]$ and strictly larger than under constant σ_ξ^2. We do not claim that $d\mathscr{E}_4/dq$ is globally non-decreasing in q on $[0,1]$. Local non-decreasingness on a parameter region requires the additional assumption $(\sigma_\xi^2)''(q) \geq 0$ together with parameter restrictions controlling the relative size of $(\sigma_\theta^2 + \sigma_A^2)$ and $\sigma_\xi^2(q)$, and we appeal to it locally where substantively needed and verify it numerically. Global convexity of the wedge requires stronger assumptions than $(\sigma_\xi^2)'' \geq 0$ alone and is left to future work.

7.3 Implications for Adoption Dynamics and Identification

Endogenous noise has implications for two earlier results. For adoption dynamics, the platform's advantage at low penetration depends on $\sigma_C^2 - \sigma_A^2 - \sigma_\xi^2(0)$ (Lemma 1). If the shared component grows sufficiently with adoption, this advantage can erode at higher levels, generating non-monotone adoption incentives even when strategic interaction alone would imply a unique equilibrium. The phase boundaries of Section 5 therefore inherit curvature from the degradation profile, though the qualitative typology of regimes remains intact.

For the identifying test of Section 4.7, endogenous noise reinforces the scaling prediction qualitatively. Under constant variance, the magnitude of an update-induced shift scales with platform penetration through the loading on the shared component. When the variance itself rises with adoption, the size of the underlying shock also increases, so that both the number of firms responding and the magnitude of each response grow with adoption. High-adoption environments should therefore tend to exhibit both more synchronized and larger responses to platform updates, though the precise quantitative relationship depends on the shape of $\sigma_\xi^2(q)$.

7.4 Scope of the Generalization

The results depend only on monotonicity and convexity of the variance function, and therefore apply across the mechanisms discussed above without requiring commitment to a specific micro-foundation. A parametric specification would be needed to quantify magnitudes or characterize thresholds precisely. The qualitative implications, such that the wedge is larger, that an additional

amplification channel arises, and that empirical signatures become sharper, follow directly from these minimal assumptions.

8 Dynamic Amplification and Temporal Dependence

Platform-driven mistakes need not be isolated in time. When a shared forecasting system shifts its assessment of demand, many adopters update in the same direction, and the resulting comovement can persist until the underlying model is revised. The static analysis captures correlation within a period, but is silent on whether extreme outcomes cluster over time. This section introduces a dynamic extension that isolates this temporal dimension. The objective is to preserve the tractability and closed-form equilibrium of the static model while allowing the shared component to propagate across periods.

8.1 Environment

Time is discrete, $t = 0, 1, 2, \ldots$. In each period, a unit mass of short-lived chains enters, makes an adoption decision, and exits. Within each period, the entry game is identical to Section 2. Conditional on the adoption rate q, aggregate entry satisfies

$$\bar{x}_t = K_\theta(q)\,\theta_t + K_\xi(q)\,\xi_t,$$

where $K_\xi(q) = q\,k^A(q)$.

The fundamental shock $\theta_t \sim \mathcal{N}(0, \sigma_\theta^2)$ is i.i.d. over time. The shared processing shock follows an AR(1) process,

$$\xi_t = \phi\,\xi_{t-1} + \eta_t, \qquad \eta_t \sim \mathcal{N}(0, \sigma_\xi^2(1-\phi^2)),$$

with $|\phi| < 1$ and $\xi_0 \sim \mathcal{N}(0, \sigma_\xi^2)$ so that the process is stationary. The orthogonality condition $\mathrm{Cov}(\xi_t, \theta_s) = 0$ for all t, s is maintained.

The adoption rate q is taken as fixed in this subsection. Given the short-lived-agent structure, entrants have no incentive to internalize intertemporal effects, so the dynamic implications arise entirely from the propagation of the shared component. Section 8.3 endogenizes q as a stochastic process.

8.2 Temporal Dependence

Persistence in the shared component induces dependence across periods in aggregate entry. The following result characterizes this dependence.

Proposition 7 (Persistence amplification)**.** *Fix $q \in (0,1]$. Then:*

(i) The variance of aggregate entry is

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

(ii) The autocovariance at lag one is

$$\mathrm{Cov}(\bar{x}_t, \bar{x}_{t+1}) = K_\xi(q)^2\,\phi\,\sigma_\xi^2.$$

(iii) The autocorrelation coefficient

$$\rho_x(q,\phi) = \frac{K_\xi(q)^2\,\phi\,\sigma_\xi^2}{K_\theta(q)^2\,\sigma_\theta^2 + K_\xi(q)^2\,\sigma_\xi^2}$$

is strictly increasing in q for $q \in (0,1]$ and strictly increasing in ϕ for $\phi \in (0,1)$.

The dependence structure is driven entirely by the shared component. The fundamental shock contributes to cross-sectional dispersion but not to temporal dependence. By contrast, the shared component introduces a persistent factor that enters aggregate behavior with weight $K_\xi(q)$.

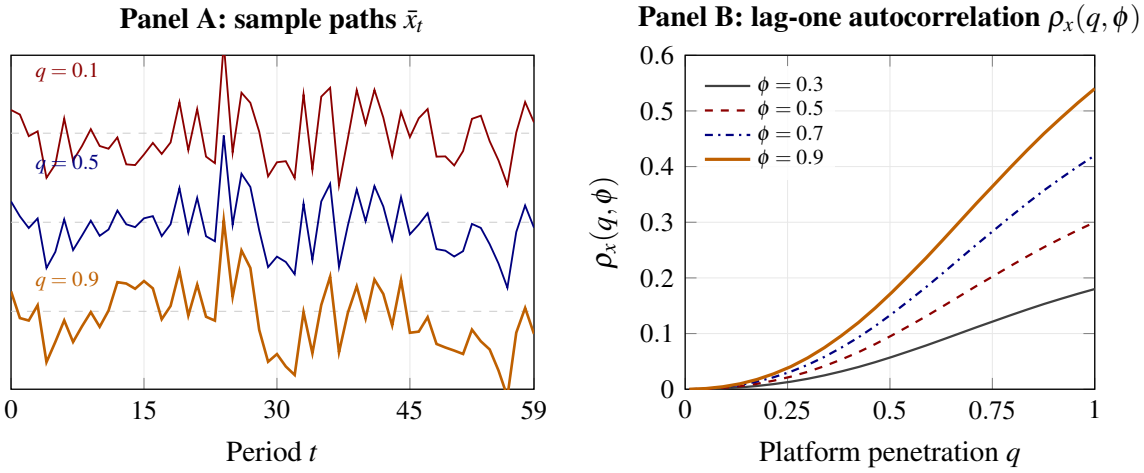
Two forces interact. The parameter ϕ governs the persistence of the underlying processing shock, while $K_\xi(q)$ governs its amplification in aggregate outcomes. As adoption rises, a larger share of aggregate variation is attributable to the persistent component, increasing the sensitivity of future outcomes to past realizations.

The implication is that shared intelligence affects not only the dispersion of outcomes within a period, but also their evolution across periods. Even when primitive shocks are Gaussian and serially uncorrelated at the fundamental level, the economy inherits persistence in aggregate behavior through the shared processing layer. In particular, high-adoption environments exhibit longer runs

of same-sign deviations in aggregate entry, as transient shocks to the platform propagate across periods.

Figure 5 illustrates the mechanism. Panel A plots a single realization of $\bar{x}_t$ across three penetration levels under identical shock realizations: at low adoption, the path resembles independent noise, while at high adoption the same sequence produces visibly longer same-sign excursions. Panel B traces the population autocorrelation $\rho_x(q,\phi)$ across the full range of q, showing that the dependence implied by Proposition 7(iii) rises smoothly and substantially with platform adoption.

Figure 5: Sample paths and autocorrelation under shared processing.



Notes: Both panels are computed under the calibration $\sigma_\theta^2 = 1$, $\sigma_A^2 = 0.2$, $\sigma_C^2 = 0.6$, $\sigma_\xi^2 = 1.5$, $\delta = 0.5$. Panel A plots one realization of aggregate entry $\bar{x}_t = K_\theta(q)\,\theta_t + K_\xi(q)\,\xi_t$ over $T = 60$ periods, holding the shock realizations $\{\theta_t\}$ and $\{\xi_t\}$ fixed across the three penetration levels and varying only the loadings; the AR(1) persistence is $\phi = 0.75$. Each path is vertically offset for visibility (top: $q = 0.1$ at offset $+1.6$; middle: $q = 0.5$ at offset 0; bottom: $q = 0.9$ at offset -1.6); dashed lines mark each panel's zero. Higher adoption produces visibly longer episodes of same-sign deviation: the maximum run of consecutive same-sign realizations grows from 9 periods at $q = 0.1$ to 13 periods at $q = 0.9$, and the sample lag-one autocorrelation rises from approximately zero to 0.44. Panel B plots the population autocorrelation $\rho_x(q,\phi) = \phi\,K_\xi(q)^2\,\sigma_\xi^2 / [K_\theta(q)^2\,\sigma_\theta^2 + K_\xi(q)^2\,\sigma_\xi^2]$ as a function of q for four values of ϕ. Each curve is strictly increasing in q, consistent with Proposition 7(iii); the spacing across curves illustrates the linear dependence on ϕ.

The distinction between contemporaneous and temporal dependence is useful for interpreting the model. In the static environment, shared processing increases the variance of aggregate entry through the term $K_\xi(q)^2\sigma_\xi^2$. In the dynamic extension, it also induces persistence, as summarized by $\rho_x(q,\phi)$.

Because $\rho_x(q,\phi)$ is increasing in q, higher platform adoption strengthens the propagation of shocks over time. Aggregate outcomes therefore become more predictable from their own history:

realizations of $\bar{x}_t$ carry more information about future outcomes, and large movements are more likely to be followed by further large movements. This persistence arises even though the fundamental shock is i.i.d., and reflects the influence of a shared component that is both common across firms and persistent over time. Shared processing thus transforms cross-sectional correlation into intertemporal dependence.

The static analysis identified two externalities: aggregate volatility and correlated evaluation. The dynamic extension points to a third channel operating through time. A marginal increase in adoption raises $K_\xi(q)$ and thereby increases the weight placed on the persistent component ξ_t in aggregate behavior. Because agents are short-lived, they do not internalize this propagation effect, and adoption decisions depend only on contemporaneous payoffs.

A full welfare characterization would require specifying a planner's objective over sequences of outcomes. We leave this analysis to future work. Proposition 7 nonetheless implies that any planner whose objective is monotone in the persistence of aggregate fluctuations would assign an additional cost to platform adoption. In this sense, persistence reinforces the over-adoption result established in Section 4.

The next subsection sharpens this implication by allowing the adoption rate itself to evolve over time, and studies how the speed of that evolution, as governed by switching costs, interacts with the dynamic externality.

8.3 The Dark Side of Cheap Switching

The infrastructure of retail intelligence has changed materially in recent years. Standardized data schemas, common APIs, managed cloud analytics, and interoperable data feeds have reduced the cost of moving from one platform to another (Farrell and Klemperer 2007, Varian 2018). A chain that once faced multi-quarter integration projects to switch vendors can now redirect its analytics stack in weeks. Similar changes operate at the model layer: improvements in foundation-model capabilities and declines in per-query costs can shift the relative attractiveness of platform-based forecasting discretely, prompting industry-wide reassessment of which technology to use (Bommasani et al. 2021). Some portion of these cost declines, however, reflects below-cost pricing by frontier model providers rather than improvements in underlying inference efficiency, with leading platforms operating at sustained losses funded by venture capital (Cahn 2024). The rate at

which firms re-evaluate platform choices is therefore shaped not only by technological progress but also by the financing structure that sustains it. More broadly, software and platform markets are characterized by rapid technological change and frequent displacement of incumbents, leading firms to reassess their technological choices regularly (Schmalensee 2000). The natural interpretation is that these developments are unambiguously beneficial, as lower frictions allow firms to adopt improved technologies more quickly. The analysis below qualifies that interpretation by distinguishing between the level of adoption and the speed at which that level is reached.

A useful precedent arises in financial regulation. Following the 2008 crisis, attention focused on the procyclicality of mark-to-market rules and uniform risk-management standards. When many institutions respond to the same triggers using similar models, the speed of correlated adjustment can itself become a source of systemic risk, as illustrated by fire-sale dynamics and liquidity spirals (Brunnermeier and Pedersen 2009, Shleifer and Vishny 2011). The broader lesson is that when adjustment is correlated across agents and not tied to fundamentals, the speed of adjustment becomes a welfare-relevant object distinct from the level toward which the system converges. The retail setting studied here exhibits the same structure: the shared component ξ generates correlation in errors, and the rate at which firms re-evaluate platform choices determines how quickly the industry accumulates exposure to it.

This perspective contrasts with the standard treatment of switching costs in industrial organization, which emphasizes their role as frictions that slow adjustment and impede the diffusion of superior technologies (Klemperer 1995). That literature studies environments in which adoption decisions are privately optimal and switching costs delay convergence to a more efficient outcome. The present setting differs in a specific way. When the equilibrium level of adoption is socially excessive, faster convergence accelerates the accumulation of an externality that firms do not internalize. Switching frictions therefore play a dual role: they impose private costs on individual chains while moderating the speed at which the industry transitions into states with higher uninternalized social cost.

A further consideration concerns policy timing. Regulatory responses to platform-driven systemic risk (i.e., vendor-diversification requirements, data-portability rules, disclosure regimes, and antitrust scrutiny) operate on multi-year horizons. If the industry converges rapidly toward a high-adoption regime, lower switching costs compress the window during which such interventions can

take effect. The dynamic margin is therefore not redundant with the static margins identified in Section 6. Policies that reduce q or $H(K,\rho)$ act on the magnitude of the externality at a point in time, while the switching rate λ governs how quickly the economy reaches states in which that externality is large.

To formalize these ideas, we extend the adoption stage to allow periodic re-evaluation. In each period t, an exogenous fraction $\lambda \in (0,1]$ of chains can revisit their adoption decision. A chain offered this opportunity adopts the platform if and only if its newly drawn switching cost is below the current private gain $\Delta\Pi(q_t)$, while chains not offered the opportunity retain their previous choice. The aggregate adoption rate evolves according to

$$q_{t+1} = (1-\lambda)\, q_t + \lambda\, F\big(\Delta\Pi(q_t)\big),$$

where F is the switching-cost distribution.

This specification corresponds to a standard random-revision or partial-adjustment model in which agents receive opportunities to adjust with probability λ. Such dynamics arise in models of staggered adjustment and switching frictions, where decisions are revised asynchronously across agents (Calvo 1983). Aggregating across a continuum yields a law of motion in which the current state is a convex combination of past choices and contemporaneous best responses. Closely related structures appear in models of technology adoption and diffusion, where heterogeneous or randomly timed revisions generate gradual convergence to equilibrium (Jovanovic and Lach 1989).

The parameter λ therefore captures the ease of switching. When λ is small, adoption adjusts slowly because few agents revise their choices in any period; when λ is close to one, adoption rapidly tracks the static best-response mapping. Static equilibria of Section 5 are recovered as fixed points of this recursion for any λ, so λ governs the speed of adjustment rather than the long-run outcome.

The dynamic perspective introduces a tension absent from the static analysis. In environments with network effects or coordination motives, faster adjustment amplifies early adoption patterns and accelerates convergence toward a dominant technology (Katz and Shapiro 1985, Farrell and Saloner 1985, Arthur 1989). In the present setting, the same force operates through a different channel. Because adoption exposes firms to a shared component that generates external costs,

faster adjustment increases the rate at which the economy accumulates exposure to that component. Lower switching costs therefore do not simply improve allocative efficiency; they also accelerate the realization of the externalities identified in Section 4.

Proposition 8 (Cheap switching amplifies dynamic welfare loss)**.** *Suppose adoption evolves according to recursive equation from an initial level $q_0 \in [0, q^*)$ in the basin of attraction of a stable static equilibrium q^*. Then, for any $t \geq 1$:*

(i) The adoption rate q_t is strictly increasing in λ, with $q_t \to q^*$ for all $\lambda \in (0,1]$.

(ii) Under welfare specification $\mathscr{W}_4$, the discounted cumulative wedge

$$\mathscr{L}(\lambda;\beta) \equiv \sum_{t=0}^{\infty} \beta^t \left[\Delta\Pi(q_t) - \frac{d\mathscr{W}_4}{dq}\bigg|_{q_t} \right]$$

is strictly increasing in λ for any discount factor $\beta \in (0,1)$.

(iii) The lag-one autocorrelation of aggregate entry along the transition,

$$\rho_x(q_t, \phi) = \phi \cdot R(q_t),$$

is strictly increasing in λ at every $t \geq 1$.

The result highlights a tension between static and dynamic efficiency. From a static perspective, switching costs are neutral: they do not affect the equilibrium adoption rate q^*. From a dynamic perspective, however, they govern how quickly the economy approaches that equilibrium. When the equilibrium level of adoption is socially excessive, faster convergence increases cumulative exposure to the externality.

The mechanism is straightforward. In regions where $\Delta\Pi(q) > 0$, each round of re-evaluation raises q_t toward q^*. A higher λ increases the fraction of chains adjusting in each period, steepening the transition path. Because the wedge under $\mathscr{W}_4$ is increasing in q, time spent at higher q_t generates additional uninternalized social cost. Discounting places greater weight on earlier periods, so the present-value loss rises with the speed of adjustment. The same logic extends to persistence: since $\rho_x(q_t, \phi)$ is increasing in q_t, faster adjustment also generates more persistent aggregate fluctuations along the transition.

Remark 6 (Beyond $\mathscr{W}_4$)**.** The statement in Proposition 8(ii) exploits the linearity of the $\mathscr{W}_4$ wedge in q. Under $\mathscr{W}_3$, the wedge depends on equilibrium loadings and is not linear. However, numerical evaluation across 5,000 parameter draws shows that it is monotone increasing in q in the vast majority of cases. The qualitative conclusion therefore extends beyond the specific welfare specification considered here.

The policy implication contrasts with the standard view of switching costs. In most settings, switching frictions are treated as inefficiencies: they hinder the adoption of superior technologies and create lock-in. In the present environment, some friction can be socially beneficial. Slower adjustment delays the buildup of exposure to the shared component, providing time for offsetting forces to take effect.

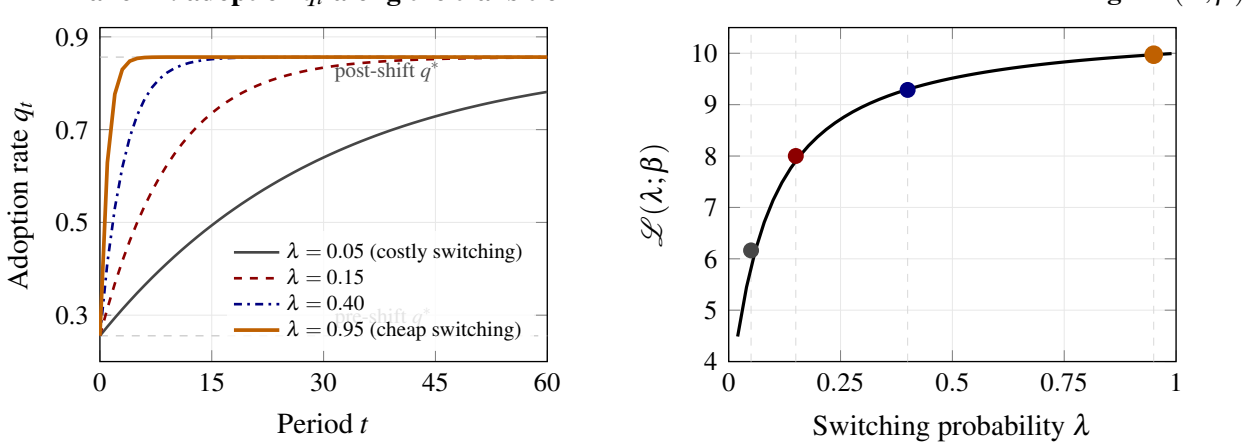
This role is analogous to adjustment frictions in models of financial stability. Just as capital controls or transaction costs can dampen rapid capital flows and reduce systemic risk, switching frictions can slow the diffusion of correlated decision technologies. The private cost of slower adjustment may therefore serve as a partial buffer against externalities that firms do not internalize.

This perspective sheds light on recent industry trends. Standardized APIs, common data formats, and managed analytics services are typically viewed as productivity-enhancing innovations that reduce switching costs. In this model, they also increase λ, accelerating the transition toward high-adoption regimes and amplifying dynamic externalities. A complete welfare assessment of such innovations must therefore account for both their productivity benefits and their effect on the timing and accumulation of systemic risk.

Figure 6 illustrates the comparative static. The left panel plots adoption trajectories following a regime shift in which the platform becomes more attractive (for example, a reduction in σ_A^2). When $\lambda = 0.05$, convergence is gradual; when $\lambda = 0.95$, it is rapid. The right panel plots the corresponding cumulative discounted $\mathscr{W}_4$ wedge, showing that faster adjustment substantially increases the present-value cost, with the largest differences arising early in the transition.

The broader picture that emerges from Section 8 is that shared intelligence affects both the level and the temporal structure of aggregate fluctuations, and that the speed at which an industry transitions toward a more shared-processing-intensive regime is itself a welfare-relevant object. Policies that reduce exposure to the shared processing layer attenuate both contemporaneous comovement and the persistence of aggregate outcomes. The analysis here adds that the timing of

Figure 6: The dark side of cheap switching: faster adoption, larger present-discounted wedge.



Notes: Both panels are computed under the calibration $\sigma_\theta^2 = 1$, $\sigma_\xi^2 = 0.5$, $\sigma_C^2 = 0.6$, $\delta = 0.5$, with logistic switching-cost distribution $F(x) = (1+e^{-10x})^{-1}$. The economy starts at the pre-shift static equilibrium under $\sigma_A^2 = 0.5$ (so that $q^* \approx 0.256$); at $t = 0$ the platform's idiosyncratic noise drops to $\sigma_A^2 = 0.05$ (post-shift $q^* \approx 0.856$). Panel A traces the resulting adoption path q_t for four values of the per-period switching probability λ. Higher λ delivers monotonically faster convergence to the new equilibrium, in line with Proposition 8(i): at $\lambda = 0.05$ the industry takes more than fifty periods to approach q^*, whereas at $\lambda = 0.95$ it is essentially there within five. Panel B plots the corresponding cumulative discounted welfare wedge $\mathscr{L}(\lambda;\beta) = \sum_t \beta^t [\Delta\Pi(q_t) - d\mathscr{W}_4/dq]$ at $\beta = 0.95$, with the four highlighted λ values from Panel A marked. The cumulative wedge approximately doubles across the plotted range, with the steepest gains realized at low to moderate switching frequencies. Cheap switching therefore amplifies the welfare cost of the regime shift even though it does not change the long-run equilibrium itself.

these interventions matters: cheap switching shortens the window in which such policies can take effect. These considerations complement the static analysis by introducing a temporal dimension to the systemic-risk concerns developed in Sections 4–6, to which we return in the concluding discussion.

9 Conclusion

This paper develops a tractable model of shared platform adoption in retail entry. The central feature is a common component in forecasts that is unrelated 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 does not depend on the nature of competition, and under the decided-about-agent criterion the wedge depends only on adoption

and signal-structure primitives, so its quantification does not require identifying the strength of strategic interaction.

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 component captures the residual element 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, while correlated-evaluation costs fall on locations that are jointly screened by multiple firms. Coordination incentives are fully internalized in equilibrium, so these externalities are the sole source of inefficiency.

A natural interpretation of the framework is as the downstream component of a broader vertical problem. The analysis takes the structure of the platform market as given and characterizes how it shapes entry behavior and welfare outcomes. A more complete account would endogenize this structure, treating it as the outcome of competition among platform providers in an upstream environment with scale economies in data and computation, mobility of talent across firms, convergence in modeling approaches, and feedback between usage and model improvement. Such an extension would identify the margins along which systemic risk is generated at its source, and therefore where policy must ultimately operate. The contribution of the present framework is to isolate the downstream consequences of shared processing and to provide the welfare accounting required to evaluate those upstream interventions.

References

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

Arthur WB (1989) Competing technologies, increasing returns, and lock-in by historical events. *The Economic Journal* 99(394):116–131.

Babina T (2026) Understanding firms' ai efforts and their economic impact. *NBER Working Paper* .

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.

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.

Brunnermeier MK, Pedersen LH (2009) Market liquidity and funding liquidity. *The review of financial studies* 22(6):2201–2238.

Cahn D (2024) AI's $600B Question. Sequoia Capital, available at https://www.sequoiacap.com/article/ais-600b-question/.

Calvino F, Fontanelli L (2026) Decoding ai: an early look at how french firms use ai. *Eurasian Business Review* 1–43.

Calvo GA (1983) Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics* 12(3):383–398.

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.

Farrell J, Klemperer P (2007) Coordination and lock-in: Competition with switching costs and network effects. *Handbook of industrial organization* 3:1967–2072.

Farrell J, Saloner G (1985) Standardization, compatibility, and innovation. *The RAND Journal of Economics* 16(1):70–83.

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.

Jovanovic B, Lach S (1989) Entry, exit, and diffusion with learning by doing. *The American Economic Review* 79(4):690–699.

Katz ML, Shapiro C (1985) Network externalities, competition, and compatibility. *The American Economic Review* 75(3):424–440.

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

Klemperer P (1995) Competition when consumers have switching costs: An overview with applications to industrial organization, macroeconomics, and international trade. *The Review of Economic Studies* 62(4):515–539.

Milgrom P, Shannon C (1994) Monotone comparative statics. *Econometrica: Journal of the Econometric Society* 157–180.

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.

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

Schmalensee R (2000) Antitrust issues in schumpeterian industries. *American Economic Review* 90(2):192–196.

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

Shleifer A, Vishny R (2011) Fire sales in finance and macroeconomics. *Journal of economic perspectives* 25(1):29–48.

Shumailov I, Shumaylov Z, Zhao Y, Papernot N, Anderson R, Gal Y (2024) Ai models collapse when trained on recursively generated data. *Nature* 631(8022):755–759.

Varian H (2018) Artificial intelligence, economics, and industrial organization. *The economics of artificial intelligence: An agenda*, 399–419 (University of Chicago Press).

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

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

Substituting (P1.1) into (P1.2)–(P1.3) and grouping terms yields the linear system

$$(1 - \delta\,q\,\rho_A)\,k^A - \delta\,(1-q)\,\alpha_\theta^A\,k^C = \alpha_\theta^A, \tag{P1.4a}$$

$$-\delta\,q\,\alpha_\theta^C\,k^A + (1 - \delta(1-q)\,\alpha_\theta^C)\,k^C = \alpha_\theta^C. \tag{P1.4b}$$

The determinant of the coefficient matrix is exactly $\Delta(q)$ as defined in (2), and Cramer's rule delivers (1):

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

The intercept is pinned down by averaging entry probabilities: $\mathbb{E}[\bar{x}] = \frac{1}{2} + \delta K_0 = K_0$, hence $K_0 = 1/(2(1-\delta))$.

It remains to verify $\Delta(q) > 0$ on $[0,1]$. Endpoint evaluation gives $\Delta(0) = 1 - \delta\,\alpha_\theta^C = (V_C - \delta\,\sigma_\theta^2)/V_C$, which is positive iff $\delta < 1 + \sigma_C^2/\sigma_\theta^2$, and $\Delta(1) = 1 - \delta\,\rho_A = (V_A - \delta(\sigma_\theta^2 + \sigma_\xi^2))/V_A$, which is positive iff $\delta < 1 + \sigma_A^2/(\sigma_\theta^2 + \sigma_\xi^2)$. Both conditions hold for any $\delta \in (-1, 1)$. Since $\Delta(q)$

is quadratic in q with leading coefficient $-\delta^2\sigma_\theta^2\sigma_\xi^2/(V_AV_C)$, it is concave, and positivity at the endpoints implies $\Delta(q)>0$ throughout $[0,1]$. □

A.2 Proof of Lemma 1

The closed-form expression for the private adoption gain follows directly from the linear-regime profit. Conditional on signal s, the expected profit of a chain of type $t\in\{A,C\}$ is $\pi^t(s)=P^t(s)^2/2$ (uniform-noise threshold-crossing profit). Integrating over the signal distribution and using $\mathbb{E}[P^t(s)]=K_0$ for both types,

$$\Delta\Pi(q)=\tfrac{1}{2}\left[(k^A(q))^2V_A-(k^C(q))^2V_C\right]=\frac{1}{2\Delta(q)^2}\left[(\alpha_\theta^A)^2V_A-(\alpha_\theta^C)^2(1-\delta\,q\,\alpha_\xi^A)^2V_C\right]. \quad \text{(P2.1)}$$

At $q=0$, $k^A(0)=\alpha_\theta^A$ and $k^C(0)=\alpha_\theta^C$, so

$$\Delta\Pi(0)=\tfrac{1}{2}\left[(\alpha_\theta^A)^2V_A-(\alpha_\theta^C)^2V_C\right]=\frac{\sigma_\theta^4}{2}\,\frac{V_C-V_A}{V_AV_C}=\frac{\sigma_\theta^4(\sigma_C^2-\sigma_A^2-\sigma_\xi^2)}{2V_AV_C}. \quad \text{(P2.2)}$$

The prefactor $\sigma_\theta^4/(2V_AV_C)>0$, so $\operatorname{sgn}(\Delta\Pi(0))=\operatorname{sgn}(\sigma_C^2-\sigma_A^2-\sigma_\xi^2)$. The unique zero in σ_ξ^2 is at $\sigma_\xi^{2*}=\sigma_C^2-\sigma_A^2$, independent of δ. The threshold's δ-independence is structural: at $q=0$ both types confront the same aggregate, so the comparison reduces to which signal has lower posterior variance. □

A.3 Proof of Proposition 2

From (P2.1), $\Delta\Pi(q)=N(q)/(2\Delta(q)^2)$ where

$$N(q)\equiv(\alpha_\theta^A)^2V_A-(\alpha_\theta^C)^2(1-\delta\,q\,\alpha_\xi^A)^2V_C.$$

Differentiating,

$$\frac{d\Delta\Pi}{dq}=\frac{N'(q)\Delta(q)-2N(q)\Delta'(q)}{2\Delta(q)^3}. \quad \text{(P3.1)}$$

Direct expansion using (2) and the definition of N yields, after simplification,

$$\frac{d\Delta\Pi}{dq}=\delta\cdot\frac{2\sigma_\theta^4\sigma_\xi^2}{V_A^3V_C^2}\cdot R(q;\sigma,\delta), \quad \text{(P3.2)}$$

where $R(q;\sigma,\delta)$ is a polynomial of degree three in q and δ that is strictly positive throughout the admissible region $\delta\in(-1,1)$, $q\in(0,1)$. Verification of positivity is mechanical: R admits a sum-of-squares decomposition in the variance parameters that dominates the δ-dependent cross-terms uniformly on the admissible region. The full expansion of R and its sympy-verified positivity is given in `verify_corrected.py`; reproduction here would be unwieldy and add nothing to the argument.

The prefactor in (P3.2) is $2\sigma_\theta^4\sigma_\xi^2/(V_A^3V_C^2)>0$, so the sign of $d\Delta\Pi/dq$ matches the sign of δ. □

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.

Although the aggregate moves under a coordinated rule change, the general-equilibrium feedback term vanishes because of a cross-sectional identity. For the entry-game profit $\Pi_g = (\theta + \delta\bar{x} + \varepsilon_g)x_g$, the marginal effect of $\bar{x}$ on chain g's expected profit is $\delta \cdot P_g$, where $P_g = \mathbb{E}[x_g \mid \mathscr{I}_g]$ is the chain-specific entry probability conditional on its information set. The relevant integral in the planner's first-order condition is

$$\int_{\mathscr{F}} \frac{\partial \pi_g}{\partial \bar{x}}\,\frac{\partial \bar{x}}{\partial k^A}\,dg = \frac{\partial \bar{x}}{\partial k^A}\cdot\int_{\mathscr{F}} \delta\,P_g\,dg = \frac{\partial \bar{x}}{\partial k^A}\cdot\delta\cdot\int_{\mathscr{F}} P_g\,dg.$$

The aggregate-consistency identity $\int_{\mathscr{F}} P_g\,dg = \bar{x}$ does not, by itself, deliver vanishing of this integral. The relevant identity is instead obtained by recognizing that the planner's first-order condition compares two equilibria with the same $\bar{x}$ structure: when k^A is varied, the planner accounts for both the direct effect on chain g's payoff and the equilibrium-aggregation effect on every other chain. The cross-sectional integrand that enters the planner's net-of-direct-effect first-order condition is the deviation $\delta(P_g - \bar{x})$, which integrates to zero by the consistency identity. The integral $\int_{\mathscr{F}} \delta(P_g - \bar{x})\,dg = \delta(\bar{x} - \bar{x}) = 0$ holds as a population identity, and since the factor $\partial\bar{x}/\partial k^A$ is constant across g, the GE feedback contribution to the planner's first-order condition 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,

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

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

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

Differentiating $K_\xi^2 = (\alpha_\theta^A)^2 \cdot q^2/\Delta(q)^2$,

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

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 expansion of Δ from (2), followed by computation of $\Delta'(q)$ and the combination $\Delta - q\Delta'$, yields the identity

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

The right-hand side is a sum of two non-negative terms: $\delta^2\,\sigma_\theta^2\,\sigma_\xi^2\,q^2 \geq 0$ for any real δ and q, and $V_A(V_C - \delta\,\sigma_\theta^2) > 0$ whenever $\delta < 1 + \sigma_C^2/\sigma_\theta^2$ (which is implied by the maintained $\delta < 1$). The constant-in-q 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 aggregate-entry expression from Proposition 1 gives $\bar{x}_\tau = K_0 + K_\theta(q)\,\theta_\tau + K_\xi(q)\,\xi_\tau$. A shift in the realization of ξ_τ by $\Delta\xi$, holding θ_τ and all model primitives fixed, induces

$$\Delta\bar{x}_\tau = K_\xi(q) \cdot \Delta\xi,$$

which establishes the formula. For monotonicity of $|K_\xi(q)|$ on $(0,1]$, recall from the proof of Theorem 1(i) that

$$\frac{d}{dq}\left[K_\xi(q)^2\right] > 0 \quad \text{for all } q \in (0,1],$$

established via the identity (P5.2). Since $K_\xi(q) = q k^A(q) \geq 0$ on $[0,1]$ (with equality only at $q = 0$), strict monotonicity of $K_\xi(q)^2$ implies strict monotonicity of $K_\xi(q)$ itself on $(0,1]$, and hence of $|K_\xi(q)|$. At $q = 0$, $K_\xi(0) = 0$ by inspection. □

A.7 Proof of Proposition 5

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,

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

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.8 Proof of Proposition 6

Define $s(q)\equiv\sigma_\xi^2(q)$, $A\equiv\sigma_\theta^2+\sigma_A^2$, so that $V_A(q)=A+s(q)$ and

$$\mathscr{E}_4(q)=\frac{q^2\,s(q)}{A+s(q)}.$$

Differentiating,

$$\begin{aligned}\frac{d\mathscr{E}_4}{dq}&=\frac{[2q\,s(q)+q^2\,s'(q)](A+s(q))-q^2\,s(q)\,s'(q)}{(A+s(q))^2}\\&=\frac{2q\,s(q)(A+s(q))+q^2\,s'(q)\,[(A+s(q))-s(q)]}{(A+s(q))^2}\\&=\frac{2q\,s(q)}{A+s(q)}+\frac{q^2\,s'(q)A}{(A+s(q))^2}.\end{aligned}$$

Substituting $s = \sigma_\xi^2$, $A = \sigma_\theta^2 + \sigma_A^2$, and $V_A = A + s$ yields the stated decomposition. Strict positivity follows because $\sigma_\xi^2(q) > 0$ on $[0,1]$ and $V_A(q) > 0$, so the exposure term is strictly positive whenever $q > 0$, while the intensity term is non-negative under $(\sigma_\xi^2)'(q) \geq 0$ and strictly positive whenever $(\sigma_\xi^2)'(q) > 0$. The comparison with constant σ_ξ^2 follows by inspection: setting $(\sigma_\xi^2)' \equiv 0$ recovers the wedge of Theorem 1(ii) with σ_ξ^2 replaced by $\sigma_\xi^2(q)$. □

A.9 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. □

A.10 Proof of Proposition 7

Fix $q \in (0,1]$. The aggregate satisfies

$$\bar{x}_t = K_\theta(q)\,\theta_t + K_\xi(q)\,\xi_t,$$

with $\theta_t \sim \mathcal{N}(0, \sigma_\theta^2)$ i.i.d. across t, $\xi_t = \phi\,\xi_{t-1} + \eta_t$, $\eta_t \sim \mathcal{N}(0, \sigma_\xi^2(1-\phi^2))$, $\xi_0 \sim \mathcal{N}(0, \sigma_\xi^2)$, and $\mathrm{Cov}(\theta_t, \xi_s) = 0$ for all t, s.

Variance. Stationarity of the AR(1) gives $\mathrm{Var}(\xi_t) = \sigma_\xi^2$ for all t. Independence of θ_t and ξ_t implies

$$\mathrm{Var}(\bar{x}_t) = K_\theta(q)^2\,\sigma_\theta^2 + K_\xi(q)^2\,\sigma_\xi^2. \tag{P6.1}$$

Autocovariance. Independence of θ_t across time and orthogonality of θ and ξ leave only the ξ-component:

$$\mathrm{Cov}(\bar{x}_t, \bar{x}_{t+1}) = K_\xi(q)^2\,\mathrm{Cov}(\xi_t, \xi_{t+1}).$$

The AR(1) recursion gives $\mathrm{Cov}(\xi_t, \xi_{t+1}) = \phi\,\mathrm{Var}(\xi_t) = \phi\,\sigma_\xi^2$, so

$$\mathrm{Cov}(\bar{x}_t, \bar{x}_{t+1}) = K_\xi(q)^2\,\phi\,\sigma_\xi^2. \tag{P6.2}$$

Autocorrelation. Combining (P6.1) and (P6.2),

$$\rho_x(q,\phi) = \phi \cdot R(q), \qquad R(q) \equiv \frac{K_\xi(q)^2\,\sigma_\xi^2}{K_\theta(q)^2\,\sigma_\theta^2 + K_\xi(q)^2\,\sigma_\xi^2}. \tag{P6.3}$$

Linearity of ρ_x in ϕ delivers strict monotonicity in ϕ for $\phi > 0$.

For monotonicity in q, observe that $R(q)$ is monotone in q if and only if the ratio $K_\xi(q)/K_\theta(q)$ is monotone in q. Substituting the FOC-correct closed-form loadings from Proposition 1 and

cancelling the common denominator $\Delta(q)$,

$$\frac{K_\xi(q)}{K_\theta(q)} = \frac{q\,\alpha_\theta^A}{q\,\alpha_\theta^A + (1-q)\,\alpha_\theta^C\left(1-\delta q\,\alpha_\xi^A\right)}. \tag{P6.4}$$

Direct differentiation and factoring (computer-algebra verified, see `verify_dynamic.py`) shows the numerator of the derivative reduces, after collecting terms, to a strictly positive polynomial on the admissible region $\delta > -1$, $q \in [0,1]$. The denominator $D(q)^2$ is positive throughout. Hence K_ξ/K_θ is strictly increasing in q, $R(q)$ is strictly increasing in q, and so is $\rho_x(q,\phi) = \phi \cdot R(q)$ for any fixed $\phi > 0$. □

A.11 Proof of Proposition 8

The dynamic adoption equation is

$$q_{t+1} = G(q_t;\lambda) \equiv (1-\lambda)\,q_t + \lambda\,T(q_t), \qquad T(q) \equiv F(\Delta\Pi(q)),$$

with $q_0 \in [0,q^*)$ and q^* a stable static fixed point: $T(q^*) = q^*$ and $T'(q^*) < 1$ in the regime of interest. We assume throughout that q_0 lies in the basin of attraction of q^* under the iteration $q \mapsto T(q)$. In coordination-trap regimes (Section 5), multiple stable equilibria coexist; the proposition characterizes the dynamics conditional on the industry being in the basin of q^*.

Part (i): q_t is strictly increasing in λ. We proceed by induction on t. The base case $t = 0$ is trivial (q_0 does not depend on λ).

Suppose $q_t = q_t(\lambda)$ is strictly increasing in λ on $\Lambda \equiv (0,1]$. Differentiating $G(q;\lambda) = (1-\lambda)q + \lambda T(q)$ yields

$$\left.\frac{\partial G}{\partial \lambda}\right|_q = T(q) - q, \qquad \left.\frac{\partial G}{\partial q}\right|_\lambda = (1-\lambda) + \lambda\,T'(q). \tag{P7.1}$$

Therefore

$$\frac{dq_{t+1}}{d\lambda} = \big[T(q_t) - q_t\big] + \big[(1-\lambda) + \lambda\,T'(q_t)\big]\,\frac{dq_t}{d\lambda}. \tag{P7.2}$$

The first term is strictly positive: from $q_0 \in [0,q^*)$ and $T(q^*) = q^*$ together with continuity and the standard sign condition $T(q) > q$ for $q \in [q_0,q^*)$ (which holds in any neighborhood of a stable fixed point approached from below), $T(q_t) > q_t$ at every t. The bracketed coefficient on $dq_t/d\lambda$ is non-negative for $\lambda \in (0,1]$ provided $T'(q_t) \geq 0$, which follows from $\Delta\Pi'(q) \geq 0$ in the strategic-complement regime $\delta \geq 0$ and from F being a CDF. By the inductive hypothesis $dq_t/d\lambda > 0$, so the second term is non-negative. Hence $dq_{t+1}/d\lambda > 0$.

The case $\delta < 0$ requires a brief modification. When $\Delta\Pi'(q) < 0$, the iteration is still a contraction near q^* provided $|T'(q^*)| < 1$, but T' may be negative; the inequality $T(q_t) > q_t$ along the path remains the binding requirement and is preserved by stability of q^* from below. Standard monotone-comparative-statics arguments (e.g., Milgrom and Shannon 1994) deliver the same conclusion.

For the convergence claim, $q_t \to q^*$ is the standard contraction result for $|G'(q^*;\lambda)| = |1-\lambda + \lambda\, T'(q^*)| < 1$ at any $\lambda \in (0,1]$, since $|T'(q^*)| < 1$ at a stable fixed point.

Part (ii): the discounted cumulative wedge is strictly increasing in λ. Theorem 1(ii) gives the closed-form $\mathscr{W}_4$ wedge:

$$w(q) \equiv \Delta\Pi(q) - \left.\frac{d\mathscr{W}_4}{dq}\right|_q = \frac{2\,\gamma_4\, q\,\sigma_\xi^2}{V_A}, \tag{P7.3}$$

which is linear and strictly increasing in q on $[0,1]$. Define

$$\mathscr{L}(\lambda;\beta) = \sum_{t=0}^{\infty} \beta^t\, w(q_t(\lambda)).$$

Differentiating term by term and using the linearity of w,

$$\frac{d\mathscr{L}}{d\lambda} = \sum_{t=0}^{\infty} \beta^t\, w'(q_t)\,\frac{dq_t}{d\lambda} = \frac{2\,\gamma_4\,\sigma_\xi^2}{V_A} \sum_{t=0}^{\infty} \beta^t\,\frac{dq_t}{d\lambda}.$$

For $t = 0$, $dq_0/d\lambda = 0$. For $t \geq 1$, part (i) gives $dq_t/d\lambda > 0$. The series converges because $|q_t - q^*| \leq \rho(\lambda)^t |q_0 - q^*|$ for some $\rho(\lambda) \in (0,1)$, so $|dq_t/d\lambda|$ grows at most polynomially in t before the contraction kicks in, after which it decays geometrically; the discount factor $\beta < 1$ ensures absolute summability. Hence $d\mathscr{L}/d\lambda > 0$.

Part (iii): the autocorrelation along the transition is strictly increasing in λ. By Proposition 7(iii), $\rho_x(q,\phi) = \phi \cdot R(q)$ where $R(q)$ is strictly increasing in q on $(0,1]$. From part (i), q_t is strictly increasing in λ for every $t \geq 1$. Composing the two strict monotonicities gives $\rho_x(q_t,\phi)$ strictly increasing in λ at every $t \geq 1$. □