Shared Intelligence and Misaligned Firm Entry

Feng Chi* Nathan Yang†

June 13, 2026

Abstract

Firms increasingly rely on shared intelligence platforms for market entry. While these systems improve forecast accuracy, they route decisions through a common layer, creating correlated errors unrelated to fundamentals. We show that platform adoption is socially excessive because firms ignore two external costs: synchronized entry that strains local markets, and correlated screening errors that misallocate sites. This inefficiency persists regardless of strategic interactions between competitors. Adding vendors reduces these costs only insofar as their underlying data and models are independent; when platforms share a common core, the vendor count is irrelevant. A case study of restaurant site-selection analytics indicates that nominal vendor competition delivers limited effective diversification: under calibrations disciplined by vendor disclosures and upstream data overlap, the observed market retains roughly three-quarters of the single-platform distortion.

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

The retail landscape is increasingly shaped by a common set of algorithmic forecasting platforms. Historically, retail entry relied on local, firm-specific information (Love et al. 1986). A regional grocer learned its trade area through direct experience, while a restaurant chain relied on proprietary market research. Today, however, national chains, regional grocers, and franchisees increasingly rely on the same third-party analytics providers (e.g., CoStar, Placer.ai, and Buxton) to forecast trade-area demand, foot traffic, and competitive density.[1] While these shared intelligence platforms can improve individual forecast accuracy by pooling data and reducing idiosyncratic noise, they also route many firms' decisions through a common processing layer. As noted by the Bank of England's Financial Policy Committee (Bank of England Financial Policy Committee 2025), the potential for AI-based participants to take increasingly correlated positions is a growing concern for systemic stability. The central tension is that when a shared processing layer contains a common error unrelated to local fundamentals, adoption not only improves individual forecasts but also makes industry-wide mistakes more correlated.

This phenomenon is readily observable in the quick-service restaurant industry. 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, even after accounting for observable demand fundamentals such as population, density, and income. Furthermore, this alignment is especially pronounced in counties with greater exposure to AI-supported analysis (Felten et al. 2021). While this co-location could stem from common demand fundamentals, agglomeration economies, or strategic imitation, the rapid proliferation of shared analytics vendors suggests an additional, distinct channel. If competing firms rely on the same forecasting infrastructure, they may enter the same markets not only because those markets are attractive, but also because the same underlying processing error points them there.

Standard models of retail entry do not readily generate this form of correlation.[2] In the existing

[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`.

[2] In canonical frameworks (Bresnahan and Reiss 1991, Seim 2006, Jia 2008, Holmes 2011), firms receive information about

literature, correlation in firm behavior is typically attributed to strategic interaction, reputational herding (Scharfstein and Stein 1990, Froot et al. 1992), or informational cascades (Banerjee 1992, Bikhchandani et al. 1992). In contrast, the channel we study operates through a shared processing shock that is orthogonal to fundamentals. Recent evidence supports the concern that algorithmic prediction can improve individual decisions while increasing correlation across firms. For instance, firms investing heavily in AI exhibit increases in systematic risk (Babina et al. 2024), and firms sharing data through common software development kits (SDKs) exhibit co-movement that survives controls for product similarity (Bian et al. 2024). 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.

To address this question, we develop a parsimonious model of firm entry in which a continuum of potential entrants chooses whether to rely on custom research or a shared forecasting platform. Custom research generates idiosyncratic forecast errors, whereas the shared platform reduces idiosyncratic noise but introduces a common processing shock shared across adopters. Our model admits a closed-form equilibrium with heterogeneous adoption and cleanly separates three objects that are often conflated: strategic interaction, shared processing, and the resulting welfare loss. The main finding of our analysis 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 (e.g., lenders, landlords, and municipalities), and one borne by locations evaluated by multiple firms through the same forecasting technology. Consequently, equilibrium adoption is excessive. Notably, this over-adoption result does not depend on whether entry decisions are strategic complements or substitutes; in the location-screening case, the welfare discrepancy depends only on the adoption rate and signal-structure primitives.[3] Furthermore, we extend the baseline model to investigate the role of vendor diversity and dynamic switching. A dynamic extension demonstrates that shared processing induces persistence in aggregate outcomes

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] A closely related concern arises in the AI-trading literature. Dou et al. (2025) show that AI speculators equipped with planning modules autonomously learn intertemporal strategies that create and burst bubbles, amplifying market fragility. In their setting, correlation is engineered: agents internalize the shared exposure and coordinate to exploit it. In our work, the correlation is incidental, as it arises precisely because adopters do not internalize the externality, and over-adoption follows from the signal structure alone, without coordination or intent. The two mechanisms are poles of the same phenomenon, and the policy levers differ accordingly.

even when fundamentals are independent over time, and that lower switching costs accelerate the realization of the externality. To quantify the practical implications of our theory, we conduct a descriptive case study of the U.S. restaurant site-selection analytics market. Although several vendors compete at the front end (e.g., Placer.ai, Tango, Buxton), adoption is concentrated and many tools rely on overlapping upstream inputs such as roadway traffic data and mobility panels. This structure is important because vendor competition need not imply independent decision rules. Recent evidence from delegated algorithmic pricing shows that firms often outsource strategic algorithm development to third-party contractors or LLMs, and that these delegates can produce systematic, economically meaningful design errors even when firms provide nontechnical instructions (Ershov and Lyons 2024). In that setting, the relevant risk is not only explicit algorithmic collusion, but also common mis-specification and reliance on similar reduced-form prediction routines. The same logic applies to site-selection analytics: multiple vendor brands may still transmit correlated errors if their models are trained on similar data, built with similar forecasting conventions, or embedded in similar outsourced decision pipelines. Under calibrations disciplined by vendor disclosures and evidence on upstream-input overlap, the observed market retains roughly three-quarters of the single-platform distortion despite multiple vendor brands. This indicates that nominal vendor competition provides limited effective diversification when front-end providers draw on overlapping upstream data sources and common algorithmic design practices.

Our paper contributes to three streams of literature. First, we contribute to the literature on the economics of data and platform market power (Kirpalani and Philippon 2020, Acemoglu et al. 2022, Choi et al. 2019, Ichihashi 2021). While prior work focuses on externalities in data collection or platform bargaining power, we highlight an externality that operates one layer downstream, in the processing that converts data into forecasts, which falls on third parties exposed to correlated decisions. Second, we contribute to the literature on the social value of information (Morris and Shin 2002, Angeletos and Pavan 2007, Hellwig and Veldkamp 2009) and algorithmic monoculture (Kleinberg and Raghavan 2021, Peng and Garg 2024). We extend this literature into an entry environment where shared evaluation shapes both aggregate outcomes and welfare. Similar to the division of labor identified by Clayton and Coppola (2026) in financial regulation, our framework separates the structural impact of policy from the predictive component of shared models, allowing the welfare wedge to depend solely on signal-structure primitives.[4] Finally, we contribute to the

[4]The invariance mechanisms differ across the two settings. In Clayton and Coppola (2026), policy invariance holds because

empirical literature on retail entry and industry dynamics (Bresnahan and Reiss 1991, Igami and Yang 2016, Blevins et al. 2018). By providing a structural interpretation of cross-firm alignment, we offer a new mechanism for correlated entry that is distinct from agglomeration, imitation, or strategic preemption.

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 whether the two chains 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, and it is not offered as causal evidence of any particular mechanism. Its purpose is 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, especially in high-AI-exposure counties.

2.1 Industry Setting and Data

We assemble a county-level snapshot of McDonald's and Burger King locations in the United States.[5] Brand locations come from OpenStreetMap, and each location is assigned to a county

private agents observe the state while the regulator does not. Here, invariance is a property of the welfare object rather than of equilibrium loadings: the decided-about-agent wedge does not load on strategic interaction because it depends only on the encounter probability and the variance share of the shared component.

[5] We focus on McDonald's and Burger King because they are large national chains in the same product segment and have broad geographic footprints. In the 2025 QSR 50, which ranks U.S. limited-service chains by 2024 system-wide sales, McDonald's is the largest chain overall and the largest burger chain, while Burger King ranks eighth overall and third among burger chains. Their common segment helps hold product category, service format, and target occasion relatively fixed, while their national footprints make county-level variation informative about broad site-selection patterns rather than the expansion strategy of a narrow regional brand. Both chains also maintain substantial in-house real-estate and analytics capabilities, so the pair is, if anything, conservative for the shared-platform channel, which Section 3 argues binds most strongly for mid-size chains and franchisees that consume platform forecasts as point estimates. The vendor disclosures assembled in Appendix E identify chains that are clients of the same platform (e.g., Shake Shack, Five Guys, and Freddy's Frozen Custard at eSite Analytics); comparing alignment among disclosed co-clients of a common vendor against alignment among non-co-client pairs is a sharper descriptive cut that we leave to future work with store-census data.

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).

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 weights these 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 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. 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 only 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 cosine similarity of the coefficient vectors, correlation of the coefficient vectors, and correlation between 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, treating fitted values as 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 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 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 indicators for at least one McDonald's location and 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. Co-location among burger chains is not new: Toivanen and Waterson (2005) document strategic interaction and co-location in U.K. hamburger-chain entry using data that predate the modern analytics platforms studied here. The mechanism we formalize therefore does not claim to originate cross-firm alignment; it predicts that shared forecasting infrastructure adds a component of alignment that is orthogonal to fundamentals and that scales with platform penetration. A single cross-sectional snapshot cannot separate this component from the longstanding channels what Toivanen and Waterson (2005) identify, which is precisely why Section B states the identifying restriction in terms of platform-side update events rather than levels of co-location.

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 whether they select markets according to similar observable characteristics. 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 county 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 coefficients can be compared across brands.

The two chains differ in scale. McDonald's has 4.85 locations per county on average, compared with 2.11 for Burger King. Conditional on scale, however, the observable market-selection rules 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. Observable fundamentals therefore rank counties in almost the same way for McDonald's and Burger King, even though McDonald's has the larger footprint.

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.

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

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.

The similarity is high in both subsamples, but tighter in high-AIGE counties. The predicted-score correlation is 0.9849 in low-AIGE counties and 0.9978 in high-AIGE counties, with coefficient-vector similarity showing the same pattern. This does not imply that AI exposure causes chains to behave more similarly. It shows that counties more exposed to AI, as measured by local industry and occupation composition, are also places where the two chains' observable market-selection rules are especially close in directional terms.

These patterns motivate the model in two ways. McDonald's and Burger King co-locate strongly across counties, and they 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 and 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.

We emphasize that the descriptive alignment is consistent with the shared-processing mechanism but does not identify it: a cosine similarity cannot separate a common processing error ξ from common fundamentals θ, agglomeration, or imitation, each of which predicts that the two

chains select similar markets. Isolating ξ requires variation that moves the processing layer without moving fundamentals which we formalize as an identifying restriction in Appendix B and leave to future work to implement. The exercise here should therefore be read as motivating the object the model formalizes, not as a structural calibration of it or a test of its central mechanism.

As an exploratory diagnostic, we decompose the high-minus-low AIGE difference in predicted-score alignment across the seven continuous fundamentals using exact Shapley attribution.[6] For each subset of fundamentals, we assign both chains the high-AIGE coefficients for those fundamentals and the low-AIGE coefficients for the remaining fundamentals, then evaluate 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. The seven contributions sum to the total high-minus-low alignment difference.

Because predicted-score alignment is close to one in both subsamples, the difference being decomposed is small. The resulting contributions are small, partially offsetting, and imprecisely estimated; bootstrap confidence intervals are wide enough 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 instead points to broad-based alignment across observables.

3 Model

The model studies retail entry when firms choose both whether to enter and which information technology to use when forecasting profitability. The environment consists of a single market, two information technologies, and a continuum of potential entrants. This sparse structure isolates the feature of interest: a shared processing layer that introduces correlated errors across adopters while preserving closed-form tractability.

A key distinction is between primitives and outcomes. The variance of the shared component is a primitive of the information structure, whereas platform adoption is an endogenous equilibrium

[6] 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.

outcome reflecting firms' responses to those primitives.[7]

3.1 Environment

A unit continuum of retail chains $f \in [0,1]$ decides whether to enter a local market. The entry 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. Positive values capture agglomeration economies, such as food halls, outlet centers, or fast-food strips where rival entry expands demand; negative values capture business stealing in markets with limited demand. The restriction $\delta > -1$ ensures a finite, positive mean entry rate $K_0 = 1/(2(1-\delta))$ in the linear equilibrium; existence of a unique linear equilibrium additionally requires the upper bound on δ collected in the admissible region $\mathscr{D}$ of Section 4.

The uniform cost shock is a payoff-smoothing device in the spirit of the global-games literature (Carlsson and van Damme 1993, Morris and Shin 1998); its role here is not equilibrium selection but linearity, as it makes entry probabilities linear in posterior beliefs. We work in the regime where this linear approximation is informative.[8] We normalize $\bar{\varepsilon} = 1/2$, so entry probability responds one-for-one to the posterior mean. This normalization is without loss of generality and does not affect the qualitative results.

3.2 Information and Processing Technologies

The two information technologies 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

[7] The analysis treats σ_ξ^2 as a fixed primitive of the information structure throughout.

[8] Formally, linearity holds whenever $|\mathbb{E}_f[\theta] + \delta \mathbb{E}_f[\bar{x}]| \leq \bar{\varepsilon}$ pointwise; a sufficient condition is $\bar{\varepsilon} > 4\sqrt{\text{Var}(\theta) + \text{Var}(\bar{x})}$. Some calibrations used in the figures lie outside this bound to illustrate qualitative regimes. The closed-form expressions of Proposition 1 remain well-defined wherever $\Delta(q) > 0$, and the qualitative comparative statics in Theorem 1 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.

across adopters because competing firms rely on similar tools. Related evidence shows that firms investing more heavily in AI exhibit higher systematic risk (Babina et al. 2024), and that firms relying on external AI providers behave differently from those building internally (Calvino and Fontanelli 2026). This literature identifies an important source of heterogeneity in firm AI use, but does not characterize the welfare implications of shared vendor solutions adopted at scale.

The signal structure below formalizes this margin. Custom research generates idiosyncratic forecast errors that differ across firms. Platform-based forecasting may reduce this idiosyncratic variation, but it also introduces a common processing component shared across adopters. A chain using custom research observes

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

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

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

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

The distinction between the two technologies is therefore 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 replaces part of the idiosyncratic error with shared noise. This trade can be privately attractive while generating aggregate effects that no individual chain internalizes.

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.

This structure differs 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 inherits a common mistake.[9]

The shared component can be interpreted as a latent processing factor. 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 unrelated to fundamentals (Baker and Wurgler 2006, 2007) and correlated belief errors (De Long et al. 1990).

A related 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. The mechanism here formalizes a different source of the gap between subjective and objective beliefs. Adopters' beliefs depart from the objective distribution not because of cognitive bias, but because forecasts are outsourced to a shared technology that introduces correlated error. The welfare consequence is similar because both mechanisms generate a common error component that firms do not internalize, but the policy levers differ. Shared algorithmic errors leave an audit trail, such as model versions, training data, and parameter updates, so interventions such as vendor diversification, model-update disclosure, or training-data audits can target the source of 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

[9] 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.

is a rational-expectations fixed point in which individual choices reproduce the aggregate:

$$q^* = F(\Delta\Pi(q^*)),$$

where $\Delta\Pi(q) \equiv \pi^A(q) - \pi^C(q)$ is the private gain from adoption. If F has bounded support $[0, \bar{c}]$, corner outcomes arise when $\Delta\Pi(0) \leq 0$ or $\Delta\Pi(1) \geq \bar{c}$. With unbounded support, such as the logistic specification used below, these cases are replaced by interior equilibria close to but bounded away from 0 or 1. This distinction does not affect the qualitative results.

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

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

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 assign it its prior mean. The platform signal's common component is summarized by

$$\chi_A \equiv \alpha_\theta^A + \alpha_\xi^A = \frac{\sigma_\theta^2 + \sigma_\xi^2}{V_A},$$

the share of signal variance attributable to components common across adopters. This object governs how strongly the shared shock enters equilibrium behavior.[10]

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 has the analytic capacity to characterize the platform's residual error structure. Such a chain recognizes that the forecast inherits noise from underlying data sources, that this noise varies across trade areas, and that this variation is independent of local demand. It therefore filters the platform signal correctly and weights it against its prior as in Proposition 1.

[10] The scalar χ_A is a variance share and is distinct from the profit functions $\pi^A(q)$ and $\pi^C(q)$ introduced in Section 5, which are equilibrium payoffs.

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 may instead observe only 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 then weighted as if it were more informative than it is, and entry decisions over-respond to the common component.

The rational-filtering benchmark therefore delivers a lower bound on platform-driven distortion. When chains under-correct for σ_ξ^2, 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 isolates the welfare wedge generated by the information structure itself, abstracting from additional distortions introduced by misperception.[11]

4 Equilibrium with Heterogeneous Adoption

This section characterizes equilibrium behavior for a fixed adoption rate q. Throughout, 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 $\Delta(q)$ in Proposition 1 is strictly positive for all $q \in [0,1]$.

4.1 Best Responses and the Linear Fixed Point

A chain enters when its posterior expected profit is nonnegative,

$$\mathbb{E}_f[\theta] + \delta \mathbb{E}_f[\bar{x}] + \varepsilon_f \geq 0.$$

[11] The binary technology choice places all exposure to the shared component on the extensive margin of adoption. Allowing firms to choose a continuous platform weight does not eliminate the externality: when custom research is costly, the privately optimal platform weight strictly exceeds the social optimum whenever $\sigma_\xi^2 > 0$, because each chain internalizes the private precision cost of relying less on the platform but not the correlation cost imposed on co-screened locations. The correlated-evaluation externality therefore reappears on the intensive margin.

With $\varepsilon_f \sim \text{Uniform}[-1/2, 1/2]$, the entry probability is

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

in the interior region. We look for an equilibrium in which the aggregate entry rate is linear in the underlying shocks:

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

In a continuum, idiosyncratic noise averages out. If k^A and k^C denote the slopes of entry probabilities with respect to signals for platform adopters and custom chains, consistency requires

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

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

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

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

where $\chi_A = \alpha_\theta^A + \alpha_\xi^A = (\sigma_\theta^2 + \sigma_\xi^2)/V_A$. The equilibrium exists and is unique on $\mathscr{D}$.[12]

The loadings differ because platform adopters respond directly to both θ and ξ, while custom chains respond only to θ. The term $-\delta q \alpha_\xi^A$ in k^C captures the indirect effect of the shared shock on custom chains through aggregate entry. Although custom chains do not observe ξ, they anticipate that adopters respond to it. They therefore offset this effect when $\delta < 0$ and reinforce it when $\delta > 0$,

[12] The two bounds in $\mathscr{D}$ ensure $\Delta(0) > 0$ and $\Delta(1) > 0$, respectively, and concavity of $\Delta(q)$ implies positivity throughout $[0,1]$. The lower bound $\delta > -1$ is sufficient for the global-games argument but not for existence of the linear equilibrium, which also requires the upper bound in $\mathscr{D}$. For $\delta \geq 1$, this requires $\sigma_C^2/\sigma_\theta^2 > \delta - 1$ 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 extend analytically to such cases but should then be interpreted as approximations to the truncated equilibrium described in Section 3.

with the adjustment becoming stronger as adoption rises or 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 private-signal entry game, with $k^C(0) = \alpha_\theta^C/(1-\delta\alpha_\theta^C)$. When $q = 1$, all chains adopt the platform and $k^A(1) = \alpha_\theta^A/(1-\delta\chi_A)$, where the denominator reflects that the platform 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

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

The intercept $K_0 = 1/(2(1-\delta))$ can exceed unity for δ close to one, in which case the linear expression for $\bar{x}$ should be read as the interior-regime approximation discussed in Section 3. None of the objects analyzed below is affected: the variance decomposition, the adoption gain $\Delta\Pi(q)$, and every welfare wedge depend only on deviations of $\bar{x}$ from its mean, so the intercept drops out of all subsequent expressions.

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 affected parties outside the chain population. The analysis shows that coordination motives are internalized in equilibrium, but the shared component generates two externalities that are not. Either 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, after 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

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

The dependence of $\Delta\Pi(q)$ on q reflects both forecast precision and strategic alignment. The platform changes the precision of a chain's information and also changes how closely its entry decision tracks rivals. The relative importance of these effects is governed by δ.

Proposition 2 (Monotonicity of the adoption gain)**.** *The function $\Delta\Pi(q)$ is independent of q when $\delta = 0$. Whenever the polynomial factor $R(q;\sigma,\delta)$ defined in Appendix F is strictly positive on $\mathscr{D}$ (i.e., a condition we verify numerically over* 200,000 *parameter draws in the replication package) $\Delta\Pi(q)$ is 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 on the signal-precision tradeoff. When $\delta > 0$, adoption becomes more attractive as it spreads because platform adopters track the shared component that increasingly shapes aggregate entry. When $\delta < 0$, the same mechanism reverses: chains benefit from differentiation, so the correlation induced by platform adoption reduces its marginal value.

The result follows from the derivative of $\Delta\Pi(q)$. Differentiating the closed-form expression derived in Appendix F gives

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

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 externalities, each affecting a different group.

The first operates at the aggregate level. It affects agents whose payoffs depend on the variance of entry around fundamentals, such as lenders exposed to correlated default risk, landlords facing synchronized vacancies, and municipalities 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 location level. When multiple chains evaluate the same location through the shared platform, their signals contain the same realization of ξ, making their evaluation errors positively correlated. As adoption rises, more evaluations share this component, reducing the effective independence of screening across firms. The corresponding welfare object is the expected pairwise covariance of evaluation errors normalized by total signal variance, which scales as $q^2\sigma_\xi^2/V_A$: the probability that two evaluators both adopt the platform times the variance share of the shared component.

This second channel is related to the externality in Kleinberg and Raghavan (2021) and Peng and Garg (2024). In those settings, the decided-about agents are individuals, and the key statistic is the probability that at least one decision-maker recognizes a favorable case. Here the decided-about agent is a location, and the cost comes from correlated evaluation errors across firms screening the same market. The mechanisms are related, but the welfare objects differ.

Remark 1 (A screening micro-environment for $\mathscr{W}_4$)**.** The baseline environment features a single

market evaluated by the full continuum, so the pairwise object in $\mathscr{W}_4$ deserves an explicit micro-environment. Suppose that, alongside the entry game, there is a unit continuum of ex-ante identical candidate locations, each screened by two chains drawn independently at random from the population, with each evaluator applying its equilibrium forecasting technology to a location-specific realization of (θ, ξ, u). For a given location, the normalized covariance of the two evaluators' forecast errors equals σ_ξ^2/V_A when both evaluators are platform adopters and zero otherwise, so the expected normalized error covariance across locations is exactly $q^2\,\sigma_\xi^2/V_A$, the penalty term in $\mathscr{W}_4$. The weight γ_4 prices the loss of an independent second opinion: when evaluation errors are uncorrelated, a location wrongly screened out by one chain retains a chance of being recognized by the other, whereas correlated errors eliminate this diversification, in the spirit of Kleinberg and Raghavan (2021). Because a location passed over by both evaluators in error and a location entered by both in error are both costly, these losses do not net out across locations even though the shared error is mean zero; the relevant social loss is therefore the covariance itself rather than its (zero) expectation.

The two externalities also differ in their dependence on equilibrium behavior. The aggregate-volatility externality depends on equilibrium loadings, while the correlated-evaluation externality depends only on the signal structure and adoption rate. Both originate from the orthogonality of ξ to θ, but they 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 3.[13]

The utilitarian benchmark, denoted $\mathscr{W}_1$, is

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

This specification captures the welfare of chains alone. Any wedge between private and social

[13] The indices 1, 3, and 4 are used for the specifications analyzed in this section, with the index 2 reserved for the fundamental-tracking criterion $\mathscr{W}_2$ of Remark 3. This numbering matches the labels used in the figures and the replication code, and is kept fixed throughout to avoid relabeling across documents.

incentives under $\mathscr{W}_1$ would reflect a failure to internalize strategic interaction, 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 $\mathscr{W}_4$, 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 normalization by V_A reflects that a location screened by multiple adopters bears the shared component relative to total signal uncertainty, not its absolute variance alone. The parameter γ_4 measures the marginal cost of correlated evaluation.

The distinction between $\mathscr{W}_3$ and $\mathscr{W}_4$ is structural. The former penalizes realized aggregate behavior and therefore depends on equilibrium loadings. The latter penalizes correlation in signals and depends only on adoption and signal-structure primitives. Hence the $\mathscr{W}_4$ wedge at any fixed q is independent of δ, while the $\mathscr{W}_3$ wedge varies with δ through equilibrium responses.[14] The penalty weights γ_3 and γ_4 set the scale of each externality but are not separately identified from the model; the identified content of the results below is therefore the sign of each wedge and ratios across regimes, not the level of the welfare loss.

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

[14] 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. More generally, any welfare penalty that depends only on the information structure and adoption rate, rather than on equilibrium responses, yields a wedge whose sign and magnitude at fixed q are independent of strategic interaction.

chains through their own profits, and Proposition 3 implies that such effects are internalized in equilibrium. The misalignment studied here instead arises from costs borne outside the chain population. A planner concerned with both objectives would use a hybrid specification, in which the externality wedge identified below adds to any fundamental-tracking distortion.

5.4 Team-Efficiency Result

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

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

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

That term vanishes by 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 $\delta(P_g - \bar{x})$ after 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 planner's condition coincides with the Nash condition. The same argument applies to changes in the adoption rate, implying $d\mathscr{W}_1/dq = \Delta\Pi(q)$.

Shared processing therefore does not create inefficiency within the chain population. Chains internalize their effects on one another through equilibrium responses, so any wedge between private and social incentives must arise from costs borne outside that population. The externalities identified in Section 5.2 operate through this channel.

Remark 2 (Robustness of the team-efficiency result)**.** The identity $d\mathscr{W}_1/dq = \Delta\Pi(q)$ relies on a continuum of agents and linear entry probabilities. The continuum assumption is mainly technical. With finitely many firms, the cross-sectional identity in Proposition 3 holds only in expectation, so the general-equilibrium term does not cancel exactly. Standard large-game approximation arguments suggest that the residual is $O(1/N)$, so the over-adoption result of Theorem 1 should survive for moderate N; a formal finite-N treatment is beyond the scope of this paper. We note that the

$\mathscr{W}_4$ wedge itself does not rely on the cancellation: the penalty $q^2\sigma_\xi^2/V_A$ is a property of the signal structure and adoption rate alone, so any finite-N residual in the team-efficiency identity perturbs the wedge's magnitude, not its source.

Linearity plays a larger role. With nonlinear entry probabilities outside the global-games regime, the marginal payoff includes a correction term that depends on the density of firms near the entry threshold, so the planner's condition need not coincide exactly with the Nash condition. Appendix D reports a numerical exercise replacing the linear rule with $P_f = \Lambda(\kappa[\mathbb{E}_f(\theta) + \delta\,\mathbb{E}_f(\bar{x})])$, which nests linearity as κ becomes small. 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 positive in the great majority of cells, with two boundary failures when the welfare object is near zero. The numerical evidence is consistent with the structural property 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 at any fixed q.

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

The two parts have the same structure: the wedge equals the marginal external cost of adoption. Under $\mathscr{W}_3$, the externality is $K_\xi(q)^2\sigma_\xi^2$, which combines the primitive variance σ_ξ^2 with the equilibrium amplification term $K_\xi(q)^2 = q^2(k^A(q))^2$. Its sign is invariant, but its magnitude varies with δ through the equilibrium loadings. Under $\mathscr{W}_4$, the externality is $q^2\sigma_\xi^2/V_A$: the probability that two evaluators both adopt the platform times the within-adopter variance share of the shared component. Equilibrium loadings do not enter this object, so the wedge at fixed q is independent of δ; strategic interaction affects it only through the equilibrium adoption level $q^*(\delta)$.

The result therefore does not rely on whether entry decisions are strategic complements or substitutes. 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 different affected populations, but both imply that equilibrium adoption is too high. The over-adoption is thus robust across additive externality penalties: it is a property of the information structure rather than of a particular such specification. It is not unconditional, however, as Remark 3 shows for the fundamental-tracking criterion $\mathscr{W}_2$, which is not of the additively separable form and leaves the sign of the wedge indeterminate.[15]

Of the two criteria, $\mathscr{W}_4$ is the cleaner one. Its sign and magnitude depend only on adoption and signal-structure primitives, so the wedge is invariant to the strategic environment and, as Appendix D shows, essentially invariant to the entry-choice rule, remaining within one percent of its closed form across every calibration tested. The aggregate-volatility wedge $\mathscr{W}_3$ delivers the same qualitative conclusion, but it inherits its magnitude from the equilibrium loadings, and that magnitude is sensitive to the curvature of the choice rule. We therefore read $\mathscr{W}_4$ as establishing the over-adoption result and $\mathscr{W}_3$ as corroborating it. The invariance under $\mathscr{W}_4$ has an additional empirical implication. The strategic parameter δ is difficult to identify because 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 3 (Misalignment under fundamental-tracking welfare $\mathscr{W}_2$ is generically indeterminate)**.** A fourth specification isolates forecast accuracy at the chain level, treating coordination payoffs

[15] 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 and aligns the decentralized condition with the planner's first-order condition $F^{-1}(q) = d\mathscr{W}_k/dq$. Any intervention that lowers the marginal private return to adoption by $\tau_k(q^*)$, such as an adoption tax or a capital requirement indexed to platform exposure, implements the constrained optimum.

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, the loading at which entry passes signals through fully; an alternative tracking criterion that penalizes the mean-squared gap between aggregate entry and the fundamental, $\mathbb{E}[(\bar{x}-\mathbb{E}[\bar{x}]-\theta)^2]$, delivers the same qualitative conclusion, because both criteria measure objects internal to the chain population and both generate indirect terms in $(k^A)'(q)$ and $(k^C)'(q)$ whose signs are calibration-dependent. Unlike $\mathscr{W}_3$ and $\mathscr{W}_4$, $\mathscr{W}_2$ is not an additive perturbation of $\mathscr{W}_1$, so 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 $(k^C)'(q)$, so its sign is calibration-dependent. This specification therefore does not deliver a determinate prediction for over- or under-adoption.

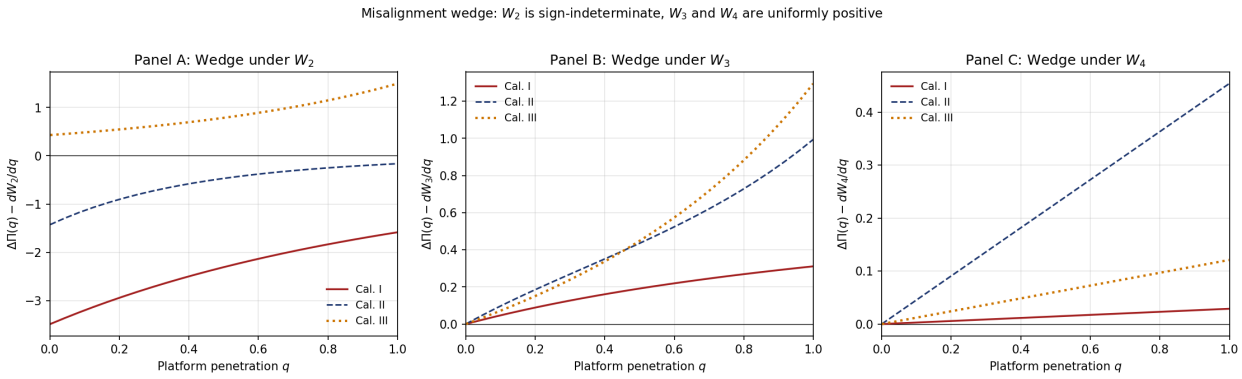
The contrast isolates the role of externalities. Under $\mathscr{W}_2$, costs are internal to the chain population and are largely internalized in equilibrium. Over-adoption arises only when the planner accounts for costs borne outside the chain population, as in $\mathscr{W}_3$ and $\mathscr{W}_4$.

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

5.6 Residual Error versus Informative Commonality

The baseline model assumes $\xi \perp \theta$ to isolate the correlated-evaluation channel from the social value of public information. A natural question is whether over-adoption depends on the common component being entirely unrelated to fundamentals. To draw this boundary, suppose the shared

Figure 1: The misalignment wedge under three welfare specifications.



Notes: Each panel plots the misalignment wedge $\Delta\Pi(q) - d\mathscr{W}_k/dq$ as a function of platform penetration q, for three calibrations sharing $\delta = 0.7$ and $\gamma_k = 1$ but differing in the relative magnitudes of σ_A^2, σ_C^2, and σ_ξ^2. Calibration I: $(\sigma_\theta^2, \sigma_\xi^2, \sigma_A^2, \sigma_C^2) = (2.79, 0.05, 0.62, 0.27)$; Calibration II: $(1.23, 0.62, 0.88, 0.47)$; Calibration III: $(2.48, 0.21, 0.78, 1.46)$. Panel A shows the wedge under fundamental-tracking welfare $\mathscr{W}_2$: it is negative throughout for Calibrations I and II and positive throughout for Calibration III, with no interior zero-crossing. 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 3). 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). Levels are in units of γ_k. Panels B and C jointly establish that the over-adoption result is robust across two structurally distinct externality penalties.

component can be decomposed as

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

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

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

with total variance

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

The informative component changes signal precision. 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 $\xi = \ell\theta + \zeta$*, where* $\zeta \perp \theta$*,* $\mathrm{Var}(\zeta) = \sigma_\zeta^2$*, and* $\ell \in [0,1)$*. Let*

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

denote the admissible region under the transformed signal, and let $V_A(\ell) \equiv (1+\ell)^2\sigma_\theta^2 + \sigma_\zeta^2 + \sigma_A^2$ *denote the unconditional variance of the transformed platform signal. Under the decided-about-agent welfare criterion* $\mathscr{W}_4$*,*

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

Consequently:

(i) For all $q \in (0,1]$ *and* $\delta \in \mathscr{D}(\ell)$*, the wedge is strictly positive whenever* $\sigma_\zeta^2 > 0$ *and zero when* $\sigma_\zeta^2 = 0$*.*

(ii) Holding residual variance σ_ζ^2 *fixed, a larger informative component* ℓ *lowers the wedge only by increasing total signal variance* $V_A(\ell)$*, and the wedge vanishes only as* $\sigma_\zeta^2 \to 0$*.*

(iii) At any fixed adoption rate, the wedge is independent of δ*.*

The proposition separates the mechanism from standard public-information effects. In public-information models (Morris and Shin 2002, Angeletos and Pavan 2007, Hellwig and Veldkamp 2009), the common signal is informative about fundamentals, so its welfare effect depends on the tradeoff between coordination benefits and reduced independent learning. Here the externality is tied to the residual error ζ, which carries no information about θ. A larger informative component dilutes the residual share, but the wedge remains positive whenever residual error remains. The orthogonal baseline is the special case $\ell = 0$, in which σ_ξ^2 itself is the relevant primitive.

Remark 4 (Information-structure robustness)**.** Proposition 4 shows that the wedge is robust to enriching the shared component with an informative part. The same logic extends to the comparison of information structures more broadly: what governs the externality is the residual common variance σ_ζ^2 that survives conditioning on fundamentals, not the joint informativeness of the signal. This is the sense in which the over-adoption result is a prediction about the information structure rather than about a particular parameterization, and it is robust across the structures a platform might plausibly implement (Bergemann and Morris 2013).

The policy implication follows directly. Improving common-signal informativeness can reduce the wedge by increasing $V_A(\ell)$, but it does not internalize the correlated-mistake channel. Mitigation requires reducing σ_ξ^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 $q^* = F(\Delta\Pi(q^*))$. The number and stability of equilibria depend on the shape of $\Delta\Pi(q)$ and on its boundary values. The first boundary admits a simple characterization.

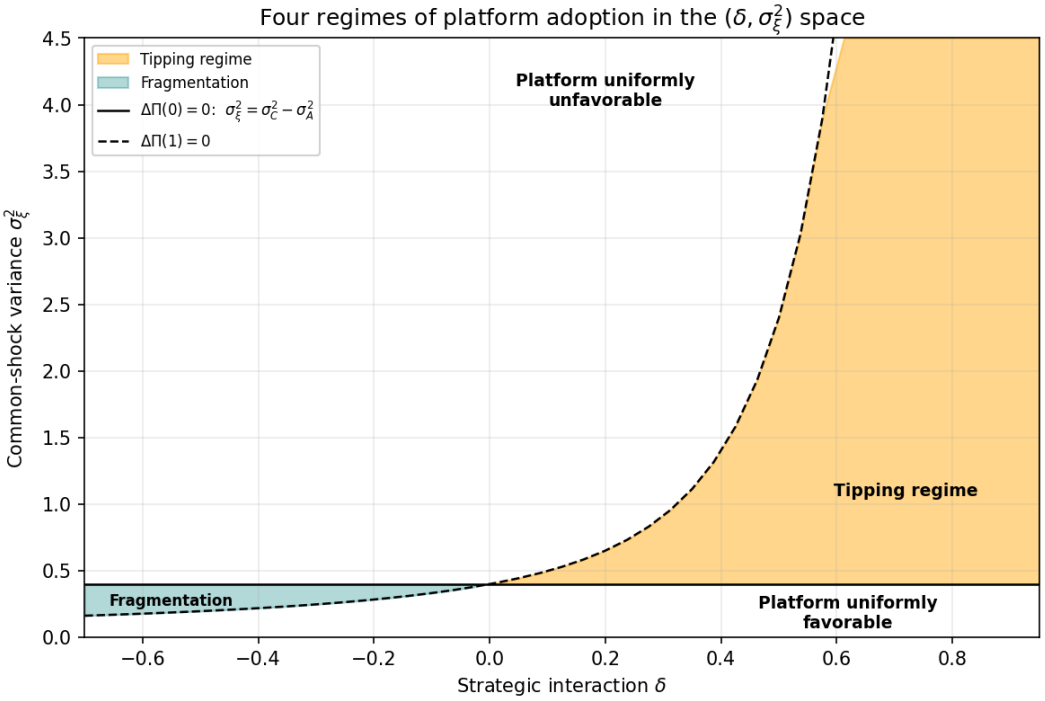
Lemma 1 (δ-independent threshold)**.** *The condition* $\Delta\Pi(0) = 0$ *holds if and only if* $\sigma_\xi^2 = \sigma_C^2 - \sigma_A^2$, *independent of* δ.

At $q = 0$, a marginal adopter faces the same aggregate as a non-adopter, so the coordination channel does not differentiate the technologies. Adoption depends only on signal quality. When $V_A = V_C$, 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 increasing in δ. Numerical evaluation shows that this locus is monotone and convex and intersects the $q = 0$ boundary at $\delta = 0$.

Figure 2 plots these two loci in (δ, σ_ξ^2) space and partitions the parameter region into four regimes. While equilibrium configurations vary across regions, the misalignment result of Theorem 1 applies throughout: strategic interaction shapes how adoption is sustained, but not whether it is excessive. The regimes therefore differ in the level and structure of adoption, not in the direction of the welfare distortion. Two cases are particularly informative: the tipping regime under agglomeration and the fragmentation regime under business stealing.

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 $\delta = 0$, and is shown only over the range of δ on which the linear equilibrium is admissible. 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. The horizontal axis is restricted to $\delta \in [-0.7, 0.95]$, inside the admissibility region $\mathscr{D}$ of Proposition 1 for this calibration; regions are not extrapolated beyond the plotted boundary loci. Calibration: $\sigma_\theta^2 = 1$, $\sigma_A^2 = 0.2$, $\sigma_C^2 = 0.6$.

6.1 Tipping Regime

When $\delta > 0$, $\Delta\Pi(0) < 0$, and $\Delta\Pi(1) > 0$, the fixed-point condition $q^* = F(\Delta\Pi(q^*))$ has at least one interior solution. If F is sufficiently curved relative to $\Delta\Pi$, so that $F'(\Delta\Pi(q))\Delta\Pi'(q) > 1$ on a non-degenerate interior interval, the equation admits three solutions: a stable low-adoption fixed point q_L^*, a stable high-adoption fixed point q_H^*, and an unstable threshold $q^\dagger$ separating their basins of attraction.[16] Figure 3 illustrates such a configuration for a logistic F with steepness parameter κ_F large enough for the curvature condition to hold.

[16] 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^*.

In this region, platform adoption is privately unattractive in isolation but attractive once adoption is sufficiently widespread. At low penetration, exposure to the shared component is limited, and chains prefer custom research. At high penetration, aggregate entry is driven more strongly by the shared component, and chains benefit from aligning with rivals who respond to the same signal. Crossing $q^\dagger$ 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, leading the market to lock in to a single technology. Here, the feedback operates through strategic interaction and shared processing rather than technological increasing returns alone. The inefficiency is also distinct: coordination may be privately valuable, but adoption remains socially excessive because the shared component generates correlated errors that firms do not internalize.

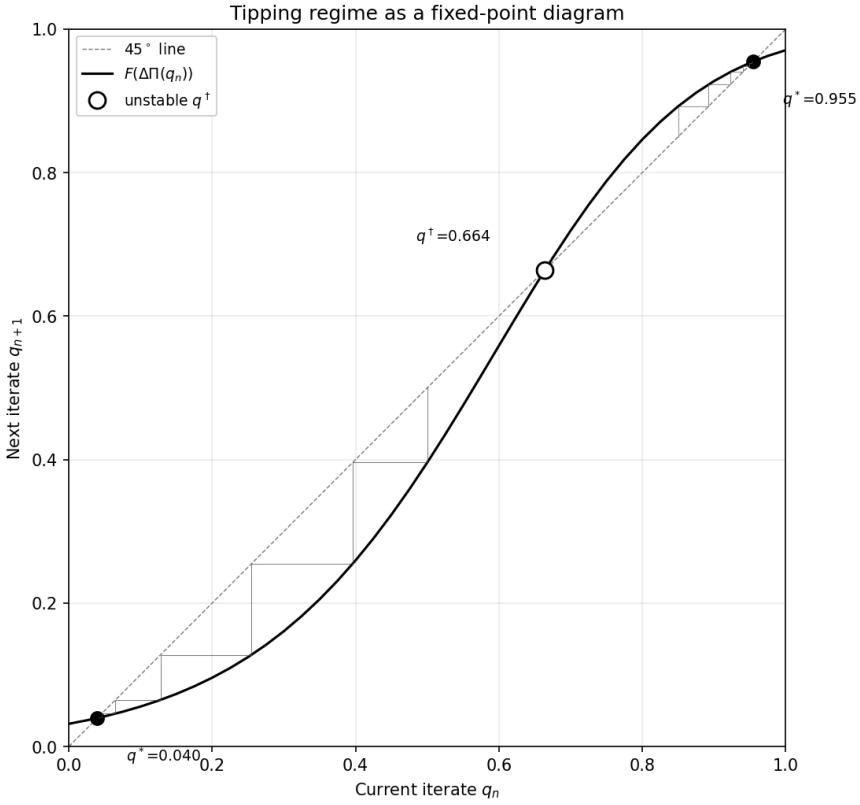
This regime provides a precise formulation of shared platforms as coordination devices. The high-adoption fixed point is sustained even though adoption is privately unattractive in isolation, because no chain has an incentive to deviate once others adopt. The mechanism is strongest in sectors with agglomeration economies. In food halls, outlet centers, and clusters of competing fast-food restaurants, forecasting rival behavior can be as valuable as forecasting demand, so a platform that reveals common signals becomes valuable 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

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. 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 $F(x) = (1 + e^{-\kappa_F x})^{-1}$ and $\kappa_F = 25$. This calibration sits inside the admissibility region $\mathscr{D}$ of Proposition 1, with $\Delta(q) > 0$ throughout $[0,1]$.

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. Since $F \circ \Delta\Pi$ 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 reverses the coordination case. The platform improves precision and is therefore attractive in isolation, but it also increases correlation across chains. When entry decisions are strategic substitutes, this correlation is costly. As adoption rises, the marginal benefit declines, and equilibrium adoption settles where precision gains are offset by the cost of moving with rivals. Sectors where differentiation is central, such as specialty grocers, boutique retail, or niche restaurant concepts, should therefore exhibit partial adoption rather than convergence to a common platform.

The welfare implication is unchanged. By Theorem 1, the 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 tipping and fragmentation regimes therefore differ in the level of q^*, not in the direction of the welfare distortion: 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 $\Delta\Pi(1) < 0$, the platform is unattractive throughout and adoption remains near zero. In both cases, strategic forces reinforce the underlying technological advantage or disadvantage rather than changing 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 matters because platforms may differ in interface or branding while relying on the same underlying data or models. The

multi-platform setup below distinguishes substantive diversity from cosmetic diversity.

7.1 Multiple Processing Technologies

Suppose adopters partition across $K \geq 1$ platforms. Throughout this section K denotes the number of platforms; it is a scalar count and is not to be confused with the aggregate loadings $K_\theta(q)$ and $K_\xi(q)$ of Section 4, which carry subscripts. Platform k generates a processing shock $\xi_k \sim \mathcal{N}(0, \sigma_\xi^2)$ and idiosyncratic noise $u_f^{A,k}$ with variance σ_A^2, so that an adopter of platform k observes

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

Let q_k denote the share of chains using platform k, with $\sum_{k=1}^K q_k = q$. The shocks $\{\xi_k\}$ are jointly Gaussian with $\text{Var}(\xi_k) = \sigma_\xi^2$ and pairwise correlation $\text{Corr}(\xi_k, \xi_{k'}) = \rho \in [0,1]$ for $k \neq k'$.[17] The single-platform model corresponds to $K = 1$.

The parameter ρ captures shared processing across platforms. When $\rho = 0$, platforms are structurally independent, with shocks arising 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.[18]

For tractability, vendor shares $\{q_k\}$ are treated as exogenous. Section 6 determines total adoption q endogenously; here we hold q fixed and isolate how its allocation across platforms affects welfare.

[17] The bare ρ in this section denotes cross-platform correlation in shared shocks. It differs from the projection coefficient $\chi_A = \alpha_\theta^A + \alpha_\xi^A$ 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.

[18] Such overlap need not be confined to data inputs. DosSantos DiSorbo et al. (2026) document that frontier language models from every major provider exhibit the same systematic rigidity in handling policy exceptions, refusing warranted deviations that humans grant; because this bias is a property of how the underlying models are trained rather than of any one vendor's interface, downstream tools built on a shared foundation-model layer inherit correlated decision errors even when separately branded. This is a behavioral instance of the high-ρ regime.

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:

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

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

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

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 and 1 when they share a common core.

Proposition 5 (Vendor diversity)**.** *Under either welfare specification, the wedge between private and social adoption incentives scales with the effective concentration index:*

(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 $\mathscr{W}_3$, 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 equals expected

correlation across adopters: with probability $1/K$, two adopters share a platform and their errors are perfectly correlated; with probability $1 - 1/K$, they use different platforms and correlation is ρ. Vendor diversity therefore reduces the externality only insofar as it reduces the likelihood that adopters share the same source of error.

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 by lowering the probability that adopters share the same error. When platforms share a common core, increasing K has no effect. Vendor diversity therefore matters through the independence of processing technologies rather than through the number of providers.

This distinction is 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. Platforms built on distinct data sources, such as transaction records, mobility traces, or satellite imagery, approximate the low-ρ regime. The relevant margin is the diversity of underlying inputs and methods, not the number of user-facing platforms.

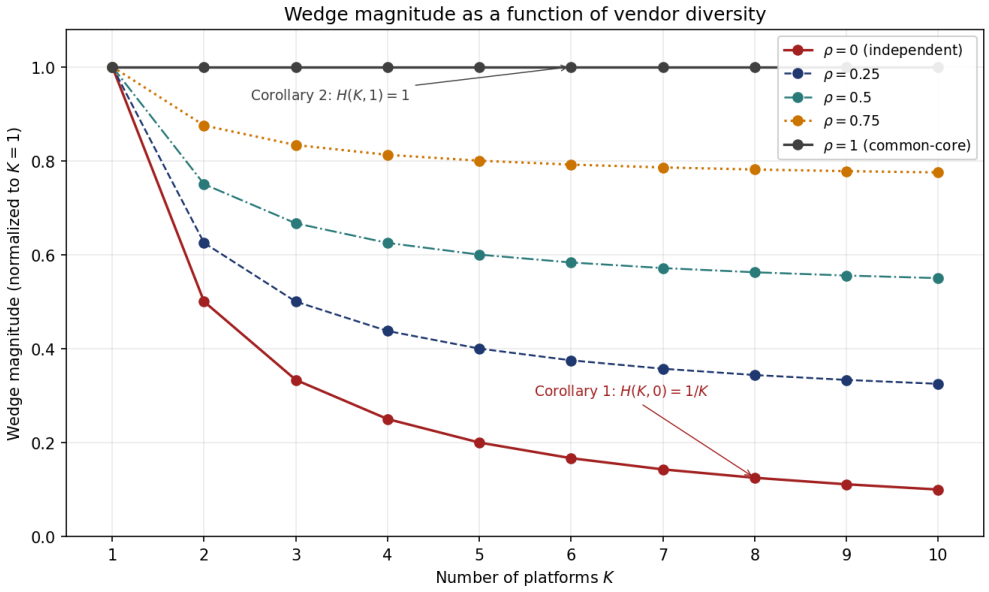
A similar pattern appears in adjacent algorithmic markets. Bommasani et al. (2022) document that fine-tuned variants of a common foundation model exhibit limited independent variation relative to structurally distinct models. Tong et al. (2023) identify systematic failure modes of CLIP's text encoder and show that the same failure inputs generate errors in Midjourney, DALL-E, and VideoFusion. In the notation of the model, distinct downstream products share an upstream processing layer, so nominal diversity leaves the externality 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. Larger and more accurate models also exhibit highly correlated errors even across distinct providers. We emphasize that error agreement is a different statistic from the correlation parameter ρ, measured in a different domain, and does not map one-for-one into our calibration; we read this evidence as establishing that structural correlation among independently branded prediction systems is large, not as a point estimate of ρ. Thus, ρ need not approach zero with provider diversity alone. When adopters converge on frontier technologies, cosmetic diversity can remain binding precisely 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.

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.

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 (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 5 (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 revises its assessment of demand, many adopters update in the same direction, and the resulting co-movement can persist until the model is revised again. The static analysis captures within-period correlation; the dynamic extension below isolates persistence across periods 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, chooses whether to adopt the platform, and exits. Within each period, the entry game is identical to Section 3. Conditional on adoption rate q, aggregate entry is

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

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

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

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

The adoption rate q is fixed in this subsection. Because agents are short-lived, they do not internalize intertemporal effects; persistence arises entirely from the shared processing component. Section C endogenizes q as a stochastic process.

8.2 Temporal Dependence

Persistence in the shared component induces persistence in aggregate entry.

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 by the shared component. The fundamental shock contributes to within-period variation but not to persistence. By contrast, ξ_t enters aggregate behavior with weight $K_\xi(q)$, so its persistence is amplified as adoption rises. Higher q increases the share of aggregate variation attributable to the persistent component, making current entry outcomes more predictive of future outcomes.

Shared intelligence therefore affects not only the dispersion of entry within a period, but also the clustering of outcomes across periods. Even when fundamentals are i.i.d., high-adoption environments can exhibit longer same-sign deviations in aggregate entry because platform shocks propagate through many firms at once.

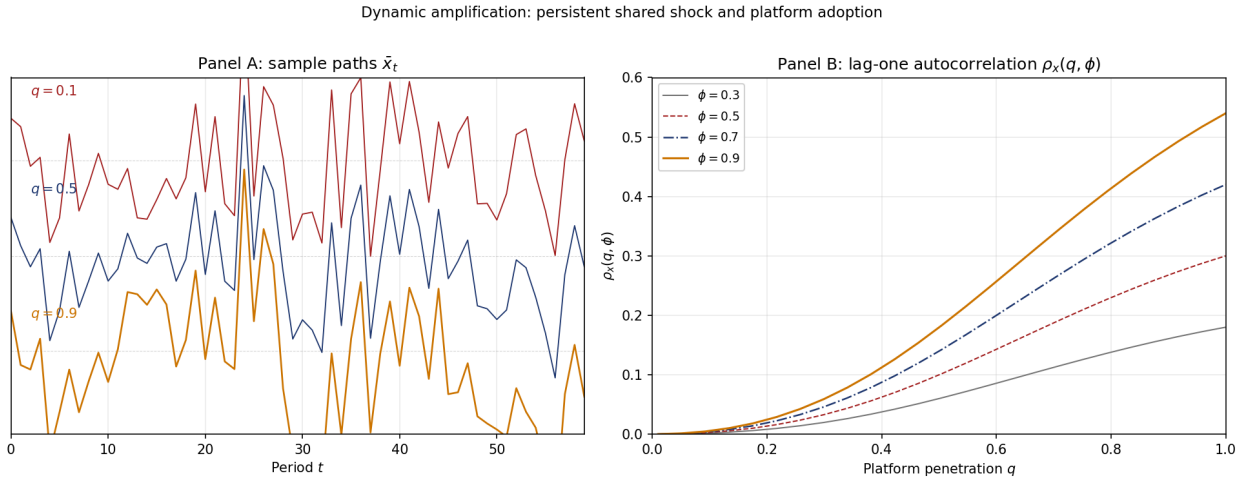
Figure 5 illustrates the mechanism. Panel A plots a single realization of $\bar{x}_t$ across three penetration levels under identical shocks: at low adoption the path resembles independent noise, while at high adoption it exhibits longer same-sign excursions. Panel B plots the population autocorrelation $\rho_x(q,\phi)$, showing that persistence rises with platform adoption.

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

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.

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.[19]

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.

[19] Short-lived agents also rule out strategic timing of entry. A two-period extension in which firms can delay entry to preempt rivals leaves the $\mathscr{W}_4$ wedge unchanged: the wedge depends only on contemporaneous adoption and signal-structure primitives, so preemptive timing shifts when exposure accrues but not its marginal level. A numerical verification of this timing-invariance is included in the replication package.

9 Case Study of Restaurant Site Selection Decision Support

Proposition 5 implies that the welfare relevance of vendor diversification depends on the effective concentration index $H(K,\rho)$, not on the number of nominal vendors. The index combines downstream adoption concentration with correlation in the underlying forecasting technologies. A market may therefore appear diversified at the interface layer while remaining concentrated in the data and modeling layer that determines forecast co-movement. Whether diversification is substantive, $H \approx 1/K$, or cosmetic, $H \approx 1$, is an empirical question about the technology stack.

This section illustrates that distinction using restaurant site-selection analytics.[20] We use the case to establish that the object of the theory, correlated platform-mediated site-selection decisions, exists at the scale the model contemplates, not to identify the mechanism. The descriptive vendor-concentration and upstream-overlap facts are consistent with the shared-error channel but do not, on their own, distinguish it from correlated fundamentals or imitation; the model supplies the interpretation, and the falsification test of Proposition 7 is what a future causal design would target. The objective is not to point-estimate H, but to show how its primitives can be disciplined using public disclosures, SEC filings, vendor case studies, and evidence on upstream data overlap. The exercise clarifies a broader implication of the model: vendor competition at the front end need not imply meaningful diversification in the forecasting technologies used by adopters.

The calibration has two components. The first is downstream vendor-share concentration, $\sum_k \omega_k^2$, the probability that two adopters use the same vendor. The second is structural correlation across vendors conditional on using different platforms, ρ, which captures shared shocks from common data inputs, model infrastructure, or upstream suppliers. The two margins enter H separately and are considered in turn.

9.1 Vendor Stack

The downstream vendor market is concentrated. Trade press, vendor disclosures, and SEC filings identify six principal providers of site-selection analytics for multi-unit restaurant chains: Placer.ai, Tango Analytics, Buxton, eSite Analytics, SiteZeus, and IdealSpot.[21] Other firms participate in the

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

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

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 smaller specialists serve single-unit operators, local businesses, or narrower use cases outside the multi-unit chain market considered here.

The model-relevant feature is that 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.[22] We note that the INRIX disclosure is the upstream vendor's own marketing material and may overstate the depth of each integration; we therefore corroborate the overlap from the demand side where possible, using the platform case studies and SEC filings assembled in Appendix E, and we treat the disclosure as evidence on the direction of upstream concentration rather than on its magnitude. The six vendor brands observed by adopters therefore need not correspond to six independent forecasting technologies. They may be competing front-end interfaces built partly on overlapping upstream inputs.

9.2 Share Concentration

Direct measurement of vendor share is not available. 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 use a revenue-proxy approach. We assemble a peer revenue pool from 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.[23] We validate relative scale and market relevance against SEC filings and earnings-call disclosures, then apply a retail-chain exposure adjustment that scales 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.

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

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

Let ω_k denote vendor k's share *among adopters* (so $\sum_k \omega_k = 1$), as distinct from the aggregate adoption rate q. Table A3 in Appendix E reports the adjusted-share estimates. The implied Herfindahl component is

$$\sum_k \omega_k^2 \approx 0.628^2 + 0.141^2 + 0.097^2 + 0.077^2 + 0.055^2 + 0.002^2 \approx 0.433.$$

The equal-share benchmark, $1/K \approx 0.17$, would be attained only if adopters were allocated evenly 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. 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. 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 under equal vendor shares. Thus, even with fully independent forecasting technologies, the market would not attain the substantive-diversification benchmark. The remaining question is whether adopters using different vendors are insulated from common forecasting shocks.

9.3 Structural Correlation

The second component of H is structural correlation across distinct vendors. This margin captures whether separately owned and branded platforms generate independent forecasts. In restaurant site selection, the relevant evidence comes less from the front-end interface than from the upstream data layer. Vendors build proprietary models, but many rely on overlapping inputs used to construct trade areas, measure traffic, characterize demand, and score candidate sites.

Four upstream input categories 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. The core inputs are therefore considerably less differentiated than the vendor interfaces.

Placer.ai is particularly informative because it occupies both sides of the stack. It is a downstream platform sold to chains, but 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.[24] 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 enter the broader analytics ecosystem.

Two chains that contract with different vendors are therefore not necessarily using independent forecasting technologies. They may rely on differentiated front-end products built over shared traffic data, point-of-interest files, visit panels, and demographic segmentation products. In the notation of the model, the shared processing shock ξ captures common errors, revisions, or methodology changes propagated through this upstream layer, including mobility-panel reweighting, point-of-interest taxonomy revisions, and demographic forecast updates.

The available data do not point-identify ρ, but they make the independence benchmark difficult to defend. A value near zero would require separate vendor ownership as well as largely independent upstream data inputs and processing pipelines. The observed stack 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 be interpreted not as a point estimate, but as an order-of-magnitude discipline imposed by the observed technology stack.

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

9.4 Discussion

Combining vendor concentration with upstream-input overlap gives

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

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

$$H \in \left[0.433 + 0.5(1 - 0.433),\ 0.433 + 0.7(1 - 0.433) \right] \approx [0.72, 0.83].$$

Table A4 reports the corresponding values 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, $1/K \approx 0.17$. 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 calibration into welfare terms. In the multi-vendor extension, both externality wedges scale linearly with H, where $H = 1$ is 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 one shared technology to fully independent vendor technologies. This ratio interpretation does not depend on γ_3 or γ_4, which determine the absolute scale of the wedge but not its proportional response to H; the γ-free ratio $H_{\text{case}}/H_{\text{single}} \approx 0.77$ is the identified object.

Figure 6 illustrates these implications over the admissible range $\delta \in [-0.95, 1.00]$. 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 closer to the cosmetic-diversity ceiling throughout the admissible region because $H_{\text{case}}/H_{\text{single}} = 0.77$ is invariant in δ, even though the level of the wedge rises with agglomeration incentives.

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 under struc-

turally independent vendors. The gap 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 with δ. The largest absolute gains from upstream-input diversification therefore arise in settings with stronger agglomeration incentives, such as quick-service restaurant markets.

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

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

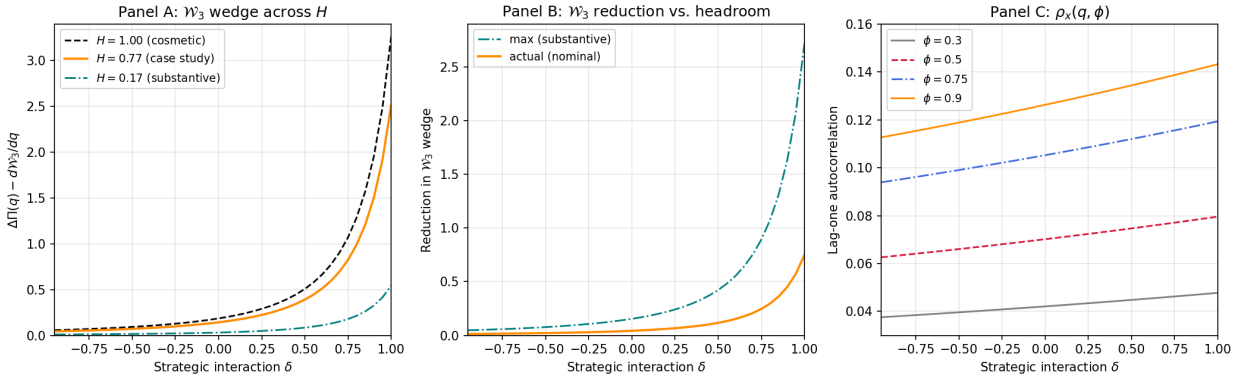
in units of γ_4, and would appear as a horizontal line. Panel C instead reports the dynamic implication by plotting the lag-one autocorrelation $\rho_x(q,\phi)$ from Proposition 6 for four values of ϕ. The autocorrelation increases mildly with δ through the ratio $K_\xi(q)/K_\theta(q)$, but most variation is governed by ϕ. The case-study calibration therefore gives a unified interpretation of the static externality in Panels A and B and the dynamic externality in Panel C, both governed by the same effective concentration index.

Policy-margin elasticities. The decomposition of H separates two intervention margins: downstream adoption shares and cross-vendor correlation. These correspond to different policy instruments and have different quantitative implications.

At $\rho = 0.6$, equalizing adoption across the six existing vendors while leaving the upstream technology stack unchanged 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.

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

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 $\mathcal{D}$ 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 $\mathcal{W}_3$ wedge (in units of γ_3) at three values of the effective concentration index H: the case-study central value $H = 0.77$ implied by Table A4; 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 $\mathcal{W}_4$ 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 $H \cdot 2q\sigma_\xi^2/V_A \approx 0.257\,\gamma_4$, 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.

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. At the observed parameter values, the upstream correlation margin is quantitatively larger than the downstream share margin, but neither margin alone eliminates the externality.

Heterogeneity across the two wedges. Although the $\mathcal{W}_3$ and $\mathcal{W}_4$ 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 $\mathcal{W}_4$ is invariant in δ at any fixed adoption rate. The aggregate-volatility wedge $\mathcal{W}_3$ instead depends on δ through equilibrium loadings.

The same calibration therefore has different interpretations depending on the welfare criterion. A regulator concerned 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 this channel non-negligible. A regulator concerned with correlated screening of locations instead reads the same value of H through the δ-independent $\mathscr{W}_4$ channel. The case study disciplines the common exposure term while leaving the relative regulatory weight on the two externalities to the welfare criterion.

10 Conclusion

In this paper, we develop a model of shared platform adoption in retail entry to study how routing firm decisions through a common forecasting layer affects market outcomes. While these platforms reduce idiosyncratic noise, they also introduce a common processing error unrelated to market fundamentals. By cleanly separating this shared error from informative common signals and strategic interaction, we show that firms internalize the private precision gains of adoption but fail to account for the external costs their correlated mistakes impose on third parties. Consequently, equilibrium platform adoption is socially excessive. Importantly, this over-adoption result holds regardless of whether firms' entry decisions are strategic complements or substitutes, and it remains robust across a broad class of welfare criteria that penalize the shared-error externality.

Beyond the theoretical characterization of the welfare distortion, our analysis provides critical insights into the effectiveness of vendor diversification. We demonstrate that nominal competition among front-end analytics providers does not automatically resolve the systemic risk of correlated errors. When competing platforms rely on overlapping upstream data inputs—such as shared traffic feeds, mobility panels, or demographic taxonomies—adopters remain exposed to the same underlying processing shocks. Our case study of the restaurant site-selection analytics market empirically illustrates this phenomenon, indicating that the observed market structure retains roughly three-quarters of the single-platform distortion. For regulators and industry practitioners, this highlights a crucial distinction: the relevant policy margin is the structural independence of the underlying forecasting technologies, rather than merely the number of competing vendor brands.

We also note the following limitations of our analysis. First, our model treats the variance of the

shared processing error as an exogenous primitive. An important avenue for future work is to endogenize this error by modeling the upstream incentives of platform providers to refine their data, or the trade-offs they face in model design and training. A preliminary step in this direction, included in the replication package, decomposes the marginal $\mathscr{W}_4$ externality under adoption-dependent noise $\sigma^2_\xi(q)$ into an exposure term (the constant-variance baseline) and an intensity term capturing marginal degradation; the decomposition suggests that capacity-strain technologies, in which noise rises with adoption, would amplify the over-adoption result rather than overturn it. Second, while our descriptive case study motivates the presence of shared processing errors, formally identifying this mechanism in the data requires exploiting exogenous variation in platform-side updates or data-source disruptions, which we leave to future empirical work. Finally, our framework focuses primarily on the entry stage; extending this analysis to post-entry dynamics—such as how shared forecasting errors propagate through ongoing pricing, inventory, or expansion decisions—could yield further insights into the long-term competitive consequences of algorithmic monoculture.

References

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.

Bergemann D, Morris S (2013) Robust predictions in games with incomplete information. *Econometrica* 81(4):1251–1308.

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.

Blevins JR, Khwaja A, Yang N (2018) Firm expansion, size spillovers, and market dominance in retail chain dynamics. *Management Science* 64(9):4070–4093.

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.

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

Clayton C, Coppola A (2026) The optimal use of ai in financial regulation. Technical report, National Bureau of Economic Research.

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.

DosSantos DiSorbo M, Ju H, Aral S (2026) Teaching ai to handle exceptions: Supervised fine-tuning with human-aligned judgment. *PNAS Nexus* pgag144.

Dou WW, Goldstein I, Ji Y (2025) Financial market fragility in the era of ai planning. *Available at SSRN 5763222* .

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

Ershov D, Lyons E (2024) Outsourcing algorithm development: Evidence from contractors and llms. *Available at SSRN* .

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.

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.

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.

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.

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.

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 An Identifying Restriction for Empirical Work **3**

C Dark Side of Cheap AI **5**

D Nonlinear Robustness of the Team-Efficiency Benchmark **8**

D.1 Equilibrium Computation 9

D.2 Calibrations 9

D.3 Results 9

D.4 Implications 11

E Additional Details for Descriptive Case Study **12**

F Proofs **16**

F.1 Proof of Proposition 1 16

F.2 Proof of Lemma 1 17

F.3 Proof of Proposition 2 17

F.4 Proof of Proposition 3 18

F.5 Proof of Theorem 1 18

F.6 Proof of Proposition 4 20

F.7 Proof of Proposition 5 21

F.8 Proofs of Corollaries 1 and 2 22

F.9 Proof of Proposition 6 22

F.10 Proof of Proposition 8 23

F.11 Proof of Proposition 7 25

A Robustness Checks for the Fast-Food Empirical Analysis

Two checks clarify the interpretation of the similarity results. First, we assess sampling sensitivity by resampling counties by state cluster 1,000 times. In each bootstrap draw, we re-estimate the chain-specific market-fundamental regressions and recompute coefficient-vector cosine similarity, coefficient correlation, and predicted-score correlation. Second, we distinguish directional alignment from exact equality. A stacked chain-by-county regression tests whether Burger King's coefficients on the standardized continuous fundamentals equal McDonald's. Exact equality is stronger 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 identical.

A.1 Full-Sample Results

The checks reject exact equality but confirm strong directional alignment. The two specifications disagree on the equality test, and the disagreement is itself informative. A Wald test on the log-count specification rejects equality of the continuous-fundamental slopes across chains ($p < 0.001$), consistent with different footprint sizes and sensitivities to scale; the Poisson count specification reported in Section 2 does not reject equality of the continuous-fundamental slopes ($p = 0.238$). The two tests evaluate the slope contrast under different functional forms for the conditional mean, and we rely on the more conservative Poisson result in the main text. The directional diagnostics are strong under both. 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. Thus, McDonald's and Burger King 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 the checks within AIGE-stratified samples shows that alignment is strongest where local employment is most exposed to AI. In the high-AIGE subsample, sign concordance is perfect, coefficient-vector cosine similarity is 0.999, and 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: cosine similarity is 0.994, 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 high-AIGE counties.

B An Identifying Restriction for Empirical Work

Strategic interaction and shared processing operate 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, so strategic interaction dictates whether adoption is self-reinforcing or self-limiting. The welfare wedge, $\Delta\Pi(q) - d\mathscr{W}_k/dq$, governs distortion: by Theorem 1, it is strictly positive whenever $\sigma_\xi^2 > 0$ 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. Multiple equilibria do not by themselves indicate a market failure, nor does uniqueness rule one out; the distortion arises only when the common component is present and marginal adoption increases exposure to it.

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

This decomposition yields a falsifiable prediction. The mechanism studied here generates a pattern that canonical coordination or network-effects models do not: 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 is a platform-side update event. Revisions to a foundation model, changes in vendor methodology, or updates to data inputs can shift 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 into entry decisions.

Proposition 7 (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 models without a shared error component, aggregate entry depends only on fundamentals and equilibrium responses to them. Such models predict no discrete movement in aggregate entry at the time of E unless fundamentals also 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.

Implementation requires three ingredients. First, the event must be identified from platform-side records, such as release notes, methodology changes, or vendor disclosures, rather than in-

ferred from entry data. Second, fundamentals must be controlled for using observable correlates, so that residual movement is not confounded with a common demand shock. Third, cross-sectional variation in platform penetration provides the scaling dimension: the residual shift should increase with adoption intensity.

Two extensions sharpen the design. A frequency-domain approach exploits regular platform-update cadence, 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 7 is a theoretical identifying restriction rather than a turnkey procedure. Its empirical implementation requires exogenous platform events, controls for fundamentals that address the possibility that platforms update in anticipation of demand shifts, and a defensible measure of adoption at the relevant unit. The scaling prediction also requires that $\Delta\xi$ be cross-sectionally constant, or at least uncorrelated with q, across the units in which adoption is measured. If platforms update differently by subscriber tier, or if update size depends on adoption through subscriber-generated data, the clean separation between loading and shock breaks down. The test is therefore best suited to update events that affect adopters uniformly, or to specifications in which differential update intensity can be measured from platform-side disclosures. Recent regulatory discourse has begun to treat platform events as possible instruments 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. The contribution here is to provide a structural restriction that maps onto observable variation in adoption; taking it to data is left to future work.

C Dark Side of Cheap AI

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 across platforms (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 of these cost declines reflect below-cost pricing by frontier model providers rather than improvements in inference efficiency, with leading platforms operating at sustained losses funded by venture capital (Cahn 2024). More broadly, software and platform markets exhibit rapid technological change and frequent displacement of incumbents, leading firms to reassess their choices regularly (Schmalensee 2000).

Lower switching frictions are usually viewed as beneficial because they allow firms to adopt improved technologies more quickly. That view requires qualification when adoption itself is socially excessive. The financial-regulation literature provides a useful precedent. Following the 2008 crisis, uniform mark-to-market rules and risk-management standards were criticized for generating procyclical adjustment: when many institutions respond to the same triggers using similar models, the speed of correlated adjustment can become a source of systemic risk (Brunnermeier and Pedersen 2009, Shleifer and Vishny 2011). The retail setting studied here has the same structure. The shared component ξ generates correlated errors, while the rate at which firms re-evaluate platform choices determines how quickly the industry accumulates exposure to them.

This perspective contrasts with the standard treatment of switching costs in industrial organization, where such costs slow adjustment and impede diffusion of superior technologies (Klemperer 1995). In that literature, adoption is typically privately and socially desirable, so switching frictions delay convergence to a better outcome. Here the long-run adoption level is excessive, so faster convergence accelerates the accumulation of an externality. Switching frictions impose private costs but may also moderate the speed at which the industry reaches states with high uninternalized social cost.

Policy timing reinforces this distinction. Regulatory responses to platform-driven systemic risk, such as 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 switching rate is therefore distinct from the static margins in Section 7: policies that reduce q or $H(K,\rho)$ affect the size of the externality at a point in time, while the switching rate governs how quickly the economy reaches states in which that externality is large.

To formalize this margin, suppose that in each period t an exogenous fraction $\lambda \in (0,1]$ of chains can revisit its adoption decision. A chain offered this opportunity adopts the platform if its newly drawn switching cost is below the current private gain $\Delta\Pi(q_t)$; chains not offered the opportunity retain their previous choice. Aggregate adoption 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 is 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 that is a convex combination of past choices and contemporaneous best responses, as in models of technology adoption and diffusion with randomly timed adjustment (Jovanovic and Lach 1989).

The parameter λ 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. The static equilibria of Section 6 are fixed points for any λ, so λ affects the speed of adjustment rather than the long-run outcome. In environments with 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 accelerates exposure to the shared component.

Proposition 8 (Cheap switching amplifies dynamic welfare loss)**.** *Suppose $\delta \geq 0$ and adoption evolves according to the recursion above from an initial level $q_0 \in [0, q^*)$ in the basin of attraction of a stable static equilibrium q^*, with the static iteration $q \mapsto F(\Delta\Pi(q))$ approaching q^* monotonically from below. Then, for any $t \geq 1$:*

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

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

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

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

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

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

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

The result highlights a tension between static and dynamic efficiency. Switching costs do not affect the static adoption equilibrium, but they govern how quickly the economy approaches it. When adoption is socially excessive, faster convergence increases cumulative exposure to the externality. 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 $\mathscr{W}_4$ wedge is increasing in q, earlier movement to higher q_t raises the discounted welfare loss. The same logic applies to persistence: since $\rho_x(q_t, \phi)$ is increasing in q_t, faster adjustment generates more persistent aggregate fluctuations along the transition.

Remark 6 (Beyond $\mathscr{W}_4$ and the strategic-substitute case)**.** Proposition 8(ii) exploits the monotonicity 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, so the qualitative conclusion extends beyond the specific welfare specification considered here.

When $\delta < 0$ and the static iteration oscillates around the stable fixed point, the path-monotonicity argument does not apply directly. If the oscillation is mild enough that q_t remains below q^*, the

proposition extends by the same argument. If the path crosses q^*, the cumulative-wedge claim holds in expectation across the oscillation cycle but not pointwise. The qualitative direction of the result is preserved: cheap switching accelerates the accumulation of welfare cost relative to the long-run static equilibrium. A complete characterization in the strategic-substitute regime is left to future work.

Remark 7 (Endogenous switching probability)**.** The analysis treats λ as exogenous. In practice, it is shaped by upstream factors such as platform pricing, integration costs, switching infrastructure, and provider competition. Below-cost pricing, standardized data formats, and managed-analytics services all tend to increase λ from the chain's perspective.

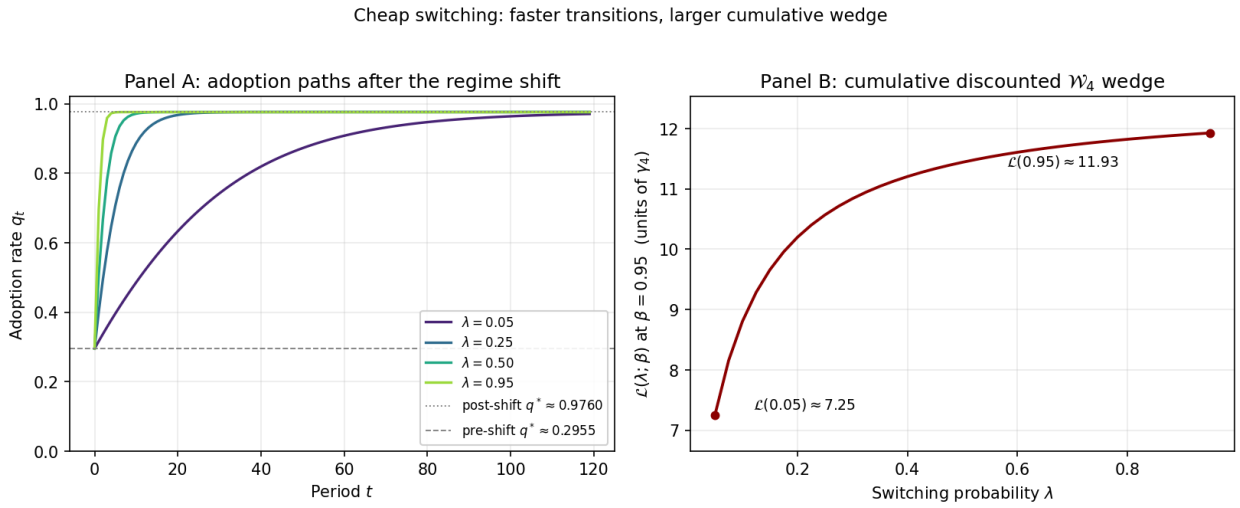
A complete model would endogenize λ and solve for a vertical equilibrium. Proposition 8 is modular with respect to this extension: any upstream model that delivers a comparative static for λ can be combined with the result above. Related work studies the upstream incentive to produce shared data inputs. Fang and Kim (2026) show that data-neutrality regulations can reduce a platform's incentive to refine its data, generating a tradeoff analogous to the one identified here. Both settings highlight that policies broadening 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 because they hinder adoption of superior technologies and create lock-in. In the present environment, some friction can be socially valuable. Slower adjustment delays the buildup of exposure to the shared component, providing time for offsetting forces such as vendor entry, diversification of data sources, or regulatory intervention to take effect. Standardized APIs, common data formats, and managed analytics services may still be productivity-enhancing, but a complete welfare assessment must also account for how they accelerate the timing and accumulation of systemic risk.

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.

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 platform idiosyncratic noise $\sigma_A^2 = 0.5$ (so that $q^* \approx 0.296$); at $t = 0$ the platform's idiosyncratic noise drops to $\sigma_A^2 = 0.05$, raising the static equilibrium to $q^* \approx 0.976$. Panel A traces the resulting adoption path q_t for four values of the per-period switching probability λ. Higher λ delivers monotonically faster convergence to the new equilibrium, in line with Proposition 8(i): at $\lambda = 0.05$ the industry takes dozens of 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 $\beta = 0.95$, in units of γ_4. The cumulative wedge rises from $\mathscr{L}(0.05) \approx 7.25$ to $\mathscr{L}(0.95) \approx 11.93$, roughly a factor of 1.6 across the plotted range of λ, 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 rises by roughly a factor of 1.6 between $\lambda = 0.05$ and $\lambda = 0.95$.

D 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 2 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.[25] 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.

D.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.

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

D.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.

[25] 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.

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.

D.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.

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

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

Estimated peer revenue is each vendor's total estimated company revenue, from third-party revenue trackers (Growjo, Latka, Owler) cross-validated against funding/valuation disclosures and customer-count claims. The exposure factor scales total revenue to the portion plausibly attributable to retail and restaurant chain analytics specifically, based on product-page coverage, public-company filings citing the vendor by use case, and vendor-disclosed customer mix. Retail-chain-adjusted share is the within-adopter share $\omega_k = (\text{revenue}_k \times \text{exposure}_k)/\sum_k(\text{revenue}_k \times \text{exposure}_k)$. Revenue-weighted shares are the appropriate weights for the model object to the extent that vendor revenue is proportional to the volume of screening decisions routed through the platform: larger chains both pay more and evaluate more candidate sites, and the model's adoption masses weight decisions rather than firms. If per-decision pricing differs systematically across vendors, revenue shares depart from decision shares; the share scenarios in Table A4 bound the resulting variation.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

F 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 $\chi_A \equiv \alpha_\theta^A + \alpha_\xi^A = (\sigma_\theta^2 + \sigma_\xi^2)/V_A$. The maintained parameter restrictions are $\delta \in \mathscr{D}$ and $\sigma_\theta^2, \sigma_A^2, \sigma_C^2, \sigma_\xi^2 > 0$.

F.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

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

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

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 $\Delta(1) = 1 - \delta\,\chi_A = (V_A - \delta(\sigma_\theta^2 + \sigma_\xi^2))/V_A$, 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 $-\delta^2\,\sigma_\theta^2\,\sigma_\xi^2/(V_A V_C)$, it is concave, and positivity at the endpoints implies $\Delta(q) > 0$ throughout $[0,1]$. □

F.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 gives $\pi^t(q) = \left[(k^t)^2 V_t + K_0^2\right]/2$; the $K_0^2/2$ terms are common to both types and cancel in the difference, so

$$\Delta\Pi(q) = \tfrac{1}{2}\left[(k^A(q))^2 V_A - (k^C(q))^2 V_C\right] = \frac{1}{2\Delta(q)^2}\left[(\alpha_\theta^A)^2 V_A - (\alpha_\theta^C)^2\,(1-\delta\,q\,\alpha_\xi^A)^2 V_C\right]. \quad \text{(L1.1)}$$

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

$$\Delta\Pi(0) = \tfrac{1}{2}\left[(\alpha_\theta^A)^2 V_A - (\alpha_\theta^C)^2 V_C\right] = \frac{\sigma_\theta^4}{2}\,\frac{V_C - V_A}{V_A V_C} = \frac{\sigma_\theta^4\,(\sigma_C^2 - \sigma_A^2 - \sigma_\xi^2)}{2\,V_A V_C}. \quad \text{(L1.2)}$$

The prefactor $\sigma_\theta^4/(2V_A V_C) > 0$, 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. □

F.3 Proof of Proposition 2

From (L1.1), $\Delta\Pi(q) = N(q)/(2\Delta(q)^2)$ where

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

Differentiating,

$$\frac{d\Delta\Pi}{dq} = \frac{N'(q)\,\Delta(q) - 2N(q)\,\Delta'(q)}{2\Delta(q)^3}. \quad \text{(P2.1)}$$

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

$$\frac{d\Delta\Pi}{dq} = \delta \cdot \frac{2\,\sigma_\theta^4\,\sigma_\xi^2}{V_A^3 V_C^2} \cdot R(q;\sigma,\delta), \quad \text{(P2.2)}$$

where $R(q;\sigma,\delta)$ is a polynomial of degree three in q and δ. On the admissible region $\delta \in \mathscr{D}$, $q \in (0,1)$, the polynomial R is strictly positive: its full expansion is reported in the replication package, where a 200,000-draw numerical check confirms strict positivity over the admissible region. We do not claim an analytic sum-of-squares certificate; the positivity statement should be read as verified numerically at this scale.

The prefactor in (P2.2) is $2\sigma_\theta^4\sigma_\xi^2/(V_A^3 V_C^2) > 0$, so the sign of $d\Delta\Pi/dq$ matches the sign of δ. □

F.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

$$\frac{\partial}{\partial k^A}\left[q\,\pi^A + (1-q)\,\pi^C\right] = q\,\frac{\partial \pi^A}{\partial k^A}\bigg|_{\bar{x}\text{ fixed}} + \int_0^1 \frac{\partial \pi_g}{\partial \bar{x}}\,\frac{\partial \bar{x}}{\partial k^A}\,dg = 0.$$

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 $\int_0^1 P_g\,dg = \bar{x}$ by aggregate consistency:

$$\int_0^1 \delta\,(P_g - \bar{x})\,dg = \delta\,(\bar{x} - \bar{x}) = 0.$$

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

$$\frac{d\mathscr{W}_1}{dq} = \pi^A(q) - \pi^C(q) + \int_0^1 \frac{\partial \pi_g}{\partial \bar{x}}\,\frac{\partial \bar{x}}{\partial q}\,dg\bigg|_{(k^A,k^C)\text{ at optimum}}.$$

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 $F^{-1}(q^*) = d\mathscr{W}_1/dq$. □

F.5 Proof of Theorem 1

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

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

$$\Delta\Pi(q) - \frac{d\mathscr{W}_3}{dq} = \gamma_3 \frac{d}{dq}\big[K_\xi(q)^2 \sigma_\xi^2\big]. \tag{T1.1}$$

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

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

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

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

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

$$V_A V_C \big[\Delta(q) - q\,\Delta'(q)\big] = V_A\,(V_C - \delta\,\sigma_\theta^2) + \delta^2\,\sigma_\theta^2\,\sigma_\xi^2\,q^2. \tag{T1.2}$$

The right-hand side is a sum of two non-negative terms: $\delta^2\,\sigma_\theta^2\,\sigma_\xi^2\,q^2 \geq 0$ for any real δ and q, and $V_A(V_C - \delta\,\sigma_\theta^2) > 0$ whenever $\delta < 1 + \sigma_C^2/\sigma_\theta^2$, which is implied by $\delta \in \mathscr{D}$. The constant-in-q term is strictly positive uniformly in q, so $\Delta(q) - q\,\Delta'(q) > 0$ throughout $[0,1]$. Combined with $V_A V_C > 0$ and $\Delta(q)^3 > 0$,

$$\frac{d}{dq}\big[K_\xi(q)^2\big] > 0 \quad \text{for all } q \in (0,1],\ \delta \in \mathscr{D}, \tag{T1.3}$$

with strict inequality whenever $q > 0$. Combining (T1.1) with (T1.3) and the positivity of γ_3 and σ_ξ^2 yields strict positivity of the wedge for all $q \in (0,1]$.

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

$$\Delta\Pi(q) - \frac{d\mathscr{W}_4}{dq} = \gamma_4 \frac{\sigma_\xi^2}{V_A} \frac{d}{dq}\big[q^2\big] = \frac{2\,\gamma_4\,q\,\sigma_\xi^2}{V_A}. \tag{T1.4}$$

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

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

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

F.6 Proof of Proposition 4

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

$$V_A(\ell) \equiv (1+\ell)^2\sigma_\theta^2 + \sigma_\zeta^2 + \sigma_A^2. \tag{P4.1}$$

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

$$\frac{\mathrm{Cov}\big(\tilde{s}_f^A - (1+\ell)\theta,\ \tilde{s}_g^A - (1+\ell)\theta\big)}{V_A(\ell)} = \frac{\mathrm{Cov}(\zeta + u_f^A,\ \zeta + u_g^A)}{V_A(\ell)} = \frac{\sigma_\zeta^2}{V_A(\ell)}, \tag{P4.2}$$

using $u_f^A \perp u_g^A$ and $u_f^A \perp \zeta$. The informative component $(1+\ell)\theta$ 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

$$\mathscr{E}_4(q) = q^2\,\frac{\sigma_\zeta^2}{V_A(\ell)}, \tag{P4.3}$$

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 $\sigma_\zeta^2/V_A(\ell)$ 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 $\sigma_\zeta^2/V_A(\ell)$ 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

$$\Delta\Pi(q) - \frac{d\mathscr{W}_4}{dq} = \gamma_4\,\frac{\sigma_\zeta^2}{V_A(\ell)}\,\frac{d}{dq}\big[q^2\big] = \frac{2\,\gamma_4\,q\,\sigma_\zeta^2}{V_A(\ell)}. \tag{P4.4}$$

Parts (i)–(iii). The right-hand side of (P4.4) is strictly positive for all $q \in (0,1]$ whenever $\sigma_\zeta^2 > 0$, and equals zero when $\sigma_\zeta^2 = 0$; the loading ℓ enters only through $V_A(\ell)$ in the denominator, establishing part (i). For fixed $q > 0$, $\partial V_A(\ell)/\partial\ell = 2(1+\ell)\sigma_\theta^2 > 0$, so the wedge is strictly decreasing in ℓ 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 $\mathscr{D}(\ell)$ by the endpoint argument of Proposition 1 with σ_θ^2 replaced by $(1+\ell)^2\sigma_\theta^2$. □

F.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

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

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

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

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

using $\operatorname{Var}(\xi_k) = \sigma_\xi^2$ and $\operatorname{Cov}(\xi_k,\xi_{k'}) = \rho\sigma_\xi^2$ for $k \neq k'$. Under symmetric vendor adoption with $q_k = q/K$, the probability that two randomly drawn adopters use the same vendor is $\sum_k (q_k/q)^2 = 1/K$, and the probability they use different vendors is $1 - 1/K$. The expected within-encounter normalized covariance is therefore

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

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

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

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

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

and

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

Both are strictly positive by (T1.3), (T1.4), 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 $\rho = 1$. □

F.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. □

F.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 $\Delta(q)$,

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

Write the ratio as $K_\xi/K_\theta = q\,\alpha_\theta^A/M(q)$ with $M(q) \equiv q\,\alpha_\theta^A + (1-q)\,\alpha_\theta^C\,(1-\delta q\,\alpha_\xi^A)$. Direct differentiation gives

$$\frac{d}{dq}\left[\frac{K_\xi(q)}{K_\theta(q)}\right] = \frac{\alpha_\theta^A\,\alpha_\theta^C\,\big(1-\delta\,q^2\,\alpha_\xi^A\big)}{M(q)^2} = \frac{\alpha_\theta^A\,\alpha_\theta^C}{M(q)^2}\cdot\frac{V_A - \delta\,q^2\,\sigma_\xi^2}{V_A},$$

using $\alpha_\xi^A = \sigma_\xi^2/V_A$. The factor $V_A - \delta\,q^2\,\sigma_\xi^2$ is strictly positive on $\mathscr{D}$: for $\delta \le 0$ it is at least $V_A > 0$, and for $\delta > 0$ the admissibility bound $\delta < 1 + \sigma_A^2/(\sigma_\theta^2+\sigma_\xi^2)$ implies $\delta\,\sigma_\xi^2 < \sigma_\xi^2 + \sigma_A^2\sigma_\xi^2/(\sigma_\theta^2+\sigma_\xi^2) < V_A$, so $\delta\,q^2\,\sigma_\xi^2 \le \delta\,\sigma_\xi^2 < V_A$ for all $q \in [0,1]$. The remaining factors $\alpha_\theta^A\,\alpha_\theta^C > 0$ and $M(q)^2 > 0$ are positive throughout ($M(q) > 0$ on $\mathscr{D}$ since $1-\delta q\alpha_\xi^A > 0$ by the same bound). A 200,000-draw numerical check in the replication package provides further confirmation. 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 $\phi > 0$. □

F.10 Proof of Proposition 8

The dynamic adoption equation is

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

with $q_0 \in [0, q^*)$ and q^* a stable static fixed point: $T(q^*) = q^*$ and $T'(q^*) < 1$ in the regime of interest. We assume throughout that q_0 lies in the basin of attraction of q^* under the iteration $q \mapsto T(q)$. In tipping regimes (Section 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

$$\left.\frac{\partial G}{\partial \lambda}\right|_q = T(q) - q, \qquad \left.\frac{\partial G}{\partial q}\right|_\lambda = (1-\lambda) + \lambda\,T'(q). \tag{P8.1}$$

Therefore

$$\frac{dq_{t+1}}{d\lambda} = \big[T(q_t) - q_t\big] + \big[(1-\lambda) + \lambda\, T'(q_t)\big]\, \frac{dq_t}{d\lambda}. \tag{P8.2}$$

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 $\delta < 0$ lies outside the hypotheses of the proposition; Remark 6 discusses the extent to which the conclusions extend to oscillating paths in the strategic-substitute regime.

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:

$$w(q) \equiv \Delta\Pi(q) - \left.\frac{d\mathscr{W}_4}{dq}\right|_q = \frac{2\,\gamma_4\, q\, \sigma_\xi^2}{V_A}, \tag{P8.3}$$

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 $dq_t/d\lambda > 0$. For absolute summability, note from (P8.2) that $D_t \equiv dq_t/d\lambda$ obeys the linear recursion $D_{t+1} = g_t + a_t D_t$ with $g_t \equiv T(q_t) - q_t$ and $a_t \equiv (1-\lambda) + \lambda\, T'(q_t)$. Convergence $q_t \to q^*$ implies $g_t \to 0$ geometrically at rate $r(\lambda) \equiv |1 - \lambda + \lambda\, T'(q^*)| < 1$, and $\limsup_t a_t = 1 - \lambda(1 - T'(q^*)) < 1$ by stability of q^*. Hence D_t is bounded and decays geometrically after finitely many periods, so the discounted series $\sum_t \beta^t D_t$ converges absolutely for any $\beta \in (0,1)$. (We write $r(\lambda)$ for this contraction factor to avoid collision with the cross-platform correlation ρ of Section 7 and the autocorrelation ρ_x.) 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 $t \geq 1$. □

F.11 Proof of Proposition 7

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 (T1.2). 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. □