Shared Intelligence and Misaligned Firm Entry

Feng Chi* Nathan Yang†

May 19, 2026

Abstract

Firm entry decisions are increasingly mediated by shared intelligence platforms whose forecasts embed a common component unrelated to 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) that persist even when firms optimally respond to one another. The wedge under correlated evaluation depends only on adoption and signal structure. A multi-vendor extension shows that diversity reduces the externality only when platforms are structurally independent.

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 have historically relied on local and firm-specific information. A regional grocer learned its trade area through years of operating stores in adjacent counties; a restaurant chain relied on market researchers who visited competing sites, counted cars, and assessed local traffic. Entry decisions were therefore shaped by information-processing routines that differed across firms. Two chains evaluating the same county could reach different conclusions, not only because they had different objectives, but because they interpreted the available information differently.

This information environment is changing. CoStar, Placer.ai, Buxton, and the analytics divisions of CBRE and JLL now provide 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 consult the same platform, receive the same model output, and update in the same direction.[1] The Bank of England's Financial Policy Committee raises an analogous concern in financial markets, noting 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). A similar issue arises in retail entry, where shared forecasting platforms increasingly mediate decisions that were once based on independent investigation.

The mechanism is simple. When firms rely on a shared system, they do not only learn from the same information; they also inherit the same errors. A platform may reduce idiosyncratic noise by pooling data and standardizing analysis, but it can also introduce a common component into users' forecasts through its data sources, model architecture, assumptions, or updates. This common component survives aggregation and induces correlated behavior unrelated to underlying market conditions. The resulting correlation matters both across firms and over time, since shared errors

[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.

can generate clusters of similar decisions across locations and persistent movements in aggregate entry.

Recent empirical work provides evidence consistent with this concern. Firms that invest more heavily in AI exhibit increases in systematic risk, becoming more correlated with aggregate outcomes in a manner not observed for other technologies (Babina et al. 2024). Bian et al. (2024) provide more direct evidence that firms sharing data through common SDKs exhibit performance comovement, return comovement, and herding in product design that cannot be explained by product similarity, supply-chain linkages, or shared analyst coverage. The introduction of Apple's App Tracking Transparency policy, which disrupted cross-firm data merging, weakened these comovements. In related work, Abis et al. (2025) show that the same policy degraded the precision of AI-derived signals, reducing analyst forecast accuracy and price efficiency. Chi (2025) show that firms relying on common data sources can make systematically misinformed entry decisions when those inputs become outdated. Taken together, these findings point to a tension: shared intelligence may improve individual decision-making while increasing correlation across firms. The welfare consequences depend on whether this correlation reflects shared information or shared error.

The relevant question is therefore not whether shared platforms produce more accurate forecasts than custom research in isolation. It is how industry-level behavior changes when forecasts that were previously independent are routed through a common processing layer. When firms rely on the same platform, forecast errors acquire a shared component originating in the platform itself. Firms may therefore not only act similarly, but err similarly.

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 information about a common underlying state. In these cases, correlation reflects economically relevant information, and forecast errors are largely independent across firms. Platform-based forecasts differ because adoption does not make the common component more informative. Instead, the same error is propagated across more decisions. Correlation is therefore technological in origin rather than informational.

This paper develops a tractable model to study this mechanism. A continuum of potential entrants chooses whether to enter a representative market and whether to rely on custom research

or on a shared platform. Custom research generates idiosyncratic forecast errors. The platform reduces idiosyncratic noise but introduces a common processing shock shared across adopters and unrelated to fundamentals. The model admits a closed-form equilibrium with heterogeneous adoption, allowing aggregate entry to be decomposed into a fundamental component and a platform-driven component.

The equilibrium has two central features. First, the influence of the shared component on aggregate entry increases with platform adoption. Idiosyncratic noise averages out in the cross section, but the shared processing shock does not. Second, the private value of adoption depends on the strategic environment. When firms benefit from aligning with rivals, adoption becomes more attractive as it spreads; when firms benefit from differentiation, adoption becomes less attractive as it spreads.

The main normative result is a robust misalignment between private and social incentives. The shared component generates two external costs that firms do not internalize: aggregate volatility and correlated evaluation across locations. The first affects parties exposed to synchronized entry, such as lenders, landlords, and municipalities. The second affects locations that are screened by multiple firms using the same forecasting technology. In both cases, equilibrium platform adoption is excessive. This result does not depend on whether entry decisions are strategic complements or strategic substitutes.

The model also clarifies how the mechanism differs from related work. A strand of research studies how shared information affects market structure through platform market power. For example, Kirpalani and Philippon (2020) show that more precise consumer data can strengthen a platform's bargaining position and reduce entry. Their mechanism operates through prices and outside options. The mechanism here operates through shared processing: adopters inherit a common signal component unrelated to fundamentals, generating correlation in errors rather than only in actions. Bian et al. (2024) document a related empirical phenomenon and develop a network model of data as a productive asset. Their framework emphasizes positive externalities from shared data and characterizes systemic importance through valuation. The present framework is complementary: it focuses on the adoption decision and on the welfare wedge generated when shared processing propagates common errors.

The paper also contributes to the economics of data. Existing work shows that externalities

in data collection can generate inefficient outcomes, including over-collection when individual data revelation affects what can be inferred about others (Acemoglu et al. 2022, Choi et al. 2019, Ichihashi 2021). The externality studied here operates one layer downstream. It arises not from data collection per se, but from the processing layer through which data are converted into forecasts. The affected parties are not other consumers whose information is revealed, but third parties exposed to correlated firm decisions.

The closest benchmarks are the social value-of-information and algorithmic-monoculture literatures. In social value-of-information models (Morris and Shin 2002, Angeletos and Pavan 2007, Hellwig and Veldkamp 2009, Myatt and Wallace 2012), common signals generate correlation that trades off coordination benefits against information loss. The signal structure studied here differs because the common component carries no information about the underlying state. In algorithmic-monoculture models (Kleinberg and Raghavan 2021, Peng and Garg 2024), agents lose independent evaluation when they rely on a common system. The present model brings this logic into an entry environment and shows how shared evaluation affects both aggregate outcomes and welfare.

Two stylized examples illustrate the mechanism. Consider quick-service restaurants that benefit from clustering along a suburban arterial. These firms value being near rivals because traffic and visibility depend on agglomeration. If many subscribe to the same trade-area platform, they learn about the same locations and inherit the same platform errors. Adoption can become self-reinforcing, and aggregate openings within a region may track the platform's assessment rather than only local fundamentals. Now consider specialty grocers, which benefit from differentiation. A shared platform is less attractive because adopters are steered toward the same neighborhoods and customers. Even so, the firms that adopt still inherit common errors: neighborhoods systematically underweighted by the platform may be passed over by all adopters at once. In both cases, firms internalize the private costs of their own entry decisions, but not the external costs created by correlated errors.

The analysis extends the baseline model in three directions. First, allowing multiple platforms shows that the magnitude of the externality depends on the independence of underlying processing technologies. Vendor diversity reduces the wedge only when platforms rely on distinct data or models; cosmetic diversity has little effect when providers share a common processing core. Second, a dynamic extension shows that shared processing induces persistence in aggregate outcomes,

even when fundamentals are independent over time. Third, when firms can switch technologies more easily, the economy accumulates exposure to the shared component more quickly. Lower switching costs therefore accelerate the realization of the externality, even though they do not change the long-run static equilibrium.

Finally, we conclude our paper with a descriptive case study of the U.S. restaurant site-selection analytics decision-support market. Six commercial vendors (e.g., Placer.ai, Tango, Buxton, eSite, SiteZeus, and IdealSpot) compete at the front end, but adoption is concentrated, with Placer.ai holding roughly two-thirds of retail-chain-adjusted share, and the same upstream data layer (e.g., INRIX roadway traffic, SafeGraph points of interest, Placer.ai mobility panels, Experian demographic segmentation) feeds all of them. Plausible parameter values place the market close to the cosmetic-diversity bound and far from substantive diversification, with the externality wedge attaining roughly three-quarters of its single-vendor magnitude despite the apparent presence of six competing platforms.

2 Model

The model studies a retail chain's entry decision when firms choose both whether to enter a representative market and which information technology to use when forecasting profitability. The environment is intentionally sparse, consisting of a single market, two information technologies, and a continuum of potential entrants. This structure isolates the feature of interest, a shared processing layer that introduces correlated errors across adopters, while retaining tractability sufficient for closed-form analysis.

A key distinction in the analysis is between primitives and outcomes. The variance of the shared component is a primitive of the information structure, whereas platform adoption is an endogenous equilibrium outcome reflecting firms' responses to that primitive.

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 work in the regime in which this 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.

2.2 Information and Processing Technologies

The two information technologies in the model correspond 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 system-

[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 lie outside this bound to illustrate qualitative regimes more clearly. The closed-form expressions of Proposition 1 remain well-defined wherever $\Delta(q) > 0$, and the qualitative comparative statics established in Theorem 1 and related results extend to a truncated-equilibrium model in which entry probabilities are clipped to $[0,1]$, with wedge sign and monotonicity preserved over the relevant parameter region.

atic 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 part 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

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

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

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

where $u_f^A \sim \mathcal{N}(0, \sigma_A^2)$ is idiosyncratic 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, but instead inherits a common mistake.[3] It is useful to interpret the shared component as a latent processing factor. This factor sum-

[3]This differs from environments such as Bergemann et al. (2015), where common signals may confound multiple sources of

marizes 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). A parallel concern arises in macroeconomics. Moll (2026) argues that rational expectations in heterogeneous-agent models require decision-makers to forecast cross-sectional distributions, a computation that is implausible even for the most advanced numerical methods, and proposes replacing the assumption with alternatives in which agents form beliefs about aggregates directly. Subjective beliefs in such environments depart from objective probabilities not because agents are biased, but because the computations rational expectations demand are infeasible. The mechanism here formalizes a different source of the same gap. Adopters' beliefs depart from the objective distribution because the forecast is outsourced to a shared technology that introduces correlated error. The departure is not a cognitive failure but a technological one, and its source can be located precisely in the properties of the platform rather than in the properties of the agents. The substantive welfare consequences are the same in equilibrium because both generate a common error component that firms do not internalize, but the policy levers differ. Algorithmic shared errors leave an audit trail (e.g., model versions, training data, parameter updates) that psychological biases do not, so interventions targeting the source of the correlation (e.g., vendor diversification, model-update disclosure, training-data audits) are feasible for shared platforms in ways they are not for sentiment-driven correlation. The model is silent on this design margin, but the framework is consistent with policy approaches that act on the technology rather than on individual decisions.

Remark 2 (Partial informativeness of the shared component)**.** The assumption that ξ is orthogonal to θ is a conservative modeling choice rather than a substantive restriction. Suppose instead that

uncertainty while remaining informative. The orthogonality of ξ to θ is a property of the processing technology itself.

the shared component is partially informative, $\xi = \lambda\,\theta + \zeta$ with $\zeta \perp \theta$, $\mathrm{Var}(\zeta) = \sigma_\zeta^2$, and loading $\lambda \in [0,1)$. The orthogonal case of the baseline model corresponds to $\lambda = 0$. The relevant primitive of the welfare analysis is the residual variance σ_ζ^2: this is the component of ξ that survives any conditioning on θ, propagates through adopters' signals as common error, and generates the externalities of Theorem 1. Decomposing ξ into its informative and residual parts and projecting the platform signal onto θ delivers an equilibrium with the same structure as Proposition 1, with σ_ξ^2 replaced by σ_ζ^2 in the externality terms. The $\mathscr{W}_4$ wedge becomes $2\,\gamma_4\,q\,\sigma_\zeta^2/V_A$, strictly positive whenever $\sigma_\zeta^2 > 0$. The over-adoption result therefore extends as long as the shared component is not fully informative ($\lambda < 1$); the orthogonal case is the limit $\lambda = 0$ that makes the residual variance maximal and simplifies exposition. The role of the orthogonality assumption is to make σ_ξ^2 itself the relevant primitive, rather than its residual component. Equivalently, orthogonality residualizes a more general signal structure to isolate the platform-driven channel; it is not load-bearing for the qualitative result.

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, with zero adoption when $\Delta\Pi(0) \leq 0$ and full adoption when $\Delta\Pi(1) \geq \bar{c}$. With unbounded support, such as the logistic specification used below, these are replaced by interior equilibria with q^* close to but bounded away from 0 or 1. This distinction does not affect any qualitative result.

2.4 Bayesian Filtering

Chains form beliefs using Bayesian filtering given their information and common knowledge of the data-generating process, including the presence and variance of the shared shock ξ. Let $V_A \equiv \sigma_\theta^2 +$

$\sigma_\xi^2 + \sigma_A^2$ and $V_C \equiv \sigma_\theta^2 + \sigma_C^2$ denote unconditional signal variances, and define projection coefficients

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

Then $\mathbb{E}_f^A[\theta] = \alpha_\theta^A \tilde{s}_f^A$, $\mathbb{E}_f^A[\xi] = \alpha_\xi^A \tilde{s}_f^A$, and $\mathbb{E}_f^C[\theta] = \alpha_\theta^C \tilde{s}_f^C$. Custom-research chains do not observe ξ and therefore assign it its prior mean.

The informational content of the platform signal is summarized by $\pi_A \equiv \alpha_\theta^A + \alpha_\xi^A = (\sigma_\theta^2 + \sigma_\xi^2)/V_A$, 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 is consequential and conservative. It corresponds to a setting in which a chain not only subscribes to the platform but also maintains the analytic capacity to characterize its residual error structure. In such a setting, a chain recognizes that a forecast inherits noise from underlying data sources, that this noise varies across the trade areas under consideration, and that this variation is independent of local demand. A chain of this type filters the platform signal correctly and weights it against its prior as in Proposition 1.

Many adopters do not resemble this benchmark. A regional grocer purchasing a CoStar subscription as part of a broader site-selection package, a franchisee relying on analytics provided by a franchisor, or a tenant receiving a platform-based recommendation through a brokerage relationship instead encounters the output as a point estimate, typically without an explicit measure of uncertainty. In such settings, the shared component is implicitly treated as if its variance were negligible. The platform signal is therefore weighted as if it were more informative than it is, and entry decisions over-respond to the common component rather than partially filtering it out.

The rational-filtering benchmark thus delivers a lower bound on platform-driven distortion. When chains under-correct for σ_ξ^2, both externalities of Section 4 are amplified. Aggregate entry inherits a larger loading on the shared component, and locations are screened with more strongly correlated errors than those implied by Theorem 1. The qualitative direction of the result is preserved under any filtering rule that assigns non-negative weight to the shared component, but the magnitude of the wedge is larger. The rational benchmark therefore isolates the component of the welfare wedge that arises from the information structure itself, abstracting from additional distortions introduced by behavioral misperception.

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. Note that throughout the equilibrium analysis we restrict attention to the *admissible region*

$$\mathscr{D} \equiv \left\{ \delta \in \mathbb{R} : -1 < \delta < \min\left(1 + \frac{\sigma_C^2}{\sigma_\theta^2},\ 1 + \frac{\sigma_A^2}{\sigma_\theta^2 + \sigma_\xi^2} \right) \right\},$$

on which the equilibrium denominator $\Delta(q)$ of Proposition 1 is strictly positive on $[0,1]$.

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 equal the population-weighted contributions of each type. Let k^A and k^C denote the slopes of entry probabilities 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

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

where $\pi_A = \alpha_\theta^A + \alpha_\xi^A = (\sigma_\theta^2 + \sigma_\xi^2)/V_A$. The equilibrium exists and is unique on the admissible region $\mathscr{D}$. The two bounds in the definition of $\mathscr{D}$ ensure $\Delta(0) > 0$ and $\Delta(1) > 0$ respectively, and concavity of $\Delta(q)$ in q implies positivity throughout $[0,1]$. The restriction $\delta > -1$ used elsewhere is sufficient for the global-games argument but not for existence of the linear equilibrium, which also requires the upper bound in $\mathscr{D}$.[4]

The two loadings differ because of how the shared shock enters 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 when $\delta < 0$ and to reinforce it when $\delta > 0$. 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 $k^A(1) = \alpha_\theta^A/(1-\delta\pi_A)$, where the denominator reflects that the signal contains both informative and uninformative common components.

[4] For $\delta \geq 1$, the bound is satisfied only when $\sigma_C^2/\sigma_\theta^2 > \delta - 1$ and $\sigma_A^2/(\sigma_\theta^2 + \sigma_\xi^2) > \delta - 1$. Some calibrations lie in regions where these bounds bind. The closed-form expressions of Proposition 1 extend analytically to such cases but should then be interpreted as approximations to the truncated equilibrium described in Section 2.

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 decomposition separates two components. The first reflects fundamental tracking: the aggregate loads on θ, and its contribution scales with σ_θ^2. The second reflects shared processing: 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, the second component appears and does not vanish under aggregation. Idiosyncratic noise averages out in the continuum, but the shared shock survives intact.

The implication is that shared platforms introduce aggregate variation that is unrelated to fundamentals. Correlation across chains is therefore no longer tied solely to information about demand, but reflects a common error embedded in the processing technology. The term $K_\xi(q)^2\sigma_\xi^2$ is the object through which this mechanism operates and is 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 internalized in equilibrium, while the shared component 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 alters forecast precision and 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. When $\delta > 0$, adoption becomes more attractive as it spreads because the aggregate becomes more exposed to the shared component ξ, and platform adopters track this component more closely than custom chains. When $\delta < 0$, the same mechanism reverses: chains benefit from differentiation, so the increased correlation induced by the platform reduces its marginal value as adoption 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 strictly positive over the admissible region $\mathscr{D}$. The sign of $d\Delta\Pi/dq$ is therefore determined by δ.

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 generates two distinct externalities, each affecting a different group outside the chain population.

The first operates at the aggregate level. It affects agents whose payoffs depend on the variance of entry around fundamentals, including lenders exposed to correlated default risk, landlords facing synchronized vacancies, and municipal authorities planning around commercial activity. The relevant object is the platform-driven component of aggregate variance, $K_\xi(q)^2\sigma_\xi^2$, which increases with adoption and does not vanish under aggregation.

The second operates at the level of individual locations. Locations are evaluated by multiple chains as potential entry targets. When chains rely on the shared platform, their signals contain the same realization of ξ, so evaluations of a given location become positively correlated. As adoption rises, a larger share of evaluations shares this common component, reducing the effective independence of screening across firms.

This mechanism is related to the externality emphasized by Kleinberg and Raghavan (2021) and Peng and Garg (2024). In those settings, the decided-about agents are individuals, and the relevant statistic is the probability that at least one decision-maker recognizes a favorable case. Here the decided-about agent is a location, and the welfare object is the expected pairwise covariance of evaluation errors normalized by total signal variance, which scales as $q^2\sigma_\xi^2/V_A$. In other words, the encounter probability times the variance share of the shared component. The two formulations capture related but distinct consequences of correlated decision technologies.

The externalities differ in the objects they act on. The aggregate-volatility externality depends on equilibrium behavior through the loadings, while the correlated-evaluation externality depends only on the signal structure and adoption rate. Both originate from the same primitive feature of the model, 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 discussed in Remark 4.

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 strategic interactions, which Proposition 3 shows does not arise.

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

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

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

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

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

The term $q^2\sigma_\xi^2/V_A$ equals the probability that two evaluators both adopt the platform times the share of an adopter's signal variance attributable to the shared component. The variance share σ_ξ^2/V_A coincides with the Bayesian projection coefficient α_ξ^A from Section 2 and admits a direct interpretation: V_A is the total uncertainty an adopter faces when evaluating a location, so σ_ξ^2/V_A is the fraction of that uncertainty arising from the shared component. A decided-about location screened by multiple adopters bears not the absolute variance of the shared shock but its weight relative to total signal noise: a platform with high σ_ξ^2 but compensating reductions in idiosyncratic noise leaves the location no worse off than under custom research with comparable total uncertainty. Normalizing by V_A captures this relative-weight intuition. The parameter γ_4 measures the marginal cost of correlated evaluation.

Because the term depends only on adoption and signal-structure primitives, the $\mathscr{W}_4$ wedge at

any fixed q is independent of equilibrium loadings, and therefore of δ. This δ-invariance is a corollary of a more general property: any signal-primitive welfare penalty (i.e., any specification of the externality term that depends only on the information structure and on the adoption rate, not on equilibrium responses) yields a wedge whose sign and magnitude are independent of strategic interaction. The $\mathscr{W}_4$ formulation is the natural representative of this class, but the invariance is a property of the welfare-object class rather than of the specific normalization.

The distinction between $\mathscr{W}_3$ and $\mathscr{W}_4$ is structural. The former penalizes realized aggregate behavior and therefore depends on equilibrium loadings, while the latter penalizes correlation in signals and depends only on the information structure. These correspond to different affected parties, agents exposed to aggregate outcomes in the first case and locations subject to correlated screening in the second, so results under one specification do not mechanically carry over to the other.[5]

A natural alternative is the fundamental-tracking criterion $\mathscr{W}_2$ studied in the social-value-of-information literature (Angeletos and Pavan 2007, Morris and Shin 2002). This specification captures costs internal to the chain population: deviations from fundamental tracking are borne by chains through their own profits, and Proposition 3 implies that such effects are fully internalized in equilibrium. The misalignment identified here arises from costs borne outside the chain population, which $\mathscr{W}_2$ does not capture. A planner concerned with both objectives would adopt a hybrid specification, in which case the externality wedge identified below adds to any fundamental-tracking distortion. The externality component is robust in sign and provides the clean source of over-adoption.

4.4 The Team-Efficiency Result

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

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

[5] 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.

The result follows from 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 a common decision rule to each type faces the same problem, except for a general-equilibrium term capturing how coordinated changes affect aggregate entry.

This term vanishes because of a cross-sectional identity. The marginal payoff with respect to $\bar{x}$ for chain g is δP_g, whose average equals $\delta \bar{x}$. The planner's first-order condition isolates the deviation $\delta(P_g - \bar{x})$ by netting out the direct effect on chain g's own payoff. This deviation integrates to zero, and because $\partial \bar{x}/\partial k$ is constant across chains, the general-equilibrium feedback term vanishes. The planner's condition therefore 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. Chains fully internalize the effect of their decisions on one another through equilibrium responses, so any wedge between private and social incentives must 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.

Remark 3 (Robustness of the team-efficiency result)**.** The identity $d\mathscr{W}_1/dq = \Delta\Pi(q)$ relies on two features of the baseline environment: a continuum of agents and linear entry probabilities. Both are simplifications, and their roles in the proof differ.

The continuum is innocuous. With finitely many firms, the cross-sectional identity in the proof of Proposition 3 holds only in expectation, so the general-equilibrium term does not cancel exactly. Standard arguments establish that the residual is $O(1/N)$ and that the over-adoption result of Theorem 1 remains intact for any $N \geq 2$.

The linearity assumption does more work. With nonlinear entry probabilities outside the global-games regime, the marginal payoff acquires a correction term that depends on the density of firms near the entry threshold. The planner's first-order condition no longer coincides exactly with the Nash condition, and any residual deviation enters the wedge between private and social adoption incentives. The qualitative externality results of Theorem 1 nonetheless survive this perturbation. Appendix B reports a numerical exercise replacing the linear entry rule with a logistic specification $P_f = \Lambda(\kappa[\mathbb{E}_f(\theta) + \delta\,\mathbb{E}_f(\bar{x})])$ that nests linearity as the small-κ limit. The $\mathscr{W}_4$ wedge is strictly positive across all calibrations and within 1% of its closed-form linear value; the $\mathscr{W}_3$

wedge is also strictly positive across the great majority of cells, with two boundary failures in regimes where the welfare object is itself near zero. The numerical evidence is therefore consistent with the structural property identified in Theorem 1(ii): the $\mathscr{W}_4$ externality depends only on signal-structure primitives and inherits no sensitivity to the choice rule.

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 $\delta \in \mathscr{D}$ 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 depends only on adoption and signal-structure primitives and is invariant in δ when evaluated at any fixed q.[6]

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

The two parts share the same structure: the wedge equals the derivative of the externality with respect to adoption. In each case, the relevant object is increasing in q, so marginal adoption raises

[6] Appendix B reports a numerical robustness exercise replacing the linear global-games rule with a logistic entry rule $P_f = \Lambda(\kappa[\mathbb{E}_f(\theta) + \delta\,\mathbb{E}_f(\bar{x})])$. Across three calibrations spanning $\delta \in \{-0.5, 0, +0.7\}$ and four choice-sharpness values $\kappa \in \{1, 2, 4, 8\}$, the $\mathscr{W}_4$ wedge remains strictly positive in every cell and within 1% of its closed-form linear value, consistent with the signal-primitive structure of the externality term. The $\mathscr{W}_3$ wedge remains positive in 34 of 36 cells.

a cost that chains do not internalize.[7]

The difference lies in what this object depends on. Under $\mathscr{W}_3$, the externality is $K_\xi(q)^2\sigma_\xi^2$, which combines the primitive variance σ_ξ^2 with an amplification term $K_\xi(q)^2 = q^2(k^A(q))^2$ that 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 $q^2\sigma_\xi^2/V_A$. 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, so the wedge at any fixed q is independent of δ. Dependence on δ arises only through the equilibrium adoption level $q^*(\delta)$.

The over-adoption result therefore does not rely on strategic interaction. It follows from the presence 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 imply the same direction.[8]

Each welfare specification corresponds to a distinct externality acting on a different population. Their agreement shows that over-adoption is a property of the information structure rather than of a particular welfare criterion.

The invariance under $\mathscr{W}_4$ has an additional implication. The strategic parameter δ is difficult to identify empirically, as 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.

Remark 4 (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):

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

[7] 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 $\tau_k(q^*)$ (e.g., adoption taxes or capital requirements indexed to platform exposure) implements the constrained optimum.

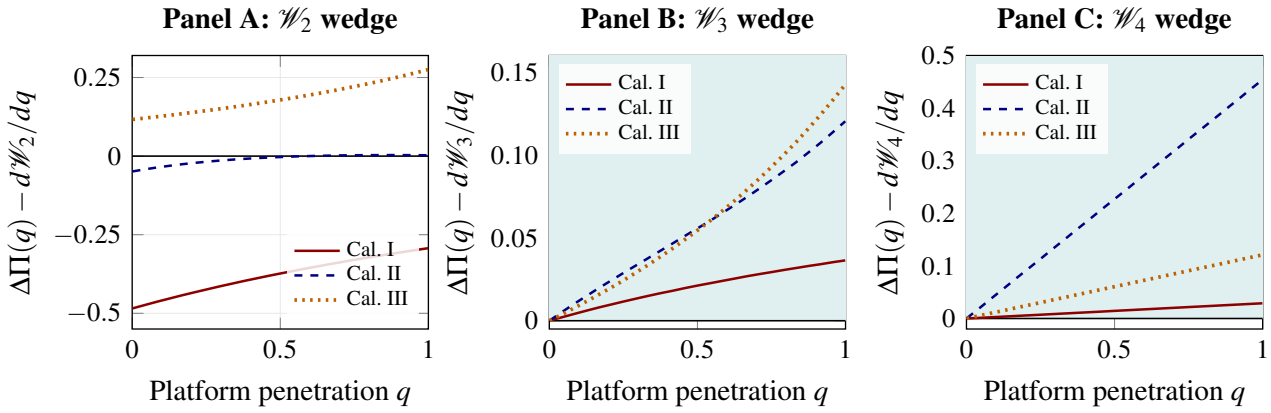
[8] See Appendix B for a numerical robustness check under an alternative logistic entry rule.

This criterion penalizes deviations of each type's signal loading from one (perfect tracking). 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 and the team-efficiency identity $d\mathscr{W}_1/dq = \Delta\Pi(q)$ does not apply.

The resulting wedge $\Delta\Pi(q) - d\mathscr{W}_2/dq$ includes indirect terms in $(k^A)'(q)$ and $(k^C)'(q)$, and its sign is therefore calibration-dependent. In contrast to Theorem 1, this specification does not deliver a determinate prediction for over- or under-adoption.

The contrast isolates the role of externalities. Under $\mathscr{W}_2$, all costs are internal to the chain population and are largely internalized in equilibrium. Over-adoption arises only when the planner accounts for costs borne outside the chain population, as in $\mathscr{W}_3$ and $\mathscr{W}_4$.

Figure 1: The misalignment wedge under three welfare specifications.



Notes: Each panel plots the misalignment wedge $\Delta\Pi(q) - d\mathscr{W}_k/dq$ as a function of platform penetration q, for three calibrations sharing $\delta = 1.5$ and $\gamma_k = 1$ but differing in the relative magnitudes of σ_A^2, σ_C^2, and σ_ξ^2. Calibration I: $(\sigma_\theta^2, \sigma_\xi^2, \sigma_A^2, \sigma_C^2) = (2.79, 0.05, 0.62, 0.27)$; Calibration II: $(1.23, 0.62, 0.88, 0.47)$; Calibration III: $(2.48, 0.21, 0.78, 1.46)$. Panel A shows the wedge under fundamental-tracking welfare $\mathscr{W}_2$: it takes negative values throughout for one calibration, positive values throughout for another, and crosses zero in the interior for the third. The same primitive parameters therefore deliver no determinate prediction about whether the platform is over- or under-adopted relative to the $\mathscr{W}_2$ planner (Remark 4). Panel B shows the wedge under aggregate-volatility welfare $\mathscr{W}_3$: it is strictly positive across all three calibrations and all $q \in (0, 1]$, with magnitude depending on δ through equilibrium loadings (Theorem 1, part i). Panel C shows the wedge under decided-about-agent welfare $\mathscr{W}_4$: it is strictly positive, linear in q, and invariant in δ (Theorem 1, part ii). Panels B and C jointly establish that the over-adoption result is robust to the choice of welfare specification across two structurally distinct externalities.

The contrast across the three panels of Figure 1 isolates the role of the welfare specification. 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 yield 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 borne outside the chain population, platform adoption is excessive.

4.6 Strategic Interaction versus Shared Error

Strategic interaction and shared processing play distinct roles in the model. The parameter δ governs how firms respond to one another and therefore 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 forces operate through separate channels.

Remark 5 (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 possible. The variance σ_ξ^2 enters private incentives through V_A and π_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. Under $\mathscr{W}_4$, the wedge at any fixed q depends only on adoption and signal-structure primitives. Strategic interaction determines equilibrium selection, while shared processing determines whether each equilibrium is distorted.

Remark 5 implies that the externality results of Theorem 1 do not require strategic interaction: under either $\mathscr{W}_3$ or $\mathscr{W}_4$, the misalignment wedge is strictly positive whenever $\sigma_\xi^2 > 0$ and $q > 0$, including at $\delta = 0$. The Bayesian entry game is nonetheless retained for four interrelated reasons. First, it provides a rigorous control environment. The team-efficiency identity of Proposition 3 uses the game structure to prove that chains fully internalize their strategic effects on one another, thereby isolating the wedge as a purely *external* cost rather than a within-population coordination failure. Second, δ governs equilibrium selection, generating the tipping, fragmentation, and path-dependence outcomes documented in Section 5, which are phenomena a non-strategic model cannot produce. Third, while the *sign* of the wedge is invariant to δ, its *magnitude* under $\mathscr{W}_3$ depends on equilibrium loadings, making strategic interaction relevant for calibrating the welfare loss under $\mathscr{W}_3$ across competitive environments. Fourth, the dynamic extensions in Section 8 build directly on the within-period equilibrium; the game structure is therefore necessary to derive

the temporal-amplification and cheap-switching results. This isolation of an external cost requires a welfare criterion that accounts for outside parties; under $\mathscr{W}_2$, which captures only chain-level forecast accuracy, the wedge sign is calibration-dependent (Remark 4).

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.

This distinction is 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. If $\sigma_\xi^2 > 0$, however, Theorem 1 implies that adoption remains excessive. The distortion arises entirely from the shared component.

When $\delta < 0$, the same conclusion 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 is unique. Differentiation incentives reduce equilibrium adoption, but do not eliminate the wedge. The equilibrium lies below the level sustained under strategic complements, yet remains above the social optimum.

When $\sigma_\xi^2 = 0$, the distinction reverses. The platform reduces idiosyncratic noise without introducing correlated error, so $\mathscr{W}_3$ and $\mathscr{W}_4$ coincide with $\mathscr{W}_1$. Even if $\delta > 0$ and multiple equilibria arise, Proposition 3 implies that all equilibria are efficient. The tipping outcome is then a feature of equilibrium selection rather than a welfare distortion.

Taken together, coordination and shared processing are separable mechanisms. Strategic interaction determines whether firms align their decisions; the shared component determines whether such alignment is costly. Multiple equilibria do not by themselves indicate a market failure, and uniqueness 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 of $\Delta\Pi(q)$ and therefore whether multiplicity is possible. The vertical axis, governed by σ_ξ^2, determines both the level of $\Delta\Pi(q)$ and whether a given level of adoption generates costs for external parties. The tipping 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 coordination motives sustain adoption even though the platform is unattractive in isolation. If $\sigma_\xi^2 = 0$, this outcome is efficient. It becomes a policy concern only when $\sigma_\xi^2 > 0$, because high adoption increases 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. The shared component instead generates correlation in errors, so comovement persists after conditioning on observable conditions. In particular, synchronized shifts following platform updates, absent coincident changes in fundamentals, identify the shared-error channel. Because the wedge under $\mathscr{W}_4$ is independent of δ, the externality can be quantified without identifying the strategic environment.

Finally, the decomposition clarifies the policy margins. Policies that alter the strategic environment operate through δ and affect equilibrium selection. Policies that alter the structure of shared processing operate through σ_ξ^2 and affect the magnitude of the externality directly. Holding q fixed, an increase in σ_ξ^2 raises the welfare wedge under either criterion. Section 6 shows that this margin can be interpreted through an effective concentration index $H(K,\rho)$ that captures the probability that adopters share the same source of error. The two policy levers are complementary.

4.7 Empirical Fingerprints

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—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.

Three 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. A fourth qualification is specific to the scaling prediction. Proposition 4 states that aggregate entry shifts by $K_\xi(q)\,\Delta\xi$; the cross-sectional scaling prediction (i.e., that the magnitude of the response increases in measured adoption) requires further that $\Delta\xi$ be cross-sectionally constant, or at least uncorrelated with q across the units in which adoption is measured. In settings where platforms update differently for different subscriber tiers, or where the size of an update is itself a function of adoption (e.g., re-estimation on subscriber-generated data), $\Delta\xi$ may co-vary with q, breaking the clean separation between the loading and the shock. The scaling test is therefore best applied to update events that affect adopters uniformly, or to specifications in which differential update intensity can be controlled for using platform-side disclosures. Recent regulatory discourse has begun to engage with the platform-event-as-instrument interpretation in financial-stability contexts, with the Bank of England's Financial Policy Committee noting that AI-induced correlation across market participants is a systemic-stability concern (Bank of England Financial Policy Committee 2025); the analogous infrastructure for retail intelligence platforms remains underdeveloped, and the financial-stability analogy is therefore suggestive rather than directly portable. 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 determine the number and stability of fixed points and yield a small set of regimes that organize the model's predictions. The key objects are the signs of $\Delta\Pi(0)$ and $\Delta\Pi(1)$. The first admits a simple characterization.

Lemma 1 (δ-independent threshold)**.** *The condition* $\Delta\Pi(0) = 0$ *holds if and only if* $\sigma_\xi^2 = \sigma_C^2 - \sigma_A^2$, *independent of* δ.

At $q = 0$, a marginal adopter faces the same aggregate as a non-adopter, 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 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.

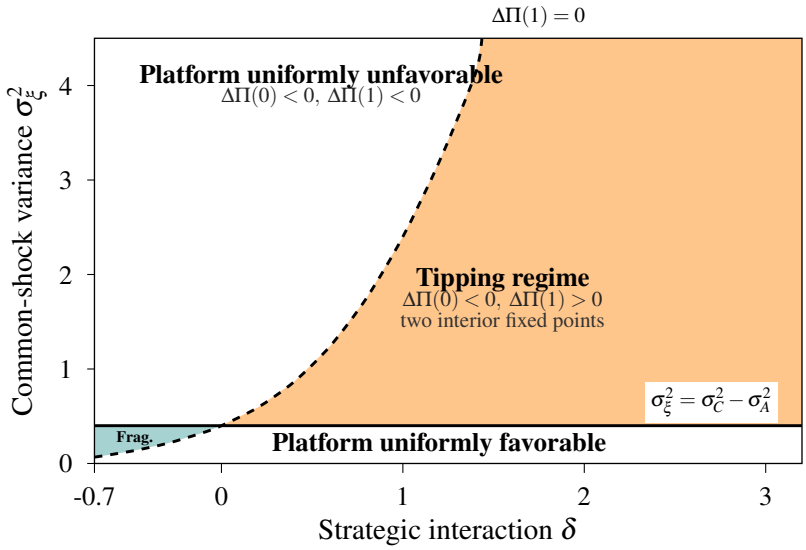
The second boundary, defined by $\Delta\Pi(1) = 0$, depends on δ. As coordination motives strengthen, deviation from the platform becomes more costly, so chains tolerate a larger shared component before switching away. The corresponding threshold in σ_ξ^2 is therefore increasing in δ. Numerical evaluation shows that this locus is monotone and convex and intersects 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 vary across regions, the misalignment result of Theorem 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, not in the direction of the welfare distortion. Two cases are particularly informative: the tipping regime under agglomeration and the fragmentation regime under business stealing.

5.1 The Tipping Regime

When $\delta > 0$, $\Delta\Pi(0) < 0$, and $\Delta\Pi(1) > 0$, the fixed-point condition $q^* = F(\Delta\Pi(q^*))$ admits at least one interior solution by the intermediate-value theorem. When F has sufficient curvature relative to $\Delta\Pi$ — concretely, when $F'(\Delta\Pi(q))\Delta\Pi'(q) > 1$ on a non-degenerate interior interval

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

— the equation admits three solutions: a stable low-adoption fixed point q_L^* near zero, a stable high-adoption fixed point q_H^* near one, and an unstable interior threshold $q^\dagger$ separating their basins of attraction.[9] Figure 3 illustrates such a configuration for a logistic F with steepness parameter λ large enough for the curvature condition to hold.

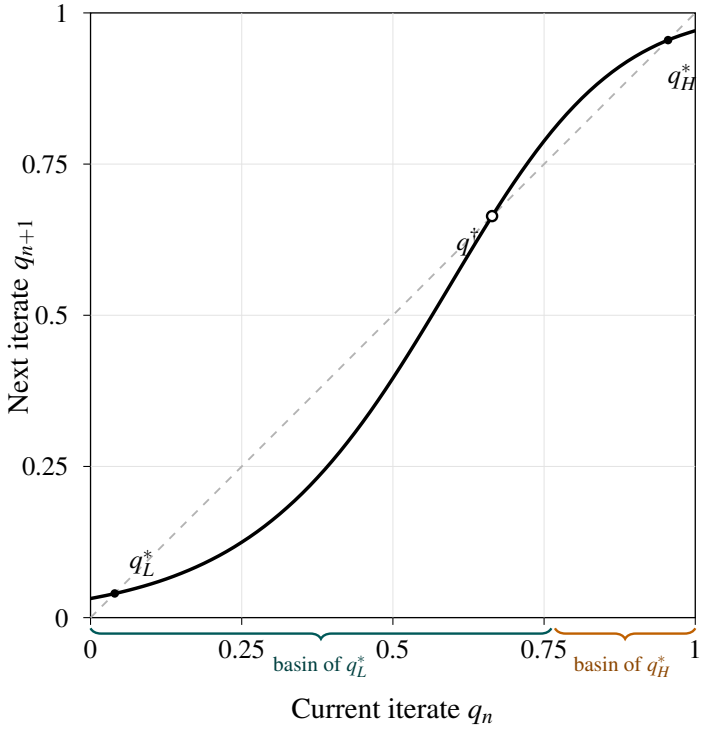
In this region, platform adoption is privately unattractive in isolation but becomes attractive once adoption is sufficiently widespread. At low penetration, exposure to the shared component is limited, and chains prefer the precision of custom research. At high penetration, aggregate entry is driven 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.

This structure parallels classic models of technology adoption under increasing returns. In

[9] Stability refers to best-response iteration on the adoption mapping: a fixed point q^* is stable if $|F'(\Delta\Pi(q^*))\,\Delta\Pi'(q^*)| < 1$, where F' is evaluated at $\Delta\Pi(q^*)$ and $\Delta\Pi'$ at q^*.

Arthur (1989), small historical events are amplified through positive feedback, so that early advantages become self-reinforcing and the market locks in to a single technology. The mechanism here is similar in its equilibrium structure but differs in its source. In Arthur's model, lock-in arises from increasing returns to adoption. Here it arises from strategic interaction operating through a shared processing layer. The inefficiency is also distinct. In the present setting, coordination can be privately optimal, yet adoption remains socially excessive because the shared component generates correlated errors that firms do not internalize.

Figure 3: The tipping regime as a fixed-point diagram.



Notes: The figure plots the next iterate $q_{n+1} = F(\Delta\Pi(q_n))$ against the current iterate q_n under best-response iteration on the static adoption map. The index n counts steps of this iteration, a graphical device used to locate the static fixed points of $q^* = F(\Delta\Pi(q^*))$ and assess their stability; it does not correspond to calendar time, which is reserved for the dynamic model of Section 8. Fixed points lie at intersections with the 45-degree line. Three intersections obtain: a stable low-adoption fixed point $q_L^* \approx 0.04$ and a stable high-adoption fixed point $q_H^* \approx 0.96$, separated by an unstable interior fixed point $q^\dagger \approx 0.66$ 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 iterate: the wide left basin captures industries that begin from any moderate level of platform adoption, while only initial values above $q^\dagger$ tip 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: $\delta = 0.7$, $\sigma_\theta^2 = 1$, $\sigma_\xi^2 = 0.6$, $\sigma_A^2 = 0.7$, $\sigma_C^2 = 0.85$, with logistic adoption distribution $F(x) = (1+e^{-\lambda x})^{-1}$ and $\lambda = 25$. This calibration sits inside the admissibility region $\mathscr{D}$ of Proposition 1, with $\Delta(q) > 0$ throughout $[0,1]$ and headroom 0.74 below the upper bound of $\mathscr{D}$.

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 $\mathscr{W}_3$ 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 distinction between marginal and global comparisons is important. The theorem implies $d\mathscr{W}_k/dq < \Delta\Pi(q)$ at every interior q, so each fixed point lies above its local social optimum. The global comparison between q_L^* and q_H^* instead depends on welfare levels, $\mathscr{W}_k(q_L^*)$ versus $\mathscr{W}_k(q_H^*)$, and therefore on the relative magnitude of coordination gains and external costs.

Under $\mathscr{W}_3$, the externality scales with $K_\xi(q)^2$ and therefore inherits curvature from δ. For moderate δ, the low-adoption equilibrium is welfare-superior because the externality dominates the coordination gain, while for sufficiently large δ the ranking can reverse. Under $\mathscr{W}_4$, the externality scales with q^2 independently of δ, so the high-adoption equilibrium is more likely to be welfare-dominated. A full global ranking is beyond the scope of the present analysis; the marginal wedge remains positive at any interior equilibrium.

A regulator seeking to reduce synchronization must therefore consider both the marginal wedge and the ranking of equilibria. When the low-adoption equilibrium 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 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 mirrors that of the tipping regime. By Theorem 1, the unique fragmentation equilibrium exhibits over-adoption under both $\mathscr{W}_3$ and $\mathscr{W}_4$. Strategic substitution limits the level of adoption but does not eliminate the externality.

The two regimes therefore differ in the level of q^*, not in the direction of the welfare distortion. For all $\delta \in \mathscr{D}$, strategic forces affect how much adoption occurs but not whether it 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 does not determine the size of the wedge. Allowing multiple competing platforms introduces a second policy margin: the diversity of processing technologies. This margin affects the size of the externality without changing its direction, and is the focus of this section.

Two related concerns motivate the extension: vendor concentration in retail intelligence markets, and the distinction between substantive and cosmetic diversity. Multiple platforms may differ in interface or branding while relying on the same underlying data or models, and the multi-platform setup below distinguishes the two cases.

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 $k \neq k'$. [10] The single-platform model corresponds to $K = 1$.

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

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

[10] The bare ρ in this section denotes cross-platform correlation in shared shocks. It differs from the projection coefficient $\pi_A = \alpha_\theta^A + \alpha_\xi^A$ in Proposition 1 and from the autocorrelation $\rho_x(q,\phi)$ in Section 8. These objects are distinct and do not appear together in a single derivation.

Under symmetric adoption, $q_k = q/K$, this expression reduces 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 ($\rho = 1$), and 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 $\delta \in \mathscr{D}$, $K \geq 1$, 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 parallel decompositions under the two welfare criteria. Under $\mathscr{W}_3$, the platform-driven component of aggregate variance factors as $K_\xi(q)^2\sigma_\xi^2$ times $H(K,\rho)$. Under $\mathscr{W}_4$, the same index governs expected correlation across adopters: with probability $1/K$ two adopters share a platform and their errors are perfectly correlated, while with probability $1 - 1/K$ they use different platforms and correlation is ρ. 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. Although the two externalities act on different objects, 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.*

When platforms are independent, increasing K reduces the externality proportionally by lowering the probability that adopters share the same error. When platforms share a common core, increasing K has no effect, as adopters remain exposed to the same underlying shock regardless of the number of intermediaries. Vendor diversity therefore matters through the independence of processing technologies rather than through the number of providers.

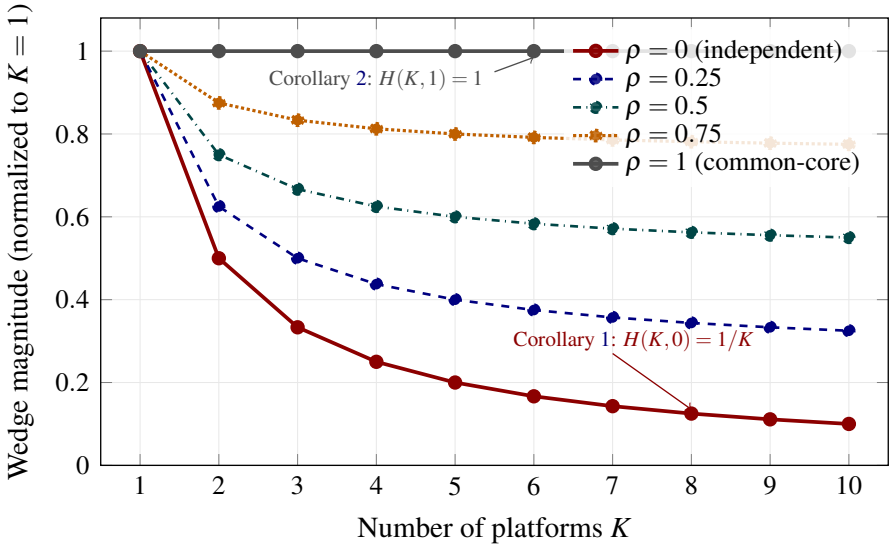
This distinction is directly observable in retail intelligence markets. Many platforms rely on common upstream data (e.g., mobile-device location pings or standardized demographic projections), so nominal variety masks high ρ. By contrast, platforms built on distinct data sources, such as transaction records, mobility traces, or satellite imagery, approximate the low-ρ regime. Bommasani et al. (2022) document an analogous pattern for foundation models, where fine-tuned variants of a common base model exhibit limited independent variation relative to structurally distinct models. The relevant margin is therefore the diversity of underlying inputs and methods, not the number of platforms.

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, while 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 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 (e.g., adoption taxes or

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.

capital requirements tied to platform exposure) operate on the first margin. Policies that reduce H (e.g., procurement diversification or data-source requirements) operate on the second.

The two margins act on distinct components of the externality. Reducing q lowers exposure but leaves the correlation structure unchanged, while reducing H lowers correlation but leaves adoption above the social optimum. Joint intervention therefore achieves the optimum with lower distortion than either margin alone. Policies targeting platform structure complement, rather than substitute for, policies targeting adoption.

Remark 6 (Asymmetric vendor shares)**.** The analysis above assumes symmetric vendor shares $q_k = q/K$. In practice, shares are uneven. With general shares, the variance decomposition 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 the qualitative results carry over: strict positivity, the distinction between substantive and cosmetic diversity, and the separation of policy margins. The symmetric case is used for transparency in comparative statics over K.

7 Endogenous Processing Noise

The analysis so far treats σ_ξ^2 as a primitive of the information structure. We now allow the variance of the shared component to depend on platform adoption, $\sigma_\xi^2 = \sigma_\xi^2(q)$, and examine how the welfare results extend.

Platform forecast quality is plausibly endogenous to usage. Capacity strain provides one channel: shared inference resources and load-induced congestion can degrade 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 increasing influence shifts incentives toward optimizing engagement or retention rather than predictive accuracy. Despite their different microfoundations, these mechanisms share a reduced-form implication that the variance of the shared component increases with adoption.

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 condition requires that adoption weakly amplifies shared error, while the second requires that this amplification is weakly convex. The constant case $(\sigma_\xi^2)' \equiv 0$ nests the baseline model. 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 does not alter the structure of the within-period equilibrium, but changes 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 channels.

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 captures endogenous amplification. An additional adopter increases the magnitude of the shared component itself, raising the distortion faced by all adopters. Its strength is governed by the slope $(\sigma_\xi^2)'(q)$ and by the share of signal variance attributable to non-shared components.

Remark 7 (Local versus global convexity)**.** Proposition 6 is a pointwise statement: under $(\sigma_\xi^2)'(q) \geq 0$, the wedge is strictly positive at each $q \in (0,1]$ and strictly larger than under constant σ_ξ^2, but $d\mathscr{E}_4/dq$ need not be globally non-decreasing. Global monotonicity requires the additional condition $(\sigma_\xi^2)''(q) \geq 0$ together with restrictions on the relative magnitude of $(\sigma_\theta^2+\sigma_A^2)$ and $\sigma_\xi^2(q)$.

7.3 Implications for Adoption Dynamics and Identification

Endogenous noise affects both adoption dynamics and identification. 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 implies a unique equilibrium. The phase boundaries of Section 5 therefore inherit curvature from the degradation profile, though the qualitative typology remains unchanged.

For the identifying test of Section 4.7, endogenous noise strengthens the scaling prediction. 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 underlying shock also grows with q, so that both the number of firms responding and the magnitude of each response increase with adoption. High-adoption environments should therefore exhibit more synchronized and larger responses to platform updates, with the precise relationship determined by the shape of $\sigma_\xi^2(q)$.

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 model is revised. The static analysis captures correlation within a period but is silent on whether outcomes cluster over time. The dynamic extension below isolates this temporal dimension while preserving the closed-form equilibrium of the static model.

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 do not internalize intertemporal effects, so the dynamics arise entirely from the persistence 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 $\phi \in (0,1)$.

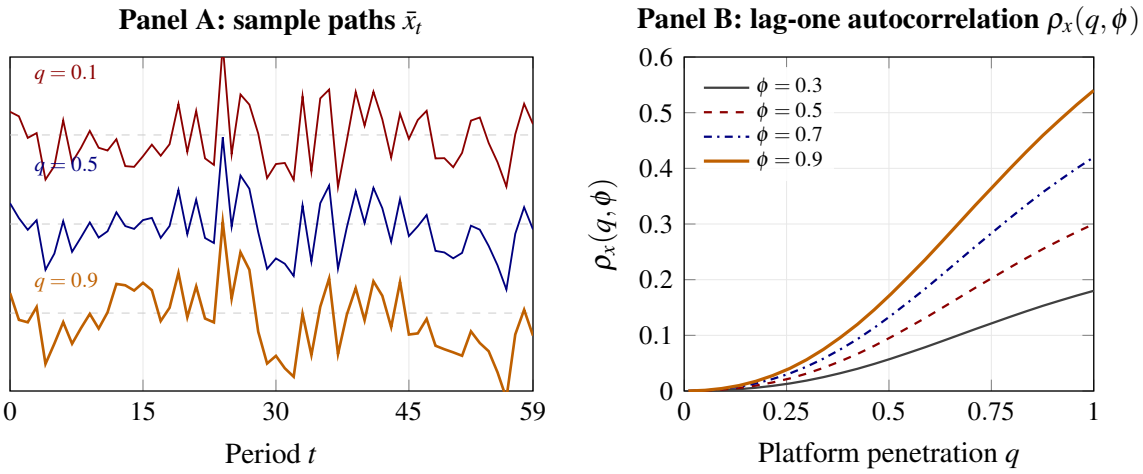
Temporal dependence is driven entirely by the shared component. The fundamental shock contributes to cross-sectional dispersion but not to persistence. By contrast, the shared component introduces a persistent factor that enters aggregate behavior with weight $K_\xi(q)$.

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 both the dispersion of outcomes within a period and their evolution across periods. Even when fundamental shocks are i.i.d., the economy inherits persistence in aggregate behavior through the shared processing layer. High-adoption environments therefore exhibit longer runs of same-sign deviations in aggregate entry, as shocks to the platform propagate over time.

Figure 5 illustrates the mechanism. Panel A plots a single realization of $\bar{x}_t$ across three penetration levels under identical shocks: at low adoption the path resembles independent noise, while at high adoption it exhibits longer same-sign excursions. Panel B plots the population autocorrelation $\rho_x(q,\phi)$, showing that persistence 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 clarifies the model's implications. In the static environment, shared processing increases the variance of aggregate entry through $K_\xi(q)^2\sigma_\xi^2$. 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 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., reflecting 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 (i.e., aggregate volatility and correlated evaluation). The dynamic extension introduces a third channel operating through time. A marginal increase in adoption raises $K_\xi(q)$ and therefore 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. 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. Persistence reinforces the over-adoption result established in Section 4.

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 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 exhibit rapid technological change and frequent displacement of incumbents, leading firms to reassess their 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. That interpretation requires qualification, however, once the level of adoption and the speed at which it is reached are distinguished.

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 is 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. Here the equilibrium level of adoption is socially excessive, so faster convergence accelerates the accumulation of an externality that firms do not internalize. Switching frictions therefore play a dual role: they impose private costs while moderating the speed at which the industry transitions into states with higher uninternalized social cost.

Policy timing provides an additional consideration. Regulatory responses to platform-driven systemic risk (e.g., vendor-diversification requirements, data-portability rules, disclosure regimes, or antitrust scrutiny) operate on multi-year horizons. If convergence toward high adoption is rapid,

lower switching costs compress the window during which such interventions can take effect. The dynamic margin is therefore distinct from the static margins identified in Section 6: policies that reduce q or $H(K,\rho)$ affect 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 revise decisions with probability λ (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. Similar structures arise in models of technology adoption and diffusion with heterogeneous or randomly timed adjustment (Jovanovic and Lach 1989).

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

The dynamic extension 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). Here, the same force operates through exposure to the shared component. Faster adjustment increases the rate at which the economy accumulates exposure to this component, so lower switching costs 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* $\delta \geq 0$ *and adoption*

evolves according to the recursion above from an initial level $q_0 \in [0, q^*)$ in the basin of attraction of a stable static equilibrium q^*, with the static iteration $q \mapsto F(\Delta\Pi(q))$ approaching q^* monotonically from below. 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 8 (Beyond $\mathscr{W}_4$ and the strategic-substitute case)**.** 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. Numerical evaluation across 5,000 parameter draws shows that it is monotone increasing in q

in the vast majority of cases. The qualitative conclusion (i.e., that cheap switching accelerates the accumulation of systemic cost) therefore extends beyond the specific welfare specification considered here.

When $\delta < 0$ and the static iteration oscillates around the stable fixed point rather than approaching it monotonically, the path-monotonicity argument in part (i) of the proposition does not apply directly. Two cases warrant separate treatment. If the oscillation is mild enough that q_t is below q^* at every t (which holds when $|F'(\Delta\Pi(q^*))\Delta\Pi'(q^*)|$ is small), the proposition extends to $\delta < 0$ via the same argument. If the oscillation crosses q^*, the cumulative-wedge claim in part (ii) holds in expectation across the oscillation cycle but not pointwise. The qualitative direction of the result—that cheap switching amplifies the accumulation of welfare cost relative to the long-run static equilibrium—is preserved in both cases. A complete characterization in the strategic-substitute regime is left to future work.

Remark 9 (Endogenous switching probability)**.** The analysis treats λ as exogenous, capturing the probability that a chain re-evaluates its adoption decision in a given period. In practice, λ is shaped by upstream factors (i.e., platform pricing, integration costs, switching infrastructure, and competitive dynamics among providers). Below-cost pricing, standardized data formats, and managed-analytics services all tend to increase λ from the chain's perspective.

A complete model would endogenize λ and solve for a vertical equilibrium. Proposition 8 is modular with respect to this extension: any upstream model that delivers a comparative static for λ can be combined with our result to assess downstream welfare. In particular, interventions that reduce switching frictions increase λ and therefore amplify the dynamic externality.

Related work studies the upstream incentive to produce shared data inputs. Fang and Kim (2026) show that data-neutrality regulations can reduce a platform's incentive to refine its data, generating a trade-off analogous to the one identified here. Both settings highlight that policies which broaden access may alter upstream supply in ways that offset their direct benefits.

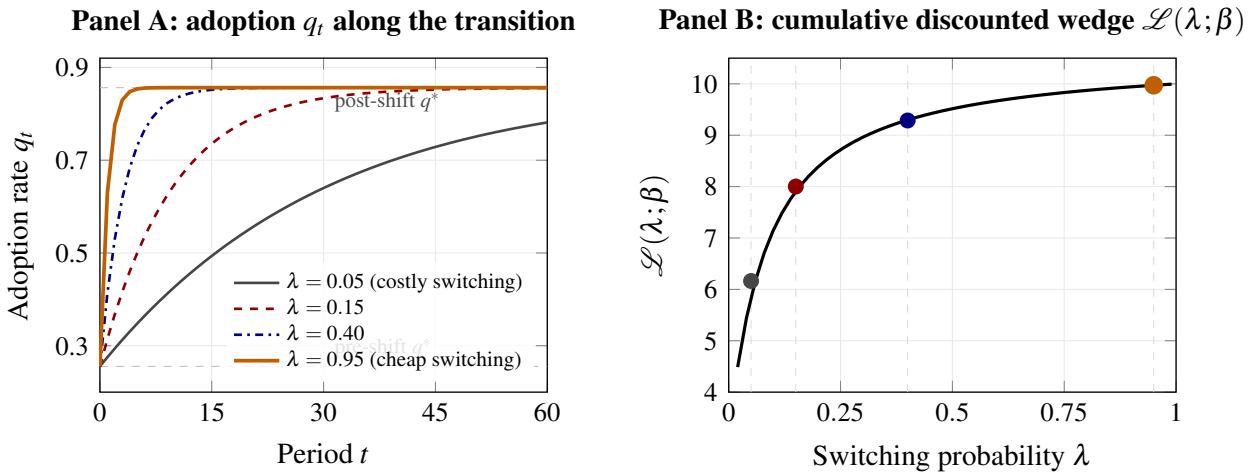
The policy implication contrasts with the standard view of switching costs. In most settings, switching frictions are treated as inefficiencies (i.e., they hinder adoption of superior technologies and create lock-in). In the present environment, some friction can be socially beneficial. Slower adjustment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

forces (i.e., vendor entry, diversification of data sources, or regulatory intervention) to take effect.

This role is analogous to adjustment frictions in models of financial stability. Just as 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: 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. The vertical units of Panel B are set by the $\mathscr{W}_4$ welfare scale (which absorbs the multiplicative constant $\gamma_4 = 1$ used here); the substantive content lies in the *ratio* of $\mathscr{L}$ across λ values, which is invariant to this scale and approximately doubles between $\lambda = 0.05$ and $\lambda = 0.95$.

Figure 6 illustrates the comparative static. The left panel plots adoption trajectories follow-

ing a regime shift in which the platform becomes more attractive (e.g., 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.

Shared intelligence therefore affects both the level and the temporal structure of aggregate fluctuations, and the speed of transition toward a shared-processing-intensive regime is itself a welfare-relevant object. Policies that reduce exposure to the shared processing layer (e.g., vendor diversity, concentration limits, or interoperability rules as in Section 6) attenuate both contemporaneous comovement and persistence. Timing matters as well: cheap switching shortens the window in which such interventions can take effect.

9 Case Study of Restaurant Site Selection Decision Support

Proposition 5 implies that the welfare relevance of vendor diversification depends not on the number of nominal vendors, but on the effective concentration index $H(K,\rho)$. This index combines two objects that are conceptually distinct: concentration in vendor adoption and correlation in the underlying forecasting technologies. A market may therefore appear diversified at the interface layer, with many competing vendor brands, while remaining highly concentrated in the data and modeling layer that determines forecast comovement. Whether a market is closer to substantive diversification ($H \approx 1/K$) or cosmetic diversification ($H \approx 1$) is consequently an empirical question about the structure of the technology stack rather than the count of observable vendors.

This section illustrates that distinction using the market for restaurant site-selection analytics[11] The objective is not to point-estimate H, but to show how the primitives entering the index can be disciplined using public disclosures, SEC filings, vendor case studies, and observed overlap in upstream data providers. The exercise clarifies a broader implication of the model: vendor competition at the front end need not imply meaningful diversification in the forecasting technology used by adopters.

The calibration has two components. The first is share concentration across downstream vendors, summarized by the Herfindahl term $\sum_k q_k^2$. This term captures the probability that two

[11] Appendix C reports the sources used to construct the vendor-share estimates and the qualitative evidence on upstream data overlap.

adopters using analytics platforms rely on the same vendor. The second is structural correlation across vendors conditional on adopters using different platforms, summarized by ρ. This term captures the extent to which distinct vendors nevertheless inherit common shocks from shared data inputs, common model infrastructure, or common upstream suppliers. The two margins enter H separately and admit separate empirical discipline. We therefore consider them in turn.

9.1 The Vendor Stack

The downstream vendor market is concentrated. Trade press, vendor disclosures, and SEC filings consistently identify six firms as the principal providers of site-selection analytics for multi-unit restaurant chains: Placer.ai, Tango Analytics, Buxton, eSite Analytics (acquired by Kalibrate in 2021), SiteZeus, and IdealSpot (acquired by Plotr in 2023). Other firms contribute to the broader location intelligence ecosystem, but play different roles. Esri's ArcGIS Business Analyst operates primarily as infrastructure rather than as a chain-facing site-selection platform, while a long tail of smaller specialists appears to serve single-unit operators, local businesses, or narrower use cases outside the multi-unit chain market considered here.

A second feature of the same market is more important for the model. The front-end vendors draw on a common upstream data layer. INRIX, the dominant supplier of roadway-traffic volume data, has disclosed that eight of the largest site-selection companies license its data, including Buxton, eSite, IdealSpot, LocateAI, Megalytics, Property Capsule, SiteZeus, and Tango, and integrate those feeds into otherwise proprietary scoring models.[12] This fact foreshadows the central interpretation of the case study: the six vendor brands observed by adopters do not correspond to six independent forecasting technologies.

9.2 Share Concentration

Direct measurement of vendor share is not available in this market. None of the six principal vendors is publicly traded as a pure-play location-intelligence firm, and none reports chain counts or revenue by restaurant-site-selection vertical. We therefore construct an approximate share measure using a revenue-proxy approach. First, we assemble a peer revenue pool from the best avail-

[12] See *QSR Magazine*, "INRIX Introduces New Dataset in Site Selection Suite" (qsrmagazine.com), and the INRIX/eSite Analytics case study (inrix.com).

able public proxies: reported ARR, third-party revenue estimates from Growjo, Owler, Latka, and Craft, employee counts, funding and valuation disclosures, and customer-count claims. Placer.ai, for example, disclosed crossing $100M ARR in February 2024 with more than 4,300 customers [13] Second, we validate scale and market relevance against SEC filings and earnings-call disclosures of public-company customers and competitors. Third, we apply a retail-chain exposure adjustment that scales each vendor's total revenue to the portion plausibly attributable to restaurant chain analytics rather than adjacent markets such as commercial real estate, civic and economic development, healthcare, and lease administration.

Table A3 in Appendix C reports the resulting adjusted-share estimates. The implied Herfindahl component is

$$\sum_k q_k^2 \approx 0.628^2 + 0.141^2 + 0.097^2 + 0.077^2 + 0.055^2 + 0.002^2 \approx 0.430.$$

The substantive-diversification benchmark, $1/K \approx 0.17$, would be attained only if adopters were allocated equally across the six platforms. Observed adoption is far from that benchmark. Even before accounting for cross-vendor correlation, share concentration alone raises the Herfindahl component to roughly 2.5 times the equal-share level. The reason is not simply that the market has few vendors; it is that adoption is highly asymmetric. Placer.ai accounts for a majority of adjusted revenue, while the remaining vendors split the residual.

This concentration is the first source of effective common exposure in the model. If two adopters draw vendors at random from the adjusted share distribution, the probability that they rely on the same platform is already much higher than it would be under equal vendor shares. Thus, even a counterfactual market in which each vendor used a fully independent forecasting technology would not deliver the substantive-diversification benchmark. The second question is whether adopters using different vendors are in fact insulated from common forecasting shocks.

9.3 Structural Correlation

The second component of H is structural correlation across distinct vendors. This margin captures whether different platforms generate independent forecasts conditional on being separately owned

[13] *TechCrunch*, "Placer.ai boosts valuation to $1.5B after quietly raising another $75M" (techcrunch.com).

and separately branded. In the restaurant site-selection market, the answer depends on the upstream data layer. The evidence indicates substantial overlap in the inputs used to construct trade areas, measure traffic, characterize nearby demand, and score potential sites.

Four categories of upstream inputs are especially relevant. First, roadway-traffic data from INRIX is licensed by multiple major site-selection platforms. Second, point-of-interest data used for competitor mapping and trade-area definition is concentrated among providers such as SafeGraph (acquired by Dewey in 2024) and Foursquare. Third, mobility and foot-traffic data come from a small set of suppliers, including Placer.ai itself, Unacast, and historically SafeGraph's mobility products. Fourth, demographic and psychographic segmentation products, such as Experian Mosaic, Claritas PRIZM, and Spectra, provide customer-profile inputs that are themselves built on overlapping consumer data sources. Vendors differ in the proprietary models they build on top of these inputs. They differ much less in the core inputs themselves.

Placer.ai is particularly informative because it occupies both sides of the stack. It is a downstream platform sold to chains, but it also functions as an upstream data provider to adjacent businesses and to other participants in the location-intelligence ecosystem. Life360's 10-K filings disclose a data-licensing relationship with Placer.ai under which Life360 contributes panel data; data-insights revenue, including the Placer.ai relationship, represented 7% of Life360's 2024 revenue and 10% of its 2022 revenue.[14] The upstream panel that feeds Placer.ai's models is therefore itself assembled through data-broker relationships with consumer-app publishers, and the resulting visit data can then enter the broader analytics ecosystem. In this sense, the dominant front-end vendor is also part of the back-end infrastructure on which competing analytics products may rely.

The implication for the model is direct. Two chains that contract with different vendors are not necessarily using independent forecasting technologies if both vendors rely on INRIX traffic data, SafeGraph or Foursquare point-of-interest data, Placer.ai visit panels, and common demographic segmentation products. They are using differently packaged front ends over a substantially shared back end. In the notation of the model, the shared processing shock ξ corresponds to common errors, revisions, or methodology changes propagated through the upstream data layer: mobility-panel reweighting, POI taxonomy revisions, demographic forecast updates, or related changes inherited across vendor models.

[14] Life360 Form 10-K for fiscal year 2024 (sec.gov, accession 158176025000008).

The available data do not point-identify ρ. They do, however, make the independence benchmark implausible. A value near zero would require not only separate vendor ownership, but also largely independent upstream data inputs and independent processing pipelines. The observed structure of the market points in the opposite direction. We therefore treat $\rho \in [0.5, 0.7]$ as a plausible range for the calibration. The upper end corresponds to synchronous propagation of dominant upstream shocks, such as INRIX traffic revisions, SafeGraph taxonomy changes, or post-ATT mobility-panel reconstitution. The lower end allows meaningful vendor-specific variation from proprietary modeling choices, brand-level training data, loyalty-program inputs, and platform-specific data partnerships. The range is therefore not an estimate in the usual sense, but an order-of-magnitude discipline imposed by the observed structure of the technology stack.

9.4 Discussion

Combining the two margins gives

$$H = \sum_k q_k^2 + \rho \left(1 - \sum_k q_k^2 \right).$$

Using the adjusted shares in Table A3 and the plausible range $\rho \in [0.5, 0.7]$ gives

$$H \in \left[0.430 + 0.5(1 - 0.430),\ 0.430 + 0.7(1 - 0.430) \right] \approx [0.72, 0.83].$$

Table A4 reports H across the full grid of ρ values and across alternative share scenarios; the central $\rho = 0.6$ point delivers $H \approx 0.77$. Across the band, the resulting value is materially closer to the cosmetic-diversity bound $H = 1$ than to the substantive-diversity benchmark $1/K \approx 0.17$. The calibration therefore suggests that the restaurant site-selection market delivers only limited effective diversification, despite the presence of multiple competing vendor brands.

Welfare implication. Proposition 5 maps this calibration into welfare. In the multi-vendor extension, both externality wedges—the aggregate-volatility wedge $\mathscr{W}_3$ and the decided-about-agent wedge $\mathscr{W}_4$—scale linearly in H, with $H = 1$ corresponding to the single-platform benchmark of Theorem 1. The ratio $H \in [0.72, 0.83]$ therefore identifies the fraction of the single-platform wedge

that the calibration retains: nominal vendor competition at the front end recovers approximately 17 to 28 percent of the diversification benefit that would arise from moving all the way from a single shared technology to fully independent vendor technologies. The ratio interpretation is invariant to the choice of γ_3 or γ_4, which set the absolute scale of the wedge but not its proportional response to H; the case study therefore disciplines proportional welfare changes across vendor structures without requiring an independent calibration of γ.

Policy-margin elasticities. The decomposition of H separates two intervention margins. The first acts on downstream adoption shares, holding the upstream technology stack fixed. The second acts on cross-vendor correlation, holding downstream shares fixed. These margins correspond to different policy instruments and have different magnitudes in this market.

Consider first an intervention that equalized adoption across the six existing vendors while leaving the shared upstream data layer unchanged. At $\rho = 0.6$, this would reduce H from 0.77 to 0.67. The associated wedge reduction is approximately 14%. This counterfactual approximates the upper bound on what front-end procurement diversification, or antitrust scrutiny focused only on Placer.ai's downstream share, could achieve without changing the common input structure.

By contrast, an intervention that eliminated structural correlation while leaving observed vendor shares unchanged would reduce H from 0.77 to 0.43. The associated wedge reduction is approximately 44%. This counterfactual isolates the value of upstream diversification: reducing dependence on common traffic data, common POI taxonomies, common mobility panels, and common segmentation products. It corresponds more closely to data-portability rules, upstream-input diversification requirements, or vertical-integration scrutiny of the INRIX/SafeGraph/Placer.ai panel layer than to policies aimed only at the number of front-end vendors.

Finally, full substantive diversification—equal vendor shares and independent forecasting technologies—would reduce H to 0.17, implying a wedge reduction of approximately 78%. The two margins are therefore not substitutes. At the observed parameter values, the upstream correlation margin is larger than the downstream share margin, but neither margin alone eliminates the externality.

Heterogeneity across the two wedges. Although $\mathscr{W}_3$ and $\mathscr{W}_4$ scale identically with H, they differ in their dependence on the strategic-interaction parameter δ. By Theorem 1(ii), the decided-

about-agent wedge $\mathscr{W}_4$ is δ-invariant, while the aggregate-volatility wedge $\mathscr{W}_3$ is not. A regulator who places greater weight on aggregate volatility will therefore view the $H \approx 0.77$ calibration as especially consequential in settings where strategic interaction amplifies correlated forecasting errors. Restaurant-chain entry into shared trade areas is a natural example, because cannibalization and local competitive spillovers make the strategic-interaction term non-negligible. A regulator concerned instead with the decided-about-agent externality reads the same H through the δ-independent channel. The case study is consistent with both interpretations; it disciplines the common exposure term rather than the relative regulatory weight placed on the two wedges.

Generalization. More generally, the exercise shows why vendor proliferation is not a sufficient statistic for substantive diversification in algorithmic decision environments. The relevant object is the effective dimensionality of the upstream information structure. A market can appear competitive at the interface layer while remaining concentrated at the data and modeling layer that governs forecast comovement. The distinction matters because the two margins identified by H—concentration in adoption and correlation conditional on adoption—are different empirical objects and respond to different interventions.

The same logic applies, with sector-specific substitutions, to other markets in which firms route decisions through algorithmic intermediaries: commercial real-estate underwriting, consumer credit scoring, advertising auction bidding, retail inventory planning, and financial-portfolio risk management. The relevant vendors and data layers differ across settings, but the analytic structure does not. In each case, the welfare relevance of diversification depends on whether observable vendor variety corresponds to independent forecasting technologies or merely to differentiated front ends built on a common data stack.

10 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 and a robust misalignment result under two distinct externalities. 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.

The framework formalizes shared intelligence as an attentional infrastructure. Shared platforms standardize how firms process information, replacing idiosyncratic variation with a common component. Because this component is unrelated to fundamentals, it generates two external costs that firms do not internalize: aggregate-volatility costs borne by agents exposed to synchronized entry and correlated-evaluation costs borne by locations that are jointly screened by adopters. Coordination incentives are internalized in equilibrium, so these externalities are the sole source of inefficiency.

A natural interpretation 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 and welfare. A more complete account would endogenize this structure in an upstream environment with scale economies in data and computation, talent mobility, and feedback between usage and model improvement. Integrating such a model with the externality structure developed here would connect the upstream production of data and models to the downstream consequences of their diffusion and identify where policy must act.

A second simplification concerns the agents that bear the externalities. In the model, lenders, landlords, municipalities, and locations enter as passive recipients of aggregate volatility and correlated evaluation. In practice, these parties respond to anticipated synchronization, partially internalizing the externality through prices and policy. Endogenizing these responses would attenuate but not eliminate the wedge and would provide a richer account of how market and regulatory adjustments interact with shared processing.

A final extension concerns preemptive motives in platform adoption. Firms may adopt not only in response to current payoffs but also to shape future behavior of rivals. Introducing such motives would add a forward-looking channel through which firms influence the transition toward high-adoption regimes. This raises the question of whether preemption accelerates convergence beyond what switching frictions imply and whether strategic anticipation amplifies or mitigates the externalities identified here. Incorporating this channel would connect the shared-processing mechanism to the literature on dynamic oligopoly and strategic entry.

References

Abis S, Tang H, Bian B (2025) Breaking the data chain: The ripple effect of data sharing restrictions on financial markets. *Available at SSRN 5334566* .

Acemoglu D, Makhdoumi A, Malekian A, Ozdaglar A (2022) Too much data: Prices and inefficiencies in data markets. *American Economic Journal: Microeconomics* 14(4):218–256.

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.

Bian B, Huang Q, Li Y, Tang H (2024) Data as a networked asset. *Available at SSRN* 5183829.

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* .

Choi JP, Jeon DS, Kim BC (2019) Privacy and personal data collection with information externalities. *Journal of Public Economics* 173:113–124.

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

Fang H, Kim SJ (2026) Data neutrality, data supply, and market competition. Technical report, Penn Institute for Economic Research, Department of Economics, University of ....

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.

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

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

Ichihashi S (2021) The economics of data externalities. *Journal of Economic Theory* 196:105316.

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.

Kirpalani R, Philippon T (2020) Data sharing and market power with two-sided platforms. Technical report, National Bureau of Economic Research.

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.

Moll B (2026) The trouble with rational expectations in heterogeneous agent models: A challenge for macroeconomics. *The Economic Journal* 136(676):1173–1216.

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

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

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

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

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).

Online Appendix

Table of Contents

A Proofs **2**

A.1 Proof of Proposition 1 . . . 2
A.2 Proof of Lemma 1 . . . 3
A.3 Proof of Proposition 2 . . . 3
A.4 Proof of Proposition 3 . . . 4
A.5 Proof of Theorem 1 . . . 4
A.6 Proof of Proposition 4 . . . 6
A.7 Proof of Proposition 5 . . . 6
A.8 Proof of Proposition 6 . . . 7
A.9 Proofs of Corollaries 1 and 2 . . . 8
A.10 Proof of Proposition 7 . . . 8
A.11 Proof of Proposition 8 . . . 9

B Nonlinear Robustness of the Team-Efficiency Benchmark **11**

B.1 Equilibrium Computation . . . 11
B.2 Calibrations . . . 12
B.3 Results . . . 12
B.4 Implications . . . 13

C Additional Details for Case Study **15**

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 $\pi_A \equiv \alpha_\theta^A + \alpha_\xi^A = (\sigma_\theta^2 + \sigma_\xi^2)/V_A$. The maintained parameter restrictions are $\delta \in \mathscr{D}$ 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

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

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

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 $\Delta(1) = 1 - \delta\, \pi_A = (V_A - \delta(\sigma_\theta^2 + \sigma_\xi^2))/V_A$, which is positive iff $\delta < 1 + \sigma_A^2/(\sigma_\theta^2 + \sigma_\xi^2)$. Both conditions hold for any $\delta \in \mathscr{D}$. 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 $\delta\in\mathscr{D}$, $q\in(0,1)$. 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 symbolic verification are provided in the replication materials.

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, so the standard atomicity argument that would deliver $\partial\bar{x}/\partial k = 0$ for an individual deviation does not apply here. The cancellation must come from a different source.

The cancellation comes from 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 δP_g, where $P_g = \mathbb{E}[x_g \mid \mathscr{I}_g]$ is the chain-specific entry probability conditional on its information set. The planner's first-order condition with respect to k^A compares the direct effect on a chain of varying its own rule against the indirect effects on all other chains' payoffs. After netting out the direct effect, the integrand that survives is the *deviation* $\delta(P_g - \bar{x})$ (i.e., the gap between chain g's individual response and the population average). This integrand vanishes when integrated across the population, since $\int_{\mathscr{F}} P_g\,dg = \bar{x}$ by aggregate consistency:

$$\int_{\mathscr{F}} \delta\,(P_g - \bar{x})\,dg = \delta\,(\bar{x} - \bar{x}) = 0.$$

Because $\partial\bar{x}/\partial k^A$ is constant across g, this identity is exactly the GE feedback contribution in the planner's first-order condition, which therefore 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 $\delta \in \mathscr{D}$. 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 $\Delta(q)^3 > 0$,

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

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. 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 variance share 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 $F^{-1}(q_k^{soc}) = d\mathscr{W}_k/dq\big|_{q_k^{soc}}$ 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 variance-share 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 covariance of signal errors normalized by the unconditional signal variance V_A is

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

using $\mathrm{Var}(\xi_k) = \sigma_\xi^2$ and $\mathrm{Cov}(\xi_k, \xi_{k'}) = \rho\sigma_\xi^2$ for $k \neq k'$. 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 normalized covariance 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 canceling the common denominator $\Delta(q)$,

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

Direct differentiation and factoring shows the numerator of the derivative reduces, after collecting terms, to a strictly positive polynomial on the admissible region $\delta \in \mathscr{D}$, $q \in [0,1]$. The symbolic verification is provided in the replication materials. 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 tipping 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 $\delta < 0$ is treated separately in Remark 8, which restricts attention to oscillation regimes where the path-monotonicity argument extends without modification.

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 $t \geq 1$. □

B Nonlinear Robustness of the Team-Efficiency Benchmark

Proposition 3 establishes that, in the linear continuum environment of Section 2, the equilibrium adoption rate q^* coincides with the social optimum under welfare specification $\mathscr{W}_1$. Remark 3 notes that this identity relies on two features of the baseline environment: a continuum of chains and linear entry probabilities. With nonlinear entry probabilities outside the global-games regime, the cross-sectional identity in the proof acquires a correction term that does not vanish exactly, so the planner's first-order condition need not coincide exactly with the Nash condition.

This appendix tests whether the qualitative externality results of Theorem 1 nevertheless survive when the linear entry rule is replaced by a smooth nonlinear alternative. We adopt a logistic specification,

$$P_f = \Lambda\big(\kappa \cdot \big[\mathbb{E}_f(\theta) + \delta \mathbb{E}_f(\bar{x})\big]\big), \qquad \Lambda(x) = \frac{1}{1+e^{-x}},$$

where $\kappa > 0$ indexes choice sharpness and the linear global-games rule of the baseline model corresponds to the small-κ limit after rescaling.[15] The exercise covers three calibrations spanning the strategic-interaction parameter δ, four values of κ spanning a range of choice sharpness, and three interior values of the adoption rate q.

B.1 Equilibrium Computation

Under the linear rule, the aggregate $\bar{x}(\theta,\xi) = K_\theta(q)\,\theta + K_\xi(q)\,\xi$ is a linear function of the underlying shocks with known closed-form loadings (Proposition 1). Under the logistic rule, $\bar{x}(\theta,\xi)$ is genuinely nonlinear in (θ,ξ), so closed-form expressions for aggregate variance, per-type expected profits, and the welfare gradients no longer obtain. Equilibrium objects must therefore be computed by simulation.

We parameterize the equilibrium by the linear-projection loadings $(\bar{K}_\theta(q), \bar{K}_\xi(q))$, defined as the OLS coefficients of the simulated aggregate $\bar{x}$ on (θ,ξ). Given these loadings, each chain's entry-probability function is $P^t(s) = \Lambda(\kappa \cdot c^t \cdot s)$, where c^t is determined by the Bayesian projection of $\theta + \delta\bar{x}$ onto s. The loadings $(\bar{K}_\theta, \bar{K}_\xi)$ and the entry-probability functions are jointly determined by a fixed-point condition that we solve by damped iteration, initialized at the linear-equilibrium loadings of Proposition 1.

Welfare gradients are computed by centered finite differences in q with step $h = 0.05$, using common random numbers across the $q \pm h$ evaluations to control Monte Carlo noise in the derivative. Aggregate variance and per-type profits are computed from Monte Carlo simulation over 200,000 (θ, ξ, u^A, u^C) draws per cell. The $\mathscr{W}_3$ externality is the actual ξ-attributable variance of the aggregate, $\mathrm{Var}(\mathbb{E}[\bar{x} \mid \xi])$, which reduces to $K_\xi^2 \sigma_\xi^2$ in the linear case. The $\mathscr{W}_4$ externality is the closed-form term $q^2\sigma_\xi^2/V_A$, which depends only on signal-structure primitives and is identical in the linear and nonlinear environments.

[15] Λ has slope $1/4$ at zero, so $\Lambda(\kappa z) \approx \frac{1}{2} + \frac{\kappa}{4} z$ for $|z|$ small. The linear-equilibrium responsiveness scale is recovered up to a factor of 4.

B.2 Calibrations

We use three calibrations of the strategic-interaction parameter, holding signal-structure primitives fixed across the three:

- Calibration A: $\delta = +0.7$ (strategic complements; agglomeration).
- Calibration B: $\delta = \ \ 0.0$ (no strategic interaction; pure precision).
- Calibration C: $\delta = -0.5$ (strategic substitutes; business stealing).

Each calibration uses $\sigma_\theta^2 = 1$, $\sigma_\xi^2 = 0.5$, $\sigma_A^2 = 0.3$, $\sigma_C^2 = 0.6$, $\gamma_3 = \gamma_4 = 1$. The choice-sharpness parameter κ takes four values $\{1, 2, 4, 8\}$, with the linear baseline reported separately for reference. Each (δ, κ) cell is evaluated at $q \in \{0.25, 0.50, 0.75\}$.

B.3 Results

Table A1 reports the ratio of the nonlinear wedge to the linear wedge at $q = 0.5$ for each (δ, κ) cell. Table A2 reports the direction of bias for each cell, defined as the relative gap $(\text{nonlinear} - \text{linear})/|\text{linear}|$ averaged across the q-grid.

Table A1: Magnitude of the nonlinear wedge relative to the linear wedge ($q = 0.5$).

	Calibration A: $\delta = +0.7$		Calibration B: $\delta = 0$	
κ	ratio $\mathscr{W}_3$	ratio $\mathscr{W}_4$	ratio $\mathscr{W}_3$	ratio $\mathscr{W}_4$
1.0	0.012	0.995	0.049	1.000
2.0	0.030	0.991	0.124	1.000
4.0	0.050	0.992	0.212	1.000
8.0	0.060	0.996	0.263	1.000
	Calibration C: $\delta = -0.5$			
κ	ratio $\mathscr{W}_3$	ratio $\mathscr{W}_4$		
1.0	0.080	0.999		
2.0	0.201	0.998		
4.0	0.361	0.996		
8.0	0.479	0.996		

Sign survival. The $\mathscr{W}_4$ wedge remains strictly positive in every (δ, κ, q) cell tested, with magnitude within 1% of the closed-form linear value across the entire range. The $\mathscr{W}_3$ wedge remains positive in 34 of 36 cells; two boundary cells fail (calibration C with $\kappa \in \{4, 8\}$ at $q = 0.25$), where the linear-baseline $\mathscr{W}_3$ wedge is itself small (≈ 0.005) and the nonlinear correction pushes it slightly below zero. We discuss these boundary cases below.

Magnitude. The $\mathscr{W}_4$ wedge is essentially indistinguishable from its linear value (ratios $\in [0.99, 1.00]$ in every cell). This reflects the structural property identified in Theorem 1(ii): the $\mathscr{W}_4$ externality term $q^2\sigma_\xi^2/V_A$ depends only on signal-structure primitives, not on equilibrium loadings, so the wedge inherits no sensitivity to the choice rule. The $\mathscr{W}_3$ wedge, by contrast, varies substantially across cells. At low choice sharpness ($\kappa = 1$), the nonlinear wedge is $\leq 10\%$ of the linear value; at high sharpness ($\kappa = 8$), the ratio rises but remains below 0.50 for all calibrations.

Table A2: Direction of bias: $(\text{nonlinear} - \text{linear})/|\text{linear}|$.

	Calibration A		Calibration B		Calibration C	
κ	bias $\mathscr{W}_3$	bias $\mathscr{W}_4$	bias $\mathscr{W}_3$	bias $\mathscr{W}_4$	bias $\mathscr{W}_3$	bias $\mathscr{W}_4$
1.0	-0.98	$+0.005$	-0.95	$+0.000$	-0.94	-0.006
2.0	-0.96	$+0.015$	-0.87	$+0.000$	-0.84	-0.015
4.0	-0.92	$+0.036$	-0.79	$+0.000$	-0.71	-0.028
8.0	-0.90	$+0.056$	-0.73	$+0.000$	-0.63	-0.041

Direction of bias. For the $\mathscr{W}_3$ wedge, linearity systematically overstates the wedge magnitude relative to the nonlinear equilibrium (negative bias in 12 of 12 cells). Economically, the logistic choice rule compresses adopters' response to extreme signals, reducing platform-driven aggregate variance, which in turn shrinks the $\mathscr{W}_3$ wedge. For the $\mathscr{W}_4$ wedge, the bias is much smaller in magnitude ($\leq 6\%$ in absolute value), and its sign tracks the sign of δ: positive in calibration A, zero in calibration B, slightly negative in calibration C.

Boundary cases. The two cells where the $\mathscr{W}_3$ wedge fails sign-survival (calibration C, $\kappa \in \{4, 8\}$, $q = 0.25$) are precisely the cells where the linear-baseline wedge is already small and the nonlinear correction is largest. These are calibrations with strategic substitutes and sharp choice, where the platform-driven aggregate variance is suppressed by two reinforcing forces: differentiation incentives reduce equilibrium loadings, and choice-rule curvature dampens response to extreme signals. In such cells the $\mathscr{W}_3$ welfare object is near zero by construction, so the over-adoption result is empirically empty: both private and social adoption rates are essentially aligned. The $\mathscr{W}_4$ wedge in the same cells is robustly positive, indicating that the over-adoption result of Theorem 1(ii) continues to hold via the screened-locations channel even when the aggregate-volatility channel weakens.

B.4 Implications

The exercise yields three findings of direct relevance to the main results. First, the over-adoption result of Theorem 1 survives the substitution of a smooth nonlinear entry-choice rule. The $\mathscr{W}_4$ wedge is robustly positive in every cell tested. Second, the $\mathscr{W}_4$ wedge is structurally insensitive to the choice rule, with magnitude within 1% of its linear value across the full range of κ and δ. This insensitivity is a direct corollary of the paper's argument that the $\mathscr{W}_4$ externality $q^2\sigma_\xi^2/V_A$

depends only on signal-structure primitives. The nonlinear exercise provides numerical confirmation of this structural property. Third, the $\mathscr{W}_3$ wedge is materially smaller under the nonlinear equilibrium than under the linear baseline, with magnitudes ranging from 1% to 50% of the linear value depending on (δ, κ). The economic intuition is that choice-rule curvature compresses extreme entry responses, reducing platform-driven aggregate variance. The qualitative direction of the over-adoption result is preserved; only the magnitude is attenuated. The two boundary cells where the $\mathscr{W}_3$ wedge fails sign-survival are precisely the cells where the welfare object is itself near zero, so the failure is empirically unimportant.

The exercise therefore provides further support for the paper's existing claim that the $\mathscr{W}_4$ criterion is the more robust externality criterion. The $\mathscr{W}_3$ criterion delivers the over-adoption result of Theorem 1(i) under the linear environment of the baseline model, but its magnitude is sensitive to the choice rule. The $\mathscr{W}_4$ criterion delivers a δ-invariant result under the linear environment, and the exercise shows that this δ-invariance extends approximately to the nonlinear environment as well.

C Additional Details for Case Study

The vendor-share estimates and qualitative characterizations in this subsection are constructed from public disclosures, third-party revenue trackers, vendor case studies, and SEC filings of restaurant and adjacent public companies that cite vendor relationships. Online sources are not academic publications and are listed here rather than in the bibliography. Sources are grouped by vendor and topic.

Market context. Fortune Business Insights, Retail Analytics Market; Grand View Research, Location Intelligence Market; DataIntelo, Store Catchment Analysis Platform Market; *QSR Magazine*, "INRIX Introduces New Dataset in Site Selection Suite"; INRIX/eSite Analytics case study, inrix.com/blog/case-study-esite/.

Placer.ai. *TechCrunch*, "Placer.ai boosts valuation to $1.5B after quietly raising another $75M" (Aug. 2024); Growjo, Placer.ai company profile; Placer.ai homepage and customer count; LinkedIn, Placer.ai company profile; Sidley Austin, "Sidley Advises Placer.ai in Its Additional US$75 Million Financing"; Calcalist/CTech, Placer.ai reaches $100M ARR coverage; Life360 Form 10-K (FY 2024), SEC accession 158176025000008; Life360 Form 10-12G (2022), SEC accession 119312522120522; Life360 Form 10-Q (Q2 2025), SEC accession 158176025000180; Life360 Form 10-Q (Q3 2025), SEC accession 158176025000224; FrontView REIT Q4 2025 earnings call transcript, The Motley Fool.

Tango Analytics. Growjo, Tango Analytics company profile; Tango, customer page; LinkedIn, Tango Analytics; Tango, retail real-estate suite product page; GlobeNewswire, "Tango Announces Growth Investment from Berkshire Partners" (Dec. 2021); Tango, 2020 growth announcement; CoStar Group Form 10-K (FY 2022), SEC accession 105735223000030; Latka, Tango Analytics company/revenue profile.

Buxton. Latka, Buxton revenue and competitor profile; Owler, Buxton company profile; Buxton, about page; Buxton, case studies; Buxton newsroom, Tom Buxton Hall of Fame induction (3,500+ clients claim); Buxton, Buxton fact sheet (PDF); CoStar Group Q1 2019 earnings call, The Motley Fool transcript; Bojangles' Form 10-K (FY 2017), SEC accession 119312517073323; Xponential Fitness Q3 2021 earnings call, The Motley Fool transcript.

eSite Analytics. eSite, homepage; about page; retail industry page; solution overview; clients and case studies; LinkedIn, eSite Analytics; Owler, eSite competitor profile; Craft, eSite company profile; CB Insights, eSite Analytics; Grocery Outlet Form S-1 (2019), SEC accession 119312519145450; Fogo de Chao Form S-1/A (2023), SEC accession 119312523073983; A.H. Belo Form 10-K (FY 2019), SEC accession 141389820000043; DallasNews Corp. Form 10-K (FY 2021), SEC accession 141389822000011.

SiteZeus. Growjo, SiteZeus company profile; Owler, SiteZeus profile and funding history; LinkedIn, SiteZeus; SiteZeus, customer insights platform and data page; PR Newswire, "SiteZeus Delivers

Record Growth and Breakthrough AI Advances in Site Selection in 2025" and "SiteZeus Debuts Intelligent Foot-Traffic Integration"; SiteZeus, client archive; SiteZeus / Raising Cane's, case study.

IdealSpot and Plotr. Growjo, IdealSpot company profile; Owler, IdealSpot profile; LinkedIn, IdealSpot; CB Insights, IdealSpot profile; Plotr, acquisition announcement, IdealSpot product page, about page.

Table A3: Major retail-chain analytics vendors for restaurant site selection.

Vendor	Disclosed restaurant clients	Peer revenue	Exposure factor	Retail-chain-adjusted share
Placer.ai	Cross-cutting; named in REIT and restaurant filings as benchmark; data partnership with Life360[a]	$184.6M[b]	50%	**62.8%**
Tango	Yum! Brands (KFC, Pizza Hut, Taco Bell); Wendy's[c]	$46.2M[d]	45%	**14.1%**
Buxton	Cited by CoStar (Q1 2019 call) as "industry leader in retail site analytics"; cited by Bojangles' 10-K; Xponential Fitness uses for TAM[e]	$31.6M[f]	45%	**9.7%**
eSite Analytics	Shake Shack, Five Guys, Freddy's Frozen Custard; cited in Grocery Outlet S-1 (400+ stores) and Fogo de Chao S-1/A (550+ stores)[g]	$15.0M[h]	75%	**7.7%**
SiteZeus	Subway, Metro Diner, Costa Vida, Raising Cane's[i]	$10.1M[j]	80%	**5.5%**
IdealSpot / Plotr	Plucker's Wing Bar, Bird's Barbershop[k]	$1.0M[l]	25%	**0.2%**
Total peer set		$288.5M		100.0%

Estimated peer revenue is each vendor's total estimated company revenue, from third-party revenue trackers (Growjo, Latka, Owler) cross-validated against funding/valuation disclosures and customer-count claims. The exposure factor scales total revenue to the portion plausibly attributable to retail and restaurant chain analytics specifically, based on product-page coverage, public-company filings citing the vendor by use case, and vendor-disclosed customer mix. Retail-chain-adjusted share is computed as $q_k = (\text{revenue}_k \times \text{exposure}_k) / \sum_k (\text{revenue}_k \times \text{exposure}_k)$.

[a] Life360 10-K (2024) discloses 7% of 2024 revenue from data-insights partnerships, including a Placer.ai agreement.

[b] Growjo estimate as of 2025; consistent with $100M+ ARR floor disclosed in Feb. 2024 plus subsequent growth.

[c] Tango case studies on Yum! Brands and Wendy's; CoStar 2022 10-K names Tango among lease-management competitors.

[d] Midpoint of Growjo ($58.9M) and Latka ($33.4M) estimates.

[e] CoStar Q1 2019 earnings call; Bojangles 2017 10-K; Xponential Fitness Q3 2021 earnings call.

[f] Latka estimate; consistent with Owler range of $25M–$100M.

[g] INRIX/eSite Analytics case study; Grocery Outlet S-1; Fogo de Chao S-1/A.

[h] Midpoint of Owler range of $5M–$25M.

[i] SiteZeus customer disclosures and PR releases; Raising Cane's case study.

[j] Growjo estimate, consistent with Owler range of $5M–$25M.

[k] INRIX/IdealSpot case study coverage; Plotr post-acquisition company page.

[l] Small placeholder following IdealSpot's 2023 acquisition by Plotr; Growjo lists $72.5K standalone, Owler range $0.1M–$5M.

Table A4: Sensitivity of effective concentration H to structural correlation.

	Structural correlation ρ					
Share scenario	0.0	0.2	0.4	0.6	0.8	1.0
Base case (Table A3 shares)	0.43	0.54	0.66	0.77	0.89	1.00
Lower bound (Placer 60%, others rebalanced)	0.41	0.53	0.64	0.76	0.88	1.00
Upper bound (Placer 65%, others rebalanced)	0.45	0.56	0.67	0.78	0.89	1.00
Equal-share counterfactual ($q_k = 1/6$)	0.17	0.33	0.50	0.67	0.83	1.00

$H = \sum_k q_k^2 + \rho(1 - \sum_k q_k^2)$. Share scenarios use the central, lower, and upper points of the bands recommended in Table A3. The equal-share row shows the substantive-diversity counterfactual in which the same six vendors are equally sized, illustrating that share concentration alone (apart from ρ) raises H from 0.17 to 0.43.