Shared Intelligence and Misaligned Firm Entry

Feng Chi* Nathan Yang†

July 23, 2026

Abstract

Firms increasingly buy forecasts from the same analytics platforms. Pooling data sharpens each forecast, but routing many decisions through one layer replaces idiosyncratic error, which averages out, with shared error, which does not. We model entry when firms choose between proprietary research and a shared platform, and price the correlation surviving conditioning on fundamentals. Adoption synchronizes entry and correlates screening errors, so a planner counting either cost prefers less adoption than one counting only profits, whatever the strategic environment. Vendor multiplicity helps only in proportion to platforms' structural independence: the restaurant site-selection market retains roughly three-quarters of the single-platform distortion.

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

*University of Illinois Urbana-Champaign, Department of Economics. Email: fchi@illinois.edu.
†University of Illinois Urbana-Champaign, Gies College of Business. Email: ncyang@illinois.edu.

1 Introduction

Retail chains have always faced the same question: where to open the next store. Historically they answered it with local, firm-specific knowledge (Love et al. 1986). A regional grocer learned its trade area through direct experience; a restaurant chain commissioned proprietary market research. The quality and timing of that information shapes entry and survival (Chi 2025). Today national chains, regional grocers, and franchisees increasingly buy the answer from the same external analytics providers, among them CoStar, Placer.ai, and Buxton.[1] Pooling data across clients can make each client's forecast more accurate. It also routes many firms' decisions through a single processing layer, so that errors in that layer are not merely inaccurate but common.

Commonality of this kind raises a question that entry models do not answer. When many firms buy their forecasts from one platform, which part of the resulting correlation in their behavior is a social cost, who bears that cost, and does the answer depend on how the firms compete? The organizing distinction separates the component of a shared signal that tracks fundamentals, which moves firms together because markets genuinely differ, from the residual component that survives conditioning on fundamentals, which moves firms together because they consult the same machine. The first is the reason platforms exist. The second is the object this paper prices, and we price it in its polar form, orthogonal to fundamentals, which bounds the channel from above.

The quick-service restaurant industry shows what such commonality would look like. Counties that appear attractive to McDonald's tend to appear attractive to Burger King as well. Conditional on population, density, income, urbanicity, and region, a county with a McDonald's is 11.9 percentage points more likely to contain a Burger King, and the two chains' observable selection rules are descriptively more similar in counties more exposed to AI-supported analysis, as measured by the AI exposure index of Felten et al. (2021). Co-location of this kind may reflect common demand fundamentals, agglomeration economies, or imitation. The spread of shared analytics vendors suggests one more possibility. Firms may enter the same markets not only because those markets are attractive, but because the same forecasting system points them there.

[1] A small number of providers are widely used across retail and commercial real estate. Esri reports that seven of the ten largest U.S. retailers rely on its geospatial analytics platform, and Placer.ai supplies location intelligence to clients spanning retail, real estate, and hospitality. These platforms are also deployed through intermediaries such as commercial real-estate advisory firms, so 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.

The channel is not a learning channel, and this places it outside standard entry models.[2] Correlation in firm behavior is usually attributed to strategic interaction, reputational herding (Scharfstein and Stein 1990, Froot et al. 1992), or informational cascades (Banerjee 1992, Bikhchandani et al. 1992), mechanisms in which firms learn from common information or from one another's actions. Here no firm observes another's action, and no firm learns anything from the shared component, because it carries no information about fundamentals. Firms act on a common error embedded in a forecasting system, and their decisions align even when each conditions only on its own forecast. Recent evidence points the same way: firms investing heavily in AI carry more systematic risk (Babina et al. 2024), and firms sharing data through common software development kits co-move even after controlling for product similarity (Bian et al. 2024). The channel is also distinct from algorithmic competition in pricing (Calvano et al. 2020, Assad et al. 2024, Brown and MacKay 2023, Raymond 2025), where shared algorithms soften rivalry within a market; shared forecasting instead correlates entry across markets, distorting where firms compete rather than how they price.

Two literatures sit on either side of the question and leave it open between them. On one side, data raise output by improving prediction, sharpening targeting, and propagating across tasks, which makes shared infrastructure an engine of growth (Farboodi and Veldkamp 2023, Farboodi et al. 2026, Farboodi and Veldkamp 2021, 2020, Jones and Tonetti 2020); that literature prices the informative component of shared signals and is silent on the residual one. On the other side, work on data externalities locates inefficiency in the correlation structure of information rather than its quantity (Acemoglu et al. 2022, Choi et al. 2019, Ichihashi 2021, Kirpalani and Philippon 2020, Fang and Kim 2026, Bergemann and Bonatti 2019, Veldkamp and Chung 2024), but the shared object there is data supplied, not an inference technology consumed. The closest antecedent is Veldkamp (2006), where increasing returns lead many agents to buy the same signal and prices co-move. That framework already contains the residual channel: noise in the commonly purchased signal moves prices together without moving fundamentals, alongside the shared-learning applications to business cycles and stock prices (Veldkamp and Wolfers 2007, Morck et al. 2000). This paper takes that residual channel and prices it for parties outside the adopting firms, locates

[2] In canonical frameworks (Bresnahan and Reiss 1991, Seim 2006, Jia 2008, Holmes 2011, Blevins et al. 2018, Hollenbeck 2017, Nishida and Yang 2020, Yang 2020), firms receive information about profitability, and correlation arises from strategic interaction or from information about a common underlying state. Correlation there reflects economically relevant information, and forecast errors are largely independent across firms. Platform forecasts differ because adoption does not make the common component more informative; it propagates the same error across more decisions.

it on a screening margin, classifies it by its invariance to equilibrium objects, and scales it by a vendor-concentration index. For the same reason the mechanism differs from information sharing in oligopoly (Clarke 1983, Gal-Or 1985, Shapiro 1986, Raith 1996): the firms here exchange no information but consume a common inference technology, and the distortion runs through that technology's residual error rather than through the content of what is shared. Holding the platform's precision gain fixed shuts down the growth channel and signs the residual channel in isolation.

We develop a model in which firms choose between proprietary research and a shared forecasting platform. Proprietary research produces forecast errors that differ across firms; the platform reduces firm-specific noise but introduces an error common to all of its users. Adoption then imposes two costs that no firm bears itself. Synchronized entry raises the volatility of commercial activity around fundamentals, which is costly to lenders, landlords, and municipalities. Correlated screening errors reduce the independence of decentralized market evaluation, so a market wrongly passed over by one firm is more likely to be passed over by the next. A planner who internalizes either cost prefers strictly less adoption than a planner who counts only firm profits, in every strategic environment. Whether the adopting firms themselves place too much or too little weight on the shared signal, the efficient-use-of-information question of the social-value literature, is a separate internal margin we do not sign. The screening criterion also prices the correlation of evaluations holding each evaluator's precision fixed, so when the platform's precision gain is large the net effect on the screened markets themselves can favor adoption.

The two costs separate along a single structural line: whether the welfare object carries an equilibrium quantity. The screening cost does not. Screening errors are realized conditional on fundamentals and before any firm acts, so no equilibrium object can enter them, and the associated wedge is invariant to whether entry decisions are strategic complements or substitutes, to the functional form of the entry rule, and to how firms map signals into beliefs. The volatility cost does carry the equilibrium loadings, so its sign is invariant but its magnitude moves with the strategic environment. A proposed cost of shared forecasting therefore has a determinate policy direction exactly when its welfare object is external to the adopting firms, and its robustness is governed by whether that object contains equilibrium quantities.

Three conclusions cut against natural intuitions. First, reducing a platform's firm-specific noise at fixed shared error raises its external liability, gross of the precision benefit: the screening cost scales with the share of an adopter's signal variance attributable to the common component, so low-

ering idiosyncratic noise raises the correlated-evaluation externality, and it is precisely the accurate common signal that makes coincident mistakes damaging. Accuracy achieved instead by shrinking the shared error lowers the externality. Second, the mechanism requires no communication: no herding, no observation of rivals, and no public announcement enters it, so remedies aimed at information flows do not touch the channel. Third, vendor multiplicity is not diversification: competition among platforms reduces the distortion only in proportion to the structural independence of their data and models, and a descriptive case study of the restaurant site-selection analytics market suggests that under plausible cross-vendor correlations the observed market retains between three-fifths and nine-tenths of the distortion a single platform would generate, roughly three-quarters at the central calibration.

Shared processing operates alongside the classical entry externality of Mankiw and Whinston (1986) and Spence (1976). Firms ignore their effect on aggregate entry, so the market is already inefficient before any platform exists, and whether it ultimately over-adopts or under-adopts depends on how the two forces net. The boundary admits an exact characterization: under strategic substitutes the market locally over-adopts whenever adoption raises the dispersion of aggregate entry, a condition that fails only where adoption lowers the aggregate's loading on fundamentals by enough to dominate an always-positive shared-error term, and the comparison reverses only when the classical force is positive and large enough to dominate, with strong agglomeration the leading case.

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. It does not distinguish among common demand fundamentals, agglomeration economies, imitation, franchising constraints, real-estate constraints, or shared analytics. The object formalized in the model, cross-firm alignment in market selection, is visible in a familiar retail setting before any structure is imposed: McDonald's and Burger King exhibit highly similar mappings from observable market characteristics into store counts, and this similarity is descriptively stronger 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), it combines occupation-level AI exposure scores with local employment composition, so counties with higher AIGE 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 shape 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, so we 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. Appendix A benchmarks state-level counts against company-reported store totals and excludes states with evident collection gaps; whatever coverage error survives that check attenuates the associations reported below. After restricting to counties with both brands collected and non-missing covariates, the sample contains 1,987 counties.

[3] These two chains are large national brands in the same product segment with 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 holds 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 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, and ask 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, which 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). Three measures summarize cross-chain alignment: 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). Two features bound what these measures can show. They are bounded above by one, and both coefficient vectors load most heavily on the same component, population, so they lie near one for any two chains that place stores where people live. Absent a placebo pair whose alignment is known to be weak, their levels are not interpretable, and only their ordering across subsamples is informative. We read them as directional throughout. 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, which 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 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, which asks whether market-selection alignment is more pronounced in counties whose employment structure is more exposed to AI. The exercise documents an empirical pattern that motivates

the theoretical mechanism developed below, in the descriptive spirit of Einav and Levin (2014).

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, and it 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, and Igami and Yang (2016) show that hamburger-chain entry in Canada reflects preemption and unobserved market heterogeneity, two of the confounds this exercise leaves unmodeled. Our mechanism does not seek to explain all cross-firm alignment. 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 that channel, and the evidence here is descriptive rather than causal.

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_{im} denote the number of stores operated by chain $i \in \{\text{McDonald's}, \text{Burger King}\}$ in county m. We estimate chain-specific relationships between store counts and county fundamentals:

$$y_{im} = \alpha_i + X_m' \beta_i + \varepsilon_{im},$$

where y_{im} is either $\log(1+n_{im})$ or n_{im} in a Poisson specification. The vector X_m 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 chains. 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 rank counties in almost the same way for both chains, even though McDonald's has the larger footprint. A stacked chain-by-county specification provides a complementary test.

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.

Equality of slopes is rejected in the log-count and presence specifications, largely because the two chains differ in scale and in their sensitivity to population. In the Poisson count specification, equality of the continuous-fundamental slopes is not rejected ($p = 0.238$), although equality is rejected once interactions with urbanicity and region are included. Poisson is the natural specification

for a store count with substantial mass at zero, 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. Exact coefficient equality is too strong a benchmark for shared market selection, and the test is not powered against the alternative of interest, so the non-rejection is not evidence of equality. The descriptive content is that the same observable fundamentals rank markets in nearly the same way for both chains.

2.5 AI Exposure as a Stratifying Variable

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 use it as a stratifying variable rather than a causal regressor. Counties are split at the sample median of the AI Geographic Exposure index, and the market-selection similarity measures are recomputed separately for low- and high-exposure counties. The similarity is high in both subsamples, but stronger in high-AIGE counties. The

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.

predicted-score correlation, which compares how the two models rank counties and is the preferred alignment metric for this split, is 0.9849 in low-AIGE counties and 0.9978 in high-AIGE counties. The scale-invariant coefficient measures show the same pattern, while the scale-sensitive RMS distance moves slightly in the opposite direction, reflecting larger coefficient magnitudes rather than weaker directional alignment. This does not imply that AI exposure causes chains to behave more similarly. 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 descriptive results are robust to alternative presence thresholds, specification choices, and AIGE cut points (Appendix A). 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, and no individual fundamental accounts for a statistically distinguishable share of the difference, though the decomposition has limited power against small contributions. The stronger alignment observed in high-AIGE counties is therefore not attributable to any single observable characteristic in these data. The two families of evidence document different channels, and each maps to a distinct welfare object. The coefficient and score similarity of Tables 2 and 3 measures how similarly the two chains weight observable fundamentals, the fundamentals-tracking channel, which the model treats as efficient and does not price. The moment with model-relevant content is the residual association of Table 1: 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. This residual co-presence, together with county-level co-location, is where the cross-market error covariance of the screening environment of Appendix E would appear. Candidate markets evaluated by co-adopting chains inherit positively correlated assessments, and the residual association contains that covariance, alongside unobserved fundamentals and strategic interaction, rather than measuring it. With store-census data and controls rich enough to absorb unobserved fundamentals, the residual co-presence moment could bound the normalized error covariance and so discipline the model's calibrations quantitatively. Synchronized entry within a market corresponds to the platform-driven component of aggregate variance, the object priced by the aggregate-volatility criterion. The model's falsifiable signature lies in variation the cross-section cannot supply: differential entry responses across high- and low-adoption markets after platform updates, the revision-response ratio of Appendix B, would separate shared processing from common fundamentals and agglomeration. With the empirical objects in hand, we now build the environment that prices them.

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. Shared platforms can improve forecast accuracy by reducing idiosyncratic noise, but they also route decisions through a com-

[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.

mon processing layer that may generate correlated errors across adopters. The model isolates this tradeoff in the simplest environment that permits heterogeneous adoption and strategic interaction. The baseline environment is a single market evaluated by a continuum of chains; the multi-market counterpart of the Section 2 evidence, in which many candidate counties are each screened by a small number of chains, is the screening environment of Appendix E. Primitives and outcomes are distinct throughout. 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] The assumptions divide according to the results they carry. The screening wedge of Theorem 1(ii) needs only two substantive ingredients: adopters' forecasts share a component that survives conditioning on fundamentals, and adoption raises exposure to that component. Delete either and the phenomenon dissolves; delete anything else and the result survives, because the wedge contains no equilibrium object. Its minimal environment is two evaluators screening one candidate market, developed in Appendix E. The remaining structure serves the volatility wedge of Theorem 1(i) and the classical strategic term, whose magnitudes run through equilibrium loadings, and each convenience there is paired with a named robustness exercise. The uniform cost shock delivers a linear entry rule, relaxed in Appendix D; orthogonality of the shared component to fundamentals is the polar isolating case, relaxed in Section 5.6; and rational filtering with known σ_ξ^2 is a conservative benchmark delimited in the filtering discussion below.

3.1 Environment

A unit continuum of retail chains $i \in [0,1]$ decides whether to enter a local market. The entry decision is binary, $x_i \in \{0,1\}$, and profits are

$$\Pi_i = (\theta + \delta\,\bar{x} + \varepsilon_i)\,x_i,$$

where $\theta \sim \mathcal{N}(0,\sigma_\theta^2)$ is unobserved local profitability, $\bar{x} \equiv \int_0^1 x_i\,di$ is the aggregate entry rate, and $\varepsilon_i \sim \text{Uniform}[-\bar{\varepsilon},\bar{\varepsilon}]$ 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

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

markets with limited demand. The lower bound $\delta > -1$ plays no role in equilibrium existence, which is governed by the upper bound on δ collected in the admissible region $\mathscr{D}$ of Section 4; it instead bounds the strength of business stealing so that the mean entry rate $K_0 = 1/(2(1-\delta))$ stays above $1/4$. Interior validity at high complementarity, where K_0 can exceed one and the linear rule becomes an interior-regime approximation, is treated in the linearity footnote below and in 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 a choice of units and does not affect the qualitative results.

3.2 Information and Processing Technologies

Firms can generate forecasts internally using proprietary research or rely on a shared external platform. This make-versus-buy margin 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. Babina (2026) identify the margin 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 this literature has largely left open: what are the welfare consequences when many firms rely on the same forecasting infrastructure? The signal structure formalizes the 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

[6] The linear rule is exact on the event $|\mathbb{E}_i[\theta] + \delta \mathbb{E}_i[\bar{x}]| \leq \bar{\varepsilon}$; because posterior means are Gaussian, no finite $\bar{\varepsilon}$ makes this hold pointwise, but the event holds with probability arbitrarily close to one when $\bar{\varepsilon}$ is large relative to the dispersion of posterior means, and the failure probability is small whenever $\bar{\varepsilon} \geq |\delta| K_0 + 4\sqrt{\mathrm{Var}(\theta) + \mathrm{Var}(\bar{x})}$. 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 $\Omega(q) > 0$, and the comparative statics in Theorem 1 extend to a saturating entry rule in which entry probabilities are smoothly compressed into $[0,1]$. Appendix D computes this saturating-rule equilibrium under a logistic entry probability: the wedge under the screening criterion $\mathscr{W}_4$ is exactly invariant to the replacement because it contains no equilibrium object, while the wedge under the aggregate-volatility criterion $\mathscr{W}_3$ retains its sign and its monotonicity in q, with a rule-sensitive magnitude that converges to the linear closed form in the interior-regime limit.

shared across adopters. A chain using custom research observes

$$\tilde{s}_i^C = \theta + u_i^C,$$

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

$$\tilde{s}_i^A = \theta + \xi + u_i^A,$$

where $u_i^A \sim \mathcal{N}(0, \sigma_A^2)$ 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 two technologies therefore differ not only in precision but in 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. Holding σ_A^2 and σ_ξ^2 fixed as adoption varies shuts down data-pooling feedback, the scale economy by which a larger client base sharpens each client's forecast. Two consequences of the fixed-technology benchmark deserve note. If σ_A^2 instead declined with q, the adoption gain $\Delta\Pi$ could increase in q even at $\delta = 0$, so the tipping regime of Section 6 could arise without strategic complementarity; the taxonomy there attributes tipping to coordination motives conditional on fixed technology. And because the screening exposure σ_ξ^2/V_A rises as σ_A^2 falls, data feedback would enlarge rather than shrink the screening externality, so the fixed-technology benchmark errs against the mechanism. 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 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), where a common signal induces correlation because it is informative about the underlying state, so that alignment rises 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. Orthogonality to θ is an isolating device, not a claim that real processing errors are fundamental-free. It selects the polar case in which the shared component carries no productive content, so that the correlated-error channel is separated cleanly from the social value of public information. Section 5.6 relaxes it: a shared component correlated with fundamentals ($\xi = \ell\theta + \zeta$) attenuates the wedge, which attaches only to the orthogonal residual ζ, so the empirically likely case $\ell > 0$ delivers a strictly smaller wedge than the $\ell = 0$ benchmark used here. The orthogonal baseline is therefore the maximal-effect case for the mechanism, not the typical one. Calling ξ an "error" refers to the ex-post realization, not to any ex-ante suboptimality: under the rational-filtering benchmark below, adopters know σ_ξ^2 and respond to their signals optimally. No firm makes a mistake in the decision-theoretic sense; the externality is entirely a third-party effect of individually optimal behavior that happens to share a common component. When multiple markets are in view, as in Appendices B and E, the processing error of market m decomposes as $\xi_m = \xi^P + \nu_m$, where ξ^P is a platform-wide component common to every market the platform evaluates, the component shifted by model revisions, and ν_m is a market-specific application error, independent across markets. Within any single market the sum ξ_m acts exactly as the baseline's ξ, so the single-market analysis is unchanged; the decomposition matters only for objects that compare across markets.

3.3 Adoption Stage

Each chain draws a switching cost $c_i \sim F$ 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:

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

[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.

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. The distinction does not affect the qualitative results.

3.4 Bayesian Filtering

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

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

Then $\mathbb{E}_i^A[\theta] = \alpha_\theta^A \tilde{s}_i^A$, $\mathbb{E}_i^A[\xi] = \alpha_\xi^A \tilde{s}_i^A$, and $\mathbb{E}_i^C[\theta] = \alpha_\theta^C \tilde{s}_i^C$. 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, 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. Such firms may place too much weight on the platform signal and therefore react too strongly to the shared component. Over-weighting of algorithmic advice, or automation bias, is documented in clinical and managerial decision-making (Goddard et al. 2012, Keding and Meissner 2021). The rational-filtering benchmark isolates the welfare consequences of the information structure itself, abstracting from the additional distortions that arise when firms misperceive the

precision of platform forecasts. For the screening penalty as specified, the benchmark is immaterial: the screening object is the covariance of signal errors conditional on fundamentals, $q^2\sigma_\xi^2/V_A$ (Section 5; Appendix E), which contains no filtering weights, so the corresponding wedge is invariant not only to the entry rule but to belief formation entirely. Mis-filtering changes what adopters do with their signals, not the correlation of the signals themselves. Any exposure measured instead on action-relevant forecast errors is monotone in the weight adopters place on the shared component, so over-weighting of the platform signal enlarges that alternative object. We confine the claim to these exposure objects rather than the full gap between private and social adoption, because mis-filtering also moves the private adoption gain and the classical strategic term, which we do not sign jointly. Either way, correct filtering does not overstate the correlated-evaluation externality.[8]

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 $\Omega(q)$ 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,

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

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

$$P_i(\text{enter} \mid s_i) = \tfrac{1}{2} + \mathbb{E}_i[\theta] + \delta\, \mathbb{E}_i[\bar{x}],$$

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

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

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

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

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

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

Proposition 1 (Equilibrium loadings)**.** *For $q \in (0,1)$, the unique linear equilibrium has*

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

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

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

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 uninforma-

[9] The two bounds in $\mathscr{D}$ ensure $\Omega(0) > 0$ and $\Omega(1) > 0$, respectively, and concavity of $\Omega(q)$ implies positivity throughout $[0,1]$. The lower bound $\delta > -1$ plays no role in existence. For $\delta \geq 1$, existence 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 forms extend analytically to such cases but should then be read as approximations to the saturating-rule equilibrium described in Section 3.

tive common components.

4.2 Aggregate Behavior

Substituting the equilibrium loadings into the aggregate yields

$$\bar{x} = K_0 + K_\theta(q)\,\theta + K_\xi(q)\,\xi,$$

so that, using $\mathbb{E}[\theta] = \mathbb{E}[\xi] = 0$, aggregate variance is

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

At the calibration plotted in Figure 1 ($\delta = 0.7$), the intercept $K_0 = 1/(2(1-\delta)) \approx 1.67$ exceeds one, so the linear rule there is an interior-regime device rather than a globally valid probability model. The wedge algebra is unaffected: 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. The 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, aggregate outcomes also inherit variation from the shared processing component. The two sources of variation behave differently under aggregation. Idiosyncratic noise averages out in the continuum, but the shared shock survives, so 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 is 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. Two forces operate. The first is a classical entry externality: because each atomistic chain ignores its effect on the aggregate, the private equilibrium already de-

parts 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 $\delta = 0$. Moreover, there is a neighborhood of $\delta = 0$ within $\mathscr{D}$ on which, for $\delta \neq 0$, $\Delta\Pi(q)$ is strictly increasing in q when $\delta > 0$ and strictly decreasing in q when $\delta < 0$.*

When $\delta = 0$, chains care only about matching θ, so adoption depends on the signal-precision tradeoff. When $\delta > 0$, adoption becomes more attractive as it spreads because platform adopters track the shared component that increasingly shapes aggregate entry. When $\delta < 0$, the 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 $\sigma \equiv (\sigma_\theta^2, \sigma_\xi^2, \sigma_A^2, \sigma_C^2)$ collects the variance primitives and Γ is established in closed form to be positive near $\delta = 0$ (Appendix G). Wherever $\Gamma > 0$, the sign of $d\Delta\Pi/dq$ is determined by δ. Numerical evaluation indicates that the sign pattern extends to the full admissible region: Γ is strictly positive over 200,000 parameter draws spanning $\mathscr{D}$ in the replication package, so that

$\operatorname{sgn}(d\Delta\Pi/dq) = \operatorname{sgn}(\delta)$ throughout. This global claim is a numerical finding rather than a theorem, and results that rely on it away from $\delta = 0$, such as the transition dynamics of Appendix C, invoke it as an explicit hypothesis. Away from $\delta = 0$, Γ is a polynomial in (q, δ) and the variance primitives whose monomials do not share a sign, and neither a factorization into positive terms nor a sum-of-squares certificate has been found; the $\delta = 0$ slice, where Γ collapses to a sum of positive terms (Appendix G), is the analytically tractable case.

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 operates at the aggregate level. When many firms rely on the same forecasting platform, aggregate entry becomes more sensitive to the shared processing component, which 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 concern parallels the systemic-risk literature, in which correlated exposures and homogeneous behavior amplify aggregate shocks (Acemoglu et al. 2015, Brunnermeier and Pedersen 2009, Shleifer and Vishny 2011), with AI-driven homogenization recently modeled as a distinct source of fragility (Danielsson and Uthemann 2023, Dou et al. 2025). This constituency is not hypothetical: financial-stability authorities have repeatedly flagged the possibility that widespread reliance on a small number of AI models and data providers induces correlated behavior with consequences for parties outside the adopting firms (Board 2024, Bank of England Financial Policy Committee 2025, Cecchetti et al. 2025). The decomposition $\xi_m = \xi^P + \nu_m$ of Section 3 refines who bears how much: parties exposed to a single market's aggregate, such as a landlord or a municipality, bear the full within-market variance σ_ξ^2, while parties diversified across markets, such as a lender with a national book, bear only the platform-wide component, so for them σ_ξ^2 should be read as $\sigma_{\xi^P}^2$, attenuating the wedge's magnitude while preserving its sign and its monotonicity in q. 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 operates at the level of individual candidate markets. When multiple firms evaluate the same candidate market through a shared platform, their forecasts inherit the same processing error, so evaluation mistakes

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, which scales as $q^2\sigma_\xi^2/V_A$, reflecting both the probability that two firms adopt the platform and the importance of the shared component in platform forecasts. These errors are determined once fundamentals and signals are drawn, before any entry decision is taken, so equilibrium behavior has no channel into them. That timing is the source of the δ-invariance established below, and it is a property of the screening environment rather than of the functional form: Appendix E derives the exposure. The invariance is scoped to harms attached to evaluations, the site scores and shortlists that platforms directly produce, rather than to the entry outcomes that follow them. A loss defined on realized entry, the object the co-presence evidence of Section 2 measures, would reintroduce the strategic parameter through entry probabilities, and the invariance would then hold only approximately. 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 candidate market, and the distortion arises because multiple firms evaluate the same market through correlated signals. We inherit the lost-independent-evaluation mechanism of Kleinberg and Raghavan (2021), and the contribution is to make exposure endogenous through the adoption equilibrium in q and to classify welfare objects by whether they carry an equilibrium quantity, which yields a strategy-invariance property absent in the fixed-population monoculture setting. The pairwise object priced below corresponds to an explicit encounter structure. Alongside the entry game, suppose there is a unit continuum of ex-ante identical candidate markets, each screened by two chains drawn independently at random from the population, with each evaluator applying its equilibrium forecasting technology to a market-specific realization of (θ, ξ, u). For a given market, the normalized covariance of the two evaluators' signal errors $e \equiv \tilde{s} - \theta$, the evaluation errors conditional on market quality, equals σ_ξ^2/V_A when both evaluators are platform adopters and zero otherwise, so the expected normalized error covariance across markets is exactly $q^2\,\sigma_\xi^2/V_A$. The pairwise structure is a normalization: if each market is instead screened by $n \geq 2$ evaluators drawn independently, the expected number of co-adopting pairs is $\binom{n}{2}q^2$, each contributing a normalized error covariance σ_ξ^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 mul-

tiplicative constant changes, and it is absorbed into the penalty weight. 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.

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 at the end of this section.[10] The term complete welfare designates a ledger with announced omissions. The complete-welfare planner counts chain profits and one named external cost. It does not count the precision benefit that adoption confers on the screened markets themselves, which the screening penalty holds out by fixing marginals (Appendix E), and it does not count consumers, the constituency priced by the entry apparatus of Mankiw and Whinston (1986) and Spence (1976), whose exposure to the technology margin is unsigned in this environment. Excluding unsigned terms is an isolating device: the planner's preference for less adoption is a statement about the named external costs, not a full accounting of all affected parties. 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 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:

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

[10] The indices 1, 3, and 4 are used for the specifications analyzed here, with the index 2 reserved for the fundamental-tracking criterion $\mathscr{W}_2$.

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 δ. The parameter γ_3 measures the marginal social cost of platform-induced volatility. The penalty prices only part of aggregate variance, and the restriction is deliberate. A party whose payoff is concave in aggregate entry, such as a risk-averse lender exposed to correlated default or a landlord facing convex vacancy costs, bears a loss proportional to total $\mathrm{Var}(\bar{x})$ and would also price the fundamental component $K_\theta(q)^2\sigma_\theta^2$. That component's derivative in q is unsigned, so a full-variance wedge would carry a sign-indeterminate fundamental term. The ξ-only penalty instead adopts a nonfundamental-volatility benchmark: synchronization that tracks θ moves commercial activity with genuine demand and is treated as efficient, and only variation orthogonal to fundamentals is priced. This parallels Hébert and La'O (2023), in which nonfundamental volatility is precisely the welfare-relevant object that decentralized information choices can generate. Timing supplies a microfoundation for the ξ-only penalty. Fundamentals are unobserved ex ante, which is why forecasting platforms exist, but they are revealed ex post as realized demand. Contracts that settle on realized fundamentals, such as turnover-linked rents or performance-indexed credit terms, therefore transmit to the external party only the nonfundamental deviation $K_\xi(q)\,\xi$, and a quadratic loss in that residual delivers the penalty $K_\xi(q)^2\sigma_\xi^2$. Exactness requires indexation on the fundamental-warranted level of activity $\mathbb{E}[\bar{x} \mid \theta]$, whose slope $K_\theta(q)$ re-prices with penetration. Fixed-slope indexation, such as a turnover-linked rent with slope b, delivers the penalty approximately, leaving a residual fundamental term of order $(K_\theta(q) - b)^2\sigma_\theta^2$. The derivative of that residual term in q is unsigned, so for a fixed-slope contract holder the total external effect of adoption is unsigned and the corroboration supplied by $\mathscr{W}_3$ attaches to the ξ-component of exposure alone. The decided-about-agent specification, denoted $\mathscr{W}_4$ and referred to below as the screening criterion, penalizes correlated evaluation at the level of individual candidate markets:

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

This penalty is not an object of the single-market baseline. It is the second moment of the multi-market screening environment of Appendix E, in which a continuum of candidate markets is each screened by co-adopting chains, so $\mathscr{W}_3$ and $\mathscr{W}_4$ are two external accounting exercises rather than a single ledger. 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. Concretely, σ_ξ^2/V_A is the covariance of two adopters' signal errors conditional on fundamentals, normalized by total signal variance, so a noisier platform (higher σ_A^2) dilutes that share and reduces the per-encounter loss. The apparently paradoxical implication, that a platform with less firm-specific noise carries a larger correlated-evaluation externality, is the substance of the mechanism: it is precisely the accurate common signal that makes coincident mistakes damaging. The parameter γ_4 prices the loss of an independent second opinion. When evaluation errors are uncorrelated, a market wrongly screened out by one chain retains a chance of being recognized by the other; correlated errors eliminate this diversification, in the spirit of Kleinberg and Raghavan (2021). Appendix E embeds the penalty in an explicit screening environment, in which a candidate market is recognized when at least one of its evaluators forecasts favorably, and shows that positive error covariance strictly lowers the probability that a good market is recognized by *some* evaluator. The load-bearing assumption behind the penalty's sign is a social loss convex in the number of coincident evaluation errors of either sign, the class under which Slepian's inequality makes both tails costly: a market passed over by both evaluators in error and a market entered by both in error are both costly, so these losses do not net out across markets even though the shared error is mean zero. Convexity is weaker than a free choice within a symmetric menu. An additive loss that charges each wrong entry its own cost is exactly correlation-neutral, correlation helps only under a loss concave in coincident wrong entries, and the miss tail of a best-of-n recognition rule is convex by construction. The penalty prices correlation holding each evaluator's precision fixed, so because adoption also sharpens individual evaluations when $\sigma_A^2 < \sigma_C^2$, the net effect on candidate markets can favor adoption when the platform's precision gain is large. One convention in $\mathscr{W}_4$ deserves comment. The penalty divides a covariance computed conditional on fundamentals by the unconditional signal variance V_A. Because an adopter's error $\xi + u^A$ has variance $\sigma_\xi^2 + \sigma_A^2$, the conditional normalization $\sigma_\xi^2/(\sigma_\xi^2 + \sigma_A^2)$ is equally available and removes σ_θ^2 from an object that conditions on θ. Nothing in Theorem 1 turns on the choice: under either convention the wedge is strictly positive, linear in q, invariant in δ, and increasing in platform precision, and lowering σ_A^2 raises the shared variance share under both. The choice does affect Proposition 4(ii), where the conditional normalization would make the wedge exactly invariant to the informative loading ℓ rather than merely decreasing in it. We retain V_A for continuity with the filtering coefficients of Section 3, and read the comparative static in ℓ as normalization-specific. 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 the equilibrium loading $k^A(q)$; the $\mathscr{W}_4$ criterion penalizes correlation in signals and is independent of that loading, depending instead on the within-adopter variance share σ_ξ^2/V_A. The two are genuinely independent objects: one uses an equilibrium quantity the other does not, and each attaches to a different affected population, parties exposed to aggregate fluctuations under $\mathscr{W}_3$ and screened markets under $\mathscr{W}_4$. Because the same directional conclusion, less adoption, emerges under both, it is a property of the primitive that both criteria detect, namely a common component orthogonal to fundamentals whose exposure rises with adoption, rather than an artifact of either functional form. As a corollary of the same structure, 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 comparison across the two criteria instantiates a general classification, recorded next as an accounting identity.

Lemma 1 (Criterion classification)**.** *Fix $\gamma > 0$ and let complete welfare take the form $\mathscr{W}(q) \equiv \mathscr{W}_1(q) - \gamma\mathscr{P}(q)$ for a differentiable penalty $\mathscr{P}$. The external wedge satisfies $\mathscr{X}(q) \equiv d\mathscr{W}_1/dq - d\mathscr{W}/dq = \gamma\mathscr{P}'(q)$, and therefore:*

(i) if $\mathscr{P}$ depends only on the adoption rate and the primitives of the information structure, the wedge at any fixed q is invariant to the strategic parameter, to the entry rule, and to belief formation, and is strictly positive wherever $\mathscr{P}$ is strictly increasing;

(ii) if $\mathscr{P}$ depends on q also through equilibrium objects, the wedge remains strictly positive wherever the penalized exposure is strictly increasing in q, and its magnitude varies with the strategic environment through those objects.

The screening criterion $\mathscr{W}_4$ is of class (i) and the aggregate-volatility criterion $\mathscr{W}_3$ of class (ii). A criterion that is not an additive perturbation of $\mathscr{W}_1$ falls outside both classes: its costs fold into the same profits the planner trades off, so its wedge carries indirect equilibrium terms of indeterminate sign, and the fundamental-tracking criterion discussed below is of this kind. A reader confronting a newly proposed cost of shared forecasting can locate it in this classification before writing down a model: the class of the welfare object, not the functional form of the penalty, determines whether the policy direction is determinate and how robust it is. Determinacy in each

[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.

class is carried by the assumed monotonicity of the penalized exposure, and the substantive case that a natural environment generates a class-(i) penalty is Proposition 9. 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 is therefore the sign of each external wedge and its variation across environments rather than the absolute magnitude of welfare losses. Any welfare *level* reported below in units of γ_k is ordinal, not cardinal. The model's cardinal objects are the index ratios in which the weight cancels: the effective concentration index $H(J,\rho)$ of Section 7, the revision-response ratio of Appendix B, and the discounted-exposure index ratio of Appendix C.

5.4 Marginal Chain Welfare and the Classical Entry Externality

The linear-Gaussian continuum does not deliver team efficiency. The atomistic chain takes the law of motion of the aggregate $\bar{x}$ as given and therefore does not internalize how the population's technology mix moves the equilibrium loadings, and with them the strategic component $\delta\,\mathbb{E}[x_i\bar{x}]$ of every chain's profit. The result is a strategic externality whenever $\delta \neq 0$. Planner problems are defined uniformly throughout. The chain-welfare planner chooses q to maximize $\widetilde{\mathscr{W}_1}(q) \equiv \mathscr{W}_1(q) - \int_0^q F^{-1}(s)\,ds$, chain welfare net of the aggregate switching costs incurred to reach adoption rate q, and the complete-welfare planner maximizes $\widetilde{\mathscr{W}_k}(q) \equiv \mathscr{W}_k(q) - \int_0^q F^{-1}(s)\,ds$ for $k \in \{3,4\}$. At any interior maximizer the first-order condition $F^{-1}(q) = d\mathscr{W}_k/dq$ is necessary; we use it as a marginal characterization, never as a definition of the optimum. The symbols q_1^{soc} and q_k^{soc} denote maximizers of these net objectives.

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 $d\mathscr{W}_1/dq = \Delta\Pi(q)$ for all q when $\delta = 0$, while for $\delta \neq 0$ the two differ at every q at which $\frac{d}{dq}\mathrm{Var}(\bar{x}) \neq 0$. The private equilibrium condition $F^{-1}(q^*) = \Delta\Pi(q^*)$ and the chain-welfare planner's interior first-order condition $F^{-1}(q) = d\mathscr{W}_1/dq$ differ by exactly Ψ: the private equilibrium coincides with an interior chain-welfare optimum only if $\Psi(q^*) = 0$, and whenever $\Psi(q^*) \neq 0$ the two conditions cannot hold at the same point. 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 term $\Psi(q)$ has an envelope interpretation. Each chain's decision rule is a best response to the aggregate, so a marginal change in q affects a chain's expected profit only through the equilibrium response of $\bar{x}$: as adoption rises, the loadings (k^A, k^C) adjust, the dispersion of aggregate entry changes, and the strategic component $\delta\,\mathbb{E}[x_i\bar{x}]$ of every chain's profit moves with it. Summed across the population, these uninternalized responses equal exactly $\Psi(q)$. Because firms choose forecasting technologies on private profits alone, this feedback enters no chain's decision, and the chain-welfare planner and the decentralized market generally choose different adoption levels. This is the entry externality of Mankiw and Whinston (1986) and Spence (1976), translated from the entry margin to the technology-adoption margin, and it is present even when $\sigma_\xi^2 = 0$. The classical term separates across the two shocks, which makes the boundary of the over-adoption result transparent. Write $\Psi(q) = \Psi_\theta(q) + \Psi_\xi(q)$, where $\Psi_j(q) \equiv \frac{\delta}{2}\frac{d}{dq}[K_j(q)^2\,\sigma_j^2]$ for $j \in \{\theta, \xi\}$. The component Ψ_θ survives at $\sigma_\xi^2 = 0$: it is the Mankiw and Whinston analog term of a platform-free economy, and it carries the calibration-dependent sign. The component Ψ_ξ shares the object $\frac{d}{dq}[K_\xi(q)^2\sigma_\xi^2]$ with the external wedge $\mathscr{X}_3$ of Theorem 1, so under $\mathscr{W}_3$ the shared shock's total contribution to the market-versus-complete-planner gap $\mathscr{X}_3 - \Psi$ carries the coefficient $\gamma_3 - \delta/2$; this net coefficient, which can take either sign, is distinct from the planner-versus-planner wedge $\mathscr{X}_3$ itself, which is always positive. Because $\frac{d}{dq}[K_\xi(q)^2\sigma_\xi^2] > 0$ throughout (proof of Theorem 1), two consequences follow. First, failure of the over-adoption guarantee of Theorem 1 requires the fundamental term $\frac{d}{dq}[K_\theta(q)^2\sigma_\theta^2]$ to be negative and large enough to dominate the always-positive ξ-term. Second, under business stealing ($\delta < 0$) the component Ψ_ξ is strictly negative, so shared processing pushes toward over-adoption through the classical channel as well as through the external wedge. 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. Because $\mathscr{W}_k$ differs from $\mathscr{W}_1$ only by the penalty, $\mathscr{W}_1$ cancels in the wedge, so finite firm counts, entry-rule curvature, and team-efficiency arguments touch the strategic term Ψ and the magnitude of $\mathscr{X}_3$, but never the wedges themselves. The substance lies in the penalties. The penalty $\mathscr{P}_4 = q^2\sigma_\xi^2/V_A$ contains no equilibrium object and is exactly rule-invariant, and Appendix E derives it from a screening environment, so its δ-invariance is a property of that environment rather than of a convenient functional form. The penalty $\mathscr{P}_3 = K_\xi(q)^2\sigma_\xi^2$ does carry the loadings, and $\mathscr{X}_3$ correspondingly retains its sign and its monotonicity in q with a rule-sensitive

magnitude (Appendix D). The scope of the linear rule follows the same division. The equilibrium loadings $K_\theta(q)$ and $K_\xi(q)$ that enter $\mathscr{X}_3$ and the calibrated figures are interior-regime objects, and Appendix D quantifies what saturation does to them: at the same degree of complementarity as Figure 1 ($\delta = 0.7$), under the appendix calibration of Table A1, the saturating-rule wedge is roughly 3.5 percent of the linear closed form, more than an order of magnitude smaller, while sign and monotonicity in q survive. Linear $\mathscr{X}_3$ levels therefore carry sign and q-monotonicity content only, while $\mathscr{X}_4$ contains no loading and is exactly rule-invariant. The saturation robustness certified in the appendix covers the planner-versus-planner wedges $\mathscr{X}_k$; the market-versus-planner net comparison $\mathscr{X}_k - \Psi$, and with it the substitutes over-adoption guarantee and the 96.6 percent share discussed below, additionally involve $\Delta\Pi$ and Ψ, which are interior-regime objects, so those statements are within-linear-model results.

5.5 Misalignment Theorem

Internalizing either external cost calls for strictly less adoption. The relevant object is the wedge between the chain-welfare planner and the complete-welfare planner, which isolates the part of the distortion attributable to shared processing.

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 the screening criterion $\mathscr{W}_4$,

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

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

Consequently every maximizer of $\mathscr{W}_k$ lies weakly below every maximizer of $\mathscr{W}_1$, and the ordering is strict under interiority: $q_k^{soc} < q_1^{soc}$ for $k \in \{3,4\}$ whenever the complete-welfare optimum satisfies $q_k^{soc} \in (0,1)$. The same ordering holds verbatim for the net-of-switching-costs objectives $\widetilde{\mathscr{W}_k}$ and $\widetilde{\mathscr{W}_1}$, whose common cost term cancels in the comparison exactly as $\mathscr{W}_1$ does. Relative to the private equilibrium, the total gap decomposes as

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

At any stable interior adoption equilibrium q^*, complete welfare net of switching costs rises as adoption is reduced marginally below q^* whenever $\mathscr{X}_k(q^*) > \Psi(q^*)$; we refer to this as local over-adoption. In particular, whenever $\Psi(q^*) \le 0$, which under strategic substitutes holds exactly when $\frac{d}{dq}\mathrm{Var}(\bar{x}) \ge 0$, the market locally over-adopts for any $\gamma_k > 0$. The comparison reverses only when the classical term is positive and exceeds the external wedge ($\Psi > \mathscr{X}_k$).

The ordering of optima is a monotone-comparative-statics conclusion rather than a first-order-condition comparison: $\mathscr{W}_k$ differs from $\mathscr{W}_1$ by the subtraction of a penalty that is strictly increasing in q, so the two maximizer sets are ordered in the strong set order, in the sense of Milgrom and Shannon (1994), without any concavity assumption on either objective, and strictness requires only that the complete-welfare optimum be interior (Appendix G). The market-versus-planner comparison, in turn, is a local statement at a stable equilibrium: in the tipping regime of Section 6, where multiple stable equilibria coexist, the decomposition applies equilibrium-by-equilibrium at each stable fixed point. 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 invariance concerns the wedge at a common q, not the realized equilibrium externality: the equilibrium adoption rate $q^*(\delta)$ depends on δ (Proposition 2), so the externality actually incurred, $\mathscr{X}_4(q^*(\delta))$, still varies with strategic interaction through the adoption rate it induces. With that qualification, the external result does not depend on whether entry decisions are strategic complements or substitutes. It follows from two features of the envi-

ronment: 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 whenever adoption raises the dispersion of aggregate entry, with both directions reversing when it lowers it. 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]

Corollary 1 (Sign boundary for the market comparison)**.** *At any stable interior adoption equilibrium q^*, the market locally over-adopts under criterion $k \in \{3,4\}$ if and only if $\mathscr{X}_k(q^*) > \Psi(q^*)$, and locally under-adopts if and only if the inequality reverses. Since $\mathscr{X}_k > 0$ throughout, under-adoption requires $\Psi(q^*) \geq \mathscr{X}_k(q^*) > 0$ and therefore $\delta \frac{d}{dq}\mathrm{Var}(\bar{x})\big|_{q^*} > 0$. In particular, $\delta \frac{d}{dq}\mathrm{Var}(\bar{x})\big|_{q^*} \leq 0$ is sufficient for local over-adoption at every $\gamma_k > 0$, and because $\frac{d}{dq}[K_\xi(q)^2\sigma_\xi^2] > 0$ throughout, under strategic substitutes the reversal can occur only where adoption lowers the aggregate's loading on fundamentals by enough for the fundamental term to dominate.*

The reversal region is therefore pinned by two observable objects, the sign of strategic interaction and the response of aggregate dispersion to adoption, rather than by the unidentified penalty weights. The three calibrations of Figure 1 illustrate the wedges under moderate complementarity ($\delta = 0.7$); the parameter sweeps in the replication package extend the evaluation across the admissible region, including strategic substitutes. In the strategic-substitutes region, $\Psi \leq 0$ whenever $\frac{d}{dq}\mathrm{Var}(\bar{x}) \geq 0$, in which case $\mathscr{X}_4 - \Psi \geq \mathscr{X}_4 > 0$ and over-adoption is guaranteed, a within-linear-model statement. Because the variance derivative is not uniformly signed, the guarantee is conditional: across 30,000 random admissible draws with $\delta \leq 0$ and $\gamma_4 = 1$, the inequality $\mathscr{X}_4 - \Psi > 0$ holds in 96.6 percent of draws, and every failure occurs at a calibration in which adoption lowers the dispersion of aggregate entry, so that $\Psi > 0$. Since $\frac{d}{dq}[K_\xi(q)^2\sigma_\xi^2] > 0$ always, Ψ can turn positive under substitutes only when adoption lowers the aggregate's loading on fundamentals by enough for the θ-term to dominate, and the 3.4 percent of failing draws are

[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.

exactly such calibrations. The transportable statement is the condition rather than the share: $\Psi \leq 0$ implies over-adoption for every $\gamma_k > 0$, whereas the 96.6 percent figure is conditional on the unit normalization $\gamma_4 = 1$, weakly increasing in γ_4, and conditional on the sampling measure over $\mathscr{D}$, which we choose. Of the two criteria, $\mathscr{W}_4$ provides the cleaner benchmark. Its external wedge depends only on adoption and signal-structure primitives and is invariant to the entry-choice rule by construction, since $\mathscr{W}_1$ cancels in the wedge and the penalty contains no equilibrium object. The aggregate-volatility criterion $\mathscr{W}_3$ yields the same qualitative conclusion, with a magnitude that inherits the equilibrium loadings and is rule-sensitive. We therefore view $\mathscr{W}_4$ as establishing the external-wedge result and $\mathscr{W}_3$ as corroborating it. The sign of the $\mathscr{W}_4$ wedge rests on the convex-loss class of Appendix E, and its invariance to δ, to the entry rule, and to belief formation follows from pricing evaluations realized before any firm acts. The invariance also has an empirical advantage. The strategic parameter δ is difficult to identify because agglomeration and business-stealing effects often coexist and vary across settings, and δ enters both the classical wedge Ψ and the magnitude of $\mathscr{X}_3$. Under $\mathscr{W}_4$, no such identification is required. Appendix E derives the screening penalty rather than positing it. 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, and states the wedge as Proposition 9. The same analysis shows why the δ-invariance separates the two channels. 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-agent 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. A natural alternative criterion 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):

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

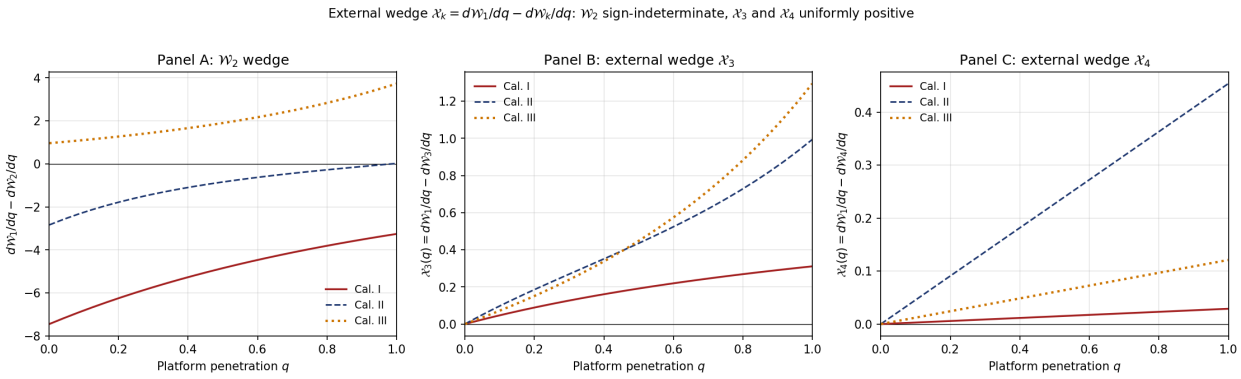
This criterion penalizes deviations of each type's signal loading from one, the loading at which entry passes signals through fully. The loading-of-one target is illustrative rather than derived: the payoff-based efficient-use-of-information benchmark of Angeletos and Pavan (2007), under

which the efficient loadings generally differ from one and depend on δ, is the correct internal comparison, and we leave it open. 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 those indirect loading terms. It therefore delivers no determinate prediction for over- or under-adoption. The sign ambiguity parallels the unresolved debate in the social-value-of-information literature over whether public disclosure raises or lowers welfare (Svensson 2006, Morris et al. 2006, Hellwig 2005), including settings where agents learn from prices or from one another (Amador and Weill 2010, Vives 1997). The contrast isolates the role of externalities: under $\mathscr{W}_2$, costs are internal to the chain population, and a determinate external wedge arises only when the planner accounts for costs borne outside it. A planner concerned with both objectives would combine the criteria, in which case the external wedge would operate alongside any distortion associated with imperfect fundamental tracking. 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. Conditional on pricing external correlated-error exposure at all, the sign of the wedge is invariant to whether that exposure is measured through realized aggregate volatility ($\mathscr{W}_3$) or through signal covariance ($\mathscr{W}_4$), two objects that differ in whether they carry an equilibrium loading and in the population they price. It is not invariant to switching the planner's target from external correlated-error costs to an internal fundamental-tracking objective, which is indeterminate. Determinacy is thus a property of internalizing external correlated-error costs, robust to how one measures them, rather than of any single functional form.

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. Whether the external wedge depends on the common component being entirely unrelated to fundamentals is a separate question. To draw the boundary, suppose the

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 $\mathscr{W}_2$ wedge $d\mathscr{W}_1/dq - d\mathscr{W}_2/dq$: 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 $\mathscr{W}_2$ planner. Panel B shows the external wedge $\mathscr{X}_3(q) = d\mathscr{W}_1/dq - d\mathscr{W}_3/dq$ 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 $\mathscr{X}_4(q)$ 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 positive under both structurally distinct external penalties, the equilibrium-loading-based $\mathscr{W}_3$ and the signal-covariance-based $\mathscr{W}_4$, while the internal fundamental-tracking criterion $\mathscr{W}_2$ in Panel A is indeterminate. The net private-versus-social comparison additionally subtracts the classical strategic term $\Psi(q)$ of Proposition 3, which is not plotted.

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 $\tilde{s}_i^A = (1+\ell)\theta + \zeta + u_i^A$, with total variance $V_A(\ell) \equiv (1+\ell)^2\sigma_\theta^2 + \sigma_\zeta^2 + \sigma_A^2$. The informative component changes signal precision. The residual component ζ is the part that remains after conditioning on fundamentals and is therefore the part to which the correlated-evaluation externality attaches.

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

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

denote the admissible region under the transformed signal, which reduces to $\mathscr{D}$ at $\ell = 0$, and let $V_A(\ell) \equiv (1+\ell)^2\sigma_\theta^2 + \sigma_\xi^2 + \sigma_A^2$. Under the screening 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. The wedge therefore prices residual commonality, not shared informativeness: a larger informative component ℓ dilutes the residual share and shrinks the wedge, but the wedge remains positive whenever residual error remains. Part (ii) is the one claim that depends on the normalization discussed above; under the conditional normalization the wedge would not depend on ℓ at all, so the separation from the public-information benchmark is if anything sharper than stated. The orthogonal baseline is the special case $\ell = 0$, in which σ_ξ^2 itself is the relevant primitive and the wedge is largest. What governs the externality is thus the residual common variance that survives conditioning on fundamentals, not the joint informativeness of the signal. Within the one-factor Gaussian family $\xi = \ell\theta + \zeta$, the wedge is pinned by σ_ζ^2 alone; this is not robustness over all information structures (cf. Bergemann and Morris 2013), and characterizing the wedge over the full set of information structures a platform could induce, subject to a fixed marginal precision, is a stronger question we leave open. The decomposition also completes an audit of how robustly each negative result holds. The screening wedge $\mathscr{X}_4$ survives every relaxation considered in this paper: it is invariant to the entry rule (Appendix D), to belief formation (Section 3), to strategic interaction (Theorem 1), and

to the sign and size of the informative loading ℓ, because it conditions on θ and attaches only to the orthogonal residual $\sigma_\zeta^2 > 0$. The volatility wedge $\mathscr{X}_3$ is the brittle one: its magnitude is rule-sensitive, and in the complementary case $\ell < 0$, a shared component that leans against fundamentals and so partly stabilizes the aggregate, it shrinks and could in principle reverse sign. We restrict attention to $\ell \geq 0$ as the empirically motivated case of a processing error that leans with fundamentals, and this asymmetry is why $\mathscr{W}_4$ serves as the paper's benchmark criterion. The policy implication turns on a distinction the decomposition makes sharp. Improving common-signal informativeness lowers the wedge only mechanically, by inflating the denominator $V_A(\ell)$ and thereby shrinking the residual's variance *share*; it does not remove the residual error and so does not internalize the correlated-mistake channel. Genuine mitigation requires reducing σ_ζ^2 itself, either by increasing structural independence across vendors, limiting overlap in upstream inputs, or auditing the processing layer that converts data into forecasts.

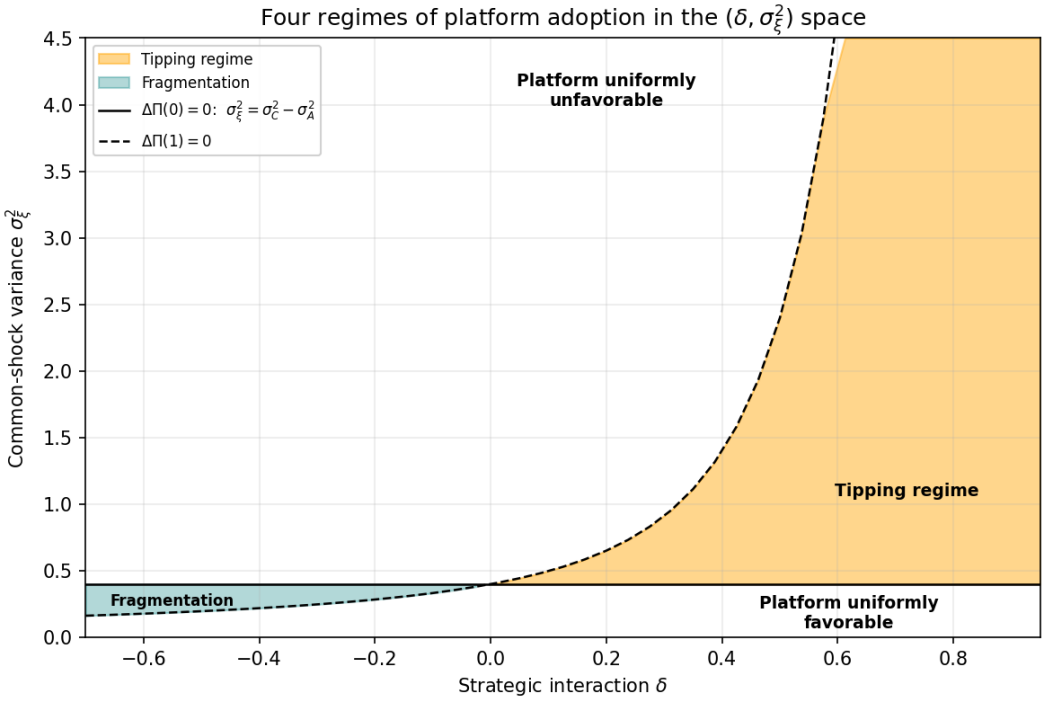
6 Four Regimes of Adoption

Adoption is determined by the fixed point $q^* = F(\Delta\Pi(q^*))$. 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 2 (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 and 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; consistent with this intuition, across the calibrations evaluated the threshold in σ_ξ^2 is increasing and convex in δ, a numerical finding rather than a proved property. That the locus intersects the $q = 0$ boundary at $\delta = 0$ is, by contrast, an immediate corollary of Proposition 2 and Lemma 2: at $\delta = 0$, $\Delta\Pi$ is constant in q, so the two loci coincide there, at $\sigma_\xi^2 = \sigma_C^2 - \sigma_A^2$. Figure 2 plots these two loci in (δ, σ_ξ^2)

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

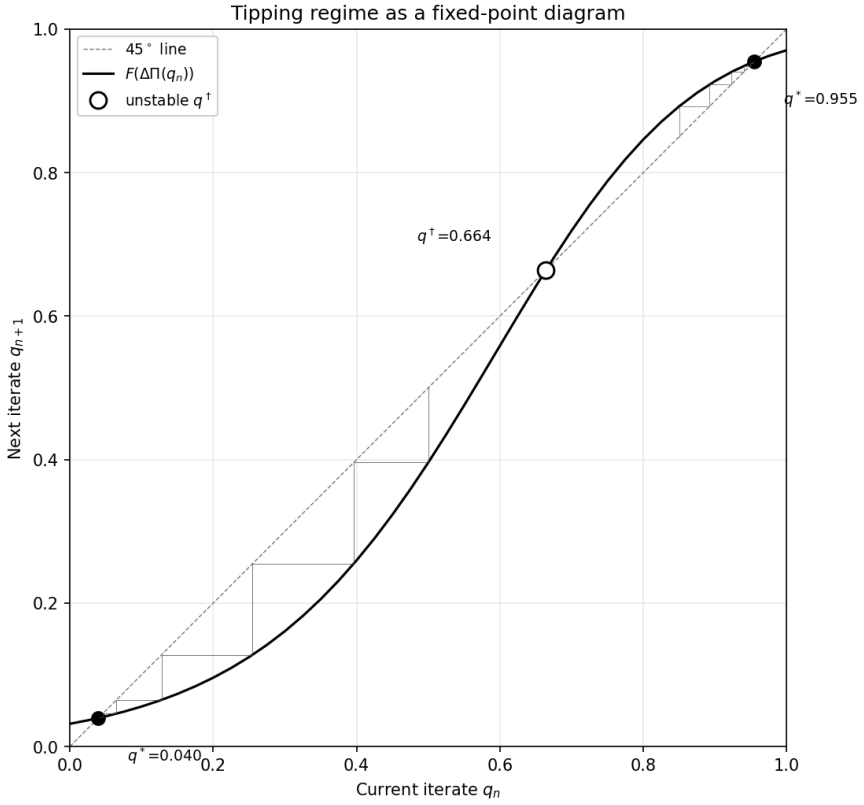
space and partitions the parameter region into four regimes. While equilibrium configurations vary across regions, the external wedge of Theorem 1 is strictly positive throughout: strategic interaction shapes how adoption is sustained and whether the *net* private distortion is over- or under-adoption, 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 at least three solutions; in the generic three-solution configuration these are 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 self-reinforcing adoption also echoes the network-effects and switching-cost classics (Katz and Shapiro 1985, Farrell and Saloner 1985, Klemperer 1995, Farrell and Klemperer 2007), though here the complementarity is informational: adopters gain from tracking the signal their rivals act on. 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 formalizes 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

[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 $\Omega(q) > 0$ throughout $[0,1]$.

point also lies above the complete planner's own optimum depends on the net sign $\mathscr{X}_k - \Psi$, which is positive whenever $\Psi \leq 0$ and can turn negative when $\Psi > 0$ is large, as under strong agglomeration with $\frac{d}{dq}\mathrm{Var}(\bar{x}) > 0$. 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, $\mathscr{W}_k(q_L^*)$ 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; under $\mathscr{W}_4$, it scales with q^2 independently of δ. In numerical examples, the low-adoption fixed point is welfare-superior

when γ_k is large relative to the coordination gain, and the ranking reverses when the coordination gain dominates; because welfare levels scale with the unidentified γ_k, no unconditional ranking is available. A full global ranking is left open; 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, an interior fixed point exists; invoking the global monotonicity reported after Proposition 2 as a hypothesis, $F \circ \Delta\Pi$ is strictly decreasing in q and the fixed point is unique. 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 locally over-adopts whenever the external wedge dominates the classical term Ψ, which is guaranteed whenever $\frac{d}{dq}\mathrm{Var}(\bar{x}) \geq 0$, so that $\Psi \leq 0$ by Proposition 3; under the screening criterion with unit penalty weight, over-adoption holds in the large majority of admissible substitutes calibrations (Section 5). 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

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 $u_i^{A,j}$ with variance σ_A^2, so that an adopter of platform j observes

$$\tilde{s}_i^{A,j} = \theta + \xi_j + u_i^{A,j}.$$

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

[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. The foundation-model paradigm makes such shared cores pervasive: many separately branded products inherit one pretrained model

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.

7.2 Effective Concentration Index

The analysis maintains the following informational device.

Assumption 1 (Vendor-anonymous filtering)**.** Each adopter conditions on its own signal knowing the total variance of its common component but not the vendor split across rival adopters: branding masks vendor identities, so an adopter forecasting the aggregate treats rival adopters' processing shocks as exchangeable with its own. Under this assumption, each adopter filters exactly as in the single-platform model.

Under Assumption 1, each adopter observes a single processing shock and the equilibrium loadings (k^A, k^C) are exactly those of 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, $q_j = q/J$, the bracketed term reduces to $q^2\,H(J,\rho)$, where

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

so that the platform-driven variance is $K_\xi(q)^2\,\sigma_\xi^2 \cdot H(J,\rho)$ with $K_\xi(q) = q\,k^A$. 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. It is the variance factor of an equally weighted portfolio of equicorrelated risks, equivalently a rescaled design effect for clustered samples (Kish 1965); the content here is its mapping to vendor concentration and the calibration of Section 9.

(Bommasani et al. 2021), whose common training pipeline, from instruction tuning (Ouyang et al. 2022) to recursive exposure to model-generated data (Shumailov et al. 2024), is itself a correlated-error source.

Proposition 5 (Vendor diversity)**.** *Maintain Assumption 1. 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 shared components coincide; 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. Vendor-anonymous filtering is not innocuous when $J > 1$, $\rho \in (0,1)$, and $\delta \neq 0$. A fully rational adopter of platform j forecasting $\bar{x}$ would load on rival platforms' shocks through $\mathrm{Corr}(\xi_j, \xi_{j'}) = \rho$, and the fixed point of Proposition 1 would replace the common-component share χ_A with $\alpha_\theta^A + H(J,\rho)\,\alpha_\xi^A$, with the equilibrium denominator adjusted to $\Omega_H(q) = (1-\delta q(\alpha_\theta^A + H\alpha_\xi^A))(1-\delta(1-q)\alpha_\theta^C) - \alpha_\theta^A\alpha_\theta^C\delta^2 q(1-q)$. Three statements survive this replacement. The $\mathscr{W}_4$ wedge is unchanged, because it contains no equilibrium loading. The $\mathscr{W}_3$ wedge retains strict positivity and monotonicity in q, because the argument of Theorem 1(i) goes through with χ_A replaced by the smaller share. What is specific to the assumption is the exact proportionality of the $\mathscr{W}_3$ wedge to $H(J,\rho)$ in Proposition 5(i): in the rational benchmark, exact H-proportionality holds at $\delta = 0$, where the loadings contain no strategic term, and at $\rho = 1$, where the two filtering rules coincide, and is otherwise approximate because H then also enters the loadings themselves.

7.3 Substantive versus Cosmetic Diversity

Two corner cases of Proposition 5 isolate the policy content.

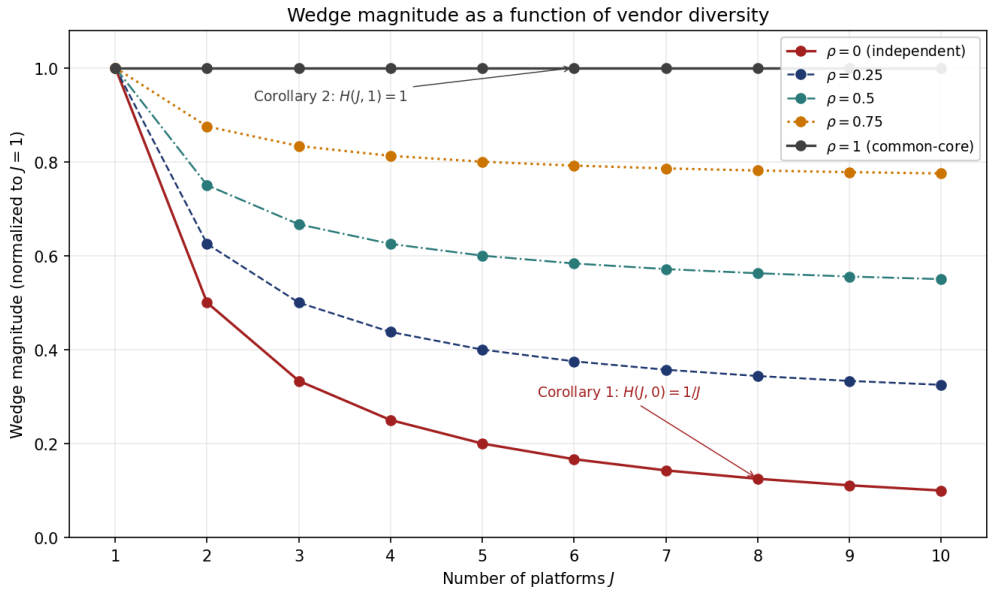
Corollary 2 (Independent platforms)**.** *When $\rho = 0$, $H(J,0) = 1/J$. The external-wedge magnitude under either welfare specification falls linearly in J.*

Corollary 3 (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. Between these corners the index is strictly decreasing in J for every $\rho < 1$, so multiplicity always helps somewhat and never fully: vendor diversity reduces the externality in proportion to the structural independence of processing technologies, and the number of providers matters only through that channel. 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; it establishes that structural correlation among independently branded prediction systems is large, not 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 2). At $\rho = 1$ (gray), platforms share a common processing core and the magnitude is invariant in J (Corollary 3). 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 and the effective concentration of processing technologies. 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 $\mathscr{W}_4$). Adoption taxes or capital requirements tied to platform exposure operate on the first margin; procurement diversification or data-source requirements operate on the second. The diversification margin is priced here through correlation alone: splitting adopters across more vendors shrinks each vendor's client pool, and if pooling improves precision this raises each vendor's σ_A^2, a cost the fixed-technology model does not count.

The comparative statics in $H(J,\rho)$ therefore measure the correlation benefit of diversification, not its net value. The two margins act on distinct components of the externality. Reducing adoption lowers exposure but leaves the correlation structure unchanged, while reducing concentration lowers correlation but leaves adoption above the complete-welfare optimum. Joint intervention therefore achieves the optimum with lower distortion than either margin alone, and policies targeting platform structure complement, rather than substitute for, policies targeting adoption. This is a comparative-statics result rather than a mechanism-design solution: the model shows how the externality scales with the concentration index $H(J,\rho)$, but J and ρ are exogenous primitives here, so reducing H describes the direction a structural intervention would take, not an equilibrium response of vendors to policy. Endogenizing vendor entry and data differentiation, so that J and ρ respond to regulation, is left to future work. Two features of observed vendor markets sit outside the symmetric partition and move the index in opposite directions. Shares are uneven in practice, and 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 strict positivity, the distinction between substantive and cosmetic diversity, and the separation of policy margins all carry over; the symmetric case is used for transparency in comparative statics over J. The partition also rules out multi-homing, and the disclosed-client compilation of Appendix F shows several chains appearing in multiple vendors' materials. Multi-homing raises rather than lowers effective concentration: a chain that combines the J vendor signals loads on the cross-vendor composite, which it shares exactly with every other multi-homer, so the pairwise correlation of shared components among multi-homers is one and multi-homing behaves like a $\rho = 1$ composite vendor. The partition-based index therefore understates effective concentration in the presence of multi-homing, and the calibrated $H = 0.76$ of Section 9 is conservative for the multi-homing subpopulation.

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 arrives, chooses whether to adopt the platform, plays the entry game, and exits. Within each period, the entry game is identical to Section 3. Arriving cohorts observe their own signal only and no lagged aggregates, so within-period filtering is exactly that of the static model. Conditional on adoption rate q, aggregate entry in deviations from its constant mean K_0 is

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

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. Appendix C endogenizes q along a deterministic partial-adjustment path. If lagged aggregate entry were public, adopters could partially strip the predictable component $\mathbb{E}[\xi_t \mid \text{history}]$ from their signals, and because lagged entry confounds θ_{t-1} and ξ_{t-1} the stripping is partial, which would attenuate $\rho_x(q, \phi)$ below. Proposition 6 is therefore the no-history-learning benchmark, an upper bound on shared-processing persistence under history-based learning. The revision-response identification of Appendix B is robust to this concern, because revisions are

innovations η_t, unpredictable from history.

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

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

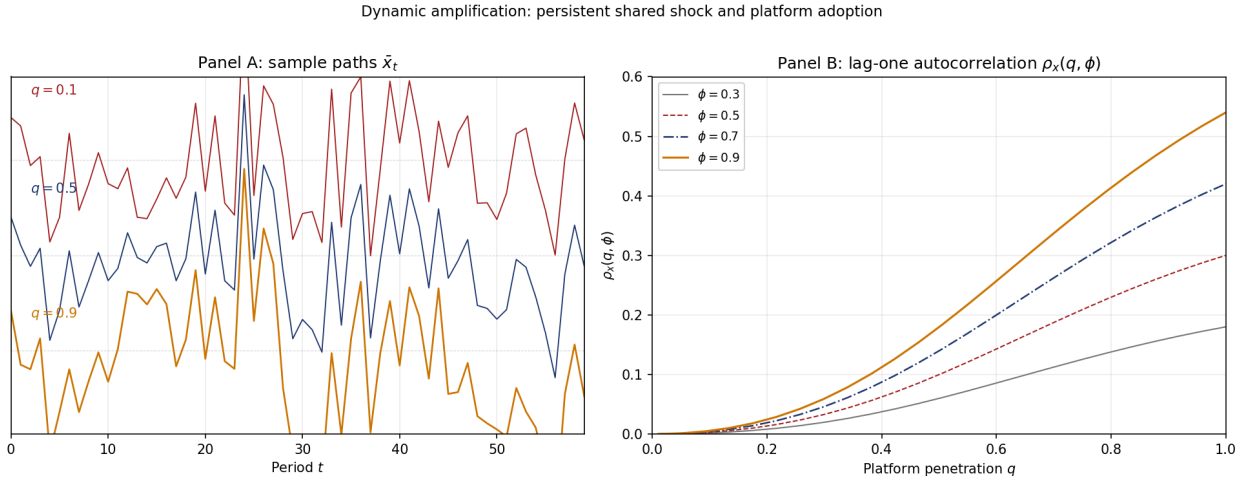
satisfies the following: for $\phi \in (0,1)$, $\rho_x(q,\phi)$ is strictly increasing in q on $(0,1]$ and strictly increasing in ϕ; for $\phi \in (-1,0)$ the monotonicity in q reverses; and $|\rho_x(q,\phi)|$ is strictly increasing in q for every $\phi \neq 0$.

In particular $\rho_x(q,\phi) = \phi\,R(q)$, where $R(q) \equiv K_\xi(q)^2\,\sigma_\xi^2 / \left(K_\theta(q)^2\,\sigma_\theta^2 + K_\xi(q)^2\,\sigma_\xi^2\right)$ denotes the share of aggregate variance attributable to the shared component.

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. In the static environment, shared processing increases the

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 in deviations from its mean, $\bar{x}_t - \mathbb{E}[\bar{x}_t] = K_\theta(q)\,\theta_t + K_\xi(q)\,\xi_t$, 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.

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 for $\phi > 0$, and $|\rho_x|$ is for every $\phi \neq 0$, 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. In this re-

spect the shared component plays a role analogous to the noise and sentiment shocks that generate aggregate fluctuations without fundamentals (Lorenzoni 2009, Angeletos and La'o 2013); here the non-fundamental driver is not a belief shock but an error embedded in a forecasting technology, and its aggregate footprint scales with the technology's adoption. The dynamic extension does not add a third priced externality; it establishes an amplification result. A marginal increase in adoption raises $K_\xi(q)$ and therefore increases the weight placed on the persistent component ξ_t in aggregate behavior. Because agents are short-lived, they do not internalize this propagation effect, and adoption decisions depend only on contemporaneous payoffs. Whether persistence is itself a social cost requires a planner objective defined over sequences, which we do not impose; absent one, the result strengthens the two static externalities rather than adding a third.[15] 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.

9 A Case Study of Restaurant Site-Selection Analytics

The model's central objects, the adoption rate, the vendor count, and the cross-vendor correlation, are in principle measurable. This section provides a descriptive case study of the restaurant site-selection analytics market and uses it to discipline a calibration of the effective concentration index.

9.1 Market Structure

The market for restaurant site-selection analytics is served by a small set of identifiable vendors. Public client disclosures, case studies, and industry coverage indicate that a handful of platforms, including Placer.ai, Buxton, Esri, SiteZeus, and Kalibrate (which acquired Pitney Bowes' location-intelligence assets), account for the bulk of disclosed chain-restaurant engagements. Appendix F compiles disclosed clients by vendor. The disclosed-client evidence is consistent with a market in which a typical adopting chain relies on one of roughly four to six substantive platforms, so we take $J \in \{4,5,6\}$ as the empirically relevant range. Vendors share upstream data inputs: mobile-device location panels are concentrated among a small number of suppliers, and demographic projections

[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 is included in the replication package.

derive from a common Census base. Vendors also share methodological practice, with trade-area definition, gravity-style competition modeling, and analog-based forecasting recurring across platforms' technical documentation. Both point to positive structural correlation across nominally distinct vendors. Direct evidence on prediction systems more broadly, reviewed in Section 7, finds error correlations that are large even across independent providers.

9.2 Calibrated Concentration

Table 4 reports the effective concentration index $H(J,\rho)$ over the empirically relevant grid. At $J = 5$ and $\rho = 0.7$, the observed market retains roughly three-quarters of the single-platform distortion. Even at the most favorable corner of the grid ($J = 6$, $\rho = 0.5$), it retains 0.58. Across the grid the retained share ranges from 0.58 to 0.93; the point value 0.76 is the central calibration rather than an estimate, and at correlations below the grid, for instance $\rho = 0.25$, the retained share falls to roughly two-fifths. Vendor multiplicity as observed in this market, absent structural independence, eliminates at most about two-fifths of the correlated-error externality, and under the central calibration roughly one-quarter. Figure 6 translates the calibrated concentration into the

Table 4: Effective concentration $H(J,\rho)$ at empirically relevant values

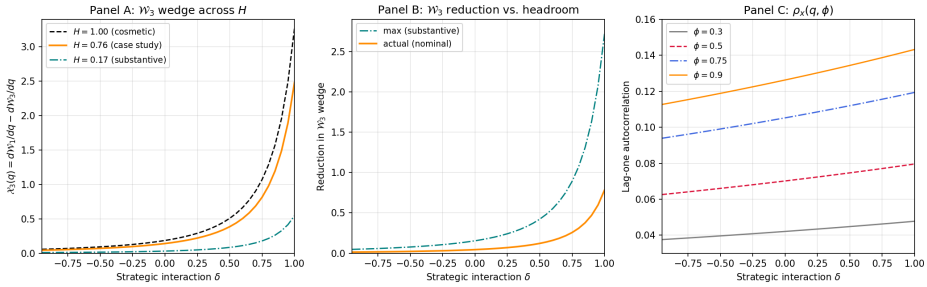
	$\rho = 0.5$	$\rho = 0.7$	$\rho = 0.9$
$J = 4$	0.625	0.775	0.925
$J = 5$	0.600	0.760	0.920
$J = 6$	0.583	0.750	0.917

Notes: Entries are $H(J,\rho) = 1/J + (1-1/J)\rho$, the share of the single-platform external wedge retained under J symmetric vendors with cross-vendor correlation ρ (Proposition 5). The empirically relevant range $J \in \{4,5,6\}$ is motivated by the disclosed-client compilation in Appendix F; $\rho \in [0.5, 0.9]$ reflects shared upstream data inputs and methodological overlap.

two external wedges across the strategic spectrum $\delta \in [-0.95, 1]$. The penetration level $q = 0.6$ is illustrative, a mid-to-high value reflecting the breadth of disclosed adoption among large chains in Appendix F rather than an estimate; both wedges are increasing in q (Theorem 1), so comparisons across concentration levels are unaffected by the choice. The $\mathscr{W}_4$ wedge needs no plot: because it contains no equilibrium object, its value at $q = 0.6$ and $\sigma_\xi^2 = 0.5$ is $0.333\,\gamma_4$ under the single-platform benchmark and $0.253\,\gamma_4$ at the case-study concentration $H = 0.76$, at every δ in the admissible region. Panel A plots the $\mathscr{W}_3$ wedge $H \cdot \mathscr{X}_3(q)$ at three concentration levels, the

single-platform benchmark, the case-study calibration, and six structurally independent vendors. While the wedge rises steeply with complementarity through the equilibrium loadings, the case-study curve is the benchmark scaled by $H = 0.76$ at every δ: concentration rescales the externality uniformly across strategic environments, an exactness that holds under the vendor-anonymous filtering of Assumption 1. The δ-steepness of the plotted linear $\mathscr{X}_3$ levels is an interior-regime feature that saturation attenuates (Appendix D), so the $H = 0.76$ scaling, constant in δ, is the panel's robust content. Panel B decomposes the reduction, showing that observed nominal diversity delivers roughly 29 percent of the reduction achievable under structural independence, a share again constant in δ. Panel C reports the lag-one autocorrelation of aggregate entry at the same calibration, computed under the no-history-learning benchmark of Section 8. The case study thereby illustrates the two policy margins jointly: vendor multiplicity as observed removes about a quarter of either externality, uniformly in δ, and the remaining $\mathscr{W}_4$ wedge is itself invariant to strategic interaction.

Figure 6: The case-study external wedge across the strategic spectrum.



Notes: All panels are evaluated at $q = 0.6$ under the case-study calibration $\sigma_\theta^2 = 1$, $\sigma_\xi^2 = 0.5$, $\sigma_A^2 = 0.3$, $\sigma_C^2 = 0.6$, $\gamma_3 = 1$, for $\delta \in [-0.95, 1]$ inside the admissible region $\mathscr{D}$ of Proposition 1. Panel A plots the external $\mathscr{W}_3$ wedge $H \cdot \mathscr{X}_3(q)$ of Theorem 1(i) at three effective concentration levels: $H = 1$ (single platform, or perfectly correlated vendors), $H = 0.76$ (the case-study calibration $J = 5$, $\rho = 0.7$ of Table 4), and $H = 1/6$ (six structurally independent vendors). The wedge varies with δ through the equilibrium loadings, but each curve is the single-platform benchmark scaled by its H, so the case-study curve is 76 percent of the benchmark at every δ. Panel B plots the wedge reduction achieved by observed nominal diversity ($H = 1$ versus $H = 0.76$) against the maximum reduction available under structural independence ($H = 1$ versus $H = 1/6$); the realized share, roughly 29 percent, is constant in δ. Panel C plots the lag-one autocorrelation $\rho_x(q, \phi)$ of aggregate entry (Proposition 6) for $\phi \in \{0.3, 0.5, 0.75, 0.9\}$. The δ-invariant $\mathscr{W}_4$ wedge is omitted from Panels A and B because it would appear as a horizontal line; its values are $0.333\,\gamma_4$ at $H = 1$ and $0.253\,\gamma_4$ at $H = 0.76$, at every δ (Theorem 1(ii)). The externality realized in equilibrium would still vary with δ through the adoption rate $q^*(\delta)$ it induces (Proposition 2).

9.3 Interpretation and Caveats

The case study is descriptive and the calibration is illustrative. Client disclosures are voluntary and selected; vendor shares are not observed, so the symmetric benchmark is an approximation; and ρ is not directly measured in this market, so the grid $\rho \in [0.5, 0.9]$ is a plausibility range motivated by input and method overlap rather than an estimate. The exercise shows that the model's policy-relevant objects map to observable features of a real analytics market, and that plausible values imply the distortion survives observed vendor competition largely intact. Measuring ρ directly, for instance by comparing forecast errors across vendors for common candidate sites, is a natural target for future empirical work.

10 Conclusion

When many firms buy market forecasts from the same platform, part of the correlation in their behavior reflects fundamentals and part reflects the machine. This paper separates the two and prices the second. The residual component of a shared forecast, the part that survives conditioning on fundamentals, imposes two costs on parties outside the adopting firms: entry synchronizes, raising the volatility of commercial activity, and screening errors correlate, eroding the value of independent second opinions on candidate markets. A planner who prices either cost prefers strictly less adoption than a planner who counts only firm profits, in every strategic environment. That is a comparison between two planners rather than a claim that the market itself over-enters, and the screening wedge at a given adoption rate does not depend on whether entry decisions are strategic complements or substitutes, because the harm it prices attaches to evaluations completed before any firm acts rather than to the entry outcomes that follow. The portable lesson is a classification. External criteria built only from adoption and the information structure yield wedges that are invariant to strategic interaction, to the entry rule, and to belief formation, because the errors they price are realized before equilibrium behavior can enter them. External criteria that carry equilibrium objects keep a determinate sign but inherit the strategic environment in their magnitude. Criteria internal to the adopting population are sign-indeterminate, because their costs fold into the same profits the planner is trading off. Determinacy comes from the externality, not from the functional form. Three beliefs should change. Reducing a platform's firm-specific noise, hold-

ing its shared error fixed, raises its external liability: lower idiosyncratic noise raises the weight of the common component, and it is the accurate common signal that makes coincident mistakes damaging, while accuracy achieved by shrinking the shared error lowers the liability instead. The mechanism needs no communication: unlike herding or cascades, no firm herds, observes a rival, or receives a public announcement, so remedies aimed at information flows do not touch it. And vendor multiplicity is not diversification: the distortion falls with the structural independence of platforms' data and models, not with the count of user-facing providers, so policy should target the correlation structure of forecasting infrastructure, through data-source diversity, disclosure of upstream dependencies, or audits of the processing layer, alongside any intervention on adoption itself. Three qualifications bound the conclusions. First, the classical entry externality operates alongside the new channel, so whether the market itself over-adopts depends on how the two net; the boundary is exact, with over-adoption guaranteed under strategic substitutes whenever adoption raises the dispersion of aggregate entry, and reversal confined to configurations where the classical force is positive and large, with strong agglomeration the leading case. Second, the social cost of correlated errors is not identified, so welfare levels are ordinal, and the screening criterion prices correlation at fixed evaluator precision, so a large enough precision gain can reverse the net comparison for screened markets; the interpretable magnitude is the effective concentration index, in which the unidentified cost cancels. Third, the descriptive evidence is consistent with shared processing but does not identify it, since common fundamentals and agglomeration generate similar cross-firm alignment. Several extensions follow. Letting the shared error's variance respond to adoption, through congestion or data feedback, adds an intensity margin to the exposure margin studied here; a preliminary decomposition is included in the replication package. Long-lived firms and forward-looking adoption would connect the framework to dynamic entry games. And measuring cross-vendor error correlation directly, using common candidate sites evaluated by multiple platforms, would convert the calibration exercise into estimation. The broader agenda is to treat forecasting infrastructure as an economic object in its own right, with market structure, correlated failure modes, and welfare consequences that standard entry models do not capture.

References

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

Acemoglu D, Ozdaglar A, Tahbaz-Salehi A (2015) Systemic risk and stability in financial networks. *American Economic Review* 105(2):564–608.

Amador M, Weill PO (2010) Learning from prices: Public communication and welfare. *Journal of Political Economy* 118(5):866–907.

Angeletos GM, La'o J (2013) Sentiments. *Econometrica* 81(2):739–779.

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

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

Assad S, Clark R, Ershov D, Xu L (2024) Algorithmic pricing and competition: Empirical evidence from the german retail gasoline market. *Journal of Political Economy* 132(3):723–771.

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.

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

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

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.

Brown ZY, MacKay A (2023) Competition in pricing algorithms. *American Economic Journal: Microeconomics* 15(2):109–156.

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

Calvano E, Calzolari G, Denicolo V, Pastorello S (2020) Artificial intelligence, algorithmic pricing, and collusion. *American Economic Review* 110(10):3267–3297.

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

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

Cecchetti S, Lumsdaine RL, Peltonen T, Sánchez Serrano A (2025) Artificial intelligence and systemic risk. Advisory Scientific Committee Report No. 16, European Systemic Risk Board, URL https://www.esrb.europa.eu/pub/pdf/asc/esrb.asc2512~ai_and_systemic_risk.en.pdf.

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

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

Clarke RN (1983) Collusion and the incentives for information sharing. *The Bell Journal of Economics* 14(2):383–394.

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.

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

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

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

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 (2020) Long-run growth of financial data technology. *American Economic Review* 110(8):2485–2523.

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.

Gal-Or E (1985) Information sharing in oligopoly. *Econometrica* 53(2):329–343.

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

Goddard K, Roudsari A, Wyatt JC (2012) Automation bias: a systematic review of frequency, effect mediators, and mitigators. *Journal of the American Medical Informatics Association* 19(1):121–127.

Hébert B, La'O J (2023) Information acquisition, efficiency, and nonfundamental volatility. *Journal of Political Economy* 131(10):2666–2723.

Hellwig C (2005) Heterogeneous information and the welfare effects of public information disclosures. Technical report, UCLA mimeo.

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

Hollenbeck B (2017) The economic advantages of chain organization. *The RAND Journal of Economics* 48(4):1103–1135.

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

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

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

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

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

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

Keding C, Meissner P (2021) Managerial overreliance on ai-augmented decision-making processes: How the use of ai-based advisory systems shapes choice behavior in r&d investment decisions. *Technological Forecasting and Social Change* 171:120970.

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

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

Kish L (1965) *Survey Sampling* (New York: John Wiley & Sons).

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.

Lorenzoni G (2009) A theory of demand shocks. *American Economic Review* 99(5):2050–2084.

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.

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

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

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

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

Morris S, Shin HS, Tong H (2006) Social value of public information: Morris and shin (2002) is actually pro-transparency, not con: Reply. *American Economic Review* 96(1):453–455.

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

Nishida M, Yang N (2020) Threat of entry and organizational-form choice: The case of franchising in retailing. *Journal of Marketing Research* 57(5):810–830.

Ouyang L, Wu J, Jiang X, Almeida D, Wainwright C, Mishkin P, Zhang C, Agarwal S, Slama K, Ray A, et al. (2022) Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems* 35:27730–27744.

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.

Raith M (1996) A general model of information sharing in oligopoly. *Journal of Economic Theory* 71(1):260–288.

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

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

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.

Shapiro C (1986) Exchange of cost information in oligopoly. *The Review of Economic Studies* 53(3):433–446.

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

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

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

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

Svensson LEO (2006) Social value of public information: Comment: Morris and shin (2002) is actually pro-transparency, not con. *American Economic Review* 96(1):448–452.

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.

Veldkamp L, Chung C (2024) Data and the aggregate economy. *Journal of Economic Literature* 62(2):458–484.

Veldkamp L, Wolfers J (2007) Aggregate shocks or aggregate information? costly information and business cycle comovement. *Journal of Monetary Economics* 54:37–55.

Veldkamp LL (2006) Information markets and the comovement of asset prices. *The Review of Economic Studies* 73(3):823–845.

Vives X (1997) Learning from others: A welfare analysis. *Games and Economic Behavior* 20(2):177–200.

Yang N (2020) Learning in retail entry. *International Journal of Research in Marketing* 37(2):336–355.

Online Appendix

A Robustness of the Motivating Evidence

First, the co-occurrence pattern of Table 1 is robust to alternative presence thresholds. Requiring at least two locations of each brand, rather than one, lowers the joint-presence share mechanically but leaves the conditional association intact: among counties with two or more McDonald's, 54.7 percent contain two or more Burger King locations, against 2.1 percent among the remainder. Second, the similarity measures of Table 2 are stable across specification choices. Dropping the Rural-Urban Continuum Code fixed effects changes the cosine similarity by less than 0.001 in the log-count specification. Winsorizing store counts at the 99th percentile, to guard against dense urban counties driving the fit, changes the score correlation from 0.9964 to 0.9958. Restricting to counties with at least one location of either brand, so that joint absences do not inflate similarity, yields a score correlation of 0.9902. Third, the AIGE split of Table 3 is robust to alternative cut points. Splitting at the 40th and 60th percentiles rather than the median produces the same ordering, with high-exposure similarity exceeding low-exposure similarity in every case. Using terciles, the score correlation rises monotonically from 0.9822 in the bottom tercile to 0.9981 in the top tercile. Fourth, the residual conditional co-presence of Table 1, unlike the fundamentals-tracking similarity of Tables 2 and 3, carries model-relevant content, and it survives sampling uncertainty. A linear probability model of Burger-King presence on McDonald's presence, conditioning on population, density, income, urbanicity, and region with standard errors clustered by state, recovers the 11.9-point association of Section 2 as a coefficient of 0.121, with a state-cluster bootstrap 95 percent confidence interval of $[0.068, 0.178]$ over 1,000 resamples. The association is 0.092 ($[0.023, 0.166]$) in below-median AIGE counties and 0.117 ($[0.046, 0.188]$) in above-median counties. Their difference is 0.025, with a bootstrap interval of $[-0.077, 0.122]$ that contains zero (two-sided $p = 0.67$); a permutation test that reshuffles the AIGE-high label across counties, fixing the number of high-exposure counties, places the difference near the center of its null over 2,000 reshuffles, at a two-sided p-value near 0.62 whether the label is permuted freely or within Census region. The conditional co-presence is thus robustly positive, while any AIGE gradient in it is too small to sign. Common fundamentals, agglomeration, and firm learning would each generate a gradient of the sign shared processing predicts, so the split identifies no single channel, and the revision-response design of Appendix B carries that burden instead. Finally, excluding states with known collection gaps in OpenStreetMap, identified by comparing state-level counts against company-reported store totals, leaves all qualitative patterns unchanged.

B An Identification Exercise

The model implies a distinctive response of aggregate entry to platform revisions, formalized below as a partial-identification result. Suppose the shared platform issues a revision at date τ: a change in its processing layer that shifts the common component by $\Delta\xi \neq 0$ while fundamentals are unchanged. Examples include a model update, a data-source migration, or a recalibration of trade-area weights. In the decomposition of Section 3, the revision shifts the platform-wide component ξ^P, which is what makes the same revision common to both markets compared below. The revision is observed by the econometrician, since platforms announce major updates, but the magnitude $\Delta\xi$ is not.

Proposition 7 (Identification from platform revisions)**.** *Consider two markets that share the strategic and signal-structure primitives (δ, σ), indexed by their adoption rates $q_1 < q_2$, which differ through market-specific switching-cost distributions F_m or through equilibrium selection in the tipping regime, so that the mapping $K_\xi(\cdot)$ is common to both markets. Both markets are subject to the same platform revision $\Delta\xi$ at date τ, with fundamentals unchanged at τ in both. Then the model implies:*

(i) The revision shifts aggregate entry by $\Delta\bar{x}_\tau(q) = K_\xi(q)\,\Delta\xi$ in a market with adoption rate q.

(ii) The ratio of entry responses across the two markets, the revision-response ratio*,*

$$\frac{\Delta\bar{x}_\tau(q_2)}{\Delta\bar{x}_\tau(q_1)} = \frac{K_\xi(q_2)}{K_\xi(q_1)},$$

does not depend on the unobserved revision magnitude $\Delta\xi$ and is strictly greater than one.

If the econometrician also observes the adoption rates (q_1, q_2) and the signal-structure parameters up to scale, the ratio in part (ii) is informative about the strategic parameter: it identifies δ on $\mathscr{D}$ whenever the mapping $\delta \mapsto K_\xi(q_2;\delta)/K_\xi(q_1;\delta)$ is injective. Injectivity is not proved analytically; it is verified numerically across the admissible region, so point identification of δ is a numerical finding of generic identification rather than a theorem of global point identification. The obstruction is explicit: the ratio equals $q_2\,\Omega(q_1;\delta)/(q_1\,\Omega(q_2;\delta))$, and the derivative of $\Omega(q_1;\delta)/\Omega(q_2;\delta)$ in δ combines terms of both signs whose balance depends on the variance primitives, so global monotonicity in δ is not available. The result turns platform revisions into a source of identifying variation. Part (ii) is the practical content: a differential entry response across high- and low-adoption markets following a platform update is a signature of the shared-processing channel, and its magnitude is informative about the equilibrium exposure $K_\xi(q)$ without requiring the revision itself to be measured. The identifying assumption is that fundamentals do not shift discontinuously at revision dates, which is testable using pre-trends in markets with near-zero adoption. Adoption variation correlated with primitives across markets would confound the ratio, since K_ξ would then differ between markets for reasons other than q; the design requires adoption variation from switching costs or equilibrium selection, not from fundamentals. Implementation requires three data objects: entry or opening counts at high frequency around announced platform updates, a market-level measure of adoption intensity, for instance from procurement records or web-traffic data, and enough revisions to separate platform-driven shifts from coincident demand shocks. None is trivial, but each has precedents in the empirical platform literature.

C Cheap Switching and Transition Dynamics

The static adoption equilibrium of Section 6 treats the technology choice as frictionless within a single decision. In practice, chains revisit their forecasting arrangements episodically: contracts renew, analytics teams turn over, and procurement cycles create windows in which switching is cheap. This appendix embeds the static adoption map in a partial-adjustment process and asks how the speed of re-evaluation affects a discounted index of external exposure. Time is discrete. In each period, a fraction $\lambda \in (0,1]$ of chains may revisit their technology choice; the remainder hold their current technology. Adjustment is myopic: revisiting chains evaluate the current-

period adoption gain only, adopting the platform when their switching-cost draw is below it. This adaptive-adjustment assumption replaces the short-lived-agents justification of Section 8, and the no-history-learning assumption of that section carries over. The adoption rate then follows the partial-adjustment recursion

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

with q_0 given. The parameter λ indexes the cheapness of switching: high λ corresponds to frictionless re-evaluation, low λ to long procurement cycles. The object we track is a discounted exposure index along the transition path,

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

the discounted sum of the per-period marginal external cost of adoption under $\mathscr{W}_4$. It is an index, not a welfare loss, since it discounts the marginal wedge, which is no party's flow loss.

Proposition 8 (Cheap switching)**.** *Suppose $\delta \geq 0$, the platform-favorable condition $q_0 < q^*$ holds with $F(\Delta\Pi(q)) > q$ on $[q_0, q^*)$, q^* is the unique stable fixed point on $[q_0, 1]$, and $\Delta\Pi$ is non-decreasing on $[q_0, q^*]$. The last hypothesis holds analytically for δ in a neighborhood of zero (Proposition 2) and numerically throughout $\mathscr{D}$. Then:*

(i) The adoption path q_t is strictly increasing in λ at every $t \geq 1$, and converges to q^* for every $\lambda \in (0,1]$.

(ii) The discounted exposure index $\mathscr{L}(\lambda;\beta)$ is strictly increasing in λ for any $\beta \in (0,1)$.

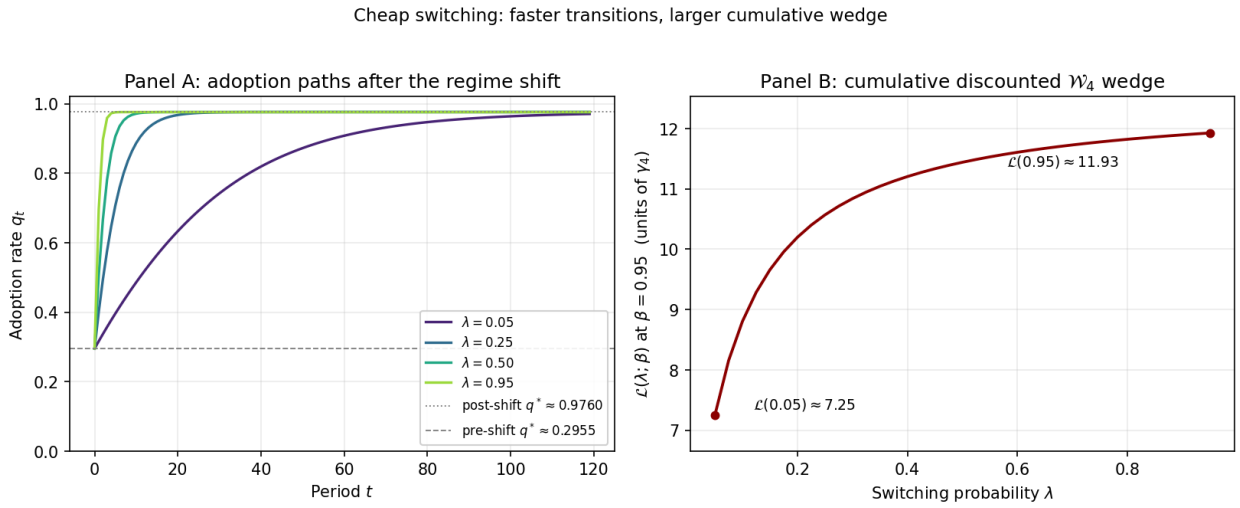
(iii) For $\phi \in (0,1)$, the stationary autocorrelation evaluated at the current adoption rate, $\rho_x(q_t,\phi) = \phi\, R(q_t)$, is strictly increasing in t; along the transition the entry process is nonstationary, so $\rho_x(q_t,\phi)$ is the fixed-q benchmark value rather than the path's realized lag-one correlation, and this benchmark builds as adoption diffuses.

The choice of w is a normalization within a class: because the adoption paths $q_t(\lambda)$ are pointwise ordered in λ by part (i), any per-period loss strictly increasing in q_t, including the level penalty $\gamma_4 q_t^2\, \sigma_\xi^2 / V_A$, inherits that order, so part (ii) holds verbatim with the level penalty in place of w. The result formalizes a policy-relevant asymmetry: institutional frictions that slow technology re-evaluation, such as long procurement cycles, delay the accumulation of external exposure. Cheap switching accelerates convergence to the high-adoption equilibrium and therefore frontloads the externality, raising its discounted total. Figure A1 illustrates with a precision-improving regime shift.

D Robustness to a Nonlinear Entry Rule

The linear equilibrium of Proposition 1 rests on the interior-regime entry rule $x_i = \frac{1}{2} + m_i$, with m_i the posterior mean. This appendix replaces it with the saturating logistic rule $x_i = \Lambda(4m_i)$, $\Lambda(z) = 1/(1+e^{-z})$, which matches the linear rule in level and slope at $m_i = 0$ and differs only through curvature: large posterior means are compressed into $[0,1]$ rather than extrapolated. For

Figure A1: Cheap switching after a precision-improving regime shift.



Notes: Both panels use the calibration $\sigma_\theta^2 = 1$, $\sigma_\xi^2 = 0.5$, $\sigma_C^2 = 0.6$, $\delta = 0.5$, $\gamma_4 = 1$, discount factor $\beta = 0.95$, and logistic switching-cost distribution $F(x) = (1+e^{-10x})^{-1}$. At $t = 0$ the platform's idiosyncratic noise falls from $\sigma_A^2 = 0.5$ to $\sigma_A^2 = 0.05$, raising the static adoption equilibrium from $q^*_{\text{pre}} \approx 0.296$, where the economy starts, to $q^*_{\text{post}} \approx 0.976$. Panel A plots the adoption path q_t under the partial-adjustment recursion $q_{t+1} = (1-\lambda)\,q_t + \lambda\,F(\Delta\Pi(q_t))$ for $\lambda \in \{0.05, 0.25, 0.50, 0.95\}$; higher λ steepens the transition toward q^*_{post} without changing its limit. Panel B plots the discounted cumulative external $\mathcal{W}_4$ wedge $\mathcal{L}(\lambda;\beta)$ over a grid of λ, in units of γ_4; it is strictly increasing in λ, from $\mathcal{L}(0.05) \approx 7.25$ to $\mathcal{L}(0.95) \approx 11.93$, a ratio of roughly 1.65, consistent with Proposition 8(ii).

each adoption rate q the aggregate $\bar{x}(\theta,\xi)$ solves a functional fixed point, which we compute by damped iteration on a tensor grid with Gauss-Hermite quadrature for all posterior integrals; effective loadings $\widehat{K}_\theta, \widehat{K}_\xi$ are the population projections of $\bar{x}$ on (θ,ξ), and $\widehat{\mathcal{X}_3}$ is obtained by a centered difference in q (`verify_nonlinear_robustness.py`). The two Theorem 1 objects behave differently under this replacement. The $\mathcal{W}_4$ external wedge $\mathcal{X}_4(q) = 2\gamma_4\, q\, \sigma_\xi^2 / V_A$ contains no equilibrium object: the penalty $q^2\sigma_\xi^2/V_A$ depends only on adoption and the signal structure, and $\mathcal{W}_1$ cancels in the wedge. It is therefore identical under any entry rule, exactly and by construction. The $\mathcal{W}_3$ wedge $\mathcal{X}_3(q) = \gamma_3 \frac{d}{dq}[K_\xi(q)^2\sigma_\xi^2]$ does depend on the rule through the equilibrium loadings, and it is the object the tables below stress-test. Table A1 reports $\mathcal{X}_3$ under the linear closed form and $\widehat{\mathcal{X}_3}$ under the logistic rule at five calibrations spanning strategic substitutes and complementarity ($\sigma_\theta^2 = 1$, $\sigma_A^2 = 0.2$, $\sigma_C^2 = 0.6$, $\gamma_3 = 1$). The wedge remains strictly positive in every cell, and $\widehat{K}_\xi(q)^2$ remains strictly increasing in q on $[0.1, 0.9]$ at both extreme calibrations, so the sign and monotonicity content of Theorem 1(i) survives saturation. The magnitude does not: the bounded rule attenuates the loadings, and the attenuation deepens with δ because for $\delta \neq 0$ the linear intercept $K_0 = 1/(2(1-\delta))$ places the average entry probability away from $1/2$, so the logistic rule operates on a flatter segment of Λ at any signal dispersion. Table A2 verifies that the attenuation is a saturation effect and that the linear closed form is the correct interior-regime limit. At $\delta = 0$ the operating point sits exactly at $1/2$, and scaling every variance by a common factor c shrinks the dispersion of posterior means relative to the saturation scale; the logistic-to-linear ratio

Table A1: External $\mathscr{W}_3$ wedge under linear and logistic entry rules

δ	σ_ξ^2	q	Linear $\mathscr{X}_3$	Logistic $\widehat{\mathscr{X}_3}$	Ratio
-0.5	0.3	0.5	0.0669	0.0207	0.310
-0.2	0.5	0.5	0.1273	0.0332	0.261
0.0	0.5	0.5	0.1730	0.0353	0.204
0.3	0.5	0.5	0.3000	0.0362	0.121
0.7	0.6	0.6	1.2334	0.0435	0.035

Notes: Calibration: $\sigma_\theta^2 = 1$, $\sigma_A^2 = 0.2$, $\sigma_C^2 = 0.6$, $\gamma_3 = 1$. The logistic rule is $\Lambda(4m)$, level- and slope-matched to the linear rule at $m = 0$. The logistic column is computed from the numerical fixed point by centered differences in q (`verify_nonlinear_robustness.py`).

rises monotonically toward one as $c \downarrow 0$, reaching 0.991 at $c = 0.002$. The qualitative content of

Table A2: Scale convergence of the $\mathscr{W}_3$ wedge at $\delta = 0$

Scale c	Linear $\mathscr{X}_3$	Logistic $\widehat{\mathscr{X}_3}$	Ratio
1.000	0.173010	0.035255	0.204
0.100	0.017301	0.012026	0.695
0.020	0.003460	0.003168	0.915
0.002	0.000346	0.000343	0.991

Notes: All variances are scaled by c: $\sigma_\theta^2 = c$, $\sigma_\xi^2 = 0.5c$, $\sigma_A^2 = 0.2c$, $\sigma_C^2 = 0.6c$; $\delta = 0$, $q = 0.5$, $\gamma_3 = 1$ (`verify_nonlinear_robustness.py`).

Theorem 1 therefore survives replacing the linear rule: $\mathscr{X}_4$ is exactly rule-invariant, $\mathscr{X}_3$ retains its sign and its monotonicity in q, and only the $\mathscr{X}_3$ magnitude is rule-sensitive, converging to the linear closed form in the interior-regime limit.

E A Screening Foundation for the Decided-About-Market Externality

The $\mathscr{W}_4$ penalty is the exact second moment of the screening environment developed here. The decided-about agents are candidate markets.

E.1 Unified Environment

Extend the baseline environment as follows. Alongside the entry game of Section 3, there is a unit continuum of candidate markets indexed by $m \in [0,1]$, each characterized by an independent draw $(\theta_m, \xi_m, \{u_{im}\})$ from the same distributions as the baseline. Each market is screened by $n \geq 2$ chains drawn independently at random from the population, with each chain applying its equilibrium forecasting technology: a platform adopter evaluating market m observes $\theta_m + \xi_m +$

u^A_{im}, and a custom chain observes $\theta_m + u^C_{im}$. The evaluation outcome for market m is the profile of evaluation errors of its n evaluators, defined next. The independent-across-markets draw of ξ_m is a normalization on the application error ν_m in the decomposition $\xi_m = \xi^P + \nu_m$ of Section 3: the screening exposure requires only within-market commonality of ξ_m across co-adopting evaluators, so it is invariant to the split. Evaluators' compensation for screening is separable from entry profits and identical across forecasting technologies, a per-evaluation fee independent of the outcome, so screening returns cancel in $\Delta\Pi$ and the adoption margin of the main text is unchanged; technology-dependent screening returns, which would add a term to $\Delta\Pi$, are an open extension. The welfare of decided-about markets is measured by the expected aggregate error covariance across evaluator pairs:

$$S(q) \equiv \mathbb{E}\left[\sum_{i<j} \frac{\mathrm{Cov}(e_{im}, e_{jm} \mid \text{technologies})}{V_A}\right],$$

where $e_{im} \equiv \tilde{s}_{im} - \theta_m$ is evaluator i's signal error in market m, equivalently its evaluation error conditional on market quality, and the expectation is over the random assignment of evaluators to markets. The signal error is the relevant object because the recommendation event below operates on the signal itself and the welfare object conditions on θ_m.

E.2 The Exposure Is an Exact Second Moment

For a platform adopter, the signal error is $e_{im} = \xi_m + u^A_{im}$; for a custom chain it is $e_{im} = u^C_{im}$. Conditional on θ_m and technologies, two adopters' signal errors therefore have covariance σ^2_ξ through the shared component ξ_m (normalized covariance σ^2_ξ / V_A), and every pair involving a custom chain has zero conditional covariance. Posterior-mean forecast errors would instead add a common prior-shrinkage covariance through θ_m to every evaluator pair, including custom-custom pairs; that component is fundamentals-tracking, and the criterion deliberately does not price it. Because σ^2_ξ / V_A contains no filtering weights, $S(q)$ is also invariant to how evaluators map signals into beliefs. With n evaluators drawn independently, the number of adopter pairs is binomial, so

$$S(q) = \binom{n}{2} q^2 \frac{\sigma^2_\xi}{V_A},$$

exactly, for every $n \geq 2$. The baseline penalty is the case $n = 2$ with the constant $\binom{2}{2} = 1$; general n rescales the constant, which is absorbed into γ_4. The normalization by V_A rather than by the conditional error variance $\sigma^2_\xi + \sigma^2_A$ is the convention discussed in Section 5; the two differ by a factor free of q, so every statement in this appendix holds under either.

E.3 Why the Covariance Is a Cost

A mean-zero covariance is not yet a welfare loss until a decision maps evaluation errors into outcomes. We attach the minimal such decision and show that the covariance is a monotone index of the resulting loss, so that $\gamma_4 > 0$ is a property of the screening problem rather than a free sign. Endow each candidate market with a latent quality and a social value $v > 0$ that is realized if at least one of its n evaluators forecasts favorably, that is, if that evaluator's signal clears a common rec-

ommendation threshold κ. For a good market, an evaluator's signal error is what pushes its signal below κ and causes a missed recommendation, so the market is overlooked only when *every* evaluator errs in the same downward direction. Write the coverage probability, the chance a good market is recognized by some evaluator, as $C(q) = 1 - \mathbb{P}(\text{all } n \text{ evaluators fall short})$. Among two platform adopters, the signal errors are jointly Gaussian with covariance $\sigma_\xi^2 \geq 0$ at fixed marginal variances, and by Slepian's inequality the probability that both fall below a common threshold is increasing in that covariance at fixed marginals. Hence $C(q)$ is decreasing in the error covariance, and the expected coverage loss $v[C^{\text{indep}} - C(q)]$ is increasing in it. The argument is two-sided, and stating the loss function explicitly makes the class precise: take the social loss to be convex in the number of coincident evaluation errors of either sign, so that it penalizes both the event that every evaluator misses a good market and the event that every evaluator wrongly favors a bad one. Slepian's inequality applies to both tails: at fixed marginals, higher error covariance raises the probability that all evaluators' signals fall below the threshold for a good market and the probability that all clear the threshold for a bad one. One movement is countervailing under other loss functions: for a bad market, higher correlation lowers the probability that at least one evaluator wrongly enters, which is a benefit if the loss attaches to any single wrong entry rather than to coincident error. The convex-in-coincident-errors class is exactly the class under which both tail movements are costly, and it is the class we impose. It is less restrictive than a symmetric menu of signs. Under the natural additive loss, in which each wrongly entered market costs its own capital, expected loss is a sum over evaluators of marginal error probabilities, depends on marginals alone, and is exactly invariant to the error correlation; for correlation to help on the entry tail, the loss must be strictly concave in the number of coincident wrong entries, so that one wasted store and two are nearly equivalent. On the miss tail no discretion arises: when a market's value is realized if at least one evaluator recognizes it, the induced loss is one minus the coverage probability, which is convex in coincident misses by construction. Weak convexity therefore delivers $\gamma_4 > 0$. The covariance $S(q)$ is the leading-order index of this loss: to first order in the small-correlation regime, the reduction in coverage probability is proportional to the pairwise error covariance, with a positive constant that a full specification of (v, κ) would pin down and that is absorbed into γ_4. Correlated evaluation is therefore costly because it destroys the option value of an independent second opinion, exactly the diversification motive of Kleinberg and Raghavan (2021), and $S(q)$ measures how much of that option value adoption removes. This delivers the direction of the penalty from primitives; its scale, like every welfare level in the paper, remains ordinal in γ_4. The claim carries two scope restrictions. First, $S(q)$ is exact as a second moment for every n, but its mapping to the coverage loss is derived to leading order in the small-correlation regime; the exactness attaches to the moment, not to the level of the loss. Second, the argument isolates the correlation channel at fixed marginals: it holds each evaluator's error variance fixed and prices only the loss of independence, so it does not net the higher per-evaluator precision that adopters enjoy when $\sigma_A^2 < \sigma_C^2$ against the second opinion they forgo. A full best-of-n comparison between a pair of precise-but-correlated adopters and a pair of noisier-but-independent custom evaluators would net these forces; we isolate the correlation force because it is the one the shared component introduces. The comparison is left uncomputed for a stated reason: the coverage probabilities it requires are bivariate normal orthant probabilities without closed form, and their difference depends on the recognition threshold κ and the market value v, objects the entry model does not restrict, so the netted boundary in $(\sigma_C^2 - \sigma_A^2, \sigma_\xi^2)$ space is a calibration exercise for work that pins down (v, κ); it is also substantially the computation of

Kleinberg and Raghavan (2021) in their setting.

E.4 The Wedge as a Proposition

Proposition 9 (Screening wedge)**.** *In the unified environment, define complete welfare as* $\mathscr{W}_1(q) - \gamma_4 S(q)$ *with* $\gamma_4 > 0$ *pricing the social cost of correlated screening errors. Then the external wedge is*

$$\frac{d\mathscr{W}_1}{dq} - \frac{d}{dq}\Big[\mathscr{W}_1(q) - \gamma_4 S(q)\Big] = \gamma_4 S'(q) = \gamma_4 \binom{n}{2} \frac{2q\,\sigma_\xi^2}{V_A} > 0$$

for all $q \in (0,1]$*, every* $n \geq 2$*, and every* $\delta \in \mathscr{D}$*. The wedge contains no equilibrium loading and is therefore invariant to* δ *at any fixed* q*.*

The screening externality runs through the residual covariance of signals, which is a property of the information structure alone. Strategic interaction affects how firms act on their forecasts, and would re-enter the screening problem only through cross-market linkages in payoffs, at which point it would operate through the same equilibrium-loading channel as the aggregate-volatility wedge $\mathscr{X}_3$. The two externalities are therefore structurally distinct, and the dependence on δ is precisely what separates them. The invariance is a property of evaluation-attached harm; a loss defined on realized entry outcomes would carry the entry rule and with it δ.

F Analytics Vendors and Disclosed Restaurant Clients

This appendix compiles publicly disclosed restaurant-chain clients of major site-selection analytics vendors, drawn from vendor case studies, press releases, and industry coverage. Disclosure is voluntary and selected, so the compilation indicates market presence rather than market shares. The

Table A3: Disclosed restaurant clients by analytics vendor

Vendor	Disclosed restaurant clients (selected)
Placer.ai	Wendy's, Chipotle, Dave's Hot Chicken, First Watch
Buxton	Chick-fil-A (historical), Checkers, Marco's Pizza, Freddy's
Esri	Starbucks (documented case), Subway (via partners)
SiteZeus	Burger King franchisees, Jersey Mike's, Checkers
Kalibrate	KFC (international), Pilot Flying J food service

Notes: Compiled from vendor case studies, press releases, and trade coverage through 2025. Client lists are illustrative, not exhaustive; disclosure is voluntary and may lag or overstate current engagements. Several chains appear in multiple vendors' materials over time, consistent with switching or multi-homing. The compilation motivates the empirically relevant range $J \in \{4,5,6\}$ used in Section 9.

same handful of vendors recurs across disclosed engagements, consistent with a concentrated market for substantive platforms. Several chains also appear in more than one vendor's materials, and vendors advertise overlapping data inputs (mobile-location panels, Census-derived demographics) and similar methodologies (trade-area gravity models, analog forecasting), consistent with positive cross-vendor correlation ρ.

G Proofs

This appendix collects proofs of the formal results. Throughout, $\sigma \equiv (\sigma_\theta^2, \sigma_\xi^2, \sigma_A^2, \sigma_C^2)$ denotes the vector of variance primitives, $V_A = \sigma_\theta^2 + \sigma_\xi^2 + \sigma_A^2$, $V_C = \sigma_\theta^2 + \sigma_C^2$, and $\alpha_\theta^A, \alpha_\xi^A, \alpha_\theta^C, \chi_A$ are the filtering coefficients of Section 3.

Proof of Proposition 1

A platform adopter with signal $\tilde{s}_i^A$ enters with probability $\frac{1}{2} + \alpha_\theta^A \tilde{s}_i^A + \delta \mathbb{E}_i^A[\bar{x}]$, and a custom chain with $\frac{1}{2} + \alpha_\theta^C \tilde{s}_i^C + \delta \mathbb{E}_i^C[\bar{x}]$. Conjecture $\bar{x} = K_0 + K_\theta \theta + K_\xi \xi$. An adopter computes $\mathbb{E}_i^A[\bar{x}] = K_0 + K_\theta \mathbb{E}_i^A[\theta] + K_\xi \mathbb{E}_i^A[\xi] = K_0 + (K_\theta \alpha_\theta^A + K_\xi \alpha_\xi^A)\tilde{s}_i^A$, and a custom chain $\mathbb{E}_i^C[\bar{x}] = K_0 + K_\theta \alpha_\theta^C \tilde{s}_i^C$, since a custom chain assigns ξ its prior mean of zero. Matching the coefficient on each own-signal gives the slopes

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

Consistency in the continuum requires $K_\theta = qk^A + (1-q)k^C$ and $K_\xi = qk^A$, since idiosyncratic noise averages out and only adopters load on ξ. Substituting K_θ and K_ξ into the two slope equations yields a linear system in (k^A, k^C). Solving,

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

with $\Omega(q) = (1 - \delta q \chi_A)(1 - \delta(1-q)\alpha_\theta^C) - \alpha_\theta^A \alpha_\theta^C \delta^2 q(1-q)$, where $\chi_A = \alpha_\theta^A + \alpha_\xi^A$. The intercept satisfies $K_0 = \frac{1}{2} + \delta K_0$, so $K_0 = 1/(2(1-\delta))$. For existence and uniqueness on $\mathscr{D}$, note Ω is quadratic in q with leading coefficient $\delta^2 \alpha_\theta^C(\alpha_\theta^A - \chi_A) = -\delta^2 \alpha_\theta^C \alpha_\xi^A \leq 0$, so Ω is concave in q. At the endpoints, $\Omega(0) = 1 - \delta\alpha_\theta^C$ and $\Omega(1) = 1 - \delta\chi_A$. The two bounds defining $\mathscr{D}$ are equivalent to $\delta\alpha_\theta^C < 1$ and $\delta\chi_A < 1$: for the first, $\delta < 1 + \sigma_C^2/\sigma_\theta^2 = V_C/\sigma_\theta^2 = 1/\alpha_\theta^C$; for the second, $\delta < 1 + \sigma_A^2/(\sigma_\theta^2 + \sigma_\xi^2) = V_A/(\sigma_\theta^2 + \sigma_\xi^2) = 1/\chi_A$. Hence $\Omega(0) > 0$ and $\Omega(1) > 0$, and concavity gives $\Omega(q) > 0$ on all of $[0,1]$. Positivity of the denominator makes the linear system nonsingular, so the linear equilibrium is unique. □

Proof of Proposition 2

Expected profit for a type with entry probability $p_i = \frac{1}{2} + m_i$, where m_i is the posterior mean of $\theta + \delta\bar{x}$ and p_i denotes the probability rather than the binary decision x_i, equals $\pi = \frac{1}{2}\mathbb{E}[p_i^2]$. To see this, condition on the signal: with $\varepsilon_i \sim \text{Uniform}[-\frac{1}{2}, \frac{1}{2}]$ and entry whenever $m_i + \varepsilon_i \geq 0$, expected profit conditional on m_i is $\int_{-m_i}^{1/2} (m_i + \varepsilon)\, d\varepsilon = \frac{1}{2}(m_i + \frac{1}{2})^2 = \frac{1}{2}p_i^2$, using that the payoff-relevant surprise is mean-independent of ε_i. Taking expectations over signals gives $\pi = \frac{1}{2}\mathbb{E}[p_i^2] = \frac{1}{2}(K_0^2 + \text{loading}^2 \cdot \text{variance})$, where the loading-variance term is $(k^A)^2 V_A$ for adopters and $(k^C)^2 V_C$ for custom chains, since an adopter's signal has variance V_A and a custom chain's has variance V_C. The common intercept K_0 cancels in the difference, so

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

Substituting the closed forms from Proposition 1 and simplifying over the common denominator $\Omega(q)^2$,

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

where Γ is a continuous function of (q,δ,σ) built from the primitives and from $\Omega(q) > 0$. At $\delta = 0$, $\Omega(q) \equiv 1$ and k^A, k^C are constants, so $\Delta\Pi$ is independent of q; moreover $\Gamma(q;\sigma,0)$ evaluates in closed form to the q-free expression

$$\Gamma(q;\sigma,0) \;=\; \frac{V_A\left[\sigma_\theta^2\left(\sigma_A^2-\sigma_C^2\right)^2 \;+\; \sigma_\xi^2\left(\sigma_C^4+\sigma_\theta^2\left(2\sigma_A^2+\sigma_\theta^2+\sigma_\xi^2\right)\right)\right]}{2\,\sigma_\xi^2},$$

in which the first bracketed term is a nonnegative square scaled by σ_θ^2 and the second is a sum of products of strictly positive primitives, so $\Gamma(q;\sigma,0) > 0$. By continuity of Γ in δ, uniformly over $q \in [0,1]$ at fixed σ, there is a neighborhood of $\delta = 0$ within $\mathscr{D}$ on which $\Gamma > 0$, and on that neighborhood the sign of $d\Delta\Pi/dq$ equals the sign of δ, which proves the proposition. Positivity of Γ over the full admissible region is a sign condition verified numerically in the replication package over 200,000 parameter draws spanning $\mathscr{D}$, and this global extension is not part of the proved statement. □

Proof of Lemma 1

Both objectives contain $\mathscr{W}_1$ identically, so $\mathscr{W}_1$ cancels in the difference of derivatives and $\mathscr{X}(q) = \gamma\mathscr{P}'(q)$. (i) If $\mathscr{P}$ is a function of (q,σ) alone, then $\mathscr{P}'$ contains no equilibrium object; the strategic parameter, the entry rule, and belief formation enter the model only through equilibrium behavior, so none of them moves the wedge at fixed q, and strict monotonicity of $\mathscr{P}$ gives $\mathscr{X}(q) > 0$. (ii) If $\mathscr{P}$ depends on equilibrium objects, the identity $\mathscr{X} = \gamma\mathscr{P}'$ still holds, so the wedge is strictly positive wherever the penalized exposure is strictly increasing in q, while its magnitude inherits the strategic environment through those objects. □

Proof of Proposition 3

Chain welfare is $\mathscr{W}_1(q) = q\pi^A(q) + (1-q)\pi^C(q)$. Using $\pi^A = \frac{1}{2}(K_0^2 + (k^A)^2 V_A)$ and $\pi^C = \frac{1}{2}(K_0^2 + (k^C)^2 V_C)$,

$$\mathscr{W}_1(q) = \tfrac{1}{2}K_0^2 + \tfrac{1}{2}\left[q(k^A)^2 V_A + (1-q)(k^C)^2 V_C\right].$$

Differentiating and grouping the direct term (the change in the technology mix at fixed loadings) against the indirect term (the equilibrium response of the loadings),

$$\frac{d\mathscr{W}_1}{dq} = \tfrac{1}{2}\left[(k^A)^2 V_A - (k^C)^2 V_C\right] + \tfrac{1}{2}\left[q\frac{d}{dq}(k^A)^2 V_A + (1-q)\frac{d}{dq}(k^C)^2 V_C\right].$$

The first bracket is $\Delta\Pi(q)$. The second bracket equals $\frac{\delta}{2}\frac{d}{dq}\mathrm{Var}(\bar{x})$. Each chain's entry rule is a best response to the law of motion of $\bar{x}$, so by the envelope theorem a marginal change in q moves a chain's expected profit only through that law of motion: $d\pi_i/dq = \delta\,\mathbb{E}[x_i\,\partial\bar{x}/\partial q]$. Integrating over

the population, $\int_0^1 \mathbb{E}[x_i\,\partial\bar{x}/\partial q]\,di = \mathbb{E}[\bar{x}\,\partial\bar{x}/\partial q] = \frac{1}{2}\,\partial\mathbb{E}[\bar{x}^2]/\partial q$, and since the intercept K_0 does not depend on q, $\partial\mathbb{E}[\bar{x}^2]/\partial q = \frac{d}{dq}\mathrm{Var}(\bar{x})$, where $\mathrm{Var}(\bar{x}) = K_\theta^2\sigma_\theta^2 + K_\xi^2\sigma_\xi^2$ with $K_\theta = qk^A + (1-q)k^C$ and $K_\xi = qk^A$. The indirect term is therefore $\frac{\delta}{2}\frac{d}{dq}\mathrm{Var}(\bar{x})$. Writing $\Psi(q) \equiv \frac{\delta}{2}\frac{d}{dq}\mathrm{Var}(\bar{x})$ gives $d\mathscr{W}_1/dq = \Delta\Pi(q) + \Psi(q)$. When $\delta = 0$, $\Psi \equiv 0$ and the two coincide. When $\delta \neq 0$, $\Psi(q)$ vanishes only where $\frac{d}{dq}\mathrm{Var}(\bar{x}) = 0$; the private equilibrium condition $F^{-1}(q^*) = \Delta\Pi(q^*)$ and the planner's interior first-order condition $F^{-1}(q) = d\mathscr{W}_1/dq$ then differ by $\Psi(q)$, so the private equilibrium can coincide with an interior chain-welfare optimum only at a point where $\Psi(q^*) = 0$, and $\mathrm{sgn}(\Psi) = \mathrm{sgn}(\delta \cdot \frac{d}{dq}\mathrm{Var}(\bar{x}))$. □

Proof of Theorem 1

For $k \in \{3,4\}$ write $\mathscr{W}_k(q) = \mathscr{W}_1(q) - \gamma_k\,\mathscr{P}_k(q)$, where $\mathscr{P}_3(q) = K_\xi(q)^2\sigma_\xi^2$ and $\mathscr{P}_4(q) = q^2\sigma_\xi^2/V_A$ are the penalty terms. By Lemma 1, the external wedge is

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

and the strategic term Ψ cancels because it appears in $d\mathscr{W}_1/dq$ on both sides of the wedge definition through $\mathscr{W}_1$. Part (i). $\mathscr{P}_3(q) = q^2(k^A(q))^2\sigma_\xi^2 = K_\xi(q)^2\sigma_\xi^2$. Since $k^A(q) = \alpha_\theta^A/\Omega(q) > 0$ on $\mathscr{D}$ and $K_\xi(q) = qk^A(q)$, we have $K_\xi(0) = 0$ and $K_\xi(q) > 0$ for $q > 0$. Differentiating, $\mathscr{P}_3'(q) = 2K_\xi(q)K_\xi'(q)\sigma_\xi^2$. Writing $\Omega(q) = a + bq + cq^2$ with $a = 1 - \delta\alpha_\theta^C$ and $c = -\delta^2\alpha_\theta^C\alpha_\xi^A$ (proof of Proposition 1), the quotient rule gives $K_\xi'(q) = \alpha_\theta^A\,[\Omega(q) - q\Omega'(q)]/\Omega(q)^2 = \alpha_\theta^A\,(a - cq^2)/\Omega(q)^2$, and substituting the filtering coefficients yields the identity

$$K_\xi'(q) \;=\; \frac{\left[V_A(V_C - \delta\sigma_\theta^2) + \delta^2\sigma_\theta^2\sigma_\xi^2\,q^2\right]\alpha_\theta^A}{V_A\,V_C\,\Omega(q)^2} \;>\; 0$$

on $\mathscr{D}$, where $V_C - \delta\sigma_\theta^2 > 0$ follows from $\delta < V_C/\sigma_\theta^2$. Hence $\mathscr{P}_3'(q) > 0$ for $q \in (0,1]$, so $\mathscr{X}_3(q) = \gamma_3\mathscr{P}_3'(q) > 0$. The sign is invariant to δ; the magnitude varies with δ through the loadings. Part (ii). $\mathscr{P}_4(q) = q^2\sigma_\xi^2/V_A$ contains no equilibrium object, so $\mathscr{P}_4'(q) = 2q\sigma_\xi^2/V_A$ and $\mathscr{X}_4(q) = 2\gamma_4 q\sigma_\xi^2/V_A > 0$ for $q \in (0,1]$, independent of δ at every fixed q. For the ordering of optima, a pointwise comparison of first-order conditions does not suffice, because neither objective is assumed concave; we argue by monotone comparative statics. By parts (i) and (ii), $\mathscr{P}_k'(q) > 0$ on $(0,1]$, so $\mathscr{P}_k$ is strictly increasing on $[0,1]$. Let q_k maximize $\mathscr{W}_k$ and q_1 maximize $\mathscr{W}_1$ on $[0,1]$, and suppose $q_k > q_1$. Optimality gives $\mathscr{W}_1(q_k) \le \mathscr{W}_1(q_1)$ and $\mathscr{W}_k(q_1) \le \mathscr{W}_k(q_k)$; summing the two inequalities and cancelling $\mathscr{W}_1$ yields $\gamma_k\,[\mathscr{P}_k(q_k) - \mathscr{P}_k(q_1)] \le 0$, contradicting strict monotonicity. Hence every maximizer of $\mathscr{W}_k$ lies weakly below every maximizer of $\mathscr{W}_1$: the two maximizer sets are ordered in the strong set order. The same argument applies verbatim to the net objectives $\widetilde{\mathscr{W}_k} = \mathscr{W}_k - \int_0^q F^{-1}(s)\,ds$ and $\widetilde{\mathscr{W}_1}$, whose maximizers are the objects denoted q_k^{soc} and q_1^{soc}: the common switching-cost term cancels when the optimality inequalities are summed, exactly as $\mathscr{W}_1$ does. For strictness under interiority, suppose $q_k \in (0,1)$ maximizes $\widetilde{\mathscr{W}_k}$. The first-order condition $\widetilde{\mathscr{W}_k}'(q_k) = 0$ gives $\widetilde{\mathscr{W}_1}'(q_k) = \gamma_k\,\mathscr{P}_k'(q_k) > 0$, so $\widetilde{\mathscr{W}_1}$ is strictly increasing at q_k and q_k cannot maximize it; combined with the weak ordering, $q_k^{soc} < q_1^{soc}$. At a common corner the two optima can coincide, which

is why the strict statement carries the interiority condition. Finally, subtracting the two first-order conditions, $\Delta\Pi(q) - d\mathscr{W}_k/dq = [\Delta\Pi - d\mathscr{W}_1/dq] + [d\mathscr{W}_1/dq - d\mathscr{W}_k/dq] = -\Psi(q) + \mathscr{X}_k(q)$, which gives the stated decomposition. At a stable interior adoption equilibrium q^*, $F^{-1}(q^*) = \Delta\Pi(q^*)$, so the complete planner's objective net of switching costs, $\mathscr{W}_k(q) - \int_0^q F^{-1}(s)\,ds$, has derivative $-[\mathscr{X}_k(q^*) - \Psi(q^*)]$ at q^*. When $\Psi(q^*) \leq 0$ this derivative is at most $-\mathscr{X}_k(q^*) < 0$, so a marginal reduction in adoption below q^* raises complete welfare for any $\gamma_k > 0$; by Proposition 3, $\Psi \leq 0$ under strategic substitutes exactly when $\frac{d}{dq}\mathrm{Var}(\bar{x}) \geq 0$. With multiple stable equilibria the same argument applies at each stable fixed point separately. The comparison reverses only if $\Psi > \mathscr{X}_k$, which requires $\Psi > 0$. □

Proof of Proposition 4

Under $\xi = \ell\theta + \zeta$ with $\zeta \perp \theta$, the platform signal is $\tilde{s}_i^A = (1+\ell)\theta + \zeta + u_i^A$ with variance $V_A(\ell) = (1+\ell)^2\sigma_\theta^2 + \sigma_\zeta^2 + \sigma_A^2$. The residual shared component across two adopters is ζ, whose normalized error covariance is $\sigma_\zeta^2/V_A(\ell)$; the informative part $\ell\theta$ is absorbed into fundamentals and does not contribute to the correlated-error penalty, which conditions on θ. Hence the $\mathscr{W}_4$ penalty becomes $\mathscr{P}_4(q) = q^2\sigma_\zeta^2/V_A(\ell)$, and by the argument of Theorem 1(ii),

$$\mathscr{X}_4(q) = \gamma_4\mathscr{P}_4'(q) = \frac{2\gamma_4 q\sigma_\zeta^2}{V_A(\ell)}.$$

Claims (i)-(iii) follow: the wedge is strictly positive iff $\sigma_\zeta^2 > 0$; holding σ_ζ^2 fixed, raising ℓ increases $V_A(\ell)$ and therefore lowers the wedge, which vanishes only as $\sigma_\zeta^2 \to 0$; and no equilibrium loading appears, so the wedge is δ-invariant at fixed q. Claim (ii) is the one statement that depends on normalizing by $V_A(\ell)$ rather than by the conditional error variance $\sigma_\zeta^2 + \sigma_A^2$; under the latter the wedge does not depend on ℓ at all, while claims (i) and (iii) are unaffected. □

Proof of Proposition 5

Under Assumption 1, each adopter filters as in the single-platform model, so the loadings are those of Proposition 1. With J platforms and shares $\{q_j\}$, only co-adopters of the same or correlated platforms share error components. The aggregate loading on platform j's shock is $K_{\xi,j} = q_j k^A$, and $\mathrm{Var}(\sum_j K_{\xi,j}\xi_j) = (k^A)^2\sigma_\xi^2[\sum_j q_j^2 + 2\rho\sum_{j<j'} q_j q_{j'}]$ using $\mathrm{Var}(\xi_j) = \sigma_\xi^2$ and $\mathrm{Corr}(\xi_j, \xi_{j'}) = \rho$. Under symmetry $q_j = q/J$, the bracket equals $q^2 H(J,\rho)$ with $H(J,\rho) = 1/J + (1-1/J)\rho$, so the platform-driven variance is $K_\xi(q)^2\sigma_\xi^2 H(J,\rho)$. The $\mathscr{W}_3$ wedge is therefore scaled by $H(J,\rho)$ relative to the single-platform case. For $\mathscr{W}_4$, the expected normalized error covariance between two randomly drawn adopters is σ_ξ^2/V_A times the expected correlation of their shared components, which equals $1/J$ (the probability of drawing the same platform) plus $(1-1/J)\rho$ (the correlation across distinct platforms), namely $H(J,\rho)\sigma_\xi^2/V_A$; multiplying by the co-adoption probability q^2 and differentiating gives the stated wedge. Both wedges are strictly positive for $\delta \in \mathscr{D}$, $J \geq 1$, $\rho \in [0,1]$, and decrease in J whenever $\rho < 1$ because H is strictly decreasing in J for $\rho < 1$. □

Proof of Lemma 2

At $q=0$, $\Omega(0)=1-\delta\alpha_\theta^C$ and $k^A(0)=\alpha_\theta^A/(1-\delta\alpha_\theta^C)$, $k^C(0)=\alpha_\theta^C/(1-\delta\alpha_\theta^C)$. From the proof of Proposition 2, $\Delta\Pi(0)=\frac{1}{2}[(k^A(0))^2V_A-(k^C(0))^2V_C]$. Now $(k^A(0))^2V_A=(\alpha_\theta^A)^2V_A/(1-\delta\alpha_\theta^C)^2=(\sigma_\theta^2)^2/[V_A(1-\delta\alpha_\theta^C)^2]$ and $(k^C(0))^2V_C=(\sigma_\theta^2)^2/[V_C(1-\delta\alpha_\theta^C)^2]$. Hence $\Delta\Pi(0)=0$ iff $1/V_A=1/V_C$, i.e. $V_A=V_C$, i.e. $\sigma_\theta^2+\sigma_\xi^2+\sigma_A^2=\sigma_\theta^2+\sigma_C^2$, i.e. $\sigma_\xi^2=\sigma_C^2-\sigma_A^2$, which does not involve δ. □

Proof of Proposition 6

With $\bar{x}_t-\mathbb{E}[\bar{x}_t]=K_\theta(q)\theta_t+K_\xi(q)\xi_t$ (the constant intercept K_0 drops from all second moments), θ_t i.i.d. and ξ_t a stationary AR(1) with $\mathrm{Var}(\xi_t)=\sigma_\xi^2$ and $\mathrm{Cov}(\xi_t,\xi_{t+1})=\phi\sigma_\xi^2$, and $\theta\perp\xi$:

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

since the i.i.d. fundamental contributes zero autocovariance. The autocorrelation is $\rho_x(q,\phi)=\phi K_\xi^2\sigma_\xi^2/(K_\theta^2\sigma_\theta^2+K_\xi^2\sigma_\xi^2)=\phi R(q)$, where $R(q)$ is the shared-variance share. $R(q)$ is strictly increasing in q on $(0,1]$. Since R is a strictly increasing transformation of the ratio K_ξ/K_θ, and both loadings are positive on $\mathscr{D}$ ($k^A>0$, and $k^C\propto 1-\delta q\alpha_\xi^A>0$ because $\delta q\alpha_\xi^A<\delta\chi_A<1$ for $\delta>0$ and is negative otherwise), it suffices that $K_\theta/K_\xi=1+(1-q)k^C/(qk^A)$ is strictly decreasing. Using the closed forms of Proposition 1, $(1-q)k^C/(qk^A)=(\alpha_\theta^C/\alpha_\theta^A)\,h(q)$ with $h(q)=(1-q)(1-\delta\alpha_\xi^A q)/q$, and $h'(q)=-1/q^2+\delta\alpha_\xi^A<0$ on $\mathscr{D}$: for $\delta\le 0$ the inequality is immediate, and for $\delta>0$ it follows from $\delta\alpha_\xi^A<\delta\chi_A<1\le 1/q^2$. Hence $R'(q)>0$, with $R(0)=0$. All claims of part (iii) then follow from the identity $\rho_x=\phi\,R(q)$ with $R'(q)>0$ and $R(q)>0$ on $(0,1]$: for $\phi>0$, ρ_x is strictly increasing in q and in ϕ; for $\phi<0$ it is strictly decreasing in q; and $|\rho_x|=|\phi|R(q)$ is strictly increasing in q for every $\phi\neq 0$. □

Proof of Proposition 7

(i) A revision shifts the shared component by $\Delta\xi$ with fundamentals fixed. Since $\bar{x}=K_0+K_\theta\theta+K_\xi\xi$, the change in aggregate entry is $\Delta\bar{x}_\tau(q)=K_\xi(q)\Delta\xi$. (ii) Taking the ratio across two markets subject to the same revision, $\Delta\bar{x}_\tau(q_2)/\Delta\bar{x}_\tau(q_1)=K_\xi(q_2)/K_\xi(q_1)$, in which $\Delta\xi$ cancels. Because $K_\xi(q)=qk^A(q)$ is strictly increasing on $(0,1]$ (proof of Theorem 1) and $q_2>q_1$, the ratio exceeds one. The injectivity underlying the point-identification discussion is verified numerically across $\mathscr{D}$ in the replication package and is not part of the proved statement. □

Proof of Proposition 8

(i) Under the recursion $q_{t+1}=(1-\lambda)q_t+\lambda F(\Delta\Pi(q_t))$ with $F(\Delta\Pi(q))>q$ on $[q_0,q^*)$, write $g(q)\equiv F(\Delta\Pi(q))-q>0$ on that interval. Then $q_{t+1}-q_t=\lambda g(q_t)$. By hypothesis, $\Delta\Pi$ is nondecreasing on $[q_0,q^*]$, so the adjustment map $q\mapsto(1-\lambda)q+\lambda F(\Delta\Pi(q))$ has derivative $(1-\lambda)+\lambda F'\Delta\Pi'\ge 1-\lambda\ge 0$ and is nondecreasing in q, while it is strictly increasing in λ at any q with $g(q)>0$. An induction on t using these two monotonicities shows q_t is strictly increasing in λ for every $t\ge 1$ while remaining below q^*, and $q_t\uparrow q^*$ for every $\lambda\in(0,1]$ because q^* is the unique

stable fixed point on $[q_0, 1]$. (ii) The per-period wedge $w(q) = 2\gamma_4 q\sigma_\xi^2/V_A$ is strictly increasing in q. Since $q_t(\lambda)$ is strictly increasing in λ at every $t \geq 1$ and equal at $t = 0$, each term $\beta^t w(q_t)$ is nondecreasing in λ and strictly increasing for $t \geq 1$, so $\mathscr{L}(\lambda;\beta) = \sum_t \beta^t w(q_t)$ is strictly increasing in λ for any $\beta \in (0,1)$. (iii) Along the increasing path q_t, $R(q_t)$ rises (proof of Proposition 6), so for $\phi \in (0,1)$ the fixed-q stationary autocorrelation $\rho_x(q_t, \phi) = \phi R(q_t)$, evaluated at the current adoption rate, is strictly increasing in t. □

Proof of Proposition 9

In the unified environment, each of $n \geq 2$ evaluators of a market is an independent draw from the population, adopting the platform with probability q. Two evaluators contribute normalized error covariance σ_ξ^2/V_A iff both are adopters, which occurs with probability q^2, and zero otherwise. Summing over the $\binom{n}{2}$ evaluator pairs, $S(q) = \binom{n}{2} q^2 \sigma_\xi^2/V_A$ exactly. With complete welfare $\mathscr{W}_1(q) - \gamma_4 S(q)$, the external wedge is $\gamma_4 S'(q) = \gamma_4 \binom{n}{2} 2q\sigma_\xi^2/V_A > 0$ for $q \in (0,1]$, every $n \geq 2$, and every $\delta \in \mathscr{D}$. No equilibrium loading appears, so the wedge is δ-invariant at fixed q. □