Shared Intelligence and Misaligned Firm Entry

Feng Chi* Nathan Yang†

May 26, 2026

Abstract

Firm entry decisions are increasingly mediated by shared intelligence platforms. These platforms can sharpen private forecasts by reducing idiosyncratic noise, but they also route many firms' decisions through a common processing layer that may introduce correlated errors unrelated to fundamentals. We develop a Bayesian entry game in which firms choose between proprietary research and a shared platform, and characterize an equilibrium with heterogeneous adoption. Adoption is excessive under two distinct externalities: synchronized aggregate entry, borne by lenders, landlords, and municipalities, and correlated evaluation, borne by locations screened through the same forecasting system. This result does not depend on whether entry decisions are strategic complements or strategic substitutes, because the framework separates shared processing errors from strategic interaction and informative common signals. A multi-vendor extension shows that vendor diversity narrows the wedge only when platforms are structurally independent. A case study of retail site-selection analytics suggests that nominal vendor competition delivers limited effective diversification.

Keywords: Firm Entry; Correlated Errors; Information Externalities; Social Value of Information; Algorithmic Monoculture; Analytics Vendor Concentration; Economics of Data

*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

Shared forecasting platforms create a simple but nonstandard welfare problem. They can be privately attractive because they pool data, reduce idiosyncratic noise, and lower the cost of analysis. At the same time, they can be socially costly because they route many firms' decisions through the same processing layer. When that layer contains a common error, adoption does not only make firms better informed; it also makes their mistakes more correlated. The relevant question is therefore not whether shared platforms produce more accurate forecasts than custom research in isolation, but how industry-level behavior changes when forecasts that were previously independent are routed through a common system.

Firm location decisions provide a natural setting for this issue. Historically, retail entry relied on local and firm-specific information (Love et al. 1986). A regional grocer learned its trade area by 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 shaped by information-processing routines that differed across firms, so two chains evaluating the same county could reach different conclusions, not only because they had different objectives but also because they interpreted the available information differently.

This 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).

The mechanism is distinct from standard sources of correlated entry. In fast food, McDonald's

[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.

and Burger King may enter the same counties because they face common demand fundamentals, benefit from agglomeration, imitate one another, operate under similar franchising constraints, or respond strategically to each other's presence. These channels are familiar in the entry literature and need not imply inefficiency. The channel studied here is different. When both firms rely on the same forecasting technology, and that technology contains a common processing component unrelated to local fundamentals, they may enter the same markets not only because those markets are attractive but also because the same forecasting error points them there. Standard models treat common signals as informative about fundamentals; here the shared component is not information about demand but a common error embedded in the processing technology.

A descriptive empirical exercise illustrates the object of interest. In county-level data, McDonald's and Burger King exhibit strong geographic alignment: counties that appear attractive to one chain also tend to appear attractive to the other after accounting for population, density, income, education, vehicle access, urbanicity, and region. This alignment is especially pronounced in counties with greater exposure to AI-supported analysis, measured using the AI Geographic Exposure index of Felten et al. (2021). The exercise does not isolate a causal channel. Its purpose is to motivate the model's central focus, cross-firm alignment in the mapping from market characteristics to entry decisions; if part of this alignment is generated by shared forecasting infrastructure, then the same technology that helps firms identify attractive markets may also propagate common errors. This concern is consistent with recent evidence that algorithmic prediction can improve individual decisions while increasing correlation across firms. Firms investing more heavily in AI exhibit increases in systematic risk (Babina et al. 2024), and firms sharing data through common SDKs exhibit co-movement and herding that survives controls for product similarity and supply-chain linkages, weakening when Apple's App Tracking Transparency policy disrupted cross-firm data merging (Bian et al. 2024); the same policy degraded the precision of AI-derived signals (Abis et al. 2025). Common errors can also originate in the inputs themselves: Chi (2025) show that firms relying on shared data sources such as the decennial U.S. Census make systematically misinformed entry decisions when those inputs become outdated, and Raymond (2025) find that digitizing U.S. housing records lowered the cost of algorithmic property valuation, induced entry by algorithmic investors, and pushed human investors toward properties harder to predict algorithmically. What remains less understood is the welfare structure that arises when many adopters route decisions through a common processing layer rather than through independent forecasting

systems.

This paper develops a parsimonious model of that 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.[2] Custom research generates idiosyncratic forecast errors. The platform reduces idiosyncratic noise but introduces a common processing shock that is shared across adopters and unrelated to fundamentals. The model admits a closed-form equilibrium with heterogeneous adoption, separating three objects that are often conflated: strategic interaction, which governs whether firms align with or differentiate from rivals; shared processing, which creates correlated errors even after strategic interaction is accounted for; and the welfare loss, which arises because these errors fall on parties outside the adopting firms.

The equilibrium has two central properties. The influence of the shared component on aggregate entry increases with adoption, because idiosyncratic noise averages out in the cross section while the shared shock does not. And the private value of adoption depends on the strategic environment: adoption grows more attractive as it spreads when firms benefit from aligning with rivals, and less attractive when they benefit from differentiation. Strategic interaction therefore shapes adoption incentives but does not eliminate the separate externality generated by shared errors.

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: one borne by agents exposed to synchronized entry, such as lenders, landlords, and municipalities, and one borne by locations evaluated by multiple firms through the same forecasting technology. Equilibrium adoption is excessive in both cases, and the result does not depend on whether entry decisions are strategic complements or substitutes; in the location-screening case the wedge depends only on adoption and signal-structure primitives.[3] This non-ambiguity is worth noting. Many information models deliver ambiguous welfare effects because common signals improve coordination while reducing

[2]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 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.

[3]Correlated investment is studied extensively in finance, where correlation typically arises from reputational herding (Scharfstein and Stein 1990, Froot et al. 1992), informational cascades (Banerjee 1992, Bikhchandani et al. 1992), or co-movement in fundamentals (Morck et al. 2000). The mechanism here is distinct: correlation arises not from agent behavior but from a structural property of the shared processing layer, which delivers a forecast component orthogonal to fundamentals and uncorrelated with strategic motives.

independent learning. Once the common component is separated from informative common signals and from strategic interaction, the welfare logic becomes one-sided.

This mechanism distinguishes the paper from three related literatures. One studies how shared information affects market structure through platform market power; Kirpalani and Philippon (2020), for instance, show that more precise consumer data strengthens a platform's bargaining position and reduces entry through prices and outside options, whereas the channel here operates through shared processing that correlates errors rather than only actions. A second concerns the economics of data, where externalities in data collection generate inefficiency such as overcollection (Acemoglu et al. 2022, Choi et al. 2019, Ichihashi 2021); the externality studied here operates one layer downstream, in the processing that converts data into forecasts, and falls on third parties exposed to correlated decisions rather than on consumers whose information is revealed. A third comprises the social value of information (Morris and Shin 2002, Angeletos and Pavan 2007, Hellwig and Veldkamp 2009, Myatt and Wallace 2012) and algorithmic monoculture (Kleinberg and Raghavan 2021, Peng and Garg 2024); the present signal structure differs from the former because the common component carries no information about the state, and brings the latter into an entry environment in which shared evaluation shapes both aggregate outcomes and welfare. Bian et al. (2024) document a related empirical phenomenon and model data as a productive asset, emphasizing positive externalities and systemic importance through valuation; the present framework is complementary in focusing on adoption incentives and the welfare wedge from propagated common errors.

Two extensions and a case study complete the analysis. A dynamic extension shows that shared processing induces persistence in aggregate outcomes even when fundamentals are independent over time, and that lower switching costs accelerate the realization of the externality without changing the long-run static equilibrium. A descriptive case study of the U.S. restaurant site-selection analytics market then quantifies the vendor-diversity margin. Although several vendors compete at the front end, including Placer.ai, Tango, Buxton, eSite, SiteZeus, and IdealSpot, adoption is concentrated and many tools rely on overlapping upstream inputs such as roadway traffic data, points-of-interest data, mobility panels, and demographic segmentation. Combining downstream concentration with upstream overlap, plausible calibrations imply that the observed market retains roughly three-quarters of the single-platform distortion despite multiple vendor brands. Reducing structural correlation across vendors produces substantially larger gains than equalizing adoption

across them, so nominal vendor competition provides limited effective diversification when front-end providers draw on overlapping upstream inputs.

2 Motivation: Fast Food Site Selection and AI Exposure

Large quick-service chains often make location decisions in geographic lockstep. A traveler who finds a McDonald's near a highway interchange often expects a Burger King nearby, and that expectation is frequently correct. This section documents that regularity in county-level location data and then asks a slightly different question: do the two chains appear to choose similar *types* of markets?

The exercise is descriptive. It does not separate demand fundamentals, agglomeration economies, imitation, franchising constraints, real-estate constraints, or shared analytics. It is not offered as causal evidence of any particular mechanism. Its purpose is narrower: to show that the object formalized in the model, cross-firm alignment in entry decisions, is visible in a familiar retail setting before any structure is imposed. The empirical pattern is not that AI exposure directly predicts more stores after controls. Rather, McDonald's and Burger King exhibit highly similar mappings from observable market fundamentals into store counts. In high-AI-exposure counties, this similarity remains especially tight.

2.1 Industry Setting and Data

We assemble a county-level snapshot of McDonald's and Burger King locations in the United States.[4] Brand locations come from OpenStreetMap, and each location is assigned to a county through point-in-polygon matching against Census county boundaries. We merge these counts to county fundamentals from the American Community Survey, the Census Gazetteer, and the USDA Rural-Urban Continuum Codes, and to the county-level AI Geographic Exposure index of Felten et al. (2021).

[4] We focus on McDonald's and Burger King for three reasons. First, they are large national chains in the same product segment, making them a natural pair for studying whether similar firms select similar markets. In the 2025 QSR 50, which ranks U.S. limited-service chains by 2024 systemwide sales, McDonald's is the largest chain overall and the largest burger chain, while Burger King ranks eighth overall and third among burger chains. Second, both chains have broad geographic footprints, so county-level variation reflects national site-selection patterns rather than the expansion strategy of a narrow regional brand. Third, because both operate in the burger quick-service segment, differences in product category, service format, and target occasion are held relatively fixed. Together, McDonald's and Burger King account for a large share of burger-chain activity in the QSR 50, making them a substantively important to focus attention on.

The AI Geographic Exposure index measures how exposed a county's local employment structure is to advances in artificial intelligence. It is derived from the authors' occupation-level AI Occupational Exposure measure, which links ten broad AI applications to workplace abilities in O*NET and then weights these ability-level exposure scores by their importance and prevalence within each occupation. Felten et al. (2021) aggregate this measure to the county level using local industry employment shares. Counties with higher AIGE therefore have industry and occupation mixes that rely more heavily on abilities exposed to AI progress.

Two features of the data discipline the exercise. First, the brand layer is a partial collection. A zero outside the subset of states for which both brands were collected reflects absence of collection rather than absence of stores. We therefore restrict attention to counties in states where both brands were collected and treat a zero within that subset as a genuine absence. Second, OpenStreetMap is volunteered geographic information rather than a store census, and the conservative brand query undercounts stores that are present but untagged. The estimates below should therefore be read as descriptive and conservative. After restricting to counties with both brands collected and non-missing fundamentals, the sample contains 1,987 counties.

2.2 Empirical Approach

The empirical exercise documents alignment in market selection rather than a causal effect of AI exposure on store entry. This distinction guides the specification. A direct regression of store counts on AI Geographic Exposure is difficult to interpret because AIGE is constructed from local industry and occupation composition, which is also related to demand, density, commuting patterns, and commercial development. We therefore use AIGE as a stratifying variable. The main object is the similarity between the two chains' mappings from observable market fundamentals to store counts.

We estimate chain-specific relationships between store counts and county fundamentals and compare the resulting coefficient vectors. This asks whether McDonald's and Burger King assign similar weights to the same observable market characteristics when sorting counties into high- and low-attractiveness markets. All continuous fundamentals are standardized before estimation so that coefficient magnitudes are comparable across variables (Gelman 2008). We summarize cross-chain alignment using three statistics: cosine similarity of the coefficient vectors, correlation of the coefficient vectors, and correlation between the predicted market scores from the two chain-specific

models. The first two statistics measure whether the slope vectors point in similar directions. The third measures whether the two models rank counties similarly in predicted attractiveness. This prediction-based comparison is useful because fitted values can be informative descriptive objects even when the exercise is not causal (Mullainathan and Spiess 2017).

As a complementary diagnostic, we estimate a stacked chain-by-county specification with chain fixed effects and interactions between a Burger King indicator and the market fundamentals. This provides a Wald test of equality of slopes across chains, in the spirit of standard coefficient-comparison tests across regression models (Paternoster et al. 1998). Exact equality is not the main benchmark. With a large county sample, modest differences in scale or in sensitivity to population can lead to rejection even when the two chains rank markets in nearly the same way. The equality test is therefore reported as a diagnostic, while the core descriptive evidence is the degree of directional alignment in coefficients and predicted market scores.

Finally, we split counties at the median of the AI Geographic Exposure index and recompute the similarity measures within low- and high-exposure subsamples. This stratification is not intended to identify the effect of AI exposure. It asks whether market-selection similarity is especially pronounced in counties whose employment structure is more exposed to AI. The empirical section therefore uses descriptive evidence in the sense emphasized by Einav and Levin (2014): it documents a robust pattern that motivates the theoretical mechanism studied below.

2.3 Synchronized Entry

The two brands are present in the same counties far more often than their separate frequencies would imply. Table 1 cross-tabulates an indicator for at least one McDonald's location against an indicator for at least one Burger King location. Among counties with a McDonald's, 67.4 percent also contain a Burger King; among counties without a McDonald's, only 6.0 percent contain a Burger King. The two presence indicators have a Pearson correlation of 0.51. Conditional on population, density, income, urbanicity, and region, a county with a McDonald's is 11.9 percentage points more likely to also contain a Burger King, with a robust standard error of 2.7 percentage points. Synchronized entry is therefore a strong feature of the data: the chains co-locate, and this co-location survives coarse controls for market fundamentals. This pattern is consistent with evidence from the U.K. fast-food industry, where Toivanen and Waterson (2005) document strategic interaction and co-location in hamburger-chain entry.

Table 1: McDonald's and Burger King at the county level

	Burger King		
McDonald's	No	Yes	Total
No	408	26	434
Yes	506	1,047	1,553
Total	914	1,073	1,987

Notes: County-level observations in states for which both brands were collected in OpenStreetMap, with non-missing fundamentals ($N = 1{,}987$). Each cell is a count of counties. "Present" denotes at least one location of the brand in the county. Among McDonald's counties, the share with a Burger King is $1{,}047/1{,}553 = 0.674$; among non-McDonald's counties, it is $26/434 = 0.060$. OpenStreetMap is a partial, volunteered source and is not a store census.

2.4 Similar Market-Selection Rules

Co-location shows that the two chains often enter the same counties. The more relevant fact for the model is that they appear to enter similar *types* of counties. Let n_{bi} denote the number of stores operated by brand $b \in \{\text{McDonald's}, \text{Burger King}\}$ in county i. We estimate brand-specific relationships between store counts and observable market fundamentals:

$$y_{bi} = \alpha_b + X_i'\beta_b + \varepsilon_{bi},$$

where y_{bi} is either $\log(1+n_{bi})$ or n_{bi} in a Poisson specification. The vector X_i includes population, population density, income, poverty, unemployment, education, vehicle access, Rural-Urban Continuum Code fixed effects, and Census-region fixed effects. Continuous fundamentals are standardized so that the estimated coefficients can be compared across brands.

The two chains differ in scale. McDonald's has more locations on average, with 4.85 locations per county compared with 2.11 for Burger King. The market rankings implied by observables, however, are nearly identical. In both the log-count and Poisson specifications, the coefficient vectors on standardized fundamentals are almost collinear, and the predicted market scores for the two chains are highly correlated. Thus, although McDonald's and Burger King differ in footprint size, observable fundamentals sort counties into high- and low-attractiveness markets in almost the same way for the two chains.

A stacked chain-by-county specification provides a complementary test. We estimate a model

Table 2: Similarity in market-selection rules

Specification	Cosine similarity	Coeff. correlation	RMS distance	Score correlation
$\log(1+\text{stores})$	0.9997	0.9997	0.0878	0.9964
Poisson count	0.9986	0.9985	0.0297	0.9885

Notes: The table compares chain-specific regressions of McDonald's and Burger King store counts on the same set of county fundamentals. Continuous fundamentals are standardized before estimation. Cosine similarity, coefficient correlation, and RMS distance are computed using the coefficient vectors on standardized continuous fundamentals. Score correlation is the correlation between fitted market scores from the two chain-specific models. Both specifications include Rural-Urban Continuum Code and Census-region fixed effects.

with chain fixed effects and interactions between the Burger King indicator and the market fundamentals. In the log-count and presence specifications, equality of slopes is rejected, largely because the two chains differ in scale and in their sensitivity to population. In the Poisson count specification, however, equality of the continuous-fundamental slopes is not rejected ($p = 0.238$), although equality is rejected once interactions with urbanicity and region are included.

This distinction is informative. Exact coefficient equality is too strong a benchmark for shared market selection. The relevant descriptive fact is directional alignment: the same observable fundamentals rank markets in nearly the same way for both chains. The two chains need not have identical coefficients for their entry rules to be highly aligned in the markets they identify as attractive.

2.5 AI Exposure as a Stratifying Variable

We do not use AI Geographic Exposure as a causal regressor in this exercise. A county's AI exposure reflects its industry and occupation mix, and is therefore related to the same market fundamentals that shape retail entry. We instead use AI exposure as a stratifying variable. Counties are split at the sample median of the AI Geographic Exposure index, and the market-selection similarity measures are recomputed separately for low- and high-exposure counties.

The similarity is high in both subsamples, but it is tighter in high-AIGE counties. The predicted-score correlation is 0.9849 in low-AIGE counties and 0.9978 in high-AIGE counties, and the coefficient-vector similarity follows the same pattern. This does not imply that AI exposure causes chains to behave more similarly. Rather, it shows that markets more exposed to AI, as measured by local industry and occupation composition, are also markets in which the two chains' observable

Table 3: Market-selection similarity by AI exposure

Sample	N	Cosine similarity	Coeff. correlation	RMS distance	Score correlation
Full sample	1,987	0.9997	0.9997	0.0878	0.9964
Low AIGE	993	0.9943	0.9941	0.0875	0.9849
High AIGE	994	0.9991	0.9991	0.1010	0.9978

Notes: The table reports similarity measures from chain-specific $\log(1+\text{stores})$ regressions estimated separately in the full sample, below-median-AIGE counties, and above-median-AIGE counties. AIGE is used only to split the sample, not as a treatment variable or direct regressor. The score correlation is the correlation between fitted market scores from the McDonald's and Burger King models.

market-selection rules are especially close in directional terms.

The descriptive evidence motivates the model in two ways. First, McDonald's and Burger King co-locate strongly across counties. Second, the two chains map observable market fundamentals into store counts in highly similar ways, especially in high-AIGE markets. This alignment is consistent with several mechanisms, including common demand, agglomeration, imitation, franchising constraints, and shared analytics. The model focuses on the last channel. It asks what happens when alignment is partly produced by shared forecasting infrastructure: when firms evaluate markets using related signals, private entry decisions may become correlated in ways that firms do not fully internalize.

As an exploratory diagnostic, we also decompose the high-minus-low AIGE difference in predicted-score alignment across the seven continuous fundamentals using an exact Shapley attribution.[5] For each subset of fundamentals, we compute an alignment value by assigning both chains the high-AIGE coefficients for those fundamentals and the low-AIGE coefficients for the remaining fundamentals, then evaluating the predicted-score correlation on a fixed set of counties. A fundamental's contribution is its average marginal effect on this alignment value across all orderings in which coefficients are switched from their low- to high-AIGE values. By construction, the seven contributions sum to the total high-minus-low alignment difference.

The exercise is descriptive. Predicted-score alignment is close to one in both subsamples, so the difference being decomposed is small. The resulting contributions are therefore small, partially offsetting, and imprecisely estimated. Bootstrap confidence intervals are wide enough

[5] Shapley-value attribution is widely used in machine-learning interpretability (Lundberg and Lee 2017). We apply the exact version of the underlying decomposition (Shapley et al. 1953): rather than approximating contributions to a single prediction, we compute each fundamental's exact contribution to the high-minus-low alignment difference.

that no individual contribution can be distinguished from zero. We therefore interpret the Shapley decomposition only as a check that the stronger high-AIGE alignment is not driven by a single market fundamental. The evidence points instead to broad-based alignment across observables.

3 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 shared intelligence platform adoption is an endogenous equilibrium outcome reflecting firms' responses to the underlying primitives.[6]

3.1 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, cap-

[6]The analysis treats σ_ξ^2 as a fixed primitive of the information structure. Appendix C relaxes this and allows the variance of the shared component to depend on adoption, $\sigma_\xi^2 = \sigma_\xi^2(q)$, capturing channels such as capacity strain, training-data feedback, and Goodhart effects. The over-adoption result is preserved and the wedge acquires an additional amplification term; the qualitative conclusions of this section are unchanged.

turing 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.[7] 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.

3.2 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) identify this make-versus-buy margin as a key organizing dimension of firm AI activity, emphasizing that vendor-provided solutions can reduce differentiation across adopters because competing firms rely on similar tools. Empirical evidence is consistent with this characterization. Firms that invest more heavily in AI exhibit higher systematic risk (Babina et al. 2024), and firms that rely on external AI providers behave differently from those that build internally (Calvino and Fontanelli 2026). This literature identifies an important source of heterogeneity in firm AI use, but it does not characterize the welfare implications 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. Platform-based forecasting reduces part of this idiosyncratic variation, but it also introduces a component shared across adopters. This shared component captures, in reduced form, the loss of differentiation emphasized in the empirical literature and is the source of the correlated behavior studied in the model.

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

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

[7] Formally, linearity holds whenever $|\mathbb{E}_f[\theta] + \delta \mathbb{E}_f[\bar{x}]| \leq \bar{\varepsilon}$ pointwise; a sufficient condition is $\bar{\varepsilon} > 4\sqrt{\mathrm{Var}(\theta) + \mathrm{Var}(\bar{x})}$. 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.

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

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

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

The specification of ξ is motivated by a central lesson from computer science: prediction systems can have similar measured accuracy while sharing latent sources of failure. Underspecification implies that different pipelines can achieve similar held-out performance while behaving differently in deployment (D'Amour et al. 2022). Production systems also accumulate risks through entanglement, data dependencies, configuration choices, feedback loops, and upstream–downstream interfaces (Sculley et al. 2015). These papers do not study shared adoption externalities, but they motivate the reduced-form object in this paper: forecast errors that originate in the processing layer rather than in market fundamentals. When many firms rely on related pipelines, such errors can become common across adopters even when each system appears accurate in isolation.

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.[8] It is useful to interpret the shared component as a latent processing factor. This factor summarizes the common distortions induced by a shared decision technology. Unlike fundamentals, which average out in the cross section, the shared processing factor persists in aggregate behavior

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

and directly shapes equilibrium outcomes.

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.

3.3 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.

3.4 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 5 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.[9]

4 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 $[0,1]$.

4.1 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

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

[9] The binary technology choice places all exposure to the shared component on the extensive margin of adoption. One might ask whether a firm could instead place partial weight on independent research and dilute its exposure to ξ enough to eliminate the externality. Appendix B shows that this does not occur in equilibrium. When the platform weight is a continuous choice and custom research is costly, the privately optimal platform weight strictly exceeds the social optimum whenever $\sigma_\xi^2 > 0$. Each chain internalizes the private precision cost of relying less on the platform, but not the correlation cost that its reliance imposes on co-screened locations. Removal is therefore feasible but privately under-supplied, and the same correlated-evaluation externality in Theorem 1 reappears on the intensive margin. Drawing on the notion of epiplexity (Finzi et al. 2026), the appendix also shows that the cost of reconstructing an independent forecast does not fall as the shared platform improves. Better shared intelligence therefore need not move the economy closer to self-protection.

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. The linear conjecture can be written as

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

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 $\mathscr{D}$.[10]

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

[10] 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 3.

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.

4.2 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.

5 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.

5.1 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 I 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.

5.2 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.

5.3 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 2.

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

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

This specification captures the welfare of chains alone. Any wedge between private and social incentives under $\mathscr{W}_1$ would reflect a failure to internalize 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 3 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.[11]

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.

5.4 Team-Efficiency Result

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

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

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

The result follows from 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 5.2 operate through this channel and drive the misalignment results that follow.

Remark 1 (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 F 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.

5.5 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.[12]

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

[12] Appendix F 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.[13]

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.[14]

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

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

This criterion penalizes deviations of each type's signal loading from one (perfect tracking). Unlike

[13] 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.

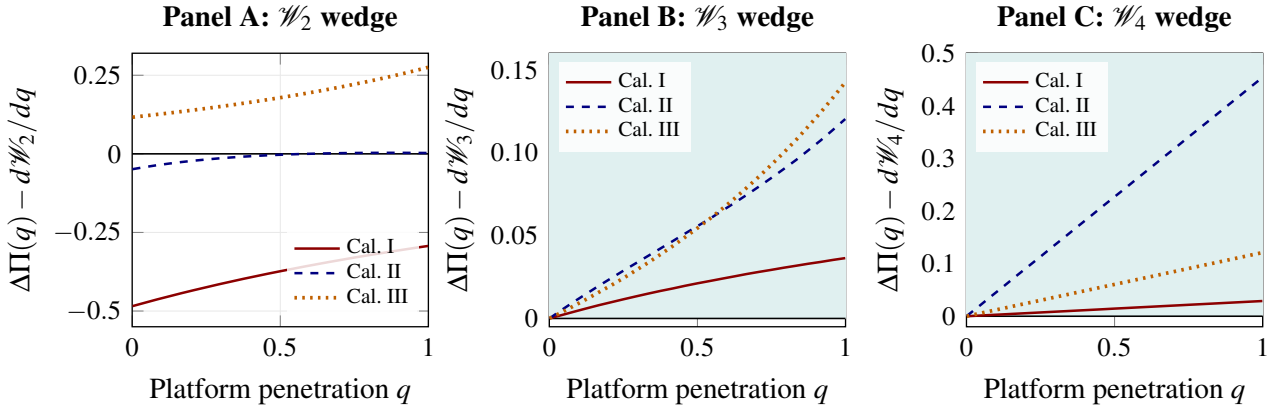
[14] See Appendix F for a numerical robustness check under an alternative logistic entry rule.

$\mathscr{W}_3$ 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 2). Panel B shows the wedge under aggregate-volatility welfare $\mathscr{W}_3$: it is strictly positive across all three calibrations and all $q \in (0, 1]$, with magnitude depending on δ through equilibrium loadings (Theorem 1, part i). Panel C shows the wedge under decided-about-agent welfare $\mathscr{W}_4$: it is strictly positive, linear in q, and invariant in δ (Theorem 1, part ii). Panels B and C jointly establish that the over-adoption result is robust to the choice of welfare specification across two structurally distinct externalities.

The contrast across the three panels of Figure 1 isolates 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 wel-

fare choice. Whenever the planner internalizes costs borne outside the chain population, platform adoption is excessive.

5.6 Residual Error versus Informative Commonality

The baseline model assumes $\xi \perp \theta$ in order to isolate the correlated-evaluation channel. This assumption separates the mechanism from the social value of public information, but it also raises a natural question: does the over-adoption result depend on the common component being entirely unrelated to fundamentals? To draw this boundary, suppose the shared component can be decomposed into an informative part and a residual part,

$$\xi = \lambda\theta + \zeta, \qquad \zeta \perp \theta, \qquad \mathrm{Var}(\zeta) = \sigma_\zeta^2,$$

where $\lambda \in [0,1)$ is the loading on fundamentals. The platform signal is then

$$\tilde{s}_f^A = (1+\lambda)\theta + \zeta + u_f^A,$$

with total variance

$$V_A(\lambda) \equiv (1+\lambda)^2\sigma_\theta^2 + \sigma_\zeta^2 + \sigma_A^2.$$

The informative component changes the precision of the common signal. The residual component, ζ, is the part that remains after conditioning on fundamentals and is therefore the part to which the correlated-evaluation externality attaches.

Proposition 4 (Residual error versus informative commonality)**.** *Suppose the shared component admits the decomposition $\xi = \lambda\theta + \zeta$, where $\zeta \perp \theta$, $\mathrm{Var}(\zeta) = \sigma_\zeta^2$, and $\lambda \in [0,1)$. Let*

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

denote the admissible region under the transformed signal. Under the decided-about-agent welfare criterion $\mathscr{W}_4$, the adoption wedge is

$$\Delta\Pi(q) - \frac{d\mathscr{W}_4}{dq} = \frac{2\gamma_4 q\,\sigma_\zeta^2}{V_A(\lambda)}.$$

Thus, for all $q \in (0,1]$ and $\delta \in \mathscr{D}(\lambda)$, the wedge is strictly positive whenever $\sigma_\zeta^2 > 0$ and is zero when $\sigma_\zeta^2 = 0$. The informative loading λ affects the wedge only through the total signal variance $V_A(\lambda)$. Holding the residual variance fixed, a larger informative component dilutes the residual variance share and reduces the magnitude of the wedge, but it does not eliminate the distortion. At any fixed adoption rate, the wedge is independent of δ.

The proposition draws the boundary between this paper's mechanism and the social value of information. In public-information models (Morris and Shin 2002, Angeletos and Pavan 2007, Hellwig and Veldkamp 2009), a common signal is informative about fundamentals. Its welfare effect is therefore ambiguous because it depends on how coordination benefits trade off against the crowding out of independent learning. The channel here operates through the residual error ζ, which carries no information about θ. Increasing the informative component λ dilutes the residual share, but the wedge remains positive as long as residual error remains. The orthogonal baseline is the special case $\lambda = 0$, in which the residual share is largest and σ_ξ^2 itself is the relevant primitive. The result therefore extends to partially informative common signals whenever the shared component contains a nonzero residual error.

The policy implication follows from this distinction. Improving the informativeness of the common signal can reduce the wedge by increasing $V_A(\lambda)$, but it does not internalize the correlated-mistake channel. Mitigation requires reducing σ_ζ^2 itself, either by increasing structural independence across vendors, limiting overlap in upstream inputs, or auditing the processing layer that converts data into forecasts.

6 Four Regimes of Adoption

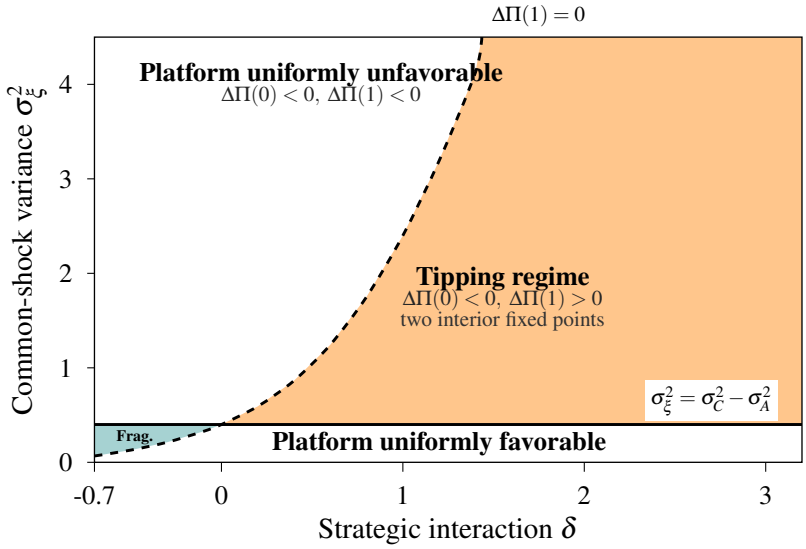
Adoption is determined by the fixed-point condition $q^* = F(\Delta\Pi(q^*))$. The equilibrium structure is therefore governed by the shape of $\Delta\Pi(q)$ together with its boundary values. These determine the number and stability of fixed points and yield a small set of regimes that organize the 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: 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: $\sigma_\theta^2 = 1$, $\sigma_A^2 = 0.2$, $\sigma_C^2 = 0.6$.

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 di-

rection of the welfare distortion. Two cases are particularly informative: the tipping regime under agglomeration and the fragmentation regime under business stealing.

6.1 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 — 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.[15] Figure 3 illustrates such a configuration for a logistic F with steepness parameter λ large enough for the curvature condition to hold.

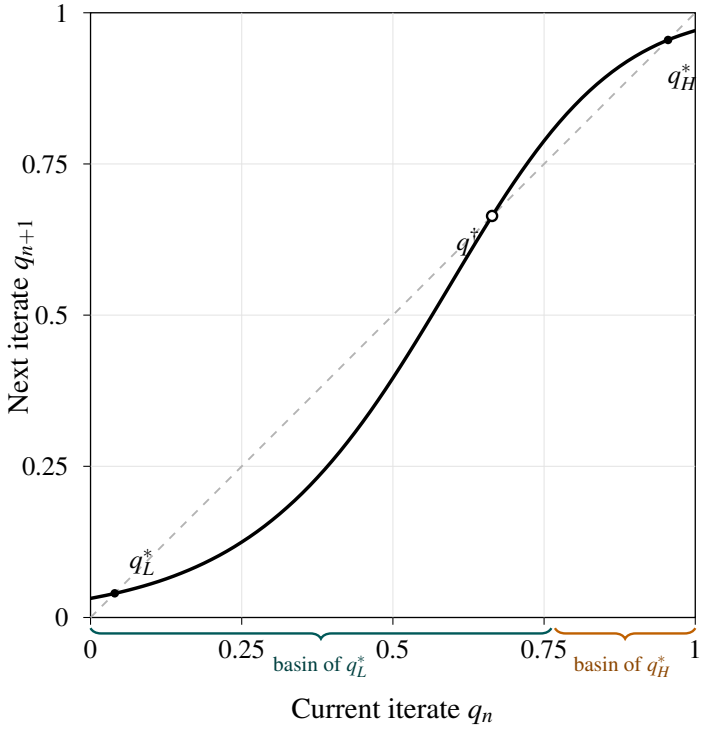
In this region, platform adoption is privately unattractive in isolation but becomes attractive once adoption is sufficiently widespread. At low penetration, exposure to the shared component is limited, and chains prefer the precision of custom research. At high penetration, aggregate entry is driven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

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-

[15] 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^*.

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 $\mathscr{D}$.

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.

6.2 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.

6.3 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.

7 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.

7.1 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 $k \neq k'$. [16] The single-platform model corresponds to $K = 1$.

The parameter ρ captures the extent of shared processing across platforms. When $\rho = 0$,

[16] 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.

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 6 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.

7.2 Effective Concentration Index

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

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

Taking variances and using the correlation structure,

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

Under symmetric adoption, $q_k = q/K$, this expression 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.

7.3 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 inputs, such as mobile-device location pings or standardized demographic projections, so nominal vendor variety can mask high ρ. By contrast, platforms built on distinct data sources, such as transaction records, mobility traces, or satellite imagery, approximate the low-ρ regime. The relevant margin is therefore the diversity of underlying inputs and methods, not the number of user-facing platforms.

A similar pattern appears in adjacent algorithmic markets. Bommasani et al. (2022) document that fine-tuned variants of a common foundation model exhibit limited independent variation relative to structurally distinct models. Direct evidence on the propagation mechanism appears in Tong et al. (2023), who identify systematic failure modes of CLIP's text encoder and show that the same failure inputs generate errors in Midjourney, DALL-E, and VideoFusion. In the notation of the model, this is a high-correlation case: distinct downstream products share an upstream processing layer, so nominal diversity leaves the externality in Theorem 1 largely intact.

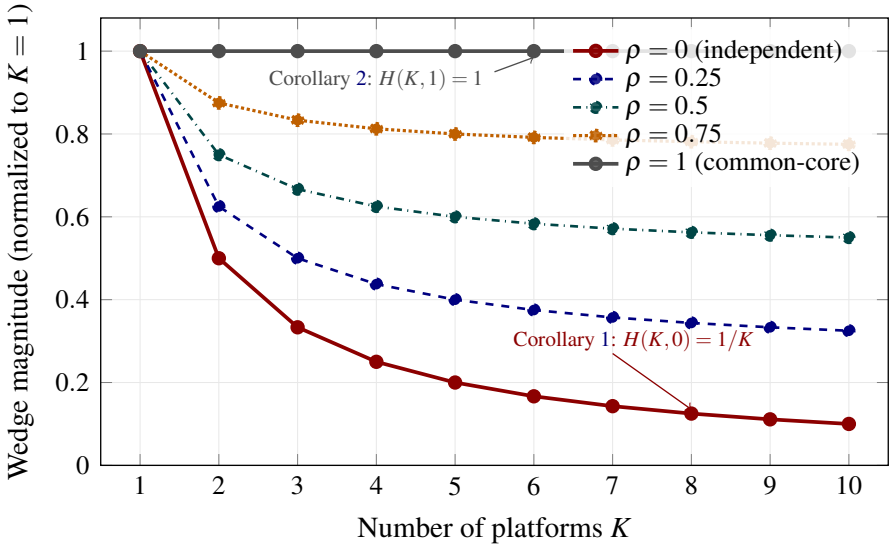
Direct measurement of cross-model error correlation is provided by Kim et al. (2025), who evaluate more than 350 LLMs on leaderboard and resume-screening tasks. They estimate error agreement of roughly 60 percent on one leaderboard dataset and show that shared architectures and providers explain part of the correlation. More importantly for the present model, larger and more accurate models exhibit highly correlated errors even across distinct providers. This finding implies that ρ need not approach zero with provider diversity alone. When adopters converge on frontier technologies, cosmetic diversity can remain binding precisely in the regime where adoption is most sophisticated.

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.

7.4 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 capital requirements tied to platform exposure) operate on the first margin. Policies that reduce H

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.

(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 3 (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.

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 co-movement 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 3. 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 E 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 6 (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 $\phi \in (0,1)$.

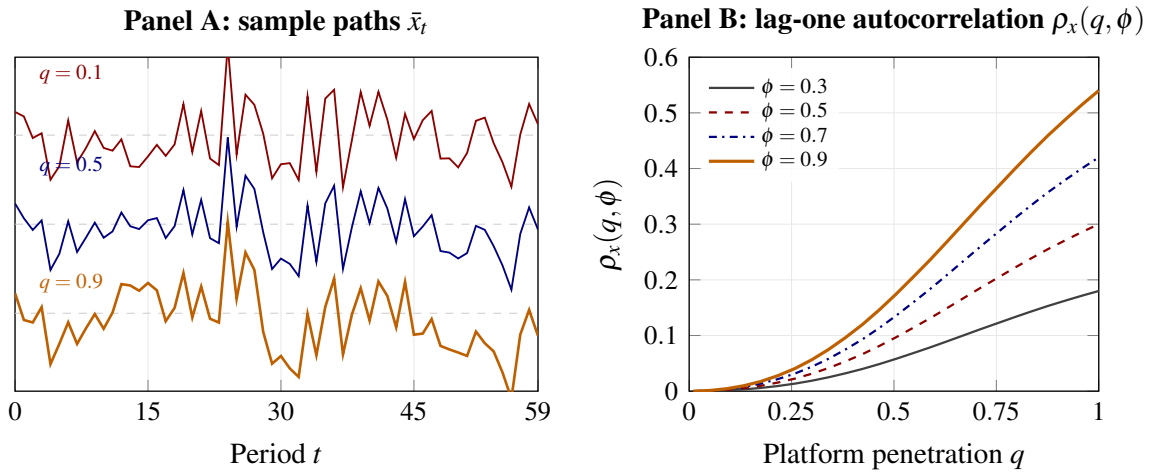
Temporal dependence is driven entirely by the shared component. The fundamental shock contributes to cross-sectional dispersion but not to persistence. By contrast, the shared component introduces a persistent factor that enters aggregate behavior with weight $K_\xi(q)$.

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 6(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 6 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 5.

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 distinct objects: concentration in downstream 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 co-movement. Whether a market is closer to substantive diversification, $H \approx 1/K$, or cosmetic diversification, $H \approx 1$, is therefore an empirical question about the structure of the technology stack rather than the number of observable vendors.

This section illustrates that distinction using the market for restaurant site-selection analytics.[17] 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 does not necessarily imply meaningful diversification in the forecasting technologies used by adopters.

The calibration has two components. The first is downstream vendor-share concentration, summarized by the Herfindahl term $\sum_k q_k^2$. This term is the probability that two adopters using analytics

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

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 can be disciplined separately. We therefore consider them in turn.

9.1 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, SiteZeus, and IdealSpot.[18] Other firms participate in the broader location-intelligence ecosystem but appear to play different roles. Esri's ArcGIS Business Analyst operates primarily as geospatial infrastructure rather than as a chain-facing site-selection platform, while a long tail of smaller specialists serves single-unit operators, local businesses, or narrower use cases outside the multi-unit chain market considered here.

A second feature of the market is more important for the model. The front-end vendors draw on a common upstream data layer. INRIX, a major 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 incorporate those feeds into otherwise proprietary scoring models.[19] This overlap is central to the case study. The six vendor brands observed by adopters do not correspond to six independent forecasting technologies. They represent competing front-end interfaces built, at least in part, on overlapping upstream inputs.

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 revenue or customer counts separately for restaurant site-selection analytics. We therefore construct an ap-

[18] eSite Analytics was acquired by Kalibrate in 2021, and IdealSpot was acquired by Plotr in 2023.

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

proximate share measure using a revenue-proxy approach. We first assemble a peer revenue pool from the best available public indicators, including 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 that it crossed $100 million in ARR in February 2024 and served more than 4,300 customers.[20] We then validate relative scale and market relevance against SEC filings and earnings-call disclosures from public-company customers and competitors. Finally, 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 A4 in Appendix H 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 raises the Herfindahl component to roughly 2.5 times the equal-share level. This is not only because the market has few vendors, but because adoption is highly asymmetric. Placer.ai accounts for a majority of adjusted revenue, while the remaining vendors divide 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 attain the substantive-diversification benchmark. The remaining question is whether adopters using different vendors are insulated from common forecasting shocks.

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

9.3 Structural Correlation

The second component of H is structural correlation across distinct vendors. This margin captures whether separately owned and separately branded platforms generate independent forecasts. In the restaurant site-selection market, this depends less on the front-end interface than 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 are licensed by multiple major site-selection platforms. Second, point-of-interest data used for competitor mapping and trade-area definition are 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 build proprietary models on top of these inputs, but the core inputs are considerably less differentiated.

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 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 percent of Life360's 2024 revenue and 10 percent of its 2022 revenue.[21] The upstream panel that feeds Placer.ai's models is therefore assembled partly 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 other analytics products may rely.

Consequently, 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 differentiated front-end products built over a substantially shared

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

back end. In the notation of the model, the shared processing shock ξ captures common errors, revisions, or methodology changes propagated through that upstream layer, including mobility-panel reweighting, point-of-interest taxonomy revisions, demographic forecast updates, and related changes inherited across vendor models.

The available data do not point-identify ρ. They do, however, make the independence benchmark difficult to defend. A value near zero would require not only separate vendor ownership, but also largely independent upstream data inputs and independent processing pipelines. The observed market structure points in the opposite direction. We therefore use $\rho \in [0.5, 0.7]$ as a disciplined calibration range. The upper end corresponds to strong 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 partnerships. The range should therefore be interpreted not as a point estimate, but as an order-of-magnitude discipline imposed by the observed structure of the technology stack.

9.4 Discussion

Combining vendor concentration with upstream-input overlap gives

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

Using the adjusted shares in Table A4 and the range $\rho \in [0.5, 0.7]$ yields

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

Table A5 reports the corresponding values of H across the full grid of ρ values and alternative share scenarios. At the central value $\rho = 0.6$, the calibration gives $H \approx 0.77$.

The implied effective concentration is much closer to the cosmetic-diversity bound, $H = 1$, than to the substantive-diversity benchmark, $1/K \approx 0.17$. The calibration therefore suggests that observed vendor variety delivers limited effective diversification. Although several vendor brands compete at the front end, their reliance on overlapping upstream inputs leaves the market substantially exposed to the same common processing component.

Welfare implication. Proposition 5 maps the case-study calibration into welfare terms. In the multi-vendor extension, both externality wedges scale linearly with the effective concentration index H, where $H = 1$ corresponds to the single-platform benchmark in Theorem 1. The calibrated range $H \in [0.72, 0.83]$ therefore measures the fraction of the single-platform wedge that remains under the observed vendor structure. Equivalently, nominal vendor competition eliminates only about 17 to 28 percent of the distortion that would be removed by moving from a single shared technology to fully independent vendor technologies. This ratio interpretation does not depend on the choice of γ_3 or γ_4, which determine the absolute scale of the wedge but not its proportional response to H. The case study therefore disciplines relative welfare changes across vendor structures without requiring an independent calibration of the marginal external cost.

Figure 6 illustrates these implications over the admissible range of the strategic-interaction parameter, $\delta \in [-0.95, 1.00]$. Panel A plots the $\mathscr{W}_3$ wedge at the case-study value $H = 0.77$, together with the single-platform ceiling $H = 1$ and the substantive-diversity floor $H = 1/6 \approx 0.17$. The case-study curve remains much closer to the cosmetic-diversity ceiling than to the substantive-diversity floor throughout the admissible region. This follows from the fact that $H_{\text{case}}/H_{\text{single}} = 0.77$ is invariant in δ, even though the level of the wedge increases as agglomeration incentives become stronger.

Panel B decomposes the welfare reduction relative to the single-vendor benchmark into the reduction generated by observed nominal vendor diversity and the maximum reduction that would be generated by structurally independent vendors. The gap between the two curves is the unrealized diversification benefit left by the observed market structure. In the calibration, roughly 72.4 percent of the achievable reduction remains unrealized at every value of δ. The proportional gap is constant because it is pinned down by H, but its level rises sharply with δ. The largest absolute welfare gains from upstream-input diversification therefore arise in settings with stronger agglomeration incentives, such as quick-service restaurant markets, even though the proportional improvement from diversification is the same across values of δ.

The $\mathscr{W}_4$ wedge is omitted from Panels A and B because, by Theorem 1(ii), it is invariant in δ at any fixed adoption rate. At $H = 0.77$ and $q = 0.6$, the wedge is

$$H \cdot \frac{2q\sigma_\xi^2}{V_A} \approx 0.257\,\gamma_4,$$

and would therefore appear as a horizontal line. Panel C instead reports the dynamic implication by plotting the lag-one autocorrelation $\rho_x(q, \phi)$ from Proposition 6 for four values of the persistence parameter ϕ. The autocorrelation increases mildly with δ through the ratio $K_\xi(q)/K_\theta(q)$, but most of the variation is governed by ϕ. The case-study calibration therefore gives a unified quantitative interpretation of both the static externality in Panels A and B and the dynamic externality in Panel C, with both governed by the same effective concentration index.

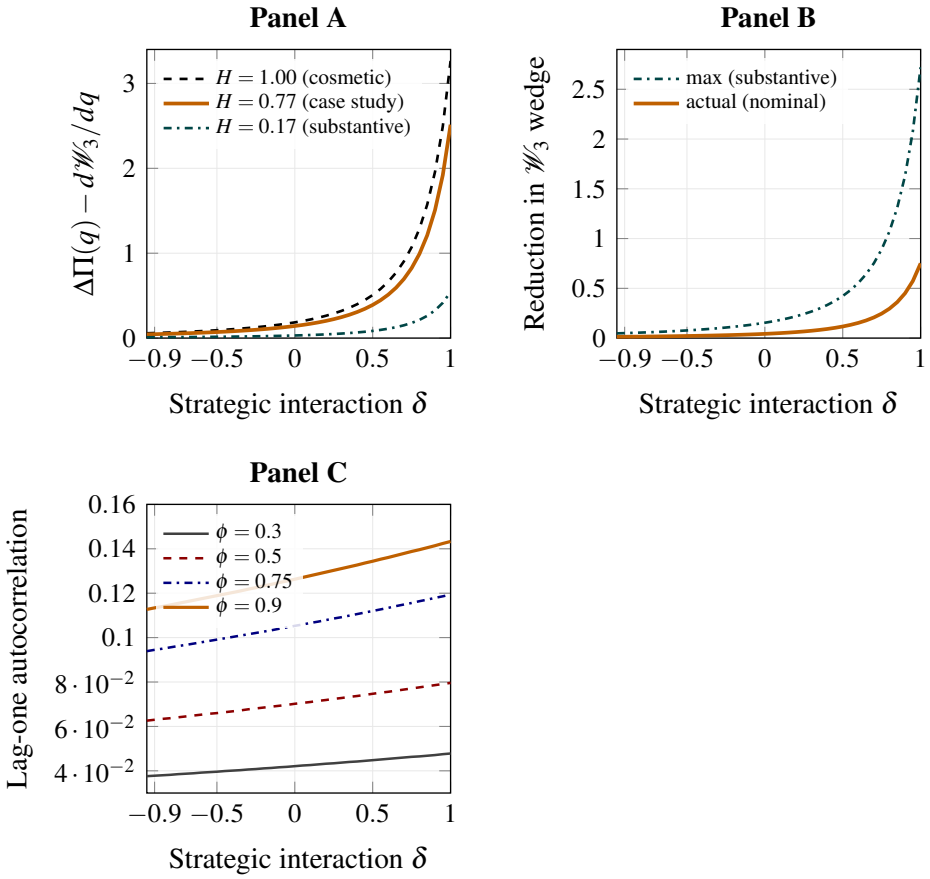
Policy-margin elasticities. The decomposition of H separates two intervention margins. The first changes downstream adoption shares while holding the upstream technology stack fixed. The second changes cross-vendor correlation while holding downstream shares fixed. These margins correspond to different policy instruments and have different quantitative implications.

Consider first a counterfactual that equalizes adoption across the six existing vendors while leaving the shared upstream data layer unchanged. At $\rho = 0.6$, this reduces H from 0.77 to 0.67, implying a wedge reduction of approximately 14 percent. This counterfactual approximates the upper bound on what can be achieved through front-end procurement diversification or antitrust scrutiny focused only on Placer.ai's downstream share, absent changes in the common input structure.

A different counterfactual eliminates structural correlation while leaving observed vendor shares unchanged. This reduces H from 0.77 to 0.43, implying a wedge reduction of approximately 44 percent. This comparison isolates the value of upstream diversification. The relevant margin is not the number of front-end vendors, but their reliance on common traffic data, common points-of-interest taxonomies, common mobility panels, and common segmentation products. The corresponding policy instruments therefore include data-portability rules, upstream-input diversification requirements, and scrutiny of vertical integration in the traffic, mobility-panel, and demographic-segmentation layers.

Full substantive diversification combines both margins: equal vendor shares and independent forecasting technologies. This reduces H to 0.17, implying a wedge reduction of approximately 78 percent. The two margins are therefore complements rather than substitutes. At the observed parameter values, the upstream correlation margin is quantitatively larger than the downstream share margin, but neither margin alone eliminates the externality.

Figure 6: Welfare implications of the case-study calibration across the strategic-interaction parameter.



Notes: All panels are evaluated at the case-study calibration $\sigma_\theta^2 = 1$, $\sigma_\xi^2 = 0.5$, $\sigma_A^2 = 0.3$, $\sigma_C^2 = 0.6$, $\gamma_3 = 1$, and at adoption level $q = 0.6$. The strategic-interaction parameter ranges over $\delta \in [-0.95, 1.00]$, the interior of the admissible region $\mathscr{D}$ of Proposition 1; the equilibrium denominator $\Delta(q)$ is strictly positive throughout and the curves diverge only at the boundary $\delta = 1.2$ where the linear equilibrium breaks down. Panel A plots the $\mathscr{W}_3$ wedge at three values of the effective concentration index H: the case-study central value $H = 0.77$ implied by Table A5; the cosmetic-diversity benchmark $H = 1$ (single vendor or six vendors with $\rho = 1$); and the substantive-diversity benchmark $H = 1/6 \approx 0.17$ (six structurally independent vendors with $\rho = 0$). The case-study curve sits much closer to the cosmetic ceiling than to the substantive floor across the entire admissible range. Panel B decomposes the welfare reduction relative to the single-vendor benchmark into the portion delivered by observed nominal vendor diversity (solid) and the maximum reduction that fully independent vendors could deliver (dashed); the gap between the curves is the regulatory headroom that the observed market leaves unrealized, and it widens sharply as strategic interaction strengthens. The $\mathscr{W}_4$ wedge is omitted from these two panels because by Theorem 1(ii) it is δ-invariant at any fixed q: at $H = 0.77$ it equals $H \cdot 2q\sigma_\xi^2/V_A \approx 0.257$, a horizontal line. Panel C plots the lag-one autocorrelation of aggregate entry from Proposition 6 for four values of AR(1) persistence ϕ; the autocorrelation increases mildly in δ through the ratio $K_\xi(q)/K_\theta(q)$ but is dominated by ϕ. Across the panels, the same H that pins down the static welfare wedge also governs the share of aggregate variance attributable to the persistent shared component, so the case-study calibration delivers a unified quantitative reading of static and dynamic externalities.

Heterogeneity across the two wedges. Although the $\mathscr{W}_3$ and $\mathscr{W}_4$ wedges both scale linearly 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 invariant in δ at any fixed adoption rate. The aggregate-volatility wedge $\mathscr{W}_3$, by contrast, depends on δ through equilibrium loadings.

This distinction affects how the same calibration should be interpreted. A regulator concerned primarily with aggregate volatility will view $H \approx 0.77$ as especially consequential in markets where strategic interaction amplifies correlated forecasting errors. Restaurant-chain entry into shared trade areas is a natural example, because agglomeration, cannibalization, and local competitive spillovers make the strategic-interaction channel non-negligible. A regulator concerned primarily with correlated screening of locations instead reads the same value of H through the δ-independent $\mathscr{W}_4$ channel. The case study is consistent with both interpretations. It disciplines the common exposure term, while leaving the relative regulatory weight on the two externalities to the welfare criterion.

10 Conclusion

This paper develops a tractable model of shared platform adoption in retail entry built around a common forecast component that is unrelated to fundamentals and therefore induces correlated decisions without improving inference. Separating this channel from common demand fundamentals and from strategic interaction yields the paper's main result: firms internalize the private precision gains from the platform and the strategic consequences of their own entry, but not the external costs their correlated errors impose on agents exposed to synchronized entry and on locations screened through the same flawed lens. Over-adoption follows regardless of whether entry decisions are strategic complements or substitutes, and in the location-screening case the wedge depends only on adoption and signal-structure primitives. This invariance is the framework's central implication: once shared processing error is separated from informative common signals and strategic interaction, the welfare distortion has a clean sign.

The analysis also clarifies why nominal vendor competition need not solve the problem. Diversification reduces the externality only when vendors rely on structurally independent data and models; when competing front-end platforms draw on the same upstream traffic data, mobility panels, points-of-interest taxonomies, or demographic inputs, adopters remain exposed to a common

processing component. The case study of restaurant site-selection analytics bears this out: despite several competing vendor brands, the calibrated effective concentration index remains close to the single-platform benchmark, so the relevant policy margin is the independence of the forecasting technologies, not the number of vendors above them.

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.

Banerjee AV (1992) A simple model of herd behavior. *The quarterly journal of economics* 107(3):797–817.

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.

Bikhchandani S, Hirshleifer D, Welch I (1992) A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of political Economy* 100(5):992–1026.

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.

D'Amour A, Heller K, Moldovan D, Adlam B, Alipanahi B, Beutel A, Chen C, Deaton J, Eisenstein J, Hoffman MD, et al. (2022) Underspecification presents challenges for credibility in modern machine learning. *Journal of Machine Learning Research* 23(226):1–61.

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

Einav L, Levin J (2014) Economics in the age of big data. *Science* 346(6210):1243089.

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.

Felten E, Raj M, Seamans R (2021) Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic management journal* 42(12):2195–2217.

Finzi M, Qiu S, Jiang Y, Izmailov P, Kolter JZ, Wilson AG (2026) From entropy to epiplexity: Rethinking information for computationally bounded intelligence. *arXiv preprint arXiv:2601.03220* .

Froot KA, Scharfstein DS, Stein JC (1992) Herd on the street: Informational inefficiencies in a market with short-term speculation. *The Journal of finance* 47(4):1461–1484.

Gelman A (2008) Scaling regression inputs by dividing by two standard deviations. *Statistics in medicine* 27(15):2865–2873.

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.

Igami M, Yang N (2016) Unobserved heterogeneity in dynamic games: Cannibalization and preemptive entry of hamburger chains in canada. *Quantitative Economics* 7(2):483–521.

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.

Kim E, Garg A, Peng K, Garg N (2025) Correlated errors in large language models. *arXiv preprint arXiv:2506.07962* .

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.

Love JFM, et al. (1986) Behind the arches. *Bantam Book: New York* .

Lundberg SM, Lee SI (2017) A unified approach to interpreting model predictions. *Advances in neural information processing systems* 30.

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

Morck R, Yeung B, Yu W (2000) The information content of stock markets: why do emerging markets have synchronous stock price movements? *Journal of financial economics* 58(1-2):215–260.

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.

Mullainathan S, Spiess J (2017) Machine learning: an applied econometric approach. *Journal of Economic Perspectives* 31(2):87–106.

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

Paternoster R, Brame R, Mazerolle P, Piquero A (1998) Using the correct statistical test for the equality of regression coefficients. *Criminology* 36(4):859–866.

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

Raymond L (2025) The market effects of algorithms. *arXiv preprint arXiv:2508.09513* .

Scharfstein DS, Stein JC (1990) Herd behavior and investment. *The American economic review* 465–479.

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

Sculley D, Holt G, Golovin D, Davydov E, Phillips T, Ebner D, Chaudhary V, Young M, Crespo JF, Dennison D (2015) Hidden technical debt in machine learning systems. *Advances in neural information processing systems* 28.

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

Shapley LS, et al. (1953) A value for n-person games. *Classics in Game Theory* .

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.

Toivanen O, Waterson M (2005) Market structure and entry: Where's the beef? *RAND Journal of Economics* 680–699.

Tong S, Jones E, Steinhardt J (2023) Mass-producing failures of multimodal systems with language models. *Advances in neural information processing systems* 36:29292–29322.

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 Robustness Checks for the Fast-Food Empirical Analysis **2**

A.1 Full-Sample Results . . . 2

A.2 Stratification by AI Exposure . . . 2

B Removal of the Shared Error and Epiplexity **2**

B.1 Removal Technology . . . 3

B.2 Blending Decision . . . 4

B.3 A Pigouvian Instrument on the Intensive Margin . . . 6

C Endogenous Processing Noise **7**

C.1 Equilibrium with Endogenous Noise . . . 7

C.2 The Wedge Under Endogenous Noise . . . 7

C.3 Implications for Adoption Dynamics and Identification . . . 8

D An Identifying Restriction for Empirical Work **9**

E Dark Side of Cheap Switching **12**

F Nonlinear Robustness of the Team-Efficiency Benchmark **17**

F.1 Equilibrium Computation . . . 17

F.2 Calibrations . . . 18

F.3 Results . . . 18

F.4 Implications . . . 19

G Strategic Timing with Preemptive Motives **21**

G.1 Two-Period Entry Game . . . 21

G.2 Wedge Decomposition . . . 22

G.3 Numerical Verification . . . 23

G.4 Implications . . . 25

H Additional Details for Descriptive Case Study **26**

I Proofs **30**

I.1 Proof of Proposition 1 . . . 30

I.2 Proof of Lemma 1 31
I.3 Proof of Proposition 2 31
I.4 Proof of Proposition 3 32
I.5 Proof of Theorem 1 32
I.6 Proof of Proposition 4 34
I.7 Proof of Proposition 5 35
I.8 Proofs of Corollaries 1 and 2 37
I.9 Proof of Proposition 6 37
I.10 Proof of Proposition 10 38
I.11 Proof of Proposition 8 42
I.12 Proof of Proposition 9 43

A Robustness Checks for the Fast-Food Empirical Analysis

Two checks clarify how the similarity results should be read. The first asks whether the alignment measures are artifacts of the particular county sample. We resample counties by state cluster 1,000 times and, in each bootstrap draw, re-estimate the chain-specific market-fundamental regressions and recompute the coefficient-vector cosine similarity, the coefficient correlation, and the predicted-score correlation. The second check separates two distinct claims that the headline numbers might be taken to support. A stacked chain-by-county regression tests whether Burger King's coefficients on the standardized continuous fundamentals are exactly equal to McDonald's. Exact equality is a stricter benchmark than the motivating fact requires: the model is motivated by whether the two chains rank markets similarly, not by whether their entry-count equations are literally identical.

A.1 Full-Sample Results

The two checks point in opposite directions on exact equality but agree on directional alignment, which is the object that matters. A Wald test rejects equality of the continuous-fundamental slopes across chains ($p < 0.001$), confirming that McDonald's and Burger King do not share an identical entry-count equation. This is unsurprising given their different footprint sizes and different sensitivities to scale. The directional diagnostics are nonetheless strong. The signs of all seven continuous-fundamental coefficients match across chains, and the rank correlation of absolute coefficient magnitudes is 0.964. The state-cluster bootstrap yields a 95 percent confidence interval of $[0.993, 0.9997]$ for cosine similarity and $[0.991, 0.997]$ for predicted-score correlation. The two chains therefore rank markets in nearly the same way on observable fundamentals even though their coefficient vectors are not identical.

A.2 Stratification by AI Exposure

Repeating both checks within the AIGE-stratified samples shows that the alignment is tightest where local employment is most exposed to AI. In the high-AIGE subsample, sign concordance is perfect, the coefficient-vector cosine similarity is 0.999, and the predicted-score correlation is 0.998, with a bootstrap confidence interval of $[0.993, 0.998]$. In the low-AIGE subsample, similarity remains high but is weaker: the cosine similarity is 0.994, the predicted-score correlation is 0.985, and sign concordance falls to six of seven fundamentals. Exact coefficient equality is rejected in every sample, but directional alignment is strongest and most stable in counties whose employment structure is more exposed to AI.

B Removal of the Shared Error and Epiplexity

The baseline model treats the forecasting-technology choice as binary: a chain either conducts custom research or relies on the platform. This places all exposure to the shared component on the extensive margin of adoption. In practice, however, a chain that subscribes to a platform need not rely on it mechanically. It can commission independent research, audit vendor predictions against field data, or combine proprietary forecasts with platform output. Each option dilutes the chain's

exposure to the common error ξ. A natural question about Theorem 1 is whether sophisticated chains may reduce their reliance on the platform until the shared component is effectively removed. This section addresses that question by giving chains an explicit, costly removal technology and allowing them to choose how intensively to use it.

The main insight is that removal is feasible but privately under-supplied. A chain that shifts weight from the platform toward independent research bears the full private cost of doing so, while capturing only the private benefit from a less noisy forecast. It does not internalize the benefit conferred on the locations it screens, whose evaluations become less correlated with those of other adopters when the chain places less weight on the shared component. The gap between the private and social reliance on shared intelligence is therefore the same correlated-evaluation externality identified in Section 5.2. Removability does not eliminate the externality. It moves the distortion from the extensive margin of *whether* to adopt to the intensive margin of *how heavily* to rely on the adopted platform.

B.1 Removal Technology

The removal technology is costly, and its cost reflects what the platform provides. A chain that subscribes to a site-selection platform is not merely buying raw data. Traffic counts, points of interest, demographic projections, and mobility panels are inputs. The platform sells the processing layer that turns those inputs into an actionable forecast: how to construct trade areas, which features predict demand, how to weight traffic against demographics, how to correct for seasonality, and where nonlinearities matter. A chain that wants to forecast without relying on the platform cannot simply avoid the subscription and retain the same information. It must reconstruct the relevant structure using its own analysts, data, and computation.

This reconstruction cost is the natural cost of self-protection. Finzi et al. (2026) formalize a related object through the notion of *epiplexity*: the length of the shortest program that a computationally bounded observer can use to extract the predictable structure in data, separate from the part that is unpredictable noise. Two features of this notion are useful for the present setting.

First, epiplexity is observer-relative. The same raw mobility panel may contain a fixed amount of statistical information, but the structure a forecaster can use depends on the computation and expertise it can deploy. A chain with a limited analytics team and a finite decision horizon is not an unbounded statistician. The usable structure in the data is the structure it can afford to extract. This is why shared platforms are valuable: the platform incurs the fixed cost of extracting and operationalizing that structure, then resells the resulting forecast technology to many subscribers. A subscriber that wants to forecast independently must re-create part of that structure itself.

Second, better platform forecasts need not make independent reconstruction cheaper. The platform's processing runs in one direction: it maps raw inputs into forecasts. Recovering a clean forecast of the fundamental from a platform output contaminated by the common component is a different problem. Having access to the forward output does not generally provide a low-cost inverse. In the entry problem, this means that a chain can consume the platform forecast cheaply, but cannot necessarily reconstruct the ξ-free fundamental without rebuilding the relevant forecasting structure. Improvements in platform accuracy may raise the value of the embedded structure without lowering the cost of reproducing it independently.

These two features pin down the cost of self-protection. It is the structural information a chain must supply on its own to forecast without routing its decision through the shared platform, and it

need not decline as the platform improves.

Definition 1 (Removal cost as conditional epiplexity)**.** Let $S_T(\theta \mid \mathscr{Z})$ denote the T-bounded epiplexity of forecasting the fundamental θ from raw market data $\mathscr{Z}$ for an observer with compute budget T, and let P_{plat} denote the platform's forecasting program. The per-unit cost of custom research is

$$\kappa \equiv \nu \cdot \left[S_T(\theta \mid \mathscr{Z}) - S_T(\theta \mid \mathscr{Z}, P_{\text{plat}}) \right],$$

where $\nu > 0$ converts structural information into the units of the profit function. The bracketed term is the conditional epiplexity of the fundamental given the platform program: the additional structural information a chain must supply to forecast θ without routing the decision through P_{plat}.

Definition 1 makes the cost a property of the forecasting technology rather than a free parameter. The next result records the bounded-computation property that prevents this cost from vanishing.

Lemma 2 (The removal cost does not vanish with platform quality)**.** *Suppose the platform program P_{plat} is efficiently evaluable within the chain's compute budget T, and that recovering the ξ-free fundamental from the platform output requires inverting a transformation whose inverse is not computable within T (Finzi et al. 2026). Then there exists $\underline{\kappa} > 0$, independent of platform forecast accuracy, such that $\kappa \geq \underline{\kappa}$. Improvements in platform quality that raise $S_T(\theta \mid \mathscr{Z})$ do not lower κ.*

The economic interpretation is that a better platform does not necessarily make it cheaper for a chain to avoid platform dependence. A more accurate vendor forecast raises the value of the structure embedded in P_{plat}, but that structure reaches the chain through an output contaminated by the common component ξ. If recovering the clean fundamental requires inverting the vendor's processing, the chain cannot appropriate the structure without also inheriting the shared error. Consuming the forecast is inexpensive; reconstructing the clean fundamental is not. Platform improvement and self-protection are therefore not substitutes, and better shared intelligence need not eliminate the cost floor that keeps chains exposed to it.[22]

B.2 Blending Decision

We now embed the removal cost in the adoption stage. A chain that has adopted the platform chooses a weight $\omega \in [0,1]$ on the platform signal relative to a custom signal, forming

$$\tilde{s}_f(\omega) = \omega(\theta + \xi + u_f^A) + (1-\omega)(\theta + u_f^C),$$

at custom-research cost $\frac{1}{2}\kappa(1-\omega)^2$.[23] The polar cases recover the baseline objects. At $\omega = 1$, the chain is a pure platform adopter; at $\omega = 0$, it is a pure custom researcher. Interior values

[22] The result is conditional on the platform's processing embedding a transformation that is hard to invert within the chain's compute budget, paralleling the conditional, one-way-function-based results that underpin the baseline information structure. This is a natural interpretation for pipelines that aggregate high-dimensional mobility panels, apply learned encoders, or join data through irreversible hashing, but it is a property of the technology that we assume rather than derive.

[23] Convexity of the custom-research cost guarantees an interior optimum and smooth comparative statics. A linear-cost specification delivers the same sign result at corners. The quadratic form is the minimal structure needed to make the intensive margin well behaved.

represent partial removal of the shared component. Exposure to ξ is therefore no longer an adoption indicator, but a continuous choice, and the aggregate loading on the shared shock becomes $K_\xi(q,\omega) = q\omega k^A$.

A chain choosing ω trades off three forces. Raising ω exploits the platform's lower idiosyncratic noise, economizes on custom research, and increases exposure to the shared component. The private optimum balances the marginal precision-and-cost benefit of platform reliance against the marginal private cost of greater exposure to ξ. The chain is not naive: it internalizes how the shared component affects its own forecast error, as in the Bayesian filtering benchmark of Section 3. The source of inefficiency is instead that the chain does not internalize how its reliance on the platform increases correlation in the evaluations faced by co-screened locations.

Proposition 7 (Insufficient self-protection)**.** *Fix an adoption rate $q \in (0,1]$ and $\delta \in \mathscr{D}$, and suppose adopters choose $\omega \in [0,1]$ to maximize private expected profit net of the custom-research cost $\frac{1}{2}\kappa(1-\omega)^2$. Let ω^* denote the private optimum and ω^{soc} the weight chosen by a planner who internalizes the decided-about-agent externality in welfare specification $\mathscr{W}_4$. Then:*

(i) ***Over-reliance.*** *Whenever $\sigma_\xi^2 > 0$, the private optimum exceeds the social optimum, $\omega^* > \omega^{soc}$.*

(ii) ***The intensive-margin wedge.*** *The gap in chosen weights satisfies*

$$\omega^* - \omega^{soc} = \frac{2\gamma_4 q\, \omega^{soc}}{\Xi} \cdot \frac{\sigma_\xi^2}{V_A} > 0,$$

where $\Xi > 0$ is the curvature of the private objective in ω. The wedge is strictly increasing in the adoption rate q, and at any fixed q it is invariant to the strategic-interaction parameter δ. It is increasing in the shared-component variance σ_ξ^2 over the range in which the platform is adopted, and declines only once σ_ξ^2 is large enough that the variance share σ_ξ^2/V_A saturates.

(iii) ***No full removal.*** *The shared component is fully removed in equilibrium, $\omega^* = 0$, only in the limits $\kappa \to 0$ or $\sigma_\xi^2 \to \infty$. For any $\kappa > 0$ and finite $\sigma_\xi^2 > 0$, the private optimum is interior or at the upper corner, and strictly exceeds ω^{soc}.*

Part (i) is the central claim. Giving chains an explicit removal technology does not eliminate the externality. A chain that lowers ω bears the private cost of independent forecasting and captures the private benefit of a less contaminated signal, but it does not capture the reduction in correlation imposed on co-screened locations. The shared error is therefore not removed in equilibrium, not because removal is infeasible, but because removal is privately under-supplied.

Part (ii) shows that the intensive-margin wedge has the same structure as the extensive-margin wedge in Theorem 1(ii). The marginal external cost is again proportional to $2\gamma_4 q\sigma_\xi^2/V_A$, now scaled by the social exposure ω^{soc} and divided by the curvature of the private problem. Its sign, its monotonicity in q, and its invariance in δ at fixed q all carry over. The reason is the same as in the baseline model: the welfare object $q^2\omega^2\sigma_\xi^2/V_A$ depends only on adoption, exposure, and signal-structure primitives, not on equilibrium loadings.

The only comparative static that differs from the extensive-margin result concerns σ_ξ^2. The wedge rises with the severity of the shared component over the range in which the platform remains

attractive. It can decline only once σ_ξ^2 is so large that the variance share σ_ξ^2/V_A saturates and the planner's chosen reliance on the platform collapses. At that point there is little remaining exposure for additional shared noise to amplify. The intensive margin therefore reproduces the δ-invariance of the $\mathscr{W}_4$ externality while adding another policy-relevant distortion: correcting adoption alone leaves excessive reliance among adopters, and correcting reliance alone leaves excessive adoption.

Part (iii) records the limiting cases. Full self-protection arises only when custom research is costless or when the shared component is so severe that firms no longer tolerate platform reliance. The first limit is precisely the one Lemma 2 rules out. Because κ is bounded below by conditional epiplexity, the cost of independent reconstruction does not vanish as platforms improve. Better shared intelligence therefore does not move the economy toward full self-protection. It may raise the value of the structure embedded in the platform without lowering the cost of reproducing that structure independently.

B.3 A Pigouvian Instrument on the Intensive Margin

Proposition 7 implies a policy instrument distinct from the adoption tax in Theorem 1. The extensive-margin instrument changes how many chains adopt the platform. The intensive-margin instrument changes how heavily adopters rely on the shared component. A per-unit subsidy to independent research of size

$$2\gamma_4 q \omega^{soc} \frac{\sigma_\xi^2}{V_A},$$

or equivalently a tax on platform weight of the same size, aligns the private choice of ω with the social optimum by offsetting the external correlation cost. The two margins are complementary. Correcting adoption alone leaves excessive reliance among adopters, while subsidizing independent research alone leaves adoption above its social level. The same caution that applies to vendor diversification applies here as well: the intervention improves welfare on the intensive margin only to the extent that the subsidized research is genuinely independent of the platform's processing core.

The connection to Section C strengthens this implication. There the variance of the shared component rises with adoption, $\sigma_\xi^2 = \sigma_\xi^2(q)$. Under endogenous noise, the intensive-margin wedge in Proposition 7(ii) inherits an additional channel through $(\sigma_\xi^2)'(q) > 0$, paralleling the adoption wedge in Proposition 8. Higher adoption both increases the number of chains relying too heavily on the platform and raises the severity of the shared error each adopter inherits. The value of self-protection therefore rises precisely when the private incentive to provide it remains incomplete.

C 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.

C.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.

C.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 8 (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 4 (Local versus global convexity)**.** Proposition 8 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 $\sigma_\xi^2(q)$.

C.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 6 therefore inherit curvature from the degradation profile, though the qualitative typology remains unchanged.

For the identifying test of Section D, 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 $\sigma_\xi^2(q)$.

D An Identifying Restriction for Empirical Work

Strategic interaction and shared processing are separable mechanisms operating through distinct channels. The slope of the adoption mapping, $d\Delta\Pi(q)/dq$, governs equilibrium selection: by Proposition 2 its sign is determined by δ alone, so δ dictates whether adoption is self-reinforcing or self-limiting and hence whether multiple equilibria arise. The welfare wedge, $\Delta\Pi(q) - d\mathscr{W}_k/dq$, governs distortion: by Theorem 1 it is strictly positive if and only if $\sigma_\xi^2 > 0$ and $q > 0$, with sign invariant to δ. Strategic interaction therefore determines whether firms align their decisions, while the shared component determines whether such alignment is costly. The limiting cases make this transparent. At $\delta = 0$ the private return to adoption is independent of penetration, yet adoption remains excessive whenever $\sigma_\xi^2 > 0$; at $\delta < 0$ differentiation incentives lower equilibrium adoption but leave it above the social optimum; and at $\sigma_\xi^2 = 0$ the welfare criteria collapse to $\mathscr{W}_1$ and all equilibria are efficient even when $\delta > 0$ generates multiplicity. Multiple equilibria thus do not by themselves indicate a market failure, nor does uniqueness rule one out: a distortion arises only when the common component is present and marginal adoption increases exposure to it.

This separation has two implications. Empirically, strategic interaction generates correlation in actions that weakens once fundamentals are controlled for, whereas the shared component generates correlation in errors that persists after such conditioning; because the wedge under $\mathscr{W}_4$ is independent of δ, the externality can be quantified without identifying the strategic environment. For policy, interventions that alter the strategic environment operate through δ and affect equilibrium selection, while those that alter shared processing operate through σ_ξ^2 and affect the magnitude of the externality directly. Section 7 develops the latter margin through an effective concentration index $H(K,\rho)$ that captures the probability that adopters share the same source of error.

This decomposition logic 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 9 (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 9 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 9 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 9 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.

E 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 7: 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 6 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 5.

Proposition 10 (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

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

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 5 (Beyond $\mathscr{W}_4$ and the strategic-substitute case)**.** Proposition 10(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 6 (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 10 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.

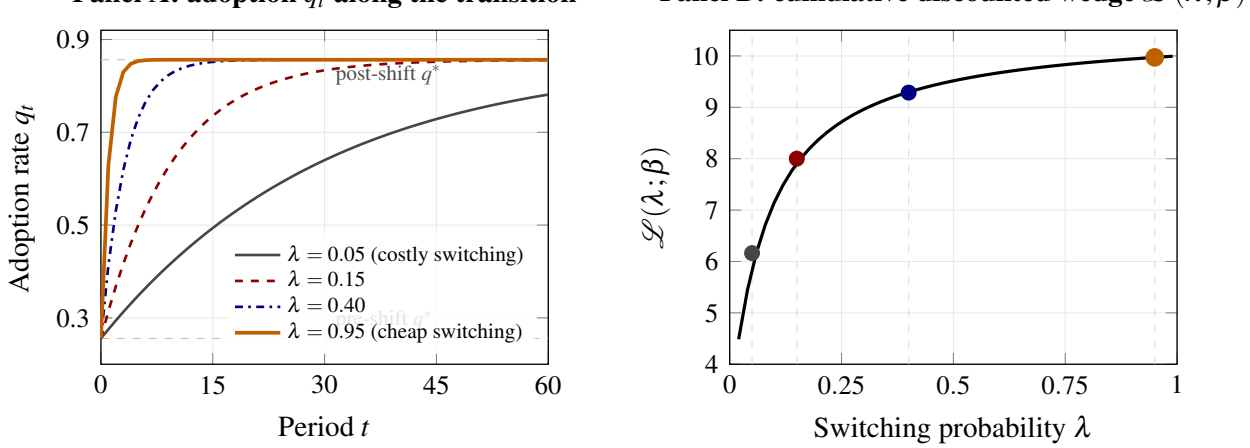
Related work studies the upstream incentive to produce shared data inputs. Fang and Kim (2026) show that data-neutrality regulations can reduce a platform's incentive to refine its data, generating a trade-off analogous to the one identified here. Both settings highlight that policies which broaden access 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 A1: 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 10(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 A1 illustrates the comparative static. The left panel plots adoption trajectories following 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 7) attenuate both contemporaneous co-movement and persistence. Timing matters as well: cheap switching shortens the window in which such interventions can take effect.

F Nonlinear Robustness of the Team-Efficiency Benchmark

Proposition 3 establishes that, in the linear continuum environment of Section 3, the equilibrium adoption rate q^* coincides with the social optimum under welfare specification $\mathscr{W}_1$. Remark 1 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.[24] 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.

F.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.

[24] 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.

F.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 $q \in \{0.25, 0.50, 0.75\}$.

F.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.

F.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.

G Strategic Timing with Preemptive Motives

The baseline model is static in its strategic structure: firms respond to platform signals without choosing when to enter relative to rivals. The empirical entry literature identifies preemption as a force that arises in dynamic timing decisions. A firm may enter early to establish a position that deters or crowds out later rivals, thereby protecting the profits it expects to earn once established (Igami and Yang 2016). This appendix embeds the shared component ξ in a two-period entry game in which preemption arises endogenously and studies how strategic timing affects the welfare wedge in Theorem 1.

The exercise focuses on the business-stealing environment most relevant for retail entry. Preemption is intrinsically tied to business stealing: a firm races to claim a position because later rival entry would erode its payoff. The result is that the equilibrium wedge contains two components. The first is the static correlated-evaluation externality from Theorem 1, which is invariant to the strategic environment. The second is a preemption component that is positive and increases as business stealing becomes stronger. The static component survives because no firm can separate ξ from the fundamental at any date. Late entrants observe early entry, but their own signals remain contaminated by the same shared component that shaped the early cohort. Strategic timing reallocates entry across cohorts; it does not remove the covariance in evaluation errors.

G.1 Two-Period Entry Game

A unit mass of ex-ante symmetric locations is evaluated by a unit mass of potential entrants over two periods, $t \in \{1,2\}$. Each location has fundamental value $\theta \sim \mathcal{N}(0, \sigma_\theta^2)$. Each adopting firm f observes a platform signal

$$s_f = \theta + \xi + u_f, \qquad \xi \sim \mathcal{N}(0, \sigma_\xi^2), \qquad u_f \sim \mathcal{N}(0, \sigma_A^2),$$

where $\xi \perp \theta$ is the shared platform component and u_f is idiosyncratic. All adopters' signals contain the same realization of ξ. Let $m_f = \mathbb{E}[\theta + \xi \mid s_f]$ denote the firm's posterior attractiveness. The firm cannot separately identify the fundamental and the shared component; it observes only their contaminated sum.

Each firm chooses whether to enter early, enter late, or stay out. Entry is irreversible. Let n_1 and n_2 denote the masses entering in periods 1 and 2, and let $N = n_1 + n_2$ denote total occupancy. Business stealing is governed by a single parameter $\delta < 0$: a firm's per-period value declines with total occupancy through the term δN. There is no separate congestion or preemption parameter; the same strategic interaction parameter governs both crowding and deterrence.

An early entrant earns a period-1 flow and a discounted period-2 flow, each evaluated at total occupancy. A late entrant earns only the discounted period-2 flow, using a sharper posterior m_f^{late} that resolves part of its idiosyncratic noise. The two continuation values are

$$V^{\text{early}}(m_f) = \big(m_f + \delta N\big) + \rho_d\big(m_f + \delta N\big), \qquad V^{\text{late}}(m_f) = \rho_d\big(m_f^{\text{late}} + \delta N\big),$$

where $\rho_d \in (0,1)$ is the discount factor. A firm enters early if

$$V^{\text{early}}(m_f) > \max\{V^{\text{late}}(m_f), 0\},$$

and enters late if it does not enter early and $V^{\text{late}}(m_f) > 0$. Equilibrium is characterized by cutoffs in posterior attractiveness, with cohort masses (n_1, n_2) determined by a fixed point.

The return to early entry arises because early entry lowers the period-2 value of the location. When $\delta < 0$, a larger early cohort increases total occupancy and reduces the payoff from late entry. Early entry therefore deters later entry, and the early entrant's period-2 position is protected by the thinner late cohort that early entry helps induce.

Following Igami and Yang (2016), who isolate preemptive motives by making rivals non-strategic, we measure preemption as the difference between two equilibria. In the strategic equilibrium, the late cohort responds to realized early entry, so deterrence is active. In the non-strategic-rivals benchmark, the late cohort's size is fixed at the level that would arise absent deterrence, so early entry provides no protection against later rivals. The benchmark fixes only the size of the late cohort, not its composition. Late entrants still evaluate locations using signals contaminated by $\theta + \xi + u_f$. Hence no firm, at any date, filters out ξ, and the static correlated-evaluation object is identical across the two regimes.

G.2 Wedge Decomposition

The welfare object is the correlated-evaluation externality from Theorem 1: the expected covariance of evaluation errors across locations attributable to the shared component, normalized by total signal variance. In the two-period timing game, this object is unchanged by strategic timing, but timing introduces an additional diagnostic channel through which the shared error affects the allocation of entry across periods.

Conjecture 1 (Strategic timing and the correlated-evaluation wedge)**.** *In the two-period entry game with shared signal $s_f = \theta + \xi + u_f$, business stealing $\delta < 0$, and irreversible timing choice, the equilibrium wedge can be represented as*

$$\mathscr{W} = \underbrace{\gamma \frac{q^2 \sigma_\xi^2}{V_A}}_{\text{static}} + T(\delta),$$

where the static term is the correlated-evaluation externality of Theorem 1. This component is invariant to the strategic environment and is identical in the strategic and non-strategic-rivals regimes. The second term, $T(\delta)$, measures the preemption channel: the additional front-loading of ξ-driven entry in the strategic equilibrium relative to the non-strategic-rivals benchmark. It is strictly positive and increasing in the intensity of business stealing, $\partial T / \partial |\delta| > 0$. The total wedge is therefore strictly positive.

The preemption term reflects the front-loading of entry induced by the shared error. As business stealing intensifies, the deterrence return to early entry increases, so firms race more aggressively to claim locations that the platform rates favorably before rivals can enter. Because the platform rating contains the shared component, the same locations inflated by ξ are also the locations into which firms enter early and together. Strategic timing therefore amplifies the allocation consequences of the shared error by concentrating them in the first period. The static term remains invariant because it depends only on adoption and signal-structure primitives, as in the baseline model.

The decomposition should be read with one qualification. The static component,

$$\gamma \frac{q^2 \sigma_\xi^2}{V_A},$$

is a welfare object derived directly from the signal structure. The comparison that defines preemption is also disciplined by the model: it is the gap between the strategic equilibrium and the non-strategic-rivals benchmark, following the logic of Igami and Yang (2016). What is not uniquely pinned down is the scalar normalization that expresses this timing gap in the same units as the static externality. The term $T(\delta)$ should therefore be interpreted as a comparably scaled diagnostic of the timing channel rather than as an independently derived welfare primitive.

This distinction does not affect the qualitative content of the result. The sign of $T(\delta)$ and its monotonicity in $|\delta|$ are invariant to any positive rescaling. What depends on normalization is only the cardinal comparison between the timing term and the static term. The decomposition is therefore best interpreted as the static correlated-evaluation wedge plus a scaled measure of the preemptive timing channel. Its purpose is to show that strategic timing does not undo the baseline externality; instead, in business-stealing environments, it adds a positive channel through which shared errors affect the timing and synchronization of entry.

G.3 Numerical Verification

Conjecture 1 is assessed by simulation. We vary the business-stealing parameter over

$$\delta \in \{-0.2, -0.4, -0.6, -0.8\},$$

holding the signal-structure primitives fixed at $\sigma_\theta^2 = 1$, $\sigma_\xi^2 = 0.5$, and $\sigma_A^2 = 0.3$. The adoption rate is fixed at $q = 0.5$, and the discount factor is $\rho_d = 0.9$. For each value of δ, we solve the strategic equilibrium and the non-strategic-rivals benchmark using common random numbers. We then measure $T(\delta)$ as the difference in the ξ-loading of early entry across the two regimes, scaled to the same units as the static externality. Equilibrium objects are computed by Monte Carlo simulation using 200,000 draws of (θ, ξ, u) per solve.

Table A3 reports the decomposition. The preemption term is positive for every value of δ and increases with the intensity of business stealing. Stronger strategic substitution raises the deterrence return to early entry, causing firms to race more aggressively on the shared signal. The cohort masses illustrate the mechanism. In the strategic equilibrium, early entry is high and late entry is thin because early entry deters later rivals. At $\delta = -0.8$, for example, the strategic equilibrium has $n_1 = 0.365$ and $n_2 = 0.046$. In the non-strategic-rivals benchmark, early entry provides no deterrence, so firms do not race in the same way: early entry falls to $n_1 = 0.309$, while late entry rises to $n_2 = 0.151$. The difference in front-loading is the preemption effect.

The static correlated-evaluation term equals $\gamma q^2 \sigma_\xi^2 / V_A = 0.069444$ at every value of δ and in both regimes, with zero numerical spread. This is the central robustness result. Preemption reallocates entry across cohorts, but it does not change the covariance object that defines the externality. The invariance is exact rather than approximate because the static term depends only on signal-structure primitives, and because the non-strategic benchmark fixes only the size of the late cohort, not the composition of the signals firms observe.

Table A3: Wedge decomposition by business-stealing intensity. Adoption rate $q = 0.5$, discount $\rho_d = 0.9$.

δ	$\mathscr{W}_4$ static	preemption $T(\delta)$	total wedge	early share (strat. vs. non)
-0.2	0.069444	$+0.00039$	0.06983	0.89 vs. 0.83
-0.4	0.069444	$+0.00086$	0.07031	0.89 vs. 0.77
-0.6	0.069444	$+0.00131$	0.07075	0.89 vs. 0.72
-0.8	0.069444	$+0.00172$	0.07117	0.89 vs. 0.67

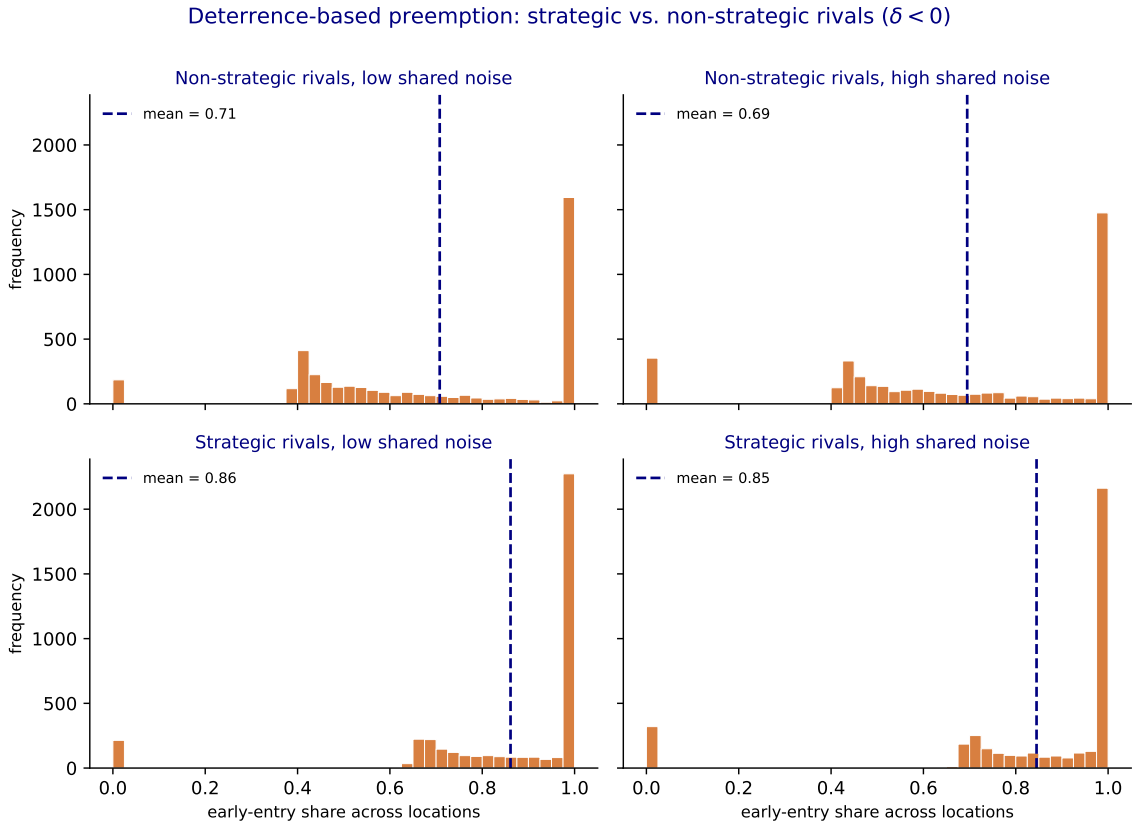


Figure A2: Distribution of the early-entry share across locations under business stealing ($\delta = -0.6$). Top row: non-strategic rivals (deterrence off). Bottom row: strategic rivals (deterrence live). Left column: low shared noise. Right column: high shared noise. Strategic rivalry concentrates entry into synchronized early waves.

The preemption gap operates through the *timing* of ξ-driven entry rather than through its overall level. The change in the ξ-loading of *total* entry between the strategic and non-strategic regimes is small and has the opposite sign of the early-entry gap, ranging from -0.0002 at $\delta = -0.2$ to -0.0010 at $\delta = -0.8$. Deterrence therefore shifts ξ-driven entry from the late cohort to the early cohort rather than materially changing total exposure to the shared component. The synchronized-wave channel, the front-loading of entry into period 1, is the dominant effect.

Figure A2 illustrates this mechanism by plotting the distribution of early-entry shares across

locations under the strategic and non-strategic-rivals regimes, separately for low and high realizations of the shared component. Under non-strategic rivals, early entry is relatively dispersed across locations. Under strategic rivals, deterrence shifts the distribution toward one, as firms front-load entry into period 1 to secure protected positions. Preemption therefore converts a common signal error into synchronized entry waves rather than dispersed mistakes, and this concentration is strongest when the shared component is large.

G.4 Implications

The exercise yields three implications for the main results.

First, the over-adoption result of Theorem 1 survives the introduction of strategic timing. The correlated-evaluation wedge remains strictly positive across the full range of business-stealing intensities considered in the simulation.

Second, the static correlated-evaluation component is invariant to the strategic environment. One might worry that forward-looking firms could anticipate that rivals read the same signal and thereby undo the correlated mistake. The exercise shows that they cannot. Preemption reallocates entry across cohorts, but it does not change the covariance of evaluation errors, because no firm at any date can filter the shared component from its signal. This identifies the precise sense in which the static analysis is robust to dynamic sophistication.

Third, preemption adds a timing channel through which shared errors affect equilibrium outcomes. When firms race to establish protected early positions on the basis of a common signal, the shared error is front-loaded into synchronized entry waves. The preemption term is positive and increases with the intensity of business stealing. This generates a comparative static absent from the static model: the welfare cost of shared intelligence is larger in industries where preemption races are important, including fast-saturating local markets with first-mover advantages and the multi-store retail setting studied by Igami and Yang (2016). The shared-intelligence externality is therefore most costly in entry environments where timing is most strategic.

H Additional Details for Descriptive 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.

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 A4: 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 A5: 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 A4 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 A4. 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.

I 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 $\sigma_\theta^2, \sigma_A^2, \sigma_C^2, \sigma_\xi^2 > 0$.

I.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 $[0,1]$. □

I.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. □

I.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. □

I.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 $F^{-1}(q^*) = d\mathscr{W}_1/dq$. □

I.5 Proof of Theorem 1

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

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

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

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

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

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

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

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

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

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

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

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,

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

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 $F^{-1}(q_k^{soc}) = d\mathscr{W}_k/dq|_{q_k^{soc}}$ for

$k \in \{3,4\}$. Since F^{-1} is strictly increasing, $\Delta\Pi(q^*) > d\mathscr{W}_k/dq|_{q^*}$ implies $q^* > q_k^{soc}$. □

I.6 Proof of Proposition 4

Under the decomposition $\xi = \lambda\theta + \zeta$ with $\zeta \perp \theta$, the platform signal is $\tilde{s}_f^A = (1+\lambda)\theta + \zeta + u_f^A$, with unconditional variance

$$V_A(\lambda) \equiv (1+\lambda)^2\sigma_\theta^2 + \sigma_\zeta^2 + \sigma_A^2. \tag{P9.1}$$

The fundamental loading is $(1+\lambda)$ rather than 1, so the projection coefficients and the equilibrium loadings of Proposition 1 are re-derived with σ_θ^2 replaced by $(1+\lambda)^2\sigma_\theta^2$ in the numerator of α_θ^A and with V_A replaced by $V_A(\lambda)$; the form is otherwise unchanged.

Encounter variance-share decomposition. The component of an adopter's signal that is orthogonal to θ is $\zeta + u_f^A$, of which only ζ is common across adopters. For two distinct platform adopters f, g, the covariance of evaluation errors normalized by the unconditional signal variance is therefore

$$\frac{\mathrm{Cov}\big(\tilde{s}_f^A - (1+\lambda)\theta,\ \tilde{s}_g^A - (1+\lambda)\theta\big)}{V_A(\lambda)} = \frac{\mathrm{Cov}(\zeta + u_f^A,\ \zeta + u_g^A)}{V_A(\lambda)} = \frac{\sigma_\zeta^2}{V_A(\lambda)}, \tag{P9.2}$$

using $u_f^A \perp u_g^A$ and $u_f^A \perp \zeta$. The informative component $(1+\lambda)\theta$ is shared signal, not shared error, and drops out of the orthogonalized covariance. Weighting by the probability q^2 that two evaluators both adopt gives the $\mathscr{W}_4$ penalty

$$\mathscr{E}_4(q) = q^2\,\frac{\sigma_\zeta^2}{V_A(\lambda)}, \tag{P9.3}$$

the decomposed counterpart of the baseline penalty $q^2\sigma_\xi^2/V_A$ in Theorem 1(ii): the encounter probability is unchanged and the within-adopter variance share is the residual share $\sigma_\zeta^2/V_A(\lambda)$ in place of σ_ξ^2/V_A.

Wedge. The decomposition preserves the continuum and linear-Gaussian structure, so the team-efficiency identity $d\mathscr{W}_1/dq = \Delta\Pi(q)$ of Proposition 3 continues to hold, with $\mathscr{W}_1$, $\Delta\Pi$, and the loadings evaluated under the transformed signal. Since $\sigma_\zeta^2/V_A(\lambda)$ does not depend on q, differentiating $\mathscr{W}_4(q) = \mathscr{W}_1(q) - \gamma_4\mathscr{E}_4(q)$ and substituting the identity gives

$$\Delta\Pi(q) - \frac{d\mathscr{W}_4}{dq} = \gamma_4\,\frac{\sigma_\zeta^2}{V_A(\lambda)}\,\frac{d}{dq}\big[q^2\big] = \frac{2\,\gamma_4\,q\,\sigma_\zeta^2}{V_A(\lambda)}. \tag{P9.4}$$

Parts (i)–(iii). The right-hand side of (P9.4) is strictly positive for all $q \in (0,1]$ whenever $\sigma_\zeta^2 > 0$, and equals zero when $\sigma_\zeta^2 = 0$; the loading λ enters only through $V_A(\lambda)$ in the denominator, establishing part (i). For fixed $q > 0$, $\partial V_A(\lambda)/\partial\lambda = 2(1+\lambda)\sigma_\theta^2 > 0$, so the wedge is strictly

decreasing in λ and vanishes only as $\sigma_\zeta^2 \to 0$, establishing part (ii). The welfare object $\mathscr{E}_4$ depends only on the adoption rate and signal-structure primitives, not on the equilibrium loadings through which δ operates, so the wedge at any fixed q is independent of δ, exactly as in the proof of Theorem 1(ii); this establishes part (iii). Positivity of $\Delta(q)$ on $[0,1]$ under the re-derived loadings holds on $\mathscr{D}(\lambda)$ by the endpoint argument of Proposition 1 with σ_θ^2 replaced by $(1+\lambda)^2\sigma_\theta^2$. □

Proof of Lemma 2

By Definition 1, $\kappa = \nu\,[S_T(\theta \mid \mathscr{Z}) - S_T(\theta \mid \mathscr{Z}, P_{\text{plat}})]$ with $\nu > 0$. It suffices to bound the bracketed conditional epiplexity below by a positive constant independent of platform accuracy.

Write the platform output as $P_{\text{plat}}(\mathscr{Z}) = G(\mathscr{Z})$ for a forecasting transformation G that is evaluable within the budget T. Recovering the ξ-free fundamental from the chain's signal requires a program that, given $G(\mathscr{Z})$, returns a forecast of θ whose residual is orthogonal to ξ. Any such program induces an inverse of the ξ-direction of G. By hypothesis this inverse is not computable within T. The conditional epiplexity of θ given $\mathscr{Z}$ and P_{plat} therefore cannot be reduced to a constant overhead on top of $|P_{\text{plat}}|$: in the language of Finzi et al. (2026, Theorem 12 and the failure of the limited-epiplexity-increase property), possession of an efficient forward program G does not yield an efficient program for the inverse, so the bound

$$S_T(\theta \mid \mathscr{Z}, P_{\text{plat}}) \;\le\; S_T(\theta \mid \mathscr{Z}) - \underline{s}$$

fails to hold with $\underline{s}$ equal to the full structural content of G. Equivalently, the conditional epiplexity retains a residual $\underline{s}_0 > 0$ that the forward program cannot absorb. Setting $\underline{\kappa} = \nu\,\underline{s}_0 > 0$ gives $\kappa \ge \underline{\kappa}$.

For the comparative static in platform quality, let a quality improvement raise $S_T(\theta \mid \mathscr{Z})$ by $\Delta > 0$, so that the forecasting task, viewed from scratch, now extracts more structure. The same improvement is embedded in G, so it raises $S_T(\theta \mid \mathscr{Z}, P_{\text{plat}})$ by at most Δ as well, because conditioning on a more informative P_{plat} cannot raise the residual structural information by more than the improvement itself. Hence the bracketed difference is non-increasing in platform quality, and κ does not fall. □

I.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,

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

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

Encounter 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

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

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 $\rho = 1$. □

I.8 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. □

I.9 Proof of Proposition 6

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 $\Delta(q)$,

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

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 $\phi > 0$. □

I.10 Proof of Proposition 10

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 6), 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 5, 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 6(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 $t \geq 1$. □

Proof of Proposition 7

We derive the two first-order conditions, establish the wedge, and treat the limits. Throughout, $q \in (0,1]$ and $\delta \in \mathscr{D}$ are fixed, and we work in the linear-Gaussian regime of Proposition 1.

Private problem. A chain placing weight ω on the platform forms $\tilde{s}_f(\omega) = \theta + \omega\,\xi + \omega\,u_f^A + (1-\omega)\,u_f^C$. Its signal has fundamental loading one and total noise variance

$$\Omega(\omega) \;\equiv\; \omega^2\,\sigma_\xi^2 + \omega^2\,\sigma_A^2 + (1-\omega)^2\,\sigma_C^2,$$

where the ξ and u^A terms enter with weight ω and the custom noise with weight $(1-\omega)$. As in the proof of Lemma 1, the chain's expected entry profit in the linear regime is, up to terms independent of ω, a strictly decreasing affine function of the posterior variance of $\theta + \delta\bar{x}$ given $\tilde{s}_f(\omega)$, which is itself increasing in $\Omega(\omega)$. Write the chain's gross private payoff as $\pi(\omega) = -\frac{1}{2}\,\psi\,\Omega(\omega) + \text{const}$, where $\psi > 0$ collects the sensitivity of profit to forecast precision. The coefficient ψ depends on (δ, q) and the variance primitives through the equilibrium loadings but not on ω. The chain solves

$$\max_{\omega \in [0,1]} \;\; \pi(\omega) - \tfrac{1}{2}\,\kappa\,(1-\omega)^2.$$

The objective is strictly concave: $\Omega''(\omega) = 2(\sigma_\xi^2 + \sigma_A^2 + \sigma_C^2) > 0$ and the cost is convex, so the second derivative of the objective is $-\psi\,(\sigma_\xi^2 + \sigma_A^2 + \sigma_C^2) - \kappa < 0$. Define the curvature

$$\Xi \;\equiv\; \psi\,(\sigma_\xi^2 + \sigma_A^2 + \sigma_C^2) + \kappa \;>\; 0,$$

so that the objective's second derivative is exactly $-\Xi$. The interior first-order condition is

$$-\tfrac{1}{2}\psi\,\Omega'(\omega) + \kappa(1-\omega) \;=\; 0, \qquad \Omega'(\omega) = 2\omega(\sigma_\xi^2 + \sigma_A^2) - 2(1-\omega)\sigma_C^2,$$

which rearranges to

$$\omega^* \;=\; \frac{\psi\,\sigma_C^2+\kappa}{\psi\,(\sigma_\xi^2+\sigma_A^2+\sigma_C^2)+\kappa}, \tag{R.1}$$

clipped to $[0,1]$. The private optimum is interior whenever $\sigma_\xi^2+\sigma_A^2>0$, which holds under $\sigma_\xi^2>0$, and reaches the upper corner $\omega^*=1$ only as $\sigma_C^2\to\infty$ or $\kappa\to\infty$.

Planner problem. The planner internalizing $\mathscr{W}_4$ values the same gross payoff and cost but subtracts the correlated-evaluation externality. From Section 5.2, two adopters placing platform weight ω each contribute to the decided-about-agent welfare object a normalized error covariance $\omega^2\sigma_\xi^2/V_A$, scaled by the encounter probability that the two screen a common location. Holding the population symmetric at weight ω, the planner's per-adopter objective adds the term $-\gamma_4\,q\,\omega^2\,\sigma_\xi^2/V_A$ to the private objective.[25] The planner solves

$$\max_{\omega\in[0,1]}\;\pi(\omega)-\tfrac{1}{2}\kappa(1-\omega)^2-\gamma_4\,q\,\omega^2\,\frac{\sigma_\xi^2}{V_A}.$$

Before differentiating, we verify that the gross payoff $\pi(\omega)$ is evaluated at loadings that remain team-optimal, so that no indirect terms in the equilibrium loadings enter the planner's first-order condition. This is the envelope step. By Proposition 3, for any fixed configuration of exposures the equilibrium loadings (k^A,k^C) maximize average chain profit under $\mathscr{W}_1$, with the cross-sectional identity $\int(P_g-\bar{x})\,dg=0$ driving the general-equilibrium feedback to zero. The blending weight ω enters the welfare objects $\mathscr{W}_1$ and $\mathscr{W}_4$ only through the per-adopter signal variance $\Omega(\omega)$ and through the exposure ω in the externality term; it does not disturb the team-optimality of the loadings at fixed ω. Differentiating the planner's objective in ω therefore does not generate indirect terms $\sum_\tau(\partial\mathscr{W}/\partial k^\tau)(k^\tau)'(\omega)$: those vanish by the same envelope argument that delivers $d\mathscr{W}_1/dq=\Delta\Pi(q)$ in Proposition 3. The planner's first-order condition is the private condition augmented by the derivative of the externality term alone:

$$-\tfrac{1}{2}\psi\,\Omega'(\omega)+\kappa(1-\omega)-\frac{2\gamma_4 q\sigma_\xi^2}{V_A}\,\omega \;=\; 0. \tag{R.2}$$

Part (i): over-reliance. Let $\Phi(\omega)\equiv-\frac{1}{2}\psi\Omega'(\omega)+\kappa(1-\omega)$ denote the left-hand side of the private first-order condition, so $\Phi(\omega^*)=0$. The planner's condition (R.2) is $\Phi(\omega)-\frac{2\gamma_4 q\sigma_\xi^2}{V_A}\omega=0$. Evaluating the planner's marginal value at the private optimum,

$$\Phi(\omega^*)-\frac{2\gamma_4 q\sigma_\xi^2}{V_A}\,\omega^* \;=\; -\frac{2\gamma_4 q\sigma_\xi^2}{V_A}\,\omega^* \;<\; 0$$

[25] The encounter probability that a given adopter shares the screening of a location with another adopter is proportional to q; the externality this adopter generates, summed over the locations it screens, scales its own weight ω against the population weight, contributing the marginal term $2\gamma_4 q\omega\sigma_\xi^2/V_A$ to the planner's first-order condition. This is the intensive-margin counterpart of the $q^2\sigma_\xi^2/V_A$ object in Theorem 1(ii), differentiated with respect to exposure rather than adoption.

whenever $\sigma_\xi^2 > 0$, $q > 0$, and $\omega^* > 0$. The planner's objective is strictly concave, since the externality term $-\gamma_4 q \omega^2 \sigma_\xi^2 / V_A$ is concave and adds to an already strictly concave private objective, so its first-order condition has a unique root and its marginal value is strictly decreasing in ω. A negative marginal value at ω^* therefore implies the planner's root lies strictly below: $\omega^{soc} < \omega^*$. The boundary case $\omega^* = 0$ requires $\sigma_C^2 = 0$ by (R.1) and is excluded under $\sigma_C^2 > 0$.

Part (ii): the wedge. The function Φ is affine in ω with slope

$$\Phi'(\omega) = -\tfrac{1}{2}\psi\,\Omega''(\omega) - \kappa = -\psi(\sigma_\xi^2 + \sigma_A^2 + \sigma_C^2) - \kappa = -\Xi.$$

The two first-order conditions read $\Phi(\omega^*) = 0$ and $\Phi(\omega^{soc}) = \frac{2\gamma_4 q \sigma_\xi^2}{V_A}\,\omega^{soc}$. Subtracting and using affineness, $\Phi(\omega^*) - \Phi(\omega^{soc}) = \Phi' \cdot (\omega^* - \omega^{soc}) = -\Xi\,(\omega^* - \omega^{soc})$, so

$$-\Xi\,(\omega^* - \omega^{soc}) \;=\; -\frac{2\gamma_4 q \sigma_\xi^2}{V_A}\,\omega^{soc} \qquad \Longrightarrow \qquad \omega^* - \omega^{soc} \;=\; \frac{2\,\gamma_4\,q\,\omega^{soc}}{\Xi} \cdot \frac{\sigma_\xi^2}{V_A},$$

which is the wedge in the statement of the proposition. Strict positivity follows from $\gamma_4, q, \omega^{soc}, \sigma_\xi^2, \Xi > 0$.

Monotonicity in q. The wedge is proportional to $q\,\omega^{soc}(q)$. The direct effect of q is positive. The induced effect operates through ω^{soc}, which is non-increasing in q because a higher encounter rate raises the marginal external cost in (R.2) and pushes the planner toward lower reliance. The product $q\,\omega^{soc}(q)$ is nonetheless increasing on the admissible region: implicit differentiation of (R.2), whose slope in ω is $-\Xi - 2\gamma_4 q \sigma_\xi^2 / V_A$, gives $d\omega^{soc}/dq = -(2\gamma_4 \sigma_\xi^2 / V_A)\,\omega^{soc} / (\Xi + 2\gamma_4 q \sigma_\xi^2 / V_A)$, so

$$\frac{d}{dq}\big(q\,\omega^{soc}\big) = \omega^{soc} + q\,\frac{d\omega^{soc}}{dq} = \omega^{soc} \cdot \frac{\Xi}{\Xi + 2\gamma_4 q \sigma_\xi^2 / V_A} > 0.$$

The factor multiplying ω^{soc} is strictly between zero and one, so the direct effect dominates and the wedge is strictly increasing in q for every admissible parameter configuration, with no auxiliary smallness condition required.

Monotonicity in σ_ξ^2. The dependence on σ_ξ^2 is not globally monotone, and characterizing it requires accounting for the fact that the precision sensitivity ψ and the total variance $V_A = \sigma_\theta^2 + \sigma_\xi^2 + \sigma_A^2$ both move with σ_ξ^2. Substituting ω^{soc} into the wedge and differentiating, the sign of $d(\omega^* - \omega^{soc})/d\sigma_\xi^2$ equals the sign of $-B(\sigma_\xi^2)$, where, treating ψ as locally constant in σ_ξ^2,

$$B(\sigma_\xi^2) \;=\; 2\psi^2\,\sigma_\xi^6 \;+\; \psi\big(2\gamma_4 q + 2\kappa + 3\psi\sigma_A^2 + 2\psi\sigma_C^2 + \psi\sigma_\theta^2\big)\,\sigma_\xi^4 \;-\; (\sigma_\theta^2 + \sigma_A^2)\big(\kappa + \psi\sigma_A^2 + \psi\sigma_C^2\big)^2,$$

in which all coefficients of positive powers of σ_ξ^2 are positive and the constant term is strictly negative. The wedge is therefore increasing in σ_ξ^2 exactly when $B(\sigma_\xi^2) < 0$. Because B has a strictly negative value at $\sigma_\xi^2 = 0$ and a strictly positive derivative for $\sigma_\xi^2 > 0$ (its derivative $6\psi^2\sigma_\xi^4 + 2\psi(\cdots)\sigma_\xi^2$ is positive and it has no linear term), B crosses zero exactly once, at a threshold $\sigma_\xi^{2\dagger} > 0$.

The wedge is strictly increasing in σ_ξ^2 on $(0, \sigma_\xi^{2\dagger})$ and strictly decreasing beyond it. A transparent sufficient condition for the increasing region, obtained by bounding the constant term below by $(\sigma_\theta^2+\sigma_A^2)\kappa^2$ and the positive part above, is

$$\psi\big(2\gamma_4 q+2\kappa+3\psi\sigma_A^2+2\psi\sigma_C^2+\psi\sigma_\theta^2\big)\,\sigma_\xi^4+2\psi^2\sigma_\xi^6 \;<\; (\sigma_\theta^2+\sigma_A^2)\,\kappa^2.$$

Economically, the wedge rises with the severity of the shared component while that severity is small or moderate relative to the rest of the signal, which is the regime in which the platform is adopted at all (Lemma 1). It can fall only once σ_ξ^2 is so large that the within-adopter variance share σ_ξ^2/V_A approaches one and saturates while the planner's chosen reliance ω^{soc} collapses toward zero; in that regime the encounter-scaled exposure stops growing because there is almost nothing left to remove. The local rather than global character of the comparative static is therefore a feature of the saturation of the variance share, not a fragility of the over-reliance result, which holds for every $\sigma_\xi^2>0$ by Part (i).

δ-invariance at fixed q. The externality term $\gamma_4 q\omega^2\sigma_\xi^2/V_A$ depends only on signal-structure primitives and the chosen exposure ω, not on the equilibrium loadings, which are the only channel through which δ enters. The marginal external cost in (R.2) is therefore independent of δ. The gross precision sensitivity ψ does depend on δ, but it enters both first-order conditions identically through Φ and cancels in the difference $\Phi(\omega^*)-\Phi(\omega^{soc})$. Hence the wedge at fixed q is independent of δ. □

Part (iii): limits. From (R.1), $\omega^*\to 0$ requires the numerator $\psi\sigma_C^2+\kappa\to 0$. With $\sigma_C^2>0$ fixed and $\psi>0$, this requires both $\kappa\to 0$ and $\psi\sigma_C^2\to 0$. The latter holds only as the precision value of forecasting collapses, which under $\sigma_\xi^2\to\infty$ drives $\psi\to 0$ because an infinitely noisy platform signal carries no precision benefit. Thus $\omega^*=0$ is approached only in the limits $\kappa\to 0$ or $\sigma_\xi^2\to\infty$. By Lemma 2, $\kappa\geq\underline{\kappa}>0$ uniformly in platform quality, so the $\kappa\to 0$ limit is not approached as platforms improve. For any $\kappa>0$ and finite $\sigma_\xi^2>0$, $\omega^*>0$ and, by Part (i), $\omega^*>\omega^{soc}$. □

I.11 Proof of Proposition 8

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 $\sigma_\xi^2(q)$. □

I.12 Proof of Proposition 9

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. □