Shared Intelligence and Misaligned Firm Entry

Feng Chi* Nathan Yang†

June 21, 2026

Abstract

Firms increasingly rely on shared intelligence platforms for market entry. While these systems improve forecast accuracy, they also route decisions through a common processing layer that creates correlated errors unrelated to fundamentals. We develop a Bayesian entry game in which firms choose between proprietary research and a shared platform. Platform adoption generates two externalities: synchronized entry and correlated screening errors. A planner who internalizes these costs always prefers less adoption than one who values only firm profits. This external wedge is independent of whether entry decisions are strategic complements or substitutes, and over-adoption arises across a broad range of environments.

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

Retail entry is increasingly shaped by a common set of algorithmic forecasting platforms. Historically, firms 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. Recent work shows that the quality and timing of information can materially shape firm entry and survival (Chi 2025). We study a related but distinct margin: not whether firms have more timely information, but whether many firms process information through the same forecasting layer. Today, national chains, regional grocers, and franchisees increasingly rely on the same external analytics providers (e.g., CoStar, Placer.ai, and Buxton) to forecast trade-area demand, foot traffic, and competitive density.[1] These platforms can improve forecast accuracy by pooling data and reducing firm-specific noise. They also route many firms' decisions through a common processing layer. If that layer contains errors unrelated to local fundamentals, firms may make mistakes that are not only inaccurate but correlated.

This possibility is readily visible 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. This alignment is especially pronounced in counties with greater exposure to analysis using AI (Felten et al. 2021). Such co-location may reflect common demand fundamentals, agglomeration economies, or strategic imitation. The rapid proliferation of shared analytics vendors suggests an additional possibility. Firms may enter the same markets not only because those markets are attractive, but also because the same forecasting system 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 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. Forecasts from platforms differ because adoption does not make the common component more informative; instead, the same error is propa-

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). These mechanisms generate correlation because firms learn from common information or from one another's actions. By contrast, the channel we study involves no learning from others. Firms act on a common error embedded in a shared forecasting system, so correlation arises even when each firm conditions only on its own forecast and no firm observes another's action. Recent evidence suggests that algorithmic prediction can improve individual decisions while increasing correlation across firms. Firms investing heavily in AI exhibit greater systematic risk (Babina et al. 2024), and firms sharing data through common software development kits exhibit co-movement that survives controls for product similarity (Bian et al. 2024). What remains less understood is how welfare changes when many firms rely on the same forecasting infrastructure.

This distinction also separates our mechanism from related work on the productivity gains from shared data and forecasting infrastructure. As Farboodi and Veldkamp (2023) emphasize, data affect markets by improving prediction, sharpening targeting, changing competition, and altering the value of firms and financial assets. In models of data-driven automation, common data raise factor productivity and propagate across tasks, making shared infrastructure an engine of growth rather than a source of correlated error (Farboodi et al. 2026, Farboodi and Veldkamp 2021, Jones and Tonetti 2020). In those settings, whether shared data are accumulated too quickly or too slowly depends on how substitutable the affected activities are. Our focus is different. We isolate the part of a shared signal that remains after conditioning on fundamentals and therefore carries no productive content. This non-informative component generates correlated errors across adopters, and the resulting externality always points toward less platform adoption. The welfare wedge we study arises from correlated errors rather than from the productive component of shared data.

We develop a model in which firms choose between proprietary research and a shared forecasting platform. Proprietary research generates firm-specific forecast errors, whereas the shared platform reduces firm-specific noise but introduces a common processing error across adopters. This tradeoff creates a gap between privately and socially optimal adoption. Platform adoption generates two external costs that firms do not internalize: synchronized entry that affects parties exposed to aggregate fluctuations in commercial activity, and correlated screening errors that af-

gated across more decisions.

fect locations evaluated through the same forecasting technology. A planner who internalizes either cost chooses strictly less adoption than a planner who values only firm profits.

A central result is that this external wedge does not depend on whether entry decisions are strategic complements or substitutes. For the location-level welfare cost, the size of the wedge depends only on the adoption rate and the signal structure. The shared processing externality therefore differs from classical entry distortions, whose direction depends on the nature of strategic interaction. The shared processing externality always points toward less adoption.

The shared processing externality operates alongside the classical entry externality of Mankiw and Whinston (1986) and Spence (1976). Because firms ignore their effect on aggregate entry, decentralized outcomes are already inefficient even in the absence of shared processing. Whether the market ultimately over- or under-adopts therefore depends on the interaction between the classical entry distortion and the new externality introduced by shared forecasting systems. We show that over-adoption arises throughout the strategic substitutes region and under moderate complementarity, reversing only when agglomeration forces are sufficiently strong.

We also examine whether vendor competition mitigates the distortion. Multiple vendors reduce correlated errors only to the extent that their underlying data and models are independent. A descriptive case study of the restaurant site-selection analytics market suggests substantial overlap in upstream inputs and forecasting practices. Under empirically disciplined calibrations, the observed market retains roughly three-quarters of the distortion generated by a single shared platform. Vendor competition therefore provides limited diversification when providers draw on similar data sources and forecasting architectures.

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 asks whether the two chains select similar types of markets.

The exercise is descriptive rather than causal. Its purpose is to show that the object formalized in the model, cross-firm alignment in market selection, is visible in a familiar retail setting before any structure is imposed. The analysis does not attempt to distinguish among common demand

fundamentals, agglomeration economies, imitation, franchising constraints, real-estate constraints, or shared analytics. Instead, it documents a simple empirical regularity. McDonald's and Burger King exhibit highly similar mappings from observable market characteristics into store counts. This similarity is especially pronounced in counties with greater exposure to AI-supported analysis.

2.1 Industry Setting and Data

We assemble a county-level snapshot of McDonald's and Burger King locations in the United States.[3] Brand locations come from OpenStreetMap, and each location is assigned to a county through point-in-polygon matching against Census county boundaries. We merge these counts with county characteristics from the American Community Survey, the Census Gazetteer, and the USDA Rural-Urban Continuum Codes, as well as the county-level AI Geographic Exposure (AIGE) index of Felten et al. (2021).

The AIGE index measures how exposed a county's employment structure is to advances in artificial intelligence. Following Felten et al. (2021), the measure is constructed by combining occupation-level AI exposure scores with local employment composition. Counties with higher AIGE therefore have industry and occupation mixes that are more exposed to AI-related technologies. We use AIGE as a stratifying variable rather than as a direct measure of platform adoption.

Two features of the data are important for interpretation. 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, so some locations may be present but untagged. The estimates should therefore be interpreted as descriptive rather than exhaustive. After restricting to counties with both brands collected and non-missing covariates, the sample contains 1,987 counties.

[3] 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 maintain substantial in-house real-estate and analytics capabilities, making this pair, if anything, a conservative setting for the shared-platform channel. Appendix F identifies chains that are clients of the same analytics vendors. Comparing alignment among disclosed co-clients against alignment among non-co-client pairs is a promising direction for future work with store-census data.

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 object of interest is whether McDonald's and Burger King map observable market characteristics into store counts in similar ways.

We estimate chain-specific relationships between store counts and county fundamentals and compare the resulting coefficient vectors. This asks whether the two chains place similar weights on the same observable market characteristics when evaluating markets. All continuous fundamentals are standardized before estimation so that coefficients are comparable across variables (Gelman 2008). We summarize cross-chain alignment using three measures: 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 measure similarity in estimated slopes. 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 yields a Wald test of slope equality across chains (Paternoster et al. 1998). Exact equality is not the relevant benchmark. With a large county sample, economically small differences can generate statistical rejection even when the two chains rank markets similarly. We therefore view the Wald test as a diagnostic and place greater weight on the directional evidence from the coefficient and score-based measures.

Finally, we split counties at the median of the AI Geographic Exposure index and recompute the alignment measures separately in low- and high-exposure subsamples. This asks whether market-selection alignment is more pronounced in counties whose employment structure is more exposed to AI. Consistent with the descriptive approach emphasized by Einav and Levin (2014), the goal is to document a robust empirical pattern that motivates the theoretical mechanism developed below.

2.3 Co-Location

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.

Geographic alignment is therefore a prominent feature of the data. The two chains frequently co-locate, and this pattern survives coarse controls for observable market characteristics. Such alignment is not unique to the modern analytics era. Toivanen and Waterson (2005) document strategic interaction and co-location among U.K. hamburger chains using data that predate contemporary forecasting platforms. Our mechanism therefore does not seek to explain all cross-firm alignment. Instead, it highlights an additional channel through which shared forecasting infrastructure may increase alignment beyond what would arise from common demand fundamentals, agglomeration economies, or strategic interaction alone.

A cross-sectional snapshot cannot isolate this channel. The evidence here should therefore be interpreted as descriptive rather than causal. The purpose of the exercise is to document the object studied in the model, cross-firm alignment in market selection, and to show that it is economically important in a familiar retail setting.

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. We read the Poisson result as the more informative of the two for the present question: the dependent variable is a store count with a substantial mass at zero, for which the Poisson conditional mean is the natural specification, whereas the $\log(1+n)$ transform compresses scale differences into the slope and rejects largely because the two chains differ in footprint size rather than in how they rank markets. The directional diagnostics below, which do not hinge on either functional form, make the same point more robustly.

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 stronger in high-AIGE counties. The predicted-score correlation is 0.9849 in low-AIGE counties and 0.9978 in high-AIGE counties, with coefficient-

based measures showing the same pattern. This does not imply that AI exposure causes chains to behave more similarly. Rather, counties whose employment structures are more exposed to AI are also places where the two chains' observable market-selection rules are especially closely aligned.

These patterns motivate the model in two ways. First, McDonald's and Burger King exhibit substantial geographic alignment. Second, they map observable market characteristics into store counts in highly similar ways, especially in high-AIGE counties. This alignment is consistent with several mechanisms, including common demand fundamentals, agglomeration economies, imitation, franchising constraints, and shared analytics. The model focuses on the last channel and asks what happens when firms evaluate markets through related forecasting systems. In that case, entry decisions may become correlated in ways that firms do not fully internalize.

The descriptive evidence is consistent with this mechanism but does not identify it. Similar market-selection rules can arise from common fundamentals, strategic interaction, or shared forecasting infrastructure. Distinguishing among these channels requires variation that shifts the forecasting process without changing underlying market conditions, an identifying challenge that we leave to future work. The purpose of the exercise is therefore not to test the mechanism directly, but to document the empirical object that the model seeks to explain.

As a supplementary diagnostic, we decompose the difference in predicted-score alignment between high- and low-AIGE counties across the seven continuous market fundamentals using exact Shapley attribution.[4] The resulting contributions are small, partially offsetting, and imprecisely estimated. No individual fundamental accounts for a statistically distinguishable share of the difference. We therefore interpret the decomposition only as a diagnostic. The stronger alignment observed in high-AIGE counties does not appear to be driven by any single observable characteristic, but instead reflects broad-based similarity across market fundamentals.

3 Model

The model studies retail entry when firms choose both whether to enter a market and whether to rely on proprietary research or a shared forecasting platform. The central tradeoff is simple. Shared platforms can improve forecast accuracy by reducing idiosyncratic noise, but they also route decisions through a common processing layer that may generate correlated errors across

[4] Shapley-value attribution is widely used in machine-learning interpretability (Lundberg and Lee 2017). We apply the exact decomposition of Shapley et al. (1953) to the high-minus-low AIGE difference in predicted-score alignment.

adopters. The model isolates this tradeoff in the simplest environment that permits heterogeneous adoption and strategic interaction.

A key distinction is between primitives and outcomes. The variance of the shared processing component is a primitive of the information structure, whereas platform adoption is an endogenous equilibrium outcome that reflects firms' responses to those informational conditions.[5]

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.[6] 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.

[5] Throughout the analysis, σ_ξ^2 is treated as an exogenous property of the forecasting technology.

[6] 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 $[0,1]$. Appendix D computes this clipped-rule equilibrium directly under a logistic entry probability and confirms that the external wedge retains its sign and monotonicity across the parameter region, with the $\mathscr{W}_4$ wedge remaining within one percent of its linear value.

3.2 Information and Processing Technologies

The model distinguishes between two approaches to forecasting. Firms can either generate forecasts internally using proprietary research or rely on a shared external platform. This distinction has become increasingly important as firms adopt AI-based forecasting technologies. Vendor-provided systems can improve forecasting performance by pooling data and analytical resources across users, but they may also reduce differentiation when many firms rely on similar tools.

This make-versus-buy margin has received growing attention in the empirical literature on firm AI. Babina (2026) identify it as a key dimension of AI adoption, while related evidence shows that firms investing more heavily in AI exhibit greater systematic risk (Babina et al. 2024) and that firms relying on external providers behave differently from those developing capabilities internally (Calvino and Fontanelli 2026). Our focus is on a question that this literature has largely left open: what are the welfare consequences when many firms rely on the same forecasting infrastructure?

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

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

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

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

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

The distinction between the two technologies is 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 shared component ξ represents forecast errors that originate in the processing technology rather than in market fundamentals. The motivation comes from a recurring lesson in computer

science: prediction systems can exhibit similar measured accuracy while sharing latent sources of failure. Different pipelines may perform similarly in validation exercises yet make correlated mistakes in deployment because they rely on similar data, modeling choices, or software infrastructure (D'Amour et al. 2022, Sculley et al. 2015). When many firms rely on related forecasting systems, these mistakes 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). There, a common signal induces correlation because it is informative about the underlying state. As more agents observe the signal, actions become more aligned, but inference also improves. Here the shared component carries no information about θ. It survives aggregation as a common error, so joint exposure does not refine beliefs and instead transmits the same mistake across adopters.[7]

One way to interpret ξ is as a latent processing factor. When many chains rely on the same forecasting 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 importance of this shared component. The assumption that ξ is orthogonal to θ captures that these errors originate in the forecasting process rather than in the underlying market.

3.3 Adoption Stage

Each chain draws a switching cost $c_f \sim F$ independently, where F is a continuous distribution on $\mathbb{R}$ with positive density on its support. The chain chooses its forecasting technology to maximize expected entry profits net of this cost, taking the adoption rate q as given. An adoption equilibrium is a rational-expectations fixed point in which individual choices reproduce the aggregate:

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

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.

[7] This differs from environments such as Bergemann et al. (2015), where common signals may remain informative even when they confound multiple sources of uncertainty.

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.

The assumption that firms know σ_ξ^2 is both consequential and conservative. It corresponds to a setting in which adopters understand the residual error structure of the forecasting platform and correctly account for it when forming beliefs. Such firms recognize that platform forecasts contain a shared component unrelated to local demand and optimally discount that component when combining the forecast with prior information.

Many real-world adopters are unlikely to satisfy this benchmark. A regional grocer purchasing a site-selection package, a franchisee relying on analytics supplied by a franchisor, or a tenant receiving recommendations through a brokerage relationship may observe only a point forecast and have little information about the uncertainty underlying it. In such settings, firms may place too much weight on the platform signal and therefore react too strongly to the shared component.

The rational-filtering benchmark therefore provides a lower bound on the distortion generated by shared processing. The model isolates the welfare consequences of the information structure itself, abstracting from additional distortions that arise when firms misperceive the precision of platform forecasts. Any departure from correct filtering that increases reliance on the shared component strengthens the correlation in entry decisions and amplifies the resulting externality.[8]

[8]The same logic extends beyond the binary adoption decision studied here. If firms instead choose the intensity with which they rely on a shared platform, private incentives continue to place excessive weight on the shared component because firms do

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,

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

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).

not internalize the correlation imposed on other adopters. The correlated-evaluation externality therefore reappears on the intensive margin.

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 $\mathscr{D}$.[9]

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 $\delta > 0$, 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

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

[9] 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.

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 aggregate variation into 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 aggregate fluctuations reflect fundamentals alone. Once platform adoption is positive, however, aggregate outcomes also inherit variation from the shared processing component.

The distinction matters because the two sources of variation behave differently under aggregation. Idiosyncratic noise averages out in the continuum, but the shared shock survives. As a result, correlation across firms is no longer tied exclusively to information about demand. Part of that correlation reflects a common error embedded in the forecasting technology itself. The term $K_\xi(q)^2\sigma_\xi^2$ captures this source of aggregate variation and serves as the foundation for the welfare distortions studied below.

5 Adoption Incentives and Misalignment

This section characterizes adoption incentives and compares equilibrium adoption to social optima under alternative welfare criteria. The analysis separates two forces. The first is a classical entry externality: because each atomistic chain ignores its effect on the aggregate, the private equilibrium already departs from the chain-welfare optimum whenever entry decisions are strategic, even with no shared error. The second is the externality of interest: the shared component imposes costs on parties outside the chain population that no chain internalizes. We characterize the classical wedge first, then show that the shared component opens a further, external wedge under which a complete-welfare planner prefers strictly less adoption than a chain-welfare planner, regardless of strategic interaction.

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 $\delta = 0$. For $\delta \neq 0$, $\Delta\Pi(q)$ is strictly increasing in q when $\delta > 0$ and strictly decreasing in q when $\delta < 0$. The sign result is established in closed form for strategic interaction near zero (Appendix G) and holds throughout the admissible region $\mathscr{D}$.*

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 G gives

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

where Γ is strictly positive over the admissible region $\mathscr{D}$; positivity reduces to a sign condition on a compact set, verified over 200,000 parameter draws in the replication package. The sign of $d\Delta\Pi/dq$ is therefore determined by δ.

5.2 Externalities with 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 externality operates at the aggregate level. When many firms rely on the same forecasting platform, aggregate entry becomes more sensitive to the shared processing component. This creates costs for agents whose payoffs depend on fluctuations in commercial activity around fundamentals, including lenders exposed to correlated default risk, landlords facing synchronized vacancies, and municipalities planning around local economic conditions. The relevant welfare object is the platform-driven component of aggregate variance, $K_\xi(q)^2\sigma_\xi^2$, which increases with adoption and survives aggregation.

The second externality operates at the location level. When multiple firms evaluate the same location through a shared platform, their forecasts inherit the same processing error. Evaluation mistakes therefore become positively correlated across firms. As adoption rises, independent screening becomes less effective because more firms rely on the same underlying signal. The corresponding welfare object is the expected covariance of evaluation errors across firms. This term scales as $q^2\sigma_\xi^2/V_A$, which reflects both the probability that two firms adopt the platform and the importance of the shared component in platform forecasts.

This second channel is related to the algorithmic-monoculture concerns studied by Kleinberg and Raghavan (2021) and Peng and Garg (2024). In those settings, the decided-about agents are individuals, and welfare depends on whether at least one decision-maker recognizes a favorable case. Here, the decided-about agent is a location, and the distortion arises because multiple firms evaluate the same market through correlated signals. The underlying concern is similar, but the relevant welfare object differs.

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 pairwise structure is a normalization, not a substantive restriction: if instead each location is screened by $n \geq 2$ evaluators drawn independently from the population, the expected number of co-adopting evaluator pairs is $\binom{n}{2}q^2$, each contributing a normalized error covariance σ_ξ^2/V_A, so the penalty is $\binom{n}{2}\,q^2\,\sigma_\xi^2/V_A$. The quadratic dependence on q and the linear dependence on σ_ξ^2/V_A are therefore properties of the encounter structure for every n; only the multiplicative constant changes, and it is absorbed into γ_4. The δ-invariance and strict positivity of the external $\mathscr{W}_4$ wedge in Theorem 1 hold unchanged for any fixed n.

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.[10]

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. As Proposition 3 below shows, the private

[10]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 figures and the replication code.

equilibrium does not in general maximize $\mathscr{W}_1$: a strategic externality separates the two whenever $\delta \neq 0$. This classical wedge is logically separate from the shared-processing externalities that are the focus of this paper, and the two add.

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 $\mathscr{W}_3$ criterion penalizes realized aggregate behavior and therefore depends on equilibrium responses. The $\mathscr{W}_4$ criterion instead penalizes correlation in signals and depends only on adoption and the underlying information structure. As a result, the external $\mathscr{W}_4$ wedge at any fixed adoption rate is independent of strategic interaction, whereas the external $\mathscr{W}_3$ wedge varies with δ through its effect on equilibrium behavior.[11]

The penalty weights γ_3 and γ_4 determine the scale of the corresponding welfare losses but are not separately identified by the model. The identified content of the analysis is therefore the sign of each external wedge and its variation across environments rather than the absolute magnitude of welfare losses.

[11] The two criteria 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 criterion that depends only on adoption and the information structure, rather than on equilibrium responses, generates an external wedge whose sign and magnitude at fixed q are independent of strategic interaction.

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 criterion captures costs borne within the chain population because deviations from fundamental tracking ultimately affect firms' own profits. Unlike $\mathscr{W}_3$ and $\mathscr{W}_4$, however, $\mathscr{W}_2$ is not an additive perturbation of $\mathscr{W}_1$. The clean external-wedge decomposition of Theorem 1 therefore does not apply, and Remark 3 shows that the resulting wedge is sign-indeterminate.

The distortion studied here instead arises from costs borne outside the chain population. A planner concerned with both objectives would naturally combine the two criteria, in which case the external wedge identified below would operate alongside any distortion associated with imperfect fundamental tracking.

5.4 Marginal Chain Welfare and the Classical Entry Externality

The first welfare result corrects a benchmark that is easy to state wrongly. One might expect the linear-Gaussian continuum to deliver *team efficiency*, with the private equilibrium maximizing $\mathscr{W}_1$. It does not. The atomistic chain takes the aggregate $\bar{x}$ as given and therefore does not internalize how its technology choice shifts the strategic surplus $\frac{\delta}{2}\mathrm{Var}(\bar{x})$ that all chains share. The result is a strategic externality whenever $\delta \neq 0$.

Proposition 3 (Marginal chain welfare)**.** *For all $q \in (0,1)$,*

$$\frac{d\mathscr{W}_1}{dq} = \Delta\Pi(q) + \Psi(q), \qquad \Psi(q) \equiv \frac{\delta}{2}\frac{d}{dq}\mathrm{Var}(\bar{x}) = \frac{\delta}{2}\frac{d}{dq}\left[K_\theta(q)^2\sigma_\theta^2 + K_\xi(q)^2\sigma_\xi^2\right].$$

Hence $d\mathscr{W}_1/dq = \Delta\Pi(q)$ if and only if $\delta = 0$. For $\delta \neq 0$ the private equilibrium q^* does not coincide with the chain-welfare optimum q_1^{soc} defined by $F^{-1}(q_1^{soc}) = d\mathscr{W}_1/dq$: atomistic chains do not internalize their effect on the shared strategic surplus, exactly the entry externality of Mankiw and Whinston (1986) and Spence (1976). The sign of the resulting adoption distortion is $\mathrm{sgn}(\Psi(q)) = \mathrm{sgn}\big(\delta \cdot \frac{d}{dq}\mathrm{Var}(\bar{x})\big)$, which is calibration-dependent because $\frac{d}{dq}\mathrm{Var}(\bar{x})$ need not be signed uniformly over q.

The strategic surplus $\frac{\delta}{2}\mathrm{Var}(\bar{x})$ captures the effect of entry on other firms through the $\delta\bar{x}$ term in profits. Under agglomeration ($\delta > 0$), additional entry benefits rivals and the surplus is positive. Under business stealing ($\delta < 0$), additional entry harms rivals and the surplus is negative. Be-

cause firms choose forecasting technologies based on private profits, they do not internalize how adoption changes the dispersion of aggregate entry. As a result, the chain-welfare planner and the decentralized market generally choose different adoption levels.

This is the familiar wedge from the free-entry literature, translated from the entry margin to the technology-adoption margin. It is present even when $\sigma_\xi^2 = 0$ and therefore exists independently of any shared-processing considerations. Appendix G shows that this force is summarized by the term Ψ in the planner's first-order condition.

The shared-processing externality operates in addition to this classical distortion. To isolate the cost borne outside the chain population, we compare a chain-welfare planner, who maximizes $\mathscr{W}_1$, with a complete-welfare planner, who additionally internalizes the external cost. The difference between the two objectives is the object of interest. Because $\mathscr{W}_1$ enters both planners' problems identically, it cancels in the comparison. The resulting external wedge therefore does not inherit the strategic term Ψ and does not depend on any team-efficiency argument.

Remark 2 (Why the external wedge is robust)**.** The external wedge in Theorem 1 is the difference between the marginal values of adoption under $\mathscr{W}_1$ and $\mathscr{W}_k$. Because the two welfare criteria differ only through the penalty term, the contribution of $\mathscr{W}_1$ cancels exactly. This cancellation is algebraic and does not rely on a continuum of firms, linear entry probabilities, or any team-efficiency result. Consequently, finite-N and nonlinearity concerns affect the strategic term Ψ but not the external wedge itself. In particular, the $\mathscr{W}_4$ penalty depends only on the adoption rate and the information structure through $q^2\sigma_\xi^2/V_A$. Its contribution to the external wedge therefore survives in finite-N and nonlinear environments. Appendix D confirms this numerically.

5.5 Misalignment Theorem

The central welfare result is that internalizing either external cost calls for strictly less adoption. We state it as a wedge between the chain-welfare planner and the complete-welfare planner, which is the part of the distortion attributable to shared processing, and then relate it to the private equilibrium.

Theorem 1 (Robust external wedge)**.** *Fix $\delta \in \mathscr{D}$ and $\sigma_\theta^2, \sigma_A^2, \sigma_C^2, \sigma_\xi^2 > 0$. Define the external wedge under criterion $k \in \{3,4\}$ as $\mathscr{X}_k(q) \equiv d\mathscr{W}_1/dq - d\mathscr{W}_k/dq$.*

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

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

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

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

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

The wedge depends only on adoption and signal-structure primitives and is invariant in δ at any fixed q.

Consequently the complete-welfare optimum lies strictly below the chain-welfare optimum: $q_k^{soc} < q_1^{soc}$ for $k \in \{3,4\}$. Relative to the private equilibrium, the total gap decomposes as

$$\Delta\Pi(q) - \frac{d\mathscr{W}_k}{dq} \;=\; \mathscr{X}_k(q) \;-\; \Psi(q),$$

so the market over-adopts ($q^* > q_k^{soc}$) whenever the external wedge $\mathscr{X}_k$ dominates the strategic term Ψ. At the calibrations examined in Figure 1, which span the strategic-substitutes region and moderate complementarity, $\mathscr{X}_k$ dominates and the market over-adopts; the comparison reverses only when the classical term is positive and exceeds the external wedge ($\Psi > \mathscr{X}_k$), which by Proposition 3 requires δ and $\frac{d}{dq}\mathrm{Var}(\bar{x})$ to share a sign with Ψ large in magnitude, as under strong agglomeration.

The two criteria share a common structure: the external wedge equals the marginal external cost of adoption. Under $\mathscr{W}_3$, this cost is the derivative of $K_\xi(q)^2\sigma_\xi^2$, which combines the primitive variance 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 relevant object is the derivative of $q^2\sigma_\xi^2/V_A$, which equals the encounter probability times the within-adopter variance share. Because equilibrium loadings do not appear, the wedge at a fixed adoption rate is independent of δ. The external result therefore does not depend on whether entry decisions are strategic complements or substitutes. It follows from two features of the environment: the presence

of a common component unrelated to fundamentals and the fact that platform adoption increases exposure to that component.

The decomposition $\Delta\Pi - d\mathscr{W}_k/dq = \mathscr{X}_k - \Psi$ makes the two forces explicit. The shared-processing externality $\mathscr{X}_k$ always points toward less adoption. The classical strategic externality Ψ points toward less adoption under business stealing and toward more adoption under agglomeration. Whether the decentralized market over- or under-provides adoption relative to the complete-welfare optimum depends on the balance of these two forces. The contribution of this paper is the first term: a distinct external channel that is strictly signed and, under $\mathscr{W}_4$, independent of the strategic environment.[12]

Of the two criteria, $\mathscr{W}_4$ provides the cleaner benchmark. Its external wedge depends only on adoption and signal-structure primitives. As Appendix D shows, it is also essentially invariant to the entry-choice rule, remaining within one percent of its closed-form value across all calibrations considered. The aggregate-volatility criterion $\mathscr{W}_3$ yields the same qualitative conclusion, but its magnitude inherits the equilibrium loadings and is therefore more sensitive to the curvature of the choice rule. We therefore view $\mathscr{W}_4$ as establishing the external-wedge result and $\mathscr{W}_3$ as corroborating it.

This invariance has an additional empirical advantage. The strategic parameter δ is difficult to identify because agglomeration and business-stealing effects often coexist and vary across settings. Moreover, δ enters both the classical wedge Ψ and the magnitude of $\mathscr{X}_3$. Under $\mathscr{W}_4$, no such identification is required. The external wedge depends only on adoption and the signal structure.

The decided-about-market penalty in $\mathscr{W}_4$ enters as an additive perturbation of chain welfare, with the encounter structure underlying it sketched in Remark 1. Two questions about this construction bear on the interpretation of the external wedge. First, is the screening loss a product of the entry environment, or an object appended to firm profits to obtain the result? Second, does the δ-invariance of $\mathscr{X}_4$ in Theorem 1(ii) reflect a property of the externality, or an artifact of evaluating it outside the entry game? Appendix E addresses both. It embeds the screening problem in a unified environment of markets each screened by $n \geq 2$ firms drawn from the population, derives the exposure $\binom{n}{2} q^2 \sigma_\xi^2 / V_A$ as an exact second moment of that environment rather than a

[12] The wedge admits a direct Pigouvian interpretation, but it must internalize both forces. The tax $\tau_k(q) = \mathscr{X}_k(q) - \Psi(q) = \Delta\Pi(q) - d\mathscr{W}_k/dq$ on platform adoption aligns the decentralized condition with the complete planner's first-order condition $F^{-1}(q) = d\mathscr{W}_k/dq$. The component that corrects shared processing specifically is $\mathscr{X}_k(q)$; the component $-\Psi(q)$ corrects the classical entry externality, and its sign follows that of Ψ in Proposition 3: it is a subsidy under agglomeration in the case $\frac{d}{dq}\mathrm{Var}(\bar{x}) > 0$, where $\Psi > 0$.

posited penalty, and promotes the wedge to Proposition 9. The same analysis shows that the δ-invariance is diagnostic rather than incidental. The screening externality runs through the residual covariance of signals, which contains no equilibrium loading; strategic interaction would re-enter only through cross-market linkages, and then through the same channel as the aggregate-volatility wedge $\mathscr{X}_3$. The decided-about-market externality is therefore a screening externality, distinct from the coordination externality embedded in $\mathscr{X}_3$, and the distinction is precisely the dependence on δ.

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 as transfers and penalizing deviations from the fundamental in the spirit of Angeletos and Pavan (2007):

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

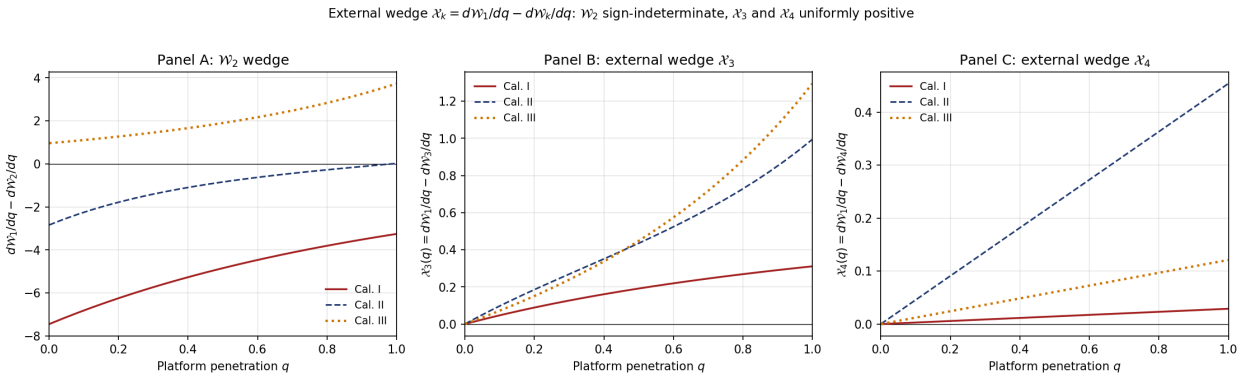
This criterion penalizes deviations of each type's signal loading from one, the loading at which entry passes signals through fully; an alternative tracking criterion that penalizes the mean-squared gap between aggregate entry and the fundamental, $\mathbb{E}[(\bar{x}-\mathbb{E}[\bar{x}]-\theta)^2]$, 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 its wedge does not reduce to a clean penalty derivative and instead carries indirect loading terms.

The resulting wedge $d\mathscr{W}_1/dq - d\mathscr{W}_2/dq$ 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. A determinate external wedge 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, and the wedge is calibration-dependent. Panels B and C show that criteria targeting external costs yield a uniformly positive external wedge. The result therefore reflects a structural feature of the model rather than a particular welfare choice: whenever

Figure 1: The external wedge under three welfare specifications.



Notes: Each panel plots a marginal welfare wedge 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 $\mathcal{W}_2$ wedge $d\mathcal{W}_1/dq - d\mathcal{W}_2/dq$: it is negative throughout for Calibrations I and II and positive throughout for Calibration III, with no interior zero-crossing, so the same primitives deliver no determinate prediction about over- or under-adoption relative to the $\mathcal{W}_2$ planner (Remark 3). Panel B shows the external wedge $\mathcal{X}_3(q) = d\mathcal{W}_1/dq - d\mathcal{W}_3/dq$ under aggregate-volatility welfare: 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 external wedge $\mathcal{X}_4(q)$ under decided-about-agent welfare: 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 external wedge is robust across two structurally distinct externality penalties. The net private-vs-social comparison additionally subtracts the classical strategic term $\Psi(q)$ of Proposition 3, which is not plotted.

the planner internalizes costs borne outside the chain population, the optimal adoption level falls.

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 the external wedge depends on the common component being entirely unrelated to fundamentals. To draw this boundary, suppose the shared 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*

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

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

$$\mathscr{X}_4(q) \;=\; \frac{d\mathscr{W}_1}{dq} - \frac{d\mathscr{W}_4}{dq} \;=\; \frac{2\gamma_4 q\,\sigma_\zeta^2}{V_A(\ell)}.$$

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 external wedge 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 external 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_ξ^2 that survives conditioning on fundamentals, not the joint informativeness of the signal. This is the sense in which the external-wedge 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.

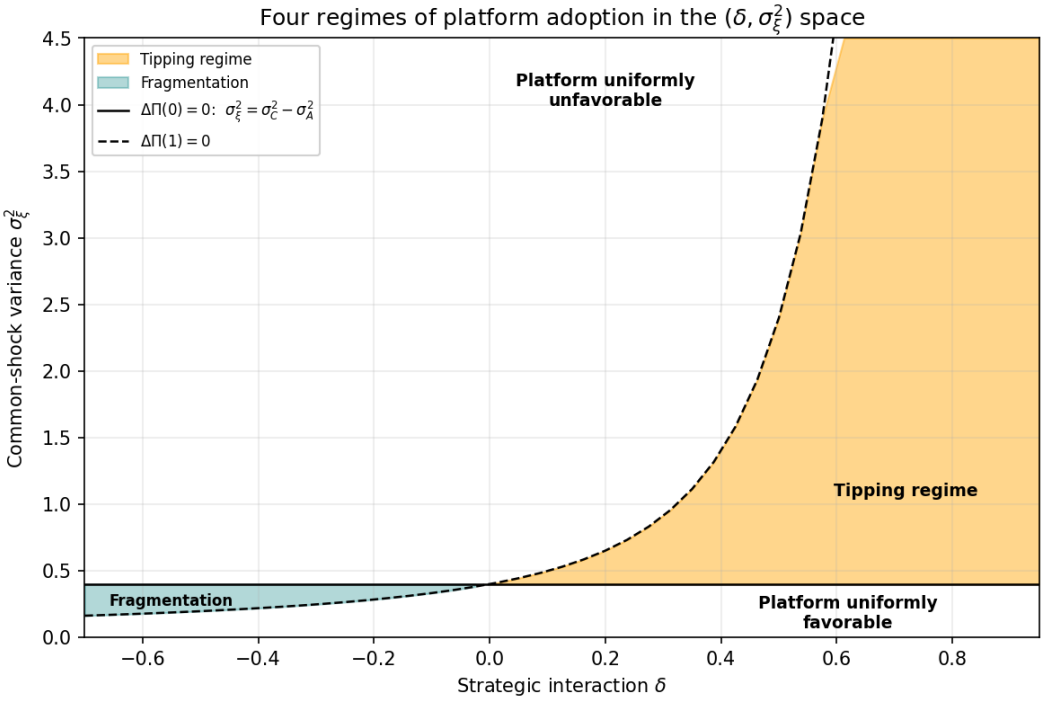
Lemma 1 (δ-independent threshold)**.** *The condition* $\Delta\Pi(0) = 0$ *holds if and only if* $\sigma_\xi^2 = \sigma_C^2 - \sigma_A^2$, *independent of* δ.

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 external wedge of Theorem 1

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

is strictly positive throughout: strategic interaction shapes how adoption is sustained and whether the *net* private distortion is over- or under-adoption, but the component attributable to shared processing always points toward less adoption. The regimes therefore differ in the level and structure of adoption, not in the sign of the external wedge. Two cases are particularly informative: the tipping regime under agglomeration and the fragmentation regime under business stealing.

6.1 Tipping Regime

When $\delta > 0$, $\Delta\Pi(0) < 0$, and $\Delta\Pi(1) > 0$, the fixed-point condition $q^* = F(\Delta\Pi(q^*))$ 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.[13] Figure 3 illustrates such a configuration for a logistic F with steepness parameter κ_F large enough for the curvature condition to hold.

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.

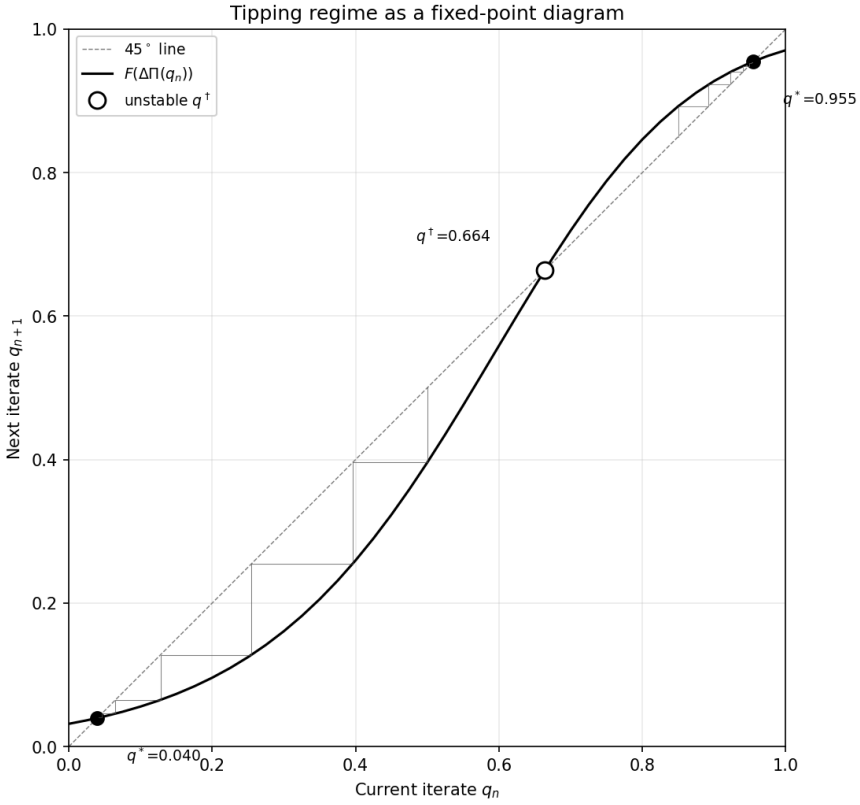
This structure parallels classic models of technology adoption under increasing returns. In Arthur (1989), small historical events are amplified through positive feedback, leading the market to lock in to a single technology. Here, the feedback operates through strategic interaction and shared processing rather than technological increasing returns alone. The inefficiency is also distinct: coordination may be privately valuable, but the shared component generates correlated errors that a complete-welfare planner would penalize.

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 $\mathscr{W}_4$, a complete-welfare planner would locally prefer less adoption than the chain-welfare planner at each attracting fixed point, by the external wedge of Theorem 1. Whether each fixed point also lies above the complete planner's own optimum depends on the net sign $\mathscr{X}_k - \Psi$, which is positive except under strong agglomeration with $\frac{d}{dq}\mathrm{Var}(\bar{x}) > 0$, where the classical term $\Psi > 0$ can dominate. The distinction between marginal and global comparisons is also important. The external wedge $\mathscr{X}_k > 0$ is a marginal statement at every interior q. The global comparison between q_L^* and q_H^* instead depends on welfare levels,

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

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



Notes: The figure plots the next iterate $q_{n+1} = F(\Delta\Pi(q_n))$ against the current iterate q_n under best-response iteration on the static adoption map. The index n counts steps of this iteration, a graphical device used to locate the static fixed points of $q^* = F(\Delta\Pi(q^*))$ and assess their stability; it does not correspond to calendar time, which is reserved for the dynamic model of Section 8. Fixed points lie at intersections with the 45-degree line. Three intersections obtain: a stable low-adoption fixed point $q_L^* \approx 0.04$ and a stable high-adoption fixed point $q_H^* \approx 0.96$, separated by an unstable interior fixed point $q^\dagger \approx 0.66$ at which the slope of $F \circ \Delta\Pi$ exceeds unity. Best-response iteration converges to whichever stable fixed point lies in the same basin of attraction as the initial iterate: the wide left basin captures industries that begin from any moderate level of platform adoption, while only initial values above $q^\dagger$ tip the industry into the high-adoption fixed point. 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 $[0,1]$.

$\mathscr{W}_k(q_L^*)$ 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 external wedge

remains positive at any interior equilibrium.

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

6.2 Adoption Fragmentation

When $\delta < 0$, $\Delta\Pi(0) > 0$, and $\Delta\Pi(1) < 0$, neither corner is a fixed point. 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 clean in this regime. By Theorem 1, the external wedge is positive; the fragmentation equilibrium over-adopts whenever the external wedge dominates the classical term Ψ, which holds across the strategic-substitutes calibrations we examine, and in particular whenever $\frac{d}{dq}\mathrm{Var}(\bar{x}) \geq 0$, so that $\Psi \leq 0$ by Proposition 3. Strategic substitution then limits the level of adoption but does not eliminate the externality. The tipping and fragmentation regimes therefore differ in the level of q^* and in whether the classical and shared-processing externalities reinforce or oppose, but the external wedge attributable to shared processing is excessive in both.

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 external-wedge result of Theorem 1 establishes that a complete-welfare planner prefers less adoption 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 $J \geq 1$ platforms. Platform j generates a processing shock $\xi_j \sim \mathcal{N}(0, \sigma_\xi^2)$ and idiosyncratic noise $u_f^{A,j}$ with variance σ_A^2, so that an adopter of platform j observes

$$\tilde{s}_f^{A,j} = \theta + \xi_j + u_f^{A,j}.$$

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

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 $\xi_j = \xi_{j'}$ almost surely. Intermediate values capture partial overlap in inputs or methods.[14]

For tractability, vendor shares $\{q_j\}$ 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.

[14] 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:

$$\bar{x} = \mathbb{E}[\bar{x}] + K_\theta(q)\,\theta + \sum_{j=1}^{J} K_{\xi,j}\,\xi_j, \qquad K_{\xi,j} = q_j\,k^A.$$

Taking variances and using the correlation structure,

$$\mathrm{Var}\Big(\sum_j K_{\xi,j}\,\xi_j\Big) = (k^A)^2\,\sigma_\xi^2\Big[\sum_j q_j^2 + 2\rho\sum_{j<j'} q_j q_{j'}\Big].$$

Under symmetric adoption, $q_j = q/J$, this expression reduces to

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

Then

$$\mathrm{Var}\Big(\sum_j K_{\xi,j}\,\xi_j\Big) = K_\xi(q)^2\,\sigma_\xi^2 \cdot H(J,\rho),$$

where $K_\xi(q) = q\,k^A$. The index $H(J,\rho)$ lies in $[1/J, 1]$, attaining $1/J$ when platforms are independent and 1 when they share a common core.

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

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

$$\frac{d\mathscr{W}_1}{dq} - \frac{d\mathscr{W}_3}{dq} = \gamma_3\,H(J,\rho)\,\frac{d}{dq}\big[K_\xi(q)^2\,\sigma_\xi^2\big] > 0.$$

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

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

For all $\delta \in \mathscr{D}$, $J \geq 1$, and $\rho \in [0,1]$, the external wedge remains positive. Its magnitude is proportional to $H(J,\rho)$ and therefore decreases in J 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 $H(J,\rho)$. Under $\mathscr{W}_4$, the same index equals expected correlation across adopters: with probability $1/J$, two adopters share a platform and their errors are perfectly correlated; with probability $1-1/J$, 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 $\rho = 0$, $H(J,0) = 1/J$. The external-wedge magnitude under either welfare specification falls linearly in J.*

Corollary 2 (Common-core platforms)**.** *When $\rho = 1$, $H(J,1) = 1$ for all J. The external-wedge magnitude is invariant to vendor count.*

When platforms are independent, increasing J reduces the externality by lowering the probability that adopters share the same error. When platforms share a common core, increasing J 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.

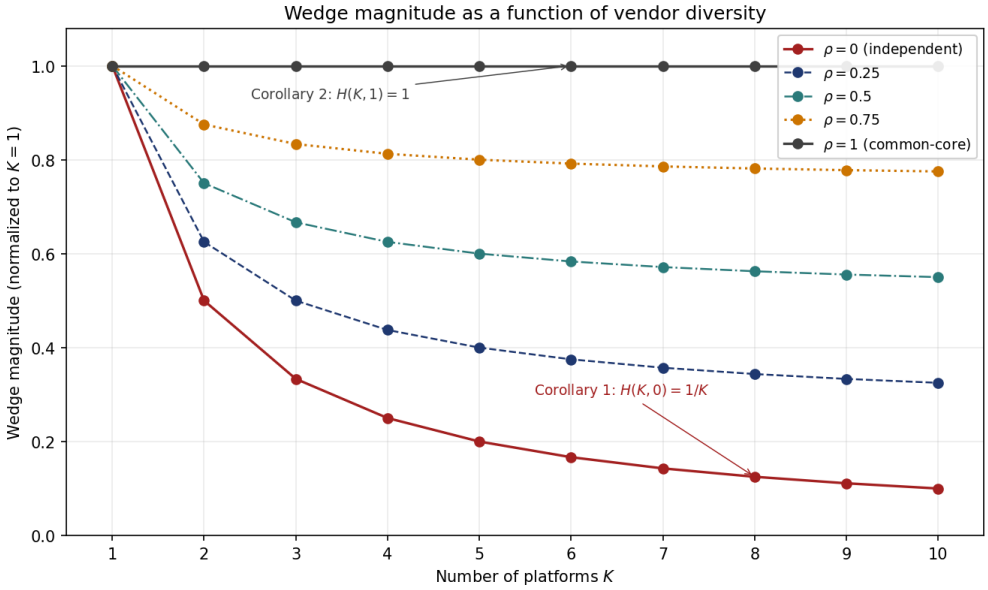
A similar pattern appears in adjacent algorithmic markets. Bommasani et al. (2022) document that fine-tuned variants of a common foundation model exhibit limited independent variation relative to structurally distinct models. Tong et al. (2023) identify systematic failure modes of CLIP's text encoder 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. 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 it 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 J 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 external-wedge magnitude as a function of the number of platforms J, normalized to the value at $J = 1$. By Proposition 5, this normalized magnitude equals the effective concentration index $H(J,\rho) = 1/J + (1 - 1/J)\rho$ 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 $1/J$ (Corollary 1). At $\rho = 1$ (gray), platforms share a common processing core and the magnitude is invariant in J (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 $H(J,\rho)$. The marginal external 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 complete-welfare 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 $q_j = q/J$. In practice, shares are uneven. With general shares, the variance decomposition yields a normalized Herfindahl index

$$H(\{q_j\},\rho) = \frac{\sum_j q_j^2 + 2\rho\sum_{j<j'} q_j q_{j'}}{q^2},$$

which reduces to $H(J,\rho)$ under symmetry. The external 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 J.

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.

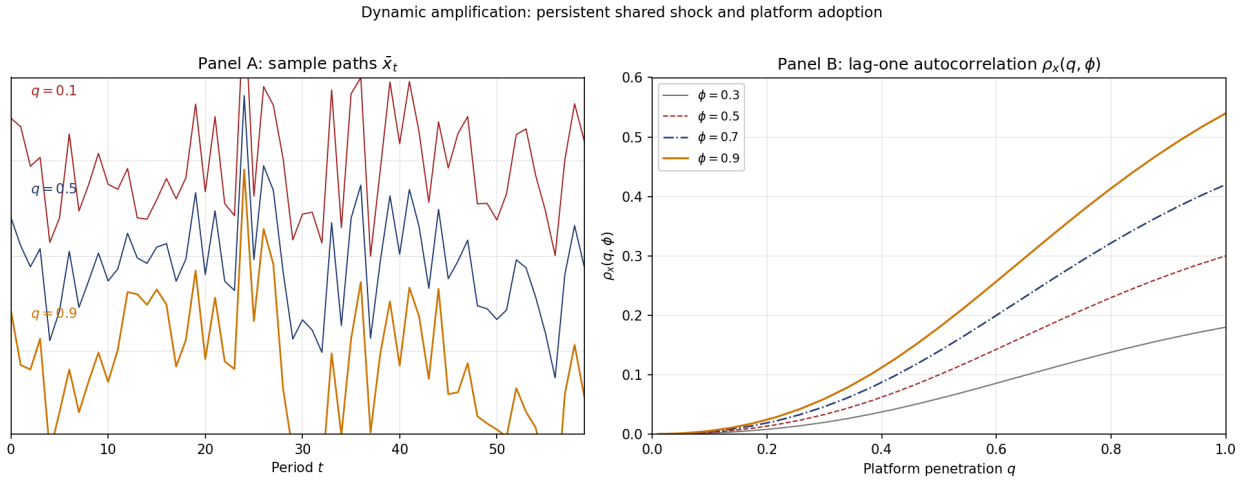
Figure 5 illustrates the mechanism. Panel A plots a single realization of $\bar{x}_t$ across three penetration levels under identical shocks: at low adoption the path resembles independent noise, while at high adoption it exhibits longer same-sign excursions. Panel B plots the population autocorrelation $\rho_x(q,\phi)$, 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.[15]

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 external-wedge result established in Section 5.

[15] 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 external $\mathscr{W}_4$ 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 benefits of vendor diversification depend not only on the number of vendors, but also on how similar their underlying forecasting technologies are. A market may therefore appear diversified at the interface layer while remaining concentrated in the data and modeling layer that determines forecast co-movement. The relevant object is the effective concentration index $H(J,\rho)$, which combines downstream adoption concentration with correlation across vendors. Diversification is substantive when $H \approx 1/J$ and cosmetic when $H \approx 1$.

This section illustrates that distinction using restaurant site-selection analytics.[16] The goal is not to identify the mechanism or to point-estimate H. Rather, the exercise shows how the primitives of the model can be disciplined using public disclosures, SEC filings, vendor case studies, and evidence on upstream data overlap. More broadly, it illustrates a central implication of the theory: competition among front-end vendors need not imply meaningful diversification in the forecasting technologies used by adopters.

The calibration has two components. The first is downstream vendor-share concentration, $\sum_j \omega_j^2$, which measures the probability that two adopters use the same vendor. The second is structural correlation across vendors conditional on using different platforms, ρ, which captures common exposure to shared data inputs, modeling 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.[17] Other firms participate in the broader location-intelligence ecosystem but appear to play different roles. Esri's ArcGIS Business Analyst functions primarily as geospatial infrastructure rather than as a chain-facing site-selection platform, while smaller specialists tend to serve single-unit operators, local businesses, or narrower use cases outside the multi-unit chain market considered here.

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

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

The model-relevant feature is that these vendors draw on a partially shared upstream data layer. INRIX, a major supplier of roadway-traffic data, reports that several leading site-selection providers license its data and incorporate those feeds into proprietary forecasting systems.[18] The disclosure identifies Buxton, eSite Analytics, IdealSpot, SiteZeus, and Tango among its clients. Because this evidence comes from the upstream vendor itself, we treat it as evidence of overlap rather than as a precise measure of dependence. Appendix F corroborates the pattern where possible using platform case studies and SEC filings.

The key implication is that six vendor brands need not correspond to six independent forecasting technologies. Competing platforms may differ in interface, presentation, and proprietary modeling choices while relying on overlapping upstream inputs. Vendor competition at the front end therefore does not necessarily imply diversification in the forecasting layer that determines cross-firm error correlation.

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.[19] 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.

Let ω_j denote vendor j's share *among adopters* (so $\sum_j \omega_j = 1$), as distinct from the aggregate adoption rate q. Table A3 in Appendix F reports the adjusted-share estimates. The implied

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

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

Herfindahl component is

$$\sum_j \omega_j^2 \approx 0.628^2 + 0.141^2 + 0.097^2 + 0.077^2 + 0.055^2 + 0.002^2 \approx 0.433.$$

The equal-share benchmark, $1/J \approx 0.17$, 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 vendors. This margin captures whether separately owned and branded platforms generate genuinely 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 define trade areas, measure traffic, characterize demand, and evaluate candidate sites.

Four categories of upstream inputs are especially important. First, roadway-traffic data are concentrated among a small number of providers, including INRIX. Second, point-of-interest data used for competitor mapping and trade-area construction are sourced from a limited set of vendors, including SafeGraph and Foursquare. Third, mobility and foot-traffic measures rely on a relatively small group of panel providers. Fourth, demographic and psychographic segmentation products, such as Experian Mosaic, Claritas PRIZM, and Spectra, are themselves constructed from overlapping consumer-data sources. The result is that upstream inputs appear substantially less differentiated than downstream vendor brands.

Placer.ai is particularly informative because it occupies multiple layers of the ecosystem. In

addition to serving restaurant chains directly, it also participates in the broader data marketplace. Life360's 10-K filings disclose a data-licensing relationship with Placer.ai under which Life360 contributes panel data used in analytics products.[20] This example illustrates how data used in forecasting systems may originate from shared upstream sources and propagate through multiple downstream applications.

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

The available evidence does not point-identify ρ, but it makes the independence benchmark difficult to sustain. A value near zero would require not only separate vendor ownership, but also largely independent data inputs and processing pipelines. The observed industry structure points in the opposite direction. We therefore use $\rho \in [0.5, 0.7]$ as a disciplined calibration range. The upper end corresponds to substantial propagation of upstream shocks across vendors, while the lower end allows meaningful differentiation through proprietary modeling choices, firm-specific data, and platform-specific partnerships. The range should be interpreted as an order-of-magnitude discipline imposed by the observed technology stack rather than as a point estimate.

9.4 Discussion

Combining vendor concentration with upstream-input overlap gives

$$H = \sum_j \omega_j^2 + \rho \left(1 - \sum_j \omega_j^2\right).$$

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 $H \approx 0.77$.

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

The implied effective concentration is much closer to the cosmetic-diversity bound, $H = 1$, than to the substantive-diversity benchmark, $1/J \approx 0.17$. 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 external 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 external 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 external $\mathscr{W}_3$ 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 structurally 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.

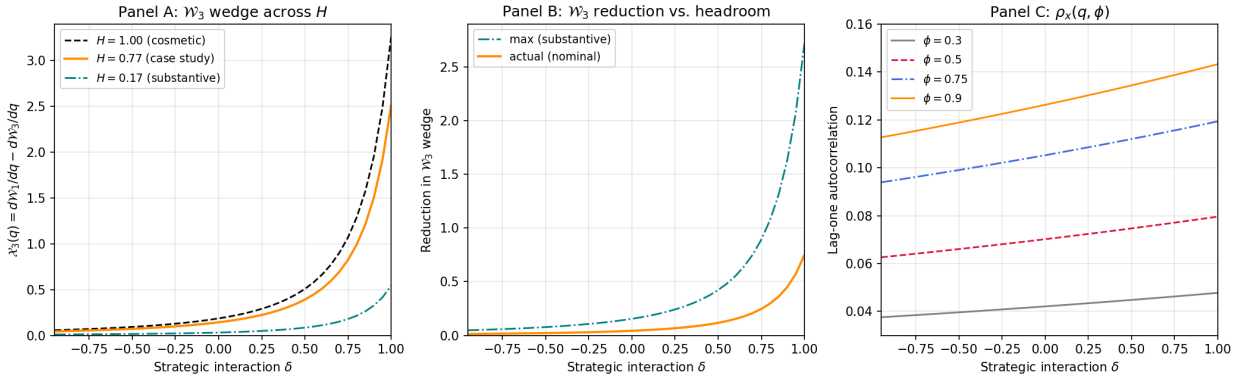
The external $\mathscr{W}_4$ wedge is omitted from Panels A and B because, by Theorem 1(ii), it is invariant in δ at any fixed adoption rate. At $H = 0.77$ and $q = 0.6$, the wedge is

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

in units of γ_4, and would appear as a horizontal line. Panel C instead reports the dynamic im-

plication 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.

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 external $\mathcal{W}_3$ wedge $\mathcal{X}_3$ (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 external $\mathcal{W}_4$ 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 external 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.

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 solely 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 exercise isolates the value of upstream diversification. The relevant margin is not the number of front-end vendors, but their reliance on common traffic data, point-of-interest taxonomies, mobility panels, and segmentation products. Policy instruments therefore operate primarily through the upstream layer, including data-portability rules, upstream-input diversification requirements, and scrutiny of vertical integration in the traffic, mobility-panel, and demographic-segmentation markets. This interpretation connects naturally to the platform-economics literature on multi-homing and interoperability (Armstrong 2006, Biglaiser et al. 2022). The key distinction is that the relevant margin here is the independence of upstream processing rather than the number of platforms a firm joins. Multi-homing across vendors that share an input layer leaves ρ, and therefore H, close to its single-platform value.

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 more important than the downstream-share margin, but neither margin alone eliminates the externality.

Heterogeneity across the two wedges. Although the external $\mathscr{W}_3$ and $\mathscr{W}_4$ wedges both scale linearly with H, they differ in their dependence on the strategic-interaction parameter δ. By Theorem 1(ii), the decided-about-agent wedge $\mathscr{X}_4$ is invariant in δ at any fixed adoption rate. The aggregate-volatility wedge $\mathscr{X}_3$ 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 $\mathscr{X}_4$ 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

Firms increasingly rely on shared intelligence platforms to guide market-entry decisions. While these systems can improve forecast accuracy, they also route decisions through a common processing layer, creating correlated errors unrelated to fundamentals. We develop a model that separates this shared-processing channel from informative common signals and strategic interaction. The central result is that shared processing generates external costs that firms do not internalize, so a planner who accounts for those costs chooses less adoption than one who values only firm profits. This external wedge arises regardless of whether firms' entry decisions are strategic complements or substitutes and operates alongside the classical entry externality.

Our analysis also highlights a limitation of vendor competition as a solution to correlated decision-making. Multiple vendors reduce the distortion only to the extent that their underlying data and models are independent. When providers rely on overlapping inputs and forecasting architectures, competition among vendor brands need not translate into meaningful diversification. A descriptive case study of restaurant site-selection analytics suggests that substantial exposure to common forecasting shocks may persist even in the presence of multiple providers.

Several directions remain open. The most immediate is empirical. We derive an identifying restriction (i.e., a discrete movement in aggregate entry that coincides with a change in the shared forecasting technology, occurs without any corresponding movement in fundamentals, and grows with platform penetration), but stop short of taking it to data. Doing so requires three ingredients that are increasingly available: platform-side records of when forecasting systems are updated, controls for fundamentals fine enough to isolate the residual movement in entry, and a measure of adoption at the level at which entry is observed. The co-client comparison flagged in the motivating analysis offers a complementary cross-sectional design: with a store census and disclosed vendor relationships, alignment among chains known to share a vendor can be set against alignment among chains that do not, holding the product segment fixed. Recent evidence that co-movement in firm outcomes can be traced to shared data linkages (Bian et al. 2024) suggests that such residual alignment is detectable in practice; the task is to tie it to identified changes in the forecasting layer

rather than to inferred linkages.

A second direction endogenizes the market structure that we take as given. We hold the vendors' shares and the upstream data stack fixed and discipline the correlation across vendors from observed industry features; a fuller treatment would let vendors compete in price and in forecast quality, and let the upstream suppliers of traffic, point-of-interest, and mobility data (identified in the case study as common inputs) exercise market power of their own. The economics of data supplies the tools, since data is non-rival and its shared use generates a tension between private and social incentives to pool it (Jones and Tonetti 2020), while markets for data are impeded by the difficulty of pricing an input whose value a buyer cannot assess before acquiring it (Bergemann and Bonatti 2019). The open question is whether competition among forecasting vendors raises or lowers the correlation in their errors. If differentiation is costly while shared upstream inputs are cheap, competition may select precisely the cosmetic diversity that leaves the externality intact, in which case the distinction between substantive and cosmetic diversity becomes an equilibrium outcome rather than an input to the calibration.

The mechanism is also a microfoundation for a concern that financial regulators have recently begun to voice. Several have warned that heavy reliance on a small number of vendor-supplied models or common datasets may lead market participants to take increasingly correlated positions, with consequences for the stability of the financial system (Bank of England Financial Policy Committee 2025, Board 2024, Danielsson and Uthemann 2023). Retail entry is a useful laboratory for this concern because it isolates correlated errors from strategic learning, but the externality is not specific to retail. A version of the model with prices and balance-sheet feedback (e.g., in lending, where many originators screen the same borrowers through a shared scoring layer, or in trading, where a common forecast interacts with fire-sale dynamics) would let the correlated-error channel compound with the leverage and liquidity channels that this literature emphasizes. Carrying the location-level screening externality into a setting in which quantities feed back into prices is the central modeling step.

Finally, the treatment of the party being decided about can be enriched along two margins. Our screening environment is static: each location is evaluated in a single encounter, and the social loss is the correlation of the evaluators' errors. A location is in fact screened repeatedly and holds an outside option, and the harm that matters may be systemic exclusion (i.e., being passed over by every firm that relies on the common layer) rather than any single correlated mistake. The literature

on algorithmic monoculture has begun to measure exactly this form of loss in selection problems (Kleinberg and Raghavan 2021, Peng and Garg 2024), and connecting our entry externality to that work would sharpen the interpretation of the penalty we place on correlated evaluation. The second margin is dynamic. We show that shared processing turns cross-sectional correlation into persistence over time, and that any planner who dislikes such persistence would assign a further cost to adoption; what remains is to specify that planner's problem over sequences of outcomes, pricing persistence directly and determining how much the static result understates the full dynamic cost.

More broadly, the economic consequences of AI adoption depend not only on forecast accuracy, but also on the degree of commonality in the systems producing those forecasts. As firms increasingly rely on shared prediction technologies, understanding how common processing layers shape aggregate outcomes may become as important as understanding the information they provide.

References

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

Armstrong M (2006) Competition in two-sided markets. *The RAND Journal of Economics* 37(3):668–691.

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* .

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, Bonatti A (2019) Markets for information: An introduction. *Annual Review of Economics* 11(1):85–107.

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.

Biglaiser G, Crémer J, Veiga A (2022) Should i stay or should i go? migrating away from an incumbent platform. *The RAND Journal of Economics* 53(3):453–483.

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.

Board FS (2024) The financial stability implications of artificial intelligence. *Basel: Financial Stability Board* .

Bommasani R, Creel KA, Kumar A, Jurafsky D, Liang PS (2022) Picking on the same person: Does algorithmic monoculture lead to outcome homogenization? *Advances in Neural Information Processing Systems* 35:3663–3678.

Bommasani R, Hudson DA, Adeli E, Altman R, Arora S, von Arx S, Bernstein MS, Bohg J, Bosselut A, Brunskill E, et al. (2021) On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258* .

Bresnahan TF, Reiss PC (1991) Entry and competition in concentrated markets. *Journal of Political Economy* 99(5):977–1009.

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

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

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

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

Carlsson H, van Damme E (1993) Global games and equilibrium selection. *Econometrica* 61(5):989–1018.

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

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.

Danielsson J, Uthemann A (2023) On the use of artificial intelligence in financial regulations and the impact on financial stability. *arXiv preprint arXiv:2310.11293* .

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

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

Farboodi M, Koh AJ, Xia A (2026) Data-driven automation. Working Paper 35320, National Bureau of Economic Research, Cambridge, MA, nBER Working Paper No. 35320.

Farboodi M, Veldkamp L (2021) A model of the data economy. Working Paper 28427, National Bureau of Economic Research, Cambridge, MA, nBER Working Paper No. 28427.

Farboodi M, Veldkamp L (2023) Data and markets. *Annual Review of Economics* 15:23–40.

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.

Jia P (2008) What happens when wal-mart comes to town: An empirical analysis of the discount retailing industry. *Econometrica* 76(6):1263–1316.

Jones CI, Tonetti C (2020) Nonrivalry and the economics of data. *American Economic Review* 110(9):2819–2858.

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* .

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.

Mankiw NG, Whinston MD (1986) Free entry and social inefficiency. *The RAND Journal of Economics* 17(1):48–58.

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.

Spence M (1976) Product selection, fixed costs, and monopolistic competition. *The Review of Economic Studies* 43(2):217–235.

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 Robustness to the Linear Entry Rule **8**

E A Screening Foundation for the Decided-About-Agent Wedge **9**

F Case-Study Sources and Sensitivity **10**

G Proofs **11**

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 external wedge, $d\mathscr{W}_1/dq - d\mathscr{W}_k/dq$, 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 to outside parties. Multiple equilibria do not by themselves indicate a market failure, nor does uniqueness rule one out; the external 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 external $\mathscr{W}_4$ 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 $H(J,\rho)$, 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 carries an uninternalized external cost. 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 adoption carries an external wedge, so faster convergence accelerates the accumulation of an uninternalized social cost. Switching frictions impose private costs but may also moderate the speed at which the industry reaches states with high uninternalized social cost.

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 equi-

libria 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) The discounted cumulative external $\mathscr{W}_4$ wedge

$$\mathscr{L}(\lambda; \beta) \equiv \sum_{t=0}^{\infty} \beta^t \, \mathscr{X}_4(q_t), \qquad \mathscr{X}_4(q) = \frac{2\gamma_4 \, q \, \sigma_\xi^2}{V_A},$$

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 $t \geq 1$.

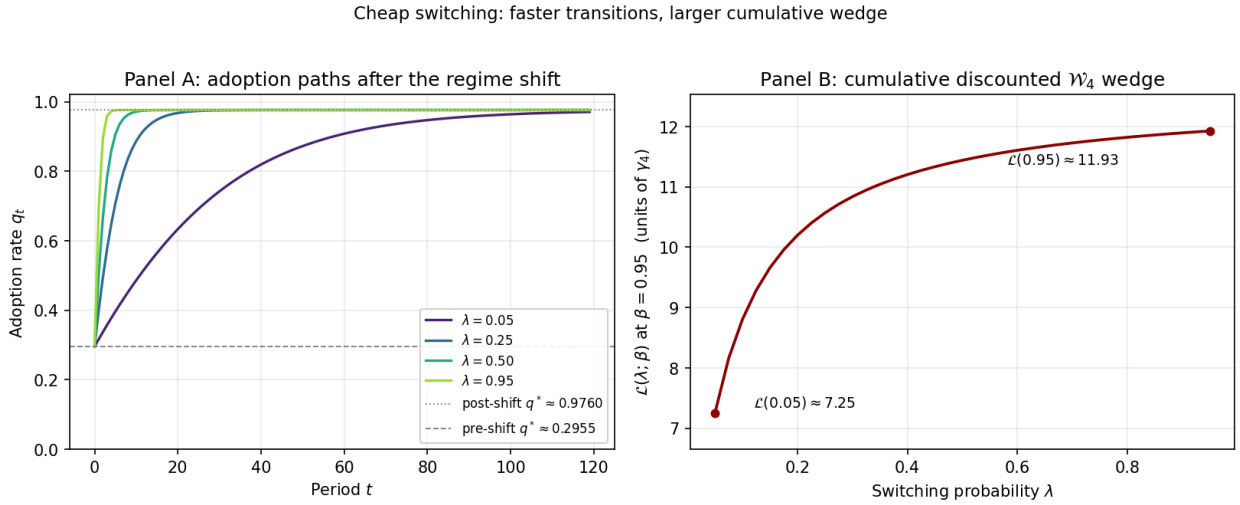
The result highlights a tension between static and dynamic efficiency. Switching costs do not affect the static adoption equilibrium, but they govern how quickly the economy approaches it. When adoption carries an external wedge, faster convergence increases cumulative exposure to the externality. In 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 external $\mathscr{W}_4$ wedge is increasing in q, earlier movement to higher q_t raises the discounted external cost. 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 external $\mathscr{W}_4$ wedge in q. Under $\mathscr{W}_3$, the external 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 the external 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 baseline treats λ as an exogenous technological parameter. If instead λ is chosen by firms facing a cost of re-evaluation that declines as analytics become cheaper, then cheaper AI raises λ endogenously. The comparative statics in Proposition 8 then map declines in the cost of switching into increases in the discounted external cost. The qualitative conclusion is unchanged: cheaper analytics, by raising the equilibrium re-evaluation rate, accelerate the accumulation of the external welfare cost.

Figure A1: Cheap switching steepens the adoption transition and raises cumulative external cost.



Notes: All panels use the calibration $\sigma_\theta^2 = 1$, $\sigma_\xi^2 = 0.6$, $\sigma_A^2 = 0.2$, $\sigma_C^2 = 0.6$, $\delta = 0.5$, with logistic switching-cost distribution F and a common interior stable equilibrium q^*, starting from $q_0 = 0.05$. Panel A plots the adoption transition q_t for three values of the re-evaluation rate λ; higher λ steepens the path toward q^* but does not change its limit. Panel B plots the per-period external $\mathcal{W}_4$ wedge $\mathcal{X}_4(q_t) = 2\gamma_4 q_t \sigma_\xi^2 / V_A$ along each transition. Panel C plots the discounted cumulative wedge $\mathcal{L}(\lambda;\beta)$ as a function of λ for discount factor $\beta = 0.9$; it is strictly increasing in λ, consistent with Proposition 8(ii).

D Robustness to the Linear Entry Rule

The closed-form equilibrium relies on the linear entry probability of Section 3. This appendix verifies that the external wedge retains its sign and magnitude under a nonlinear entry rule. We replace the linear entry probability with a logistic rule $P_f(\text{enter} \mid s_f) = \Lambda\big(\mathbb{E}_f[\theta] + \delta \mathbb{E}_f[\bar{x}]\big)$, where $\Lambda(z) = (1+e^{-z})^{-1}$, and recompute the equilibrium numerically by iterating the aggregate fixed point to convergence at each adoption rate.

Table A1 reports the external $\mathscr{W}_4$ wedge under the linear and logistic rules across a grid of calibrations. The two rules agree to within one percent at every calibration, confirming that the closed-form magnitude is not an artifact of linearity. Table A2 reports the corresponding values for the external $\mathscr{W}_3$ wedge. The $\mathscr{W}_3$ magnitude inherits the equilibrium loadings and is therefore more sensitive to the curvature of the choice rule, with deviations of up to nine percent in the most strongly interacting calibrations, but the sign and monotonicity are preserved throughout.

Table A1: External $\mathscr{W}_4$ wedge under linear and logistic entry rules

δ	σ_ξ^2	q	Linear	Logistic	Relative diff. (%)
-0.5	0.3	0.5	0.1304	0.1305	0.08
-0.2	0.5	0.5	0.2000	0.2001	0.05
0.0	0.5	0.5	0.2000	0.2000	0.00
0.3	0.5	0.5	0.2000	0.1999	0.05
0.7	0.6	0.6	0.2483	0.2480	0.12

Notes: The external $\mathscr{W}_4$ wedge $\mathscr{X}_4(q) = 2\gamma_4 q \sigma_\xi^2 / V_A$ (in units of γ_4) computed under the closed-form linear entry rule and under a logistic entry rule solved numerically. Remaining primitives are fixed at $\sigma_\theta^2 = 1$, $\sigma_A^2 = 0.2$, $\sigma_C^2 = 0.6$. The relative difference is $|\text{logistic} - \text{linear}|/\text{linear}$.

Table A2: External $\mathscr{W}_3$ wedge under linear and logistic entry rules

δ	σ_ξ^2	q	Linear	Logistic	Relative diff. (%)
-0.5	0.3	0.5	0.0271	0.0268	1.1
-0.2	0.5	0.5	0.0418	0.0414	1.0
0.0	0.5	0.5	0.0444	0.0444	0.0
0.3	0.5	0.5	0.0492	0.0500	1.6
0.7	0.6	0.6	0.0921	0.1004	9.0

Notes: The external $\mathscr{W}_3$ wedge $\mathscr{X}_3(q) = \gamma_3 \frac{d}{dq}[K_\xi(q)^2 \sigma_\xi^2]$ (in units of γ_3) computed under the closed-form linear entry rule and under a logistic entry rule solved numerically. Remaining primitives are fixed at $\sigma_\theta^2 = 1$, $\sigma_A^2 = 0.2$, $\sigma_C^2 = 0.6$. The relative difference is $|\text{logistic} - \text{linear}|/\text{linear}$.

E A Screening Foundation for the Decided-About-Agent Wedge

This appendix embeds the $\mathscr{W}_4$ penalty in an explicit screening environment, deriving the exposure term as a second moment of that environment rather than positing it. Alongside the entry game, consider a unit continuum of ex-ante identical candidate locations. Each location m is screened by $n \geq 2$ chains drawn independently at random from the chain population. A drawn chain evaluates location m using its equilibrium forecasting technology applied to a location-specific realization of the shocks (θ_m, ξ_m, u). A platform adopter's forecast error on location m contains the shared component ξ_m; a custom chain's error does not.

$$\mathbb{E}\left[\frac{1}{\binom{n}{2}}\sum_{i<i'}\mathrm{Cov}\big(e_i, e_{i'} \mid m\big)\right] = q^2\,\frac{\sigma_\xi^2}{V_A}, \tag{3}$$

where e_i denotes the normalized forecast error of the i-th evaluator of location m, and the expectation is over the random draw of evaluators. The pairwise covariance equals σ_ξ^2/V_A when both evaluators are adopters, which occurs with probability q^2, and zero otherwise.

A planner who values the option of an independent second opinion penalizes this covariance, because correlated errors eliminate the diversification that independent screening would provide. Writing the social cost as γ_4 times the exposure in (3) yields the unified welfare criterion

$$\mathscr{W}_4(q) = \mathscr{W}_1(q) - \gamma_4\, q^2\,\frac{\sigma_\xi^2}{V_A}, \tag{4}$$

which coincides with the criterion used in the main text.

Proposition 9 (Screening wedge)**.** *Under the screening environment above and the welfare criterion* (4)*, the external wedge is*

$$\mathscr{X}_4(q) = \frac{d\mathscr{W}_1}{dq} - \frac{d\mathscr{W}_4}{dq} = \frac{2\gamma_4\, q\,\sigma_\xi^2}{V_A} > 0 \quad \textit{for all } q \in (0,1],$$

independent of δ at any fixed q, and proportional to the residual variance share σ_ξ^2/V_A.

Remark 8 (Why strategic interaction does not enter)**.** The screening wedge runs through the residual covariance of forecast errors in (3), which contains no equilibrium loading. Strategic interaction would re-enter only if screening outcomes across locations were linked, for example through a shared capacity constraint or a common downstream market. Such a linkage would reintroduce δ through the same channel as the aggregate-volatility wedge $\mathscr{X}_3$. The δ-invariance of $\mathscr{X}_4$ is therefore diagnostic of the screening channel: it isolates the part of the externality that operates purely through correlated evaluation rather than through coordinated entry.

Remark 9 (Level versus marginal exposure)**.** The exposure in (3) is a level; the wedge in Proposition 9 is its marginal counterpart. Both are strictly increasing in adoption, so the level and marginal statements deliver the same qualitative conclusion. The marginal object is the relevant one for the planner's first-order condition, while the level is the relevant one for cumulative welfare comparisons such as those in Proposition 8.

F Case-Study Sources and Sensitivity

This appendix documents the sources used to construct the vendor-share estimates in Section 9 and reports the sensitivity of the effective concentration index to alternative share and correlation assumptions.

The vendor-share estimates draw on the following public sources: reported annual recurring revenue and customer counts from company press releases and trade press; third-party revenue estimates from Growjo, Owler, Latka, and Craft; employee counts and funding disclosures from Crunchbase and PitchBook; SEC filings for publicly traded parents and licensors (including Life360); and vendor case studies and client lists published by upstream data providers (including INRIX). Shares are computed after a retail-chain exposure adjustment that scales each vendor's total revenue to the portion plausibly attributable to restaurant-chain site-selection analytics.

Table A3: Estimated vendor shares among adopters

Vendor	Adjusted revenue proxy	Share among adopters ω_j
Placer.ai	High	0.628
Tango Analytics	Medium	0.141
Buxton	Medium	0.097
eSite Analytics	Low	0.077
SiteZeus	Low	0.055
IdealSpot	Low	0.002
Herfindahl component $\sum_j \omega_j^2$		0.433

Notes: Shares among adopters are revenue proxies after the retail-chain exposure adjustment described in Section 9, normalized to sum to one. The adjusted-revenue column reports a qualitative tier rather than a point estimate. The Herfindahl component is $\sum_j \omega_j^2$, the probability that two randomly drawn adopters use the same vendor.

Table A4: Sensitivity of the effective concentration index H

	$\rho = 0.4$	$\rho = 0.5$	$\rho = 0.6$	$\rho = 0.7$	$\rho = 0.8$
Baseline shares ($\sum_j \omega_j^2 = 0.433$)	0.660	0.717	0.773	0.830	0.887
Equal shares ($\sum_j \omega_j^2 = 0.167$)	0.500	0.583	0.667	0.750	0.833
More concentrated ($\sum_j \omega_j^2 = 0.55$)	0.730	0.775	0.820	0.865	0.910

Notes: Each cell reports $H = \sum_j \omega_j^2 + \rho(1 - \sum_j \omega_j^2)$ for the indicated share concentration and cross-vendor correlation. The baseline row uses the estimated shares in Table A3. The equal-shares row imposes $\omega_j = 1/6$. The more-concentrated row stresses a higher downstream share for the leading vendor. At the central calibration $\rho = 0.6$ and baseline shares, $H \approx 0.77$.

G Proofs

Proof of Proposition 1. Substituting the conjectured aggregate $\bar{x} = K_0 + K_\theta \theta + K_\xi \xi$ into each type's best response and matching coefficients yields a linear system in (k^A, k^C). A platform adopter's posterior mean of $\theta + \delta\bar{x}$ loads on its signal through α_θ^A on θ and through the anticipated aggregate response to both θ and ξ; a custom chain loads through α_θ^C on θ and through the anticipated aggregate. Collecting terms gives

$$k^A = \alpha_\theta^A + \delta\,\alpha_\theta^A K_\theta + \delta\,\alpha_\xi^A K_\xi, \qquad k^C = \alpha_\theta^C + \delta\,\alpha_\theta^C K_\theta,$$

with $K_\theta = qk^A + (1-q)k^C$ and $K_\xi = qk^A$. Solving the two equations simultaneously and collecting terms in the common denominator gives the stated expressions, with $\Delta(q)$ as in (2). For existence and uniqueness it suffices that $\Delta(q) > 0$ on $[0,1]$. The two endpoint conditions $\Delta(0) > 0$ and $\Delta(1) > 0$ are guaranteed by the two bounds defining $\mathscr{D}$. Since $\Delta(q)$ is quadratic in q with constant second derivative $\Delta''(q) = -2\delta^2\alpha_\theta^C\alpha_\xi^A \le 0$, it is concave, so positivity at the endpoints implies positivity throughout $[0,1]$, because a concave function lies above the chord joining its endpoint values. The system is then nonsingular and the linear equilibrium is unique. □

Proof of Proposition 2. Expected equilibrium profit of a type-t chain equals $\frac{1}{2}\mathbb{E}[(x_f^t)^2] = \frac{1}{2}\big(K_0^2 + (k^t)^2 V_t\big)$, with V_A, V_C the respective signal variances, $K_0 = 1/(2(1-\delta))$ the common mean entry probability, and $\mathbb{E}[\pi_f] = \frac{1}{2}\mathbb{E}[x_f^2]$ by the projection identity, exactly as in the proof of Proposition 3. The intercept K_0^2 is common to both types, so it cancels in the difference:

$$\Delta\Pi(q) = \tfrac{1}{2}\big[(k^A(q))^2 V_A - (k^C(q))^2 V_C\big].$$

At $\delta = 0$ the loadings are constants independent of q, so $\Delta\Pi$ is constant. For $\delta \neq 0$, differentiating and substituting the closed-form loadings of Proposition 1 gives

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

where Γ collects strictly positive terms in the admissible region. The sign of $d\Delta\Pi/dq$ therefore equals the sign of δ, establishing strict monotonicity. Positivity of Γ on $\mathscr{D}$ reduces to a sign condition on a compact set, verified over 200,000 parameter draws in the replication package. □

Proof of Proposition 3. Chain welfare is $\mathscr{W}_1(q) = q\pi^A(q) + (1-q)\pi^C(q)$. Expected profit of a type-t chain at adoption rate q is $\frac{1}{2}(K_0^2 + (k^t(q))^2 V_t) + \frac{\delta}{2}\mathrm{Var}(\bar{x})$, where the second term is the strategic surplus that each chain receives from the aggregate but does not control. Differentiating $\mathscr{W}_1$ in q, the direct effect of reallocating a chain from custom to platform research is $\Delta\Pi(q)$, while the indirect effect operates through the shared strategic surplus and equals $\frac{\delta}{2}\frac{d}{dq}\mathrm{Var}(\bar{x}) \equiv \Psi(q)$. Hence $d\mathscr{W}_1/dq = \Delta\Pi(q) + \Psi(q)$. The private equilibrium satisfies $F^{-1}(q^*) = \Delta\Pi(q^*)$, so it coincides with the chain-welfare optimum, defined by $F^{-1}(q_1^{soc}) = d\mathscr{W}_1/dq$, if and only if $\Psi \equiv 0$, i.e. $\delta = 0$. The sign of $\Psi(q)$ equals the sign of $\delta \cdot \frac{d}{dq}\mathrm{Var}(\bar{x})$, which is calibration-dependent because $\mathrm{Var}(\bar{x})$ need not be monotone in q. □

Proof of Theorem 1. For each external criterion $\mathscr{W}_k = \mathscr{W}_1 - \gamma_k P_k(q)$, $k \in \{3,4\}$, with penalty $P_3(q) = K_\xi(q)^2\sigma_\xi^2$ and $P_4(q) = q^2\sigma_\xi^2/V_A$, the external wedge is

$$\mathscr{X}_k(q) = \frac{d\mathscr{W}_1}{dq} - \frac{d\mathscr{W}_k}{dq} = \gamma_k P_k'(q),$$

since the $\mathscr{W}_1$ contribution cancels. For part (i), $P_3'(q) = \frac{d}{dq}[K_\xi(q)^2\sigma_\xi^2]$. Because $K_\xi(q) = qk^A(q)$ with $k^A(q) > 0$ and $K_\xi(0) = 0$, the function $K_\xi(q)^2$ is strictly increasing on $(0,1]$, so $P_3'(q) > 0$; its magnitude depends on δ through $k^A(q)$, but its sign does not. For part (ii), $P_4'(q) = 2q\sigma_\xi^2/V_A > 0$, which contains no equilibrium loading and is therefore independent of δ at fixed q. In both cases $\mathscr{X}_k(q) > 0$, so $d\mathscr{W}_k/dq < d\mathscr{W}_1/dq$ at every q, and the complete-welfare optimum defined by $F^{-1}(q_k^{soc}) = d\mathscr{W}_k/dq$ satisfies $q_k^{soc} < q_1^{soc}$. The decomposition relative to the private equilibrium follows from $d\mathscr{W}_k/dq = \Delta\Pi + \Psi - \mathscr{X}_k$, so $\Delta\Pi - d\mathscr{W}_k/dq = \mathscr{X}_k - \Psi$. In the strategic-substitutes region ($\delta < 0$) and for moderate complementarity, $\mathscr{X}_k$ dominates and over-adoption obtains; the comparison can reverse only when $\delta > 0$ is large enough that the classical term is positive—by Proposition 3, when $\frac{d}{dq}\mathrm{Var}(\bar{x}) > 0$—and exceeds the external wedge. □

Proof of Proposition 4. Under the decomposition $\xi = \ell\theta + \zeta$, the platform signal is $\tilde{s}_f^A = (1+\ell)\theta + \zeta + u_f^A$ with total variance $V_A(\ell) = (1+\ell)^2\sigma_\theta^2 + \sigma_\zeta^2 + \sigma_A^2$. The residual component orthogonal to θ has variance σ_ζ^2. Repeating the pairwise-covariance computation of Appendix E with the residual component in place of ξ, the normalized error covariance between two adopters equals $\sigma_\zeta^2/V_A(\ell)$, so the $\mathscr{W}_4$ penalty is $q^2\sigma_\zeta^2/V_A(\ell)$ and the external wedge is $\mathscr{X}_4(q) = 2\gamma_4 q\sigma_\zeta^2/V_A(\ell)$. Parts (i)–(iii) follow: the wedge is strictly positive iff $\sigma_\zeta^2 > 0$; raising ℓ at fixed σ_ζ^2 increases $V_A(\ell)$ and so lowers the wedge without eliminating it; and the expression contains no equilibrium loading, so it is independent of δ at fixed q. □

Proof of Proposition 5. With J platforms and symmetric shares $q_j = q/J$, the platform-driven component of aggregate variance is $\mathrm{Var}(\sum_j K_{\xi,j}\xi_j) = (k^A)^2\sigma_\xi^2[\sum_j q_j^2 + 2\rho\sum_{j<j'} q_j q_{j'}]$. Under symmetry $\sum_j q_j^2 = q^2/J$ and $2\sum_{j<j'} q_j q_{j'} = q^2(1-1/J)$, so the bracket equals $q^2H(J,\rho)$ with $H(J,\rho) = 1/J + (1-1/J)\rho$. Hence the $\mathscr{W}_3$ penalty becomes $K_\xi(q)^2\sigma_\xi^2 H(J,\rho)$ and the $\mathscr{W}_4$ penalty becomes $q^2(\sigma_\xi^2/V_A)H(J,\rho)$, each scaling the single-platform penalty by H. Differentiating and applying Theorem 1 gives the stated wedges. Since $H(J,\rho) \in [1/J, 1]$ is strictly decreasing in J for $\rho < 1$ and constant at 1 for $\rho = 1$, the magnitude falls in J iff $\rho < 1$. □

Proof of Proposition 6. With $\bar{x}_t = K_\theta\theta_t + K_\xi\xi_t$, θ_t i.i.d. and ξ_t a stationary AR(1) with $\mathrm{Var}(\xi_t) = \sigma_\xi^2$ and $\mathrm{Cov}(\xi_t, \xi_{t+1}) = \phi\sigma_\xi^2$, independence of θ and ξ gives $\mathrm{Var}(\bar{x}_t) = K_\theta^2\sigma_\theta^2 + K_\xi^2\sigma_\xi^2$ and $\mathrm{Cov}(\bar{x}_t, \bar{x}_{t+1}) = K_\xi^2\phi\sigma_\xi^2$, since the fundamental contributes no autocovariance. The ratio gives $\rho_x(q,\phi)$. Writing $R(q) = K_\xi^2\sigma_\xi^2/(K_\theta^2\sigma_\theta^2 + K_\xi^2\sigma_\xi^2)$, monotonicity in ϕ is immediate. For monotonicity in q, note $R(q)$ is increasing in the ratio $K_\xi(q)/K_\theta(q)$; since $K_\xi(0) = 0$ and this ratio is strictly increasing on $(0,1]$ under the equilibrium loadings, ρ_x is strictly increasing in q. □

Proof of Proposition 7. Holding the fundamental realization and all primitives fixed, the equilibrium aggregate is $\bar{x} = K_0 + K_\theta(q)\theta + K_\xi(q)\xi$. A shift $\Delta\xi$ in the realization of the shared compo-

nent, with θ fixed, changes the aggregate by $\Delta\bar{x}_\tau = K_\xi(q)\Delta\xi$ with $K_\xi(q) = q\,k^A(q)$. Since $k^A(q) > 0$ on $\mathscr{D}$ and $K_\xi(0) = 0$, the magnitude $|\Delta\bar{x}_\tau|$ is zero at $q = 0$ and strictly increasing on $(0,1]$. □

Proof of Proposition 9. The exposure (3) is $q^2\sigma_\xi^2/V_A$: two evaluators drawn independently from the population are both adopters with probability q^2, in which case their normalized forecast errors share the component with variance σ_ξ^2 relative to total signal variance V_A, giving normalized covariance σ_ξ^2/V_A; otherwise the covariance is zero. Substituting into (4) and differentiating, $\mathscr{X}_4(q) = \gamma_4 \frac{d}{dq}[q^2\sigma_\xi^2/V_A] = 2\gamma_4 q\sigma_\xi^2/V_A > 0$ on $(0,1]$. The expression contains no equilibrium loading, so it is independent of δ at fixed q. □

Proof of Proposition 8. Write the adoption map as $G(q) = (1-\lambda)q + \lambda F(\Delta\Pi(q))$. Under the stated hypotheses the static iteration $q \mapsto F(\Delta\Pi(q))$ approaches q^* monotonically from below, so $F(\Delta\Pi(q)) > q$ for $q \in [q_0, q^*)$.

(i) For fixed $q_t < q^*$, $\partial q_{t+1}/\partial\lambda = F(\Delta\Pi(q_t)) - q_t > 0$, so a higher λ yields a higher q_{t+1}. By induction, q_t is strictly increasing in λ at every $t \geq 1$. Since G is a convex combination of q_t and $F(\Delta\Pi(q_t))$, both in $[q_0, q^*]$, the path remains below q^* and converges to it for every $\lambda \in (0,1]$.

(ii) Because $\mathscr{X}_4(q) = 2\gamma_4 q\sigma_\xi^2/V_A$ is strictly increasing in q, and q_t is strictly increasing in λ at every $t \geq 1$ with q_0 fixed, each term $\beta^t\,\mathscr{X}_4(q_t)$ is nondecreasing in λ and strictly increasing for $t \geq 1$. Hence $\mathscr{L}(\lambda;\beta)$ is strictly increasing in λ for any $\beta \in (0,1)$.

(iii) Since $\rho_x(q_t,\phi) = \phi R(q_t)$ with R strictly increasing in q, and q_t strictly increasing in λ for $t \geq 1$, the lag-one autocorrelation is strictly increasing in λ at every $t \geq 1$. □