Reviews as Matchmakers: Dynamic Learning of Product Fit

Tommaso Bondi

Cornell University

February 23, 2026

Abstract

Consumer reviews are predominantly studied as quality signals – tools for *sorting* products from best to worst. But they also serve a *routing* function, helping consumers assess not whether a product is good but whether it is *good for them*. We model review text as a signal that different readers decode heterogeneously, generating personalized inferences without personalized delivery. As routing succeeds, the buyer pool homogenizes and the text signal degrades. The unique equilibrium features partial matching, sustained by a *diversity dividend* from mismatched consumers whose reviews cover features that fans' reviews miss. The binding constraint is the feature overlap between consumer types, which governs how much information a homogeneous review base can provide. The model generates implications for score-following, AI summaries that disproportionately harm minority types, and the interaction between reviews and algorithmic recommenders. An effective recommender degrades the review base by homogenizing the buyer pool, yet the two systems cover orthogonal dimensions, creating a tension platforms must manage.

1 Introduction

A neighborhood has two Thai restaurants. Restaurant A makes fiercely spicy Northern Thai food; Restaurant B makes mild, fusion-influenced Thai. Both are well-run. The difference is style, not quality.

When a review platform launches, both attract a diverse mix of diners. Restaurant A draws "incredible larb, real Isaan heat" (★★★★★) alongside "way too spicy, couldn't finish" (★★). A spice-lover reads these and knows A is her restaurant. A family diner reads the same words and reaches the opposite conclusion.

Now suppose the reviews work. Heat-seekers go to A; family diners go to B. Restaurant A's reviews become a fan-only chorus: "best som tum in the city," "never disappoints," all five stars. The two-star reviews from mismatched diners are gone. The reviews are narrower: they cover spice and authenticity but say nothing about parking, noise, or kids' menus. A newcomer can tell A is well-liked but cannot tell whether it suits *her*.

Some family diners try A and have bad experiences. Their negative reviews restore the information the system lost. The system converges to a balance: enough *routing* to improve consumer-product *matches*, but not so much that the review base loses the diversity it needs to sustain further routing.[1] This balance is the central object of the paper.

We develop a dynamic model that makes this mechanism precise. Each consumer type attends to a different subset of the feature space; the accuracy with which a reader identifies her match depends on how well the reviews cover *her* features. A diverse reviewer pool covers both subsets; a fan-only pool covers only the fans' features. We study a regime where *quantity* of reviews is not the constraint; what matters is the *variety* of the reviewer pool.[2]

Four sets of results emerge. First, review systems cannot sustain perfect matching: the diversity that makes reviews informative is what successful matching eliminates. The system

[1] We use the two terms interchangeably throughout the paper: routing is the process (directing consumers toward products), matching is the outcome (consumers end up with fitting products).

[2] The dynamics operate on the cumulative share of majority-type buyers in a product's review base. In the large-sample regime, the cumulative share approximately equals the recent flow composition.

converges to a unique equilibrium with partial matching, characterized by a *diversity equation* that decomposes the mismatch rate into three forces: the information deficit of a fan-only review base, the majority's tendency to self-sort, and the diversity value of a mixed reviewer pool. A positive fraction of mismatched consumers persist in every period, not because the system fails but because their reviews cover features that fans' reviews miss; this is the *diversity dividend*.

Second, the binding constraint on matching is the feature overlap between consumer types. Markets where types attend to similar features (hotels: cleanliness, location, and noise are near-universal) sustain high matching even with homogeneous reviewers; markets with disjoint features (books: literary fiction and thriller readers attend to entirely different dimensions) depend critically on diversity. The feature overlap governs the return on platform investments in reviewer recruitment, structured review templates, and stratified summaries, and we derive a shadow value formula that quantifies the return per marginal minority-type reviewer.

Third, the equilibrium generates implications for information design. AI-generated review summaries – such as Amazon's "Customers say" feature – erase the diversity dividend by surfacing only the features the majority mentions most; a niche hot sauce gets "great flavor, arrives well-packaged" while the fit-relevant "extremely hot, not for beginners" disappears. Score-following – purchasing the higher-rated product without reading text – has a non-monotone welfare effect (Section 4): uninformed buyers match poorly but inject diversity into the review base, and the net effect depends on how dominant the majority is. When platforms rank products by scores, a compositional barrier to entry emerges: new products score lower not because they are worse but because their buyer pool has not yet self-selected. Absent score-based ranking, the barrier does not persist: the unique equilibrium is globally stable.

Fourth, we compare reviews with algorithmic recommender systems. Both route consumers toward fitting products, but through different channels: recommenders infer types from behavioral data, reviews convey fit through natural language that different readers decode differently. We show that an effective recommender degrades the review base by homogenizing the buyer pool – the same self-undermining mechanism as in the baseline model, now triggered externally – yet covers dimensions that reviews cannot parameterize, creating a tension platforms must manage.

Roadmap. Section 1.1 reviews the literature. Section 2 develops the model. Section 3 characterizes the equilibrium and derives comparative statics. Section 4 introduces numerical scores and analyzes score-following and AI-generated summaries. Section 5 studies the interaction between reviews and algorithmic recommenders. Section 6 discusses platform implications and Section 7 concludes. All proofs are in the Appendix.

1.1 Related Literature

Online reviews have been shown to affect sales (Chevalier and Mayzlin, 2006; Luca, 2016), pricing (Sun, 2012), market structure (Godes and Mayzlin, 2004; Forman et al., 2008), and strategic manipulation (Mayzlin et al., 2014); see Pocchiari et al. (2025) and Bondi and Rossi (2026) for recent surveys. De Langhe et al. (2016) find that consumers treat average ratings as face-value quality signals even when the rating distribution is visibly polarized, making the informational content of reviews – what they convey, and to whom – a first-order question. Survey evidence confirms that consumers value text: 56% of shoppers report that review length, depth, and detail influence their purchase decisions.[3]

The theoretical literature on reviews has primarily studied the average rating as a quality signal. Acemoglu et al. (2022) characterize learning speed under endogenous selection where taste differences slow convergence to the true quality. Bondi (2025) shows that, when consumers are naive, reviewer self-selection can advantage polarizing products precisely because fit variation distorts quality signals. Chen et al. (2024) analyze the dual roles of reviews and prices in conveying quality under selection bias. Sayedi (2024) studies dynamic pricing with

[3] The Ever-Growing Power of Reviews, 2021.

observational learning, showing that prices can prevent wrong cascades. In these models, fit is a complication for quality inference; in ours, fit is the object of learning. The distinction matters for the information mechanism: in quality-learning models, endogenous selection *slows convergence* to a fixed true state (Acemoglu et al., 2022); in our model, successful matching does not slow convergence but *degrades the signal itself*, because the diversity that makes text informative is what matching eliminates. The signal is not noisy – it is shrinking.

Because we study taste heterogeneity, we naturally relate to the literature on *variance*, not just averages, of ratings. Sun (2012) shows that high variance signals horizontal differentiation; Zimmermann et al. (2018) and Bollinger et al. (2023) show that variance need not reflect just taste mismatch, and decompose it into taste and quality components. Bondi et al. (2024) document how type heterogeneity can reverse product rankings. In our model, score variance is an equilibrium outcome of the routing process.

A separate literature studies the competitive impact of fit information, but this information is not socially learned: it is either exogenous or designer-controlled. Ellison and Ellison (2014) show that improved search raises prices and welfare through better matching; Armstrong and Zhou (2022) find that consumer match information has first-order consequences for market structure; Guo and Zhang (2012) show that consumers' costly deliberation to uncover their fit has first-order effects on product line design.

The intersection of "learning from reviews" and "fit learning" is recent and predominantly empirical. Wang (2021) finds that reviews mentioning fit context reduce purchase uncertainty more effectively than quality-focused reviews. Lei et al. (2022) show experimentally that a few top-ranked text reviews can overturn the influence of average ratings, operating entirely through text content. Hong and Pavlou (2014) and Chen et al. (2021) propose methods for separating quality from fit in rating data.

The closest paper to ours is Fainmesser et al. (2021), who study a Hotelling model in which user-generated content conveys positioning information. They share our insight that detailed text reviews enable fit assessment while numerical ratings do not. The key difference is that in their framework, review informativeness is exogenous and fixed; in ours, it is an equilibrium object that degrades endogenously as the review system's own success homogenizes the reviewer pool.

Because fit learning has been primarily the domain of recommender systems, we also relate to that literature. Che and Hörner (2018) show that noise injection can improve information aggregation; our non-monotonicity result echoes this in a fit-learning context. Feldman et al. (2019) show that early adopters generate information externalities that parallel our mechanism. Rafieian and Zuo (2025) show that algorithmic dependence degrades users' independent decision-making – conceptually parallel to our finding that successful routing degrades future review informativeness. Jerath and Ren (2021) and Chakraborty et al. (2024) model selection into reviewing; our mechanism instead operates through selection into *buying*.

2 Model

2.1 Environment

Two products $j \in \{A, B\}$ have styles $\sigma_A = a$ and $\sigma_B = b$. Each period $t = 1, 2, \ldots$, a unit mass of consumers arrives.[4] Each consumer has type $\tau \in \{a, b\}$, drawn i.i.d. with $\mathbb{P}(\tau = a) = \lambda \in (1/2, 1)$; type a is the majority. The binary type represents two latent taste clusters: a consumer knows her own taste direction but not which product matches it, and holds a flat prior: $\mathbb{P}(\sigma_A = a) = \mathbb{P}(\sigma_B = b) = 1/2$.

Assumption 1 (Payoffs)**.** Consumer of type τ purchasing product j receives $u(\tau, \sigma_j) = \mathbf{1}\{\tau = \sigma_j\}$: utility 1 from a match, 0 from a mismatch. Each consumer buys exactly one product. Quality is known and equal across products.

The assumption that quality is known and equal is a key modeling choice. It isolates the

[4] We treat proportions as continuous when deriving equilibrium. The formal stability proof (Proposition 1) uses the discrete stochastic framework.

fit channel, which is the object of study. It captures settings where quality is established – a Michelin-starred restaurant, a Pulitzer-winning novelist's next book, a sequel from a trusted director – and consumers care primarily about whether the product matches *their* tastes. The entire quality-learning literature (e.g., Acemoglu et al., 2022) studies the complementary case; we study the one it abstracts from. The interaction between quality and fit uncertainty is discussed in Section 7, where we argue that the feedback loop persists when quality varies.[5]

2.2 Scores

Each buyer observes her payoff (match or mismatch) and posts a public review consisting of a numerical score and text.

Assumption 2 (Scores)**.** Each buyer posts a score $s = \mathbf{1}\{\text{match}\}$. Product j's average score at time t is $\bar{s}_j(t) = m_j(t)/n_j(t)$, where $m_j(t)$ is the number of matched buyers through period t and $n_j(t)$ is the total.

Since $\sigma_A = a$, a buyer of product A is matched if and only if she is type a, so the average score equals the fraction of A's buyers who are type a: $\bar{s}_A(t) = \phi_A(t)$, where $\phi_A(t) \equiv n_a(t)/n_A(t)$ denotes the type-a share of A's buyer pool through period t.

We focus on the text channel: consumers choose which product to buy based on their text signal, without conditioning on aggregate scores. This is the relevant model for a consumer whose goal is fit rather than quality – she needs to know whether the spice level, the pacing, or the tone matches *her*. Experimental evidence supports this: Lei et al. (2022) show that a few top-ranked text reviews can overturn the influence of average ratings, operating entirely through text content.[6] Under our equal-quality assumption, scores are clean signals of pool composition; when quality varies (Section 7), they confound quality and fit, making text the primary fit channel. We return to scores in Section 4.

[5]For example, a Goodreads review describing a novel as "beautifully written but glacially paced" is a fit signal whether the book is a masterpiece or mediocre.

[6]In a Yelp survey, 90% of consumers reported trusting reviews with text more than star ratings alone (Yelp Trust Survey, 2022); 56% of shoppers do not trust star ratings alone (PowerReviews, 2023).

2.3 Text as Feature Coverage

We now microfound the text channel. The standard approach models each review as a public signal with a common interpretation (Acemoglu et al., 2022). We instead model *heterogeneous decoding*: products have features, reviewers describe some of them, and different consumer types attend to different features.

Assumption 3 (Feature Attention)**.** Each product has K features. Type τ attends to a subset $\mathcal{F}_\tau \subset \{1, \dots, K\}$, with $|\mathcal{F}_a| = |\mathcal{F}_b| = L$. The overlap is $\mu = |\mathcal{F}_a \cap \mathcal{F}_b|$, with $0 \leq \mu < L$.

The overlap μ measures how much the types have in common: at $\mu = 0$, they attend to entirely disjoint features; near $\mu = L$, they attend to mostly the same features but disagree on preferred valences. Platform design choices affect a distinct object: the *effective* overlap $\tilde{\mu}$ in the text. Structured templates and review prompts can raise $\tilde{\mu}$ above μ by inducing reviewers to cover features outside their core interest. The comparative statics in Section 3.4 hold for both μ (cross-category variation) and $\tilde{\mu}$ (within-category design).

Each feature k carries a *style valence* $v_k \in \{a, b\}$, indicating which consumer type it appeals to.[7] The total K plays no further role beyond the constraint $K \geq 2L - \mu$; only L, μ, and ζ (defined below) affect the dynamics.

Assumption 4 (Review Content and Signals)**.** Each reviewer mentions one feature drawn uniformly and independently from her type's feature set $\mathcal{F}_\tau$. Each mentioned feature, when observed by a reader who attends to it, generates a signal about the product's style with log-likelihood ratio $\ln(\zeta/(1-\zeta))$, where $\zeta \in (1/2, 1)$.

The "one feature per review" assumption generates the coupon-collector structure that makes diversity valuable. Allowing multiple features per review would accelerate coverage but preserve the core tradeoff; the asymmetric-verbosity extension shows that the equilibrium

[7] Style valences are properties of the product ("very spicy"), not the consumer ("likes spice"). All features of a given product share the same underlying style; we assume each covered feature provides a conditionally independent signal about this common style.

structure is unchanged when mismatched reviewers mention $r > 1$ features.[8]

A consumer of type τ reads the reviews and extracts information from features in her attention set $\mathcal{F}_\tau$. This is *heterogeneous decoding*: the same reviews, read by different types, yield different observed feature sets and therefore different inferences.[9]

Heterogeneous decoding draws a distinction between *personalized information* and *personalized decoding of information*. Recommender systems personalize by showing different content to different users; reviews personalize by showing the *same* content that different users read differently. Platforms can therefore support niches and serve the long tail without algorithmic personalization – a point we develop in Section 5.

2.4 Coverage, Accuracy, and the Role of Diversity

Consider product A (style a), which has accumulated n_a reviews from type-a buyers and n_b from type-b buyers. Feature $k \in \mathcal{F}_a$ is covered if at least one reviewer mentioned it. Features exclusive to $\mathcal{F}_a$ can only be covered by type-a reviewers; shared features ($k \in \mathcal{F}_a \cap \mathcal{F}_b$) can be covered by either type. The structure is a coupon-collector problem.[10]

Lemma 1 (Expected Coverage)**.** *The expected number of features in $\mathcal{F}_a$ covered by the reviews is*

$$C_a(n_a, n_b) = (L - \mu)\Big[1 - \big(1 - \tfrac{1}{L}\big)^{n_a}\Big] + \mu\Big[1 - \big(1 - \tfrac{1}{L}\big)^{n_a + n_b}\Big]. \tag{1}$$

By symmetry, $C_b(n_a, n_b) = C_a(n_b, n_a)$.

The first term counts *exclusive* features: those in $\mathcal{F}_a$ but not $\mathcal{F}_b$. Only type-a reviewers contribute to these. The second term counts *shared* features: those in $\mathcal{F}_a \cap \mathcal{F}_b$. Both types contribute, which is why diversity helps even for shared features. This is the key

[8]Strategic review writing would add a signaling game orthogonal to our mechanism. Empirical evidence suggests most reviewers write spontaneously (Wang, 2021).

[9]Could a reviewer simply disclose her type? In practice, types are multidimensional and unlabeled. The Goodreads review "if you loved *Piranesi* for the atmosphere but found *House of Leaves* too demanding" routes by similarity without ever stating a type.

[10]The coupon-collector structure gives diminishing returns to same-type reviews, making the first reviews from a new type highly informative.

asymmetry: a fan-only review base ($n_b = 0$) provides complete coverage for the fans but only partial coverage for non-fans (only the μ shared features in $\mathcal{F}_b$).

Each covered feature carries a style valence that, when observed by an attending reader, generates a signal with log-likelihood ratio $\ln(\zeta/(1-\zeta))$ in favor of the true style, where $\zeta \in (1/2, 1)$. A reader who observes m covered features accumulates a total log-likelihood ratio of $m \ln(\zeta/(1-\zeta))$, yielding the posterior

$$\gamma(m) = \frac{\zeta^m}{\zeta^m + (1-\zeta)^m}. \tag{2}$$

The formula is strictly increasing and strictly concave in m for $m \geq 0$, with $\gamma(0) = 1/2$ (no information) and $\gamma(m) \to 1$ (perfect identification).[11] We adopt $\gamma(m)$ as the probability that the consumer correctly identifies the product's style, and therefore buys the matching product.

Definition 1 (Aggregate Text Accuracy)**.** The aggregate text accuracy is the population-weighted average: $\Gamma(n_a, n_b) = \lambda\, \gamma(C_a) + (1-\lambda)\, \gamma(C_b)$.[12]

The aggregate Γ is the expected match probability for a randomly drawn consumer, weighting each type by its population share.

Lemma 2 (Structural Determinants of Text Accuracy)**.** *Fix total reviews n and let $\phi = n_a/n$ denote the fraction from type-a reviewers.*

(i) Γ *is increasing in n (quantity). For n sufficiently large, Γ is decreasing in ϕ for $\phi > 1/2$ (increasing in variety).*[13]

(ii) *As $n \to \infty$ with $\phi \in (0,1)$: $\Gamma \to \hat{\gamma} \equiv \gamma(L)$.*

[11] Concavity follows from $\gamma'(m) = \gamma(m)(1-\gamma(m)) \ln(\zeta/(1-\zeta))$, which is maximized at $m = 0$ and decreasing. All formal results depend only on the qualitative properties: γ is increasing, concave, and bounded between 1/2 and 1.

[12] Since γ is concave, Jensen's inequality implies $\mathbb{E}[\gamma(C)] \leq \gamma(\mathbb{E}[C])$, so our formula slightly overstates accuracy. For large n, coverage concentrates and the approximation is tight.

[13] The variety result requires the exponential advantage of the minority type's marginal coverage to dominate the population weight ratio $\lambda/(1-\lambda)$. The required n grows with λ; for the large-sample regime we study (Proposition 1 onward), the condition is satisfied.

(iii) *As* $n \to \infty$ *with* $\phi = 1$*:* $\Gamma \to \bar{\gamma} \equiv \lambda\,\gamma(L) + (1-\lambda)\,\gamma(\mu)$.

(iv) $\bar{\gamma}$ *is strictly increasing in* μ*,* $\hat{\gamma}$ *is strictly increasing in* L *and* ζ*, and* $\hat{\gamma} > \bar{\gamma}$ *whenever* $\mu < L$*.*

The lemma reveals that the structural parameters (K, L, μ, ζ) reduce to two summary statistics: a ceiling $\hat{\gamma}$ (the best the system can do with a perfectly diverse pool) and a floor $\bar{\gamma}$ (the worst it can do with a completely homogeneous pool). When μ is close to L, the gap $\hat{\gamma} - \bar{\gamma}$ is small and diversity is a luxury; when μ is close to 0, the gap is large and diversity is essential.

2.5 Reduced Form for the Dynamics

As the review base grows, *quantity* ceases to be a constraint: coverage of any feature saturates. But *variety* remains binding because the state variable ϕ_t tracks the cumulative composition, dominated by recent flows.[14]

We adopt the following maintained specification for text accuracy as a function of the buyer-pool composition $\phi \equiv \phi_A(t)$:

$$\gamma(\phi) = \bar{\gamma} + (\hat{\gamma} - \bar{\gamma}) \cdot 4\phi(1-\phi). \tag{3}$$

The specification is pinned down by three requirements. First, $\gamma(1) = \bar{\gamma}$: a fully homogeneous pool achieves only the fan-only baseline. Second, $\gamma(1/2) = \hat{\gamma}$: a perfectly balanced pool achieves the maximum. Third, γ is symmetric around $\phi = 1/2$ and monotonically decreasing on $[1/2, 1]$, since increasing homogeneity reduces the pool's feature coverage regardless of which type dominates. The unique degree-two polynomial satisfying these conditions is (3); the coefficient 4 is pinned down by the boundary values.[15]

[14] Along the equilibrium path, the minority flow grows as $(1-\lambda)(1-\gamma^*)$ per period while the majority's grows as $\lambda\gamma^*$, so effective accuracy depends on flow composition even in the large-n limit.

[15] The qualitative results (uniqueness, stability, direction of comparative statics) hold for any continuous, strictly decreasing γ on $(\lambda, 1)$ with $\gamma(\lambda) > 1/2$ and $\gamma(1) < 1$. The quadratic form delivers closed-form comparative statics and the diversity equation (7).

Each consumer reads the reviews for both products and buys the one she believes matches her type.[16] The type-a fraction among new A-buyers at time t is

$$\psi(\gamma) = \frac{\lambda\gamma}{\lambda\gamma + (1-\lambda)(1-\gamma)}, \tag{4}$$

with $\psi(1/2) = \lambda$ (random assignment when text is uninformative) and $\psi(1) = 1$ (perfect matching).[17] The cumulative composition updates as

$$\phi_{t+1} = \phi_t + \frac{1}{n_t + 1}\big[\psi(\gamma_t) - \phi_t\big], \tag{5}$$

where n_t is A's total buyers through period t and $\gamma_t = \gamma(\phi_t)$. Expression (4) maps text accuracy to buyer composition; the composite map $F(\phi) = \psi(\gamma(\phi))$ gives the composition of new buyers as a function of the current stock, and a steady state satisfies $F(\phi^*) = \phi^*$.

3 Equilibrium

3.1 Partial Matching and the Diversity Dividend

Proposition 1 (Generic Impossibility)**.** *Let $\gamma : [\lambda, 1] \to (0, 1)$ be any text accuracy function satisfying:*

(i) *γ is continuous on $[\lambda, 1]$ and strictly decreasing on $(\lambda, 1)$;*

(ii) *$\gamma(\lambda) > 1/2$;*

(iii) *$\gamma(1) < 1$.*

Then the updating map $F = \psi \circ \gamma$ has a unique fixed point $\phi^* \in (\lambda, 1)$. The fixed point is globally stable. Perfect matching ($\phi = 1$) is not a steady state.

[16] Since the model has two products and a flat prior, each consumer compares her posterior for A against her posterior for B. By symmetry of the text signal across products, the choice is deterministic given the text realization.

[17] The same selection function appears in Acemoglu et al. (2022), but there γ is exogenous. Here, γ depends on buyer-pool composition, creating the feedback loop.

The conditions require only that a diverse pool is informative and a homogeneous pool cannot sustain perfect matching; both are properties of the information technology.

The impossibility result, uniqueness, and global stability hold for *any* text accuracy function satisfying conditions (i)–(iii); the closed-form diversity equation, shadow value formula, and specific thresholds in Proposition 6 use the quadratic form (3) for tractability. Under this specification, the unique steady state has $\phi^* \in (\lambda, 1)$, matching is partial ($\gamma^* \in (\bar{\gamma}, \hat{\gamma})$), and the stochastic dynamics (5) converge to ϕ^* from any initial condition.[18]

The contrast with herding models is instructive. In Banerjee (1992) and Bikhchandani et al. (1992), cascades on the wrong action can lock in permanently. Here, the fixed point ϕ^* is unique and globally attracting because successful routing is *self-undermining*: it eliminates the diversity that made routing possible, creating a restoring force.

The balance point is maintained by what we call the *diversity dividend*: the gap $\gamma^* - \bar{\gamma}$ between the equilibrium match rate and the fan-only baseline. This gap measures the efficiency contribution of reviewer-pool diversity sustained by persistent mismatch. At ϕ^*, mismatched consumers are individually worse off but serve a collective function: their reviews cover features in $\mathcal{F}_b$ that fans' reviews miss.

3.2 Characterizing the Equilibrium

The equilibrium condition $\phi^* = \psi(\gamma(\phi^*))$ does not immediately reveal how matching depends on the structural primitives. Define the *mismatch rate* $m^* = 1 - \phi^*$ and the *majority odds ratio* $R = \lambda/(1-\lambda)$.

Proposition 2 (The Diversity Equation)**.** *The equilibrium mismatch rate is characterized by:*

$$\gamma^* \;=\; \psi^{-1}(\phi^*) \;\approx\; 1 - R \cdot m^*, \qquad m^* \;\approx\; \frac{1-\bar{\gamma}}{R + 4\Delta}, \tag{6}$$

where $R = \lambda/(1-\lambda)$ and $\Delta = \hat{\gamma} - \bar{\gamma}$. The first expression is exact; the second and third are

[18] Convergence follows from a standard stochastic approximation argument (Appendix).

first-order approximations around $\phi^* = 1$.

We call (6) the *diversity equation*. A closed-form solution does not exist: ψ and γ compose into a rational function whose root cannot be expressed in elementary terms. However, the approximation is precise: with baseline parameters ($L = 10$, $\zeta = 0.56$, $\lambda = 0.55$, $\mu = 3$), the exact equilibrium is $\gamma^* = 0.8549$ while the approximation gives 0.8586, an error of 0.4%.[19]

Three structural features of the product category determine how much mismatch the market sustains:

$$m^* \approx \frac{\overbrace{1-\bar{\gamma}}^{\text{information deficit}}}{\underbrace{R}_{\text{demographic routing}} + \underbrace{4\Delta}_{\text{diversity value}}}. \tag{7}$$

The *numerator* is the information deficit: the gap between perfect text quality and the fan-only baseline $\bar{\gamma}$. When $\bar{\gamma}$ is close to 1 (high feature overlap, as in hotels), the fan-only review base is nearly sufficient and the equilibrium requires little mismatch. When $\bar{\gamma}$ is low (disjoint features, as in books), the system must sustain a substantial mismatch flow.

The *denominator* sums two compression forces: the majority's self-sorting tendency $R = \lambda/(1-\lambda)$, and the diversity value 4Δ, which converts mismatch into better future routing. The information deficit measures how much the system *needs* mismatched reviewers; demographic routing measures how quickly the majority *displaces* them; and the diversity value measures how much each mismatched reviewer *contributes*. Both the numerator and the denominator depend on μ: increasing feature overlap raises $\bar{\gamma}$ (reducing the numerator) and shrinks Δ (reducing the denominator). But the numerator effect dominates, so m^* is unambiguously decreasing in μ (Proposition 4). To make this concrete: in a book market with low feature overlap, the formula predicts roughly 11% mismatch; in a hotel market with high overlap, mismatch drops to about 7%.[20] The formula identifies the highest-priority targets

[19]Both approximations are first-order expansions in m^*; the error is quadratic in m^*. Exact bounds are $(1-\bar{\gamma})/(R+4\Delta) \le m^* \le 1-\psi(\bar{\gamma})$, with the lower bound tight for small m^*.

[20]A simple calibration illustrates the magnitudes. Fix $L = 10$, $\zeta = 0.56$, $\lambda = 0.60$. For books ($\mu = 2$): $\gamma^* = 0.845$, $m^* = 10.9\%$, diversity dividend = 4.7 pp. For hotels ($\mu = 7$): $\gamma^* = 0.896$, $m^* = 7.2\%$, diversity dividend = 0.8 pp. A frequency-based AI summary ($k = 7$) reduces minority accuracy from 92% to 62% in the book market but only from 92% to 84% in the hotel market. These are order-of-magnitude

for platform intervention: categories with low feature overlap and balanced demand.

Markets with dominant consumer segments (R large) pay a steeper efficiency cost for the same level of mismatch.

3.3 Match Efficiency

Aggregate match efficiency per period is γ^*, bounded by three benchmarks: random assignment (1/2), uninformed score-following (λ), and perfect matching (1).

Corollary 1 (Match Efficiency). *The review system closes a fraction* $(\gamma^* - \lambda)/(1 - \lambda)$ *of the gap between uninformed choice and perfect matching. This fraction is increasing in* μ *and* ζ*, and approaches 1 as* $\bar{\gamma} \to 1$*.*

In practice, the match rate γ^* maps to observable platform KPIs: product return rates, "not for me" complaints, and repeat purchase rates. The gap $\gamma^* - \bar{\gamma}$ predicts the incremental benefit of maintaining a diverse reviewer pool.

Although each type faces the same per-capita mismatch probability $1 - \gamma^*$, the *informational contribution* of mismatch is not symmetric: product A needs type-b reviewers to cover $\mathcal{F}_b$-exclusive features. Since the minority is a smaller group, each minority consumer's mismatch is more informationally pivotal. Note that "minority" here is product-specific, not demographic.

Corollary 2 (Informational Pivotality). *At the steady state, each minority reviewer's marginal contribution to text accuracy exceeds each majority reviewer's. The pivotality ratio is at least* $R = \lambda/(1 - \lambda)$ *and is increasing in* λ*.*

The asymmetry arises because the minority type's reviewer stock on its "off" product is smaller ($n_b^A/n_a^B = 1/R < 1$), and diminishing returns in the coupon-collector process make each additional review more valuable at the margin.[21]

illustrations, not estimates; the qualitative pattern – low overlap magnifies every distortion – is robust across parameterizations.

[21] The pattern echoes Waldfogel (2007)'s "tyranny of the majority": in his framework, high fixed costs cause markets to underserve preference minorities; in ours, the mechanism is informational.

Proposition 3 (Shadow Value of a Minority-Type Review)**.** *At the steady state, the marginal efficiency gain from one additional type-b review on product A is*

$$\frac{\partial \gamma^*}{\partial n_b^A} = \underbrace{(1-\lambda)}_{\text{minority weight}} \cdot \underbrace{\gamma'(m_b^*)}_{\substack{\text{signal value per} \\ \text{covered feature}}} \cdot \underbrace{\frac{(L-\mu)}{L}\left(1-\frac{1}{L}\right)^{n_b^A}}_{\text{marginal coverage probability}}, \tag{8}$$

where $n_b^A = (1-\lambda)(1-\gamma^*)n$ *is the minority reviewer stock on A and* $m_b^* = \mu + (L-\mu)[1 - (1-1/L)^{n_b^A}]$ *is the expected type-b feature coverage. The shadow value is:*

(i) *strictly decreasing in* μ *(highest in low-overlap markets),*

(ii) *strictly decreasing in* n_b^A *(highest when minority reviews are scarce),*

(iii) *for fixed L, strictly increasing in* $L - \mu$ *(highest when the minority attends to many exclusive features).*

The shadow value has a clean coupon-collector structure: the first few minority-type reviews are dramatically more valuable than the last, because the marginal coverage probability $(1-1/L)^{n_b^A}$ declines geometrically.

3.4 Comparative Statics

The comparative statics connect the equilibrium to the structural primitives of the product category. The key insight is that the binding constraint on matching is baseline text quality $\bar{\gamma}$: the informativeness of a homogeneous review base, which depends primarily on the feature overlap μ.

Proposition 4 (Feature Overlap as the Binding Constraint)**.** *At the interior steady state:*

(i) $\partial \gamma^*/\partial \mu > 0$*: greater feature overlap improves matching.*

(ii) $\partial^2 \bar{\gamma}/\partial \mu^2 < 0$*: the returns to overlap are diminishing.*

Table 1: Comparative statics of the equilibrium. Signs indicate the effect of increasing each parameter on the equilibrium outcome. Baseline values: $L = 10$, $\zeta = 0.56$, $\lambda = 0.55$.

Parameter	Interpretation	γ^*	m^*	V_A	Shadow value	$\bar{\lambda}$
μ	Feature overlap between types	+	−	−	−	+
ζ	Per-feature signal accuracy	+	−	−	−	+
L	Features per type	+	−	−	ambig.	ambig.
λ	Majority share	+	−	−	+	n/a
r	Mismatched reviewer verbosity	+	−	−	−	+

Notes: γ^* = match efficiency; $m^* = 1 - \phi^*$ = mismatch rate; $V_A = \phi^*(1 - \phi^*)$ = score variance; shadow value = marginal efficiency gain from one minority-type review (Proposition 3); $\bar{\lambda}$ = threshold above which score-following improves match efficiency (Proposition 6). Signs for μ, L, and ζ are proved in Proposition 4; signs for λ follow from the diversity equation (7), noting that R and $\bar{\gamma}$ are both increasing in λ while Δ is decreasing, and verified numerically. The effect of λ on $\bar{\lambda}$ is not defined (it *is* the threshold variable). The effect of L on shadow value and $\bar{\lambda}$ is ambiguous because L enters through both $\hat{\gamma}$ and the coupon-collector marginals. The r row refers to the asymmetric-verbosity extension ($r > 1$ features per review when mismatched).

(iii) *When the pool is sufficiently homogeneous, $\partial\gamma^*/\partial\bar{\gamma} > \partial\gamma^*/\partial\hat{\gamma}$: raising the floor matters more than raising the ceiling.*

Proposition 5 (Signal Quality and Matching)**.** *More features per type and stronger per-feature signals also improve matching: $\partial\gamma^*/\partial L > 0$ and $\partial\gamma^*/\partial\zeta > 0$.*

Increasing μ raises $\bar{\gamma}$, shifting the equilibrium toward better matching; diminishing returns follow from the concavity of $\gamma(m)$. Part (iii) of Proposition 4 says that for most empirically relevant parameter values, the platform's highest-leverage intervention is raising the floor: making fan-only reviews more informative for non-fans, rather than making diverse reviews even better.[22] This is a design principle, not just a comparative static: structured templates that induce fan reviewers to cover a few non-fan features (raising $\tilde{\mu}$) dominate investments in recruiting more diverse reviewers or improving signal precision. The intuition is that the equilibrium self-corrects toward diversity through the feedback loop, but it cannot self-correct toward a higher floor – $\bar{\gamma}$ is a structural primitive. Section 6 translates this and the remaining comparative statics into concrete platform design levers.

[22]The condition for (iii) is $8\phi^*(1-\phi^*) < 1$, which holds whenever the equilibrium pool is more than about 85% majority-type. This covers the empirically relevant range; the condition fails only when signals are very weak or the majority is barely dominant.

The equilibrium also determines score distributions. Score variance $V_A = \phi^*(1 - \phi^*)$ is decreasing in γ^* and hence in μ (by Proposition 4): better-matched products have lower score variance. The model offers a structural interpretation: low variance reflects successful *routing*, not consensus about quality, and its level is pinned down by the feature overlap μ.[23]

The score dynamics also create an informational barrier to entry. A new product enters at $\phi_0 = \lambda$ with average score $\bar{s} = \lambda$, while an established competitor at ϕ^* scores higher. The gap is compositional, not quality-based. Under the text channel alone, this barrier is transient; when platforms rank by score, it becomes persistent (cf. Hummel and Morgan, 2014; Vellodi, 2022).

4 Scores, Summaries, and Information Design

In practice, review platforms present both text and numerical scores, and increasingly layer AI-generated summaries on top of the raw review data. This section studies how these additional channels interact with text-based routing.

We first introduce *score-following*: consumers who buy the higher-rated product without reading text. Score-followers are uninformed about fit: they buy the product that the majority prefers, matching at rate λ. They still purchase, experience the product, and write reviews; what differs is only their purchase decision.[24] This creates two competing effects on the consumers who *do* read text.

Proposition 6 (The Non-Monotone Effect of Score-Following)**.** *If a fraction α of buyers follow average scores rather than reading text, the aggregate match rate is $M(\alpha) = \alpha\lambda + (1-\alpha)\gamma^*(\alpha)$. There exists a threshold $\bar{\lambda}(\mu, \zeta, L) < 1$ such that:*

(i) *When $\lambda < \bar{\lambda}$: score-following reduces match efficiency at the margin: $dM/d\alpha\big|_{\alpha=0} < 0$.*

[23] Along the convergence path ϕ_t oscillates, so score variance need not decline monotonically. The prediction is cleanest as a cross-sectional comparative static.

[24] A score-follower's type is drawn from the population (λ probability of being type a) rather than being routed by the text signal.

(ii) *When* $\lambda > \bar{\lambda}$*: score-following* improves *match efficiency at the margin:* $dM/d\alpha\big|_{\alpha=0} > 0$.

The threshold $\bar{\lambda}$ *is increasing in* μ *and in* ζ.

The *direct* effect is signal loss: uninformed consumers match at rate λ rather than γ^*. The *indirect* effect is diversity gain: score-followers who wander into the wrong product diversify the review base.[25] When λ is moderate, the direct loss dominates; when extreme, the diversity gain swamps it.

A rapidly spreading example is the AI-generated review summary. Amazon's "Customers say" feature distills hundreds of reviews into a paragraph highlighting common themes; Google Maps, Yelp, and TripAdvisor deploy similar tools. By construction, these summaries foreground themes mentioned most frequently – those mentioned by the majority – erasing fit-relevant minority perspectives.

Suppose the platform generates a summary by selecting the k most frequently mentioned features ($k \leq L$). Since shared features are mentioned by both types while exclusive features are mentioned only by one, the ranking by frequency places shared features first, majority-exclusive second, and minority-exclusive last.

Proposition 7 (Summarization Erases the Diversity Dividend)**.** *Suppose the platform generates a* k*-feature summary* ($k \leq L$) *by selecting the most frequently mentioned features.*

(i) *For* $k > \mu$*, the summary includes zero features from* $\mathcal{F}_b \setminus \mathcal{F}_a$*. A type-*$b$ *summary-reader achieves accuracy* $\gamma(\mu)$*: the fan-only baseline.*

(ii) *The effective text accuracy for summary-readers is* $\Gamma_{\text{sum}}(k) = \lambda \cdot \gamma\big(\min(k, L)\big) + (1 - \lambda) \cdot \gamma\big(\min(k, \mu)\big)$*. For* $k \geq \mu$*,* $\Gamma_{\text{sum}}(k) \leq \bar{\gamma}$*: the diversity dividend is entirely erased.*

(iii) *If a fraction* α *of consumers read only the summary, the steady-state match rate satisfies* $dM/d\alpha\big|_{\alpha=0} < 0$*: introducing summary-readers strictly reduces match efficiency at the*

[25] This connects to Che and Hörner (2018)'s result that noise injection can improve information aggregation, but in their setting the underlying attribute is common-value; in ours, the fit-learning problem is permanent.

margin.

The summary does not merely reduce information; it reduces it *asymmetrically*. The majority type loses little, because the top-k features are predominantly hers. The minority type loses almost everything, because her exclusive features are mentioned least often and summarized last.[26] The welfare loss from summaries is therefore concentrated among exactly the consumers the platform most needs to serve well: those whose preferences diverge from the majority. A summary that makes the average consumer slightly better informed makes the minority consumer substantially worse informed, and the net effect on matching is negative.

Part (iii) establishes that this is not just a distributional concern. Even accounting for the indirect diversity benefit – summary-readers match poorly and thereby inject diversity into the review base – the direct information loss dominates. The mechanism is the same as in score-following (Proposition 6), but the welfare comparison is unambiguous: unlike score-followers, who at least achieve match rate λ, summary-readers achieve $\Gamma_{\text{sum}}(k) \leq \bar{\gamma}$, which can be substantially below γ^*.

The result does not depend on the specific "top-k by frequency" rule. Any summarization procedure that (a) ranks or weights features monotonically in mention frequency and (b) is agnostic to the reader's type will produce the same qualitative distortion, because the underlying asymmetry is demographic, not algorithmic: majority-exclusive features are mentioned more often than minority-exclusive features at any interior steady state ($\phi^* > 1/2$). Modern LLM-based summaries inherit this bias to the extent that they weight input text by frequency, helpfulness votes, or representativeness – all of which correlate with majority mention share. The specific top-k rule makes the result sharpest, but the mechanism requires only majority-weighted aggregation.

[26] For example, with moderate overlap ($\mu/L = 0.3$) and a summary covering 70% of features, the majority summary-reader retains about 84% accuracy while the minority drops to roughly 67%, compared to 92% for both types under full text with a diverse review base.

5 Interaction with Recommender Systems

Reviews achieve personalization through heterogeneous decoding; recommender systems achieve it through explicit type inference. In practice, consumers use both. Aridor et al. (2024) find experimentally that recommendations' informational channel dominates their consideration-set channel; we study the reverse direction – how recommendations affect the information available in reviews.

We model pure substitution: each consumer uses one system or the other. A recommender assigns consumers to products with match rate $\gamma_R \in (1/2, 1)$, independent of the review base composition.[27] A recommender-routed consumer of type a buys product A with probability γ_R, contributing to the review base in the same way as any other buyer. Suppose a fraction β of consumers follow the recommender and $1 - \beta$ read reviews; text-readers achieve match rate $\gamma_H^*(\beta)$, which depends on the buyer pool composition.

Proposition 8 (Recommender-Induced Review Degradation). *Let $\gamma_H^*(\beta)$ denote the review system's equilibrium accuracy when a fraction β of consumers follow the recommender.*

(i) *When $\gamma_R > \gamma^*$: the recommender homogenizes the buyer pool. $\partial\gamma_H^*/\partial\beta < 0$: a more active recommender degrades the review system.*

(ii) *When $\gamma_R < \gamma^*$: the recommender diversifies the buyer pool. $\partial\gamma_H^*/\partial\beta > 0$: a more active recommender improves the review system.*

An effective recommender ($\gamma_R > \gamma^*$) routes consumers more accurately than reviews do, so recommender-routed buyers of product A are more homogeneously type a – the same self-undermining mechanism, now triggered externally. A weak recommender has the opposite effect: it injects poorly matched consumers whose reviews restore diversity.[28]

[27] The match rate γ_R captures in reduced form the recommender's ability to identify consumer types from behavioral data. The key distinction from reviews is that γ_R does not degrade as the buyer pool homogenizes.

[28] See Berman and Katona (2020) for an analogous channel in social networks, Yoganarasimhan (2020) and Holtz et al. (2020) for empirical evidence from Spotify, and Rafieian and Zuo (2025) for evidence that algorithmic dependence degrades users' independent judgment.

The degradation result has a sharp implication: a highly effective recommender cannibalizes its own data source. As the algorithm improves, the review base becomes less informative for consumers who still rely on it, echoing Fleder and Hosanagar (2009)'s finding that recommender systems can reduce sales diversity.

A natural follow-up is whether the optimal platform should deploy both systems simultaneously. In the single-dimensional match model, $M(\beta) = \beta\gamma_R + (1-\beta)\gamma_H^*(\beta)$ is a weighted average of two scalars, and the better system dominates.[29] But the complementarity that motivates platforms to maintain both channels arises from *dimensional misspecification*: the recommender covers features it can parameterize (genre, price range, popularity), while reviews cover features it cannot (pacing, prose style, emotional tone) (see also Calvano et al., 2024). These are orthogonal strengths, and the practical pattern across platforms reflects this: Amazon recommends books via collaborative filtering while prominently displaying Goodreads reviews; Netflix routes subscribers algorithmically while Letterboxd thrives as a review-based complement.[30]

6 Platform Implications

The model's parameters map to concrete platform design choices, organized around three levers: the effective feature overlap $\tilde{\mu}$, the composition of the reviewer pool, and the information architecture.

Raising effective overlap. Platforms cannot change what consumers attend to (μ), but they can change what reviewers write about ($\tilde{\mu}$). Structured templates that ask about multiple dimensions (food, service, atmosphere, value) ensure that even a fan reviewer covers features outside her core interest, raising $\tilde{\mu}$ above μ. Prompts like "Who would you *not*

[29] When $\gamma_R > \gamma^*$, $M(\beta)$ is increasing in β because the direct gain $(\gamma_R - \gamma_H^*) > 0$ dominates the degradation effect. When $\gamma_R < \gamma^*$, the review system dominates.

[30] See also Bourreau and Gaudin (2022) on strategic recommendation bias, Berman et al. (2024) on strategic algorithm design, and Ke et al. (2022) on match efficiency versus surplus extraction.

recommend this to?" serve the same function. Displaying reviewer profiles helps readers map features to their own preferences, increasing effective ζ. A fourth lever is *segmented ratings*: TripAdvisor's separate scores for families, couples, and solo travelers help each consumer type bypass the diversity constraint. All four interventions are most valuable in low-μ categories, where the gap between $\bar{\gamma} = \lambda\gamma(L) + (1-\lambda)\gamma(\mu)$ and $\hat{\gamma} = \gamma(L)$ is largest and the return to raising $\tilde{\mu}$ is highest (Proposition 4, part iii).

When type-b's feature set is nearly a subset of type-a's, the diversity problem is one-sided: type b is informationally sheltered by the majority's reviews, while type a bears the full cost of homogenization. The relevant margin is therefore *feature* diversity, not reviewer diversity per se.

Reviewer composition and the shadow value. The shadow value formula (Proposition 3) gives the marginal efficiency gain from one additional minority-type review: $\partial\gamma^*/\partial n_b^A = (1-\lambda)\cdot\gamma'(m_b^*)\cdot(L-\mu)/L\cdot(1-1/L)^{n_b^A}$. The coupon-collector term $(1-1/L)^{n_b^A}$ declines geometrically, so the first few minority-type reviews are dramatically more valuable than the last. Platforms running reviewer recruitment programs should target minority consumer segments, not just prolific reviewers: composition matters more than volume, echoing Godes and Mayzlin (2004)'s finding that WOM *dispersion* across communities predicts sales better than WOM volume. Negative reviews are informationally beneficial for the same reason: dissatisfied consumers are disproportionately mismatched, so their reviews cover features in $\mathcal{F}_b \setminus \mathcal{F}_a$ that fans' reviews miss. Platforms that suppress negative reviews through helpfulness filters lose exactly the reviews with the highest shadow value. The effect is category-dependent: in high-μ markets (hotels), suppressing negative reviews has a modest impact on match efficiency because the fan-only review base is already informative; in low-μ markets (literary fiction, niche cuisines), the same policy can substantially degrade fit information.

AI summaries and the majority bias. AI summaries that weight features by mention frequency erase the diversity dividend: for $k \geq \mu$, the summary includes zero features from

$\mathcal{F}_b \setminus \mathcal{F}_a$ (Proposition 7), and the efficiency loss $\bar{\gamma} - \gamma^*$ is concentrated in low-μ categories. Modern LLM-based summaries inherit the same majority-weighted bias to the extent that they weight features by aggregate mention frequency or helpfulness votes (cf. Dai et al., 2018). The natural remedy is *stratified summarization*: a personalized summary conditioning on the reader's inferred type would recover the lost efficiency. The staggered rollout of AI summaries across platforms and categories provides a natural experiment: our model predicts that summary adoption should disproportionately increase product return rates and "not as expected" complaints in low-μ categories (books, fragrances, wine) relative to high-μ categories (electronics, hotels).

Score-ranking and entry barriers. New products face a compositional barrier to entry: a product entering at $\phi_0 = \lambda$ scores $\bar{s} = \lambda$, below an established competitor at ϕ^*. Under the text channel alone, this barrier is transient – the entrant's diverse early pool is *more* informative for text readers. But when platforms rank by score, low scores divert traffic, making the barrier persistent (cf. Hummel and Morgan, 2014; Vellodi, 2022). Programs like Amazon's Vine, which recruit diverse early reviewers, directly address this. The model predicts that Vine-style programs should have larger effects on early sales in low-μ categories, where the compositional barrier is most severe and reviewer diversity most valuable.

7 Conclusion

This paper develops a model of social learning about idiosyncratic fit through review text. The central result is that review systems cannot sustain perfect matching: successful routing homogenizes the reviewer pool and degrades the information that enables further routing. The unique equilibrium features partial matching, sustained by a diversity dividend from mismatched consumers whose reviews cover features that fans' reviews miss. The binding constraint is the structural feature overlap between consumer types.

The routing lens reshapes how we think about market structure. When reviews function as quality signals, consumers converge on "best" products, reinforcing superstars. When reviews function as matching devices, consumers disperse across products fitting their types, supporting the long tail. The same platform can have opposite effects depending on which role dominates. Amazon acquired Goodreads in 2013; if Goodreads functions primarily as a matching device rather than a sorting device, its value lies in *demand shaping*: routing readers toward books that fit their idiosyncratic tastes redistributes demand from bestsellers toward the long tail – precisely the segment where Amazon faces the least competition (Brynjolfsson et al., 2003, 2011).

Niche products face a distinctive tradeoff: suppressing negative reviews may raise scores but degrades the fit signal sustaining the niche audience's willingness to pay.[31] More broadly, the *same* product can be a niche product or a mass-market product depending on the review base that develops around it – the information system shapes the demand structure, not just reflects it.

The model has clear limitations. The binary type space captures the essential tradeoff but cannot address richer preference structures; with more types and a matrix of pairwise overlaps, the impossibility result survives but the equilibrium involves multiple diversity equations. Noisy self-knowledge – consumers who misperceive their own type – effectively reduces per-feature signal accuracy while preserving the feedback loop, since the mechanism depends on who buys, not on how accurately consumers know their type. Quality uncertainty is more subtle: scores confound quality and fit, but text retains its fit-informational content, so the self-undermining mechanism persists – the system degrades not because the signal is noisy but because the signal is *narrowing*.

The equilibrium is identified from observables without requiring the econometrician to observe consumer types. Score variance pins down the equilibrium composition; the inverse selection function recovers match efficiency; the diversity equation recovers baseline accuracy; and the structural relationship between baseline accuracy, feature overlap, and signal

[31] In a Hotelling extension, low μ would imply strong fit signals, raising willingness to pay and softening price competition, paralleling Armstrong and Zhou (2022).

precision identifies the primitives. The number of features consumers attend to is recoverable from topic models applied to review text, and per-feature signal accuracy is estimable from within-review sentiment consistency. Crucially, the endogenous selection that complicates quality estimation is precisely what generates the equilibrium composition that identifies the model.

The deeper point is about any participatory information system. Recommendation algorithms face the same feedback loop: accurate recommendations homogenize training data, degrading the algorithm's ability to learn about users outside the modal pattern (cf. Fleder and Hosanagar, 2009). Content curation faces it too: algorithmic promotion of engaging content narrows the distribution of voices. The tension between improving outcomes and degrading the information that enables improvement arises whenever who participates determines what can be learned.

References

Acemoglu, D., Makhdoumi, A., Malekian, A., and Ozdaglar, A. (2022). Learning from Reviews: The Selection Effect and the Speed of Learning. *Econometrica*, 90(6):2857–2899.

Aridor, G., Gonçalves, D., Kluver, D., Kong, R., and Konstan, J. (2024). The Informational Role of Online Recommendations: Evidence from a Field Experiment. Working Paper.

Armstrong, M. and Zhou, J. (2022). Consumer Information and the Limits to Competition. *American Economic Review*, 112(2):534–577.

Banerjee, A. V. (1992). A Simple Model of Herd Behavior. *Quarterly Journal of Economics*, 107(3):797–817.

Berman, R. and Katona, Z. (2020). Curation Algorithms and Filter Bubbles in Social Networks. *Marketing Science*, 39(2):296–316.

Berman, R., Zhao, H., and Zhu, Y. (2024). Strategic Design of Recommendation Algorithms. Working Paper.

Bikhchandani, S., Hirshleifer, D., and Welch, I. (1992). A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades. *Journal of Political Economy*, 100(5):992–1026.

Bollinger, B., Lee, N., and Staelin, R. (2023). Vertical Versus Horizontal Variance in Online Reviews and Their Impact on Demand. *Journal of Marketing Research*, 60(3):484–503.

Bondi, T. (2025). Alone, Together: A Model of Social (Mis)Learning from Consumer Reviews. *Marketing Science*, 44(1):1–22.

Bondi, T. and Rossi, M. (2026). Online Reviews: Information Content, Biases, and Platform Design. *Working Paper*.

Bondi, T., Rossi, M., and Stevens, R. (2024). The Surprising Value of Extreme Reviews. *Working Paper*.

Bourreau, M. and Gaudin, G. (2022). Streaming Platform and Strategic Recommendation Bias. *Journal of Economics & Management Strategy*, 31(1):25–47.

Brynjolfsson, E., Hu, Y., and Smith, M. D. (2003). Consumer Surplus in the Digital Economy: Estimating the Value of Increased Product Variety at Online Booksellers. *Management Science*, 49(11):1580–1596.

Brynjolfsson, E., Hu, Y., and Simester, D. (2011). Goodbye Pareto Principle, Hello Long Tail: The Effect of Search Costs on the Concentration of Product Sales. *Management Science*, 57(8):1373–1386.

Calvano, E., Calzolari, G., Denicolò, V., and Pastorello, S. (2024). Economics of Recommender Systems. In *Proceedings of the 18th ACM Conference on Recommender Systems (RecSys '24)*, pages 1279–1280. ACM.

Chakraborty, I., Deb, J., and Öry, A. (2024). When Do Consumers Talk? Working Paper.

Che, Y.-K. and Hörner, J. (2018). Recommender Systems as Mechanisms for Social Learning. *Quarterly Journal of Economics*, 133(2):871–925.

Chen, P.-Y., Hitt, L. M., Hong, Y., and Wu, S. (2021). Measuring Product Type and Purchase Uncertainty with Online Product Ratings: A Theoretical Model and Empirical Application. *Information Systems Research*, 32(4):1470–1489.

Chen, Y., Du, J., and Lei, Y. (2024). The Interactions of Customer Reviews and Price and Their Dual Roles in Conveying Quality Information. Working Paper.

Chevalier, J. A. and Mayzlin, D. (2006). The Effect of Word of Mouth on Sales: Online Book Reviews. *Journal of Marketing Research*, 43(3):345–354.

Dai, W., Jin, G., Lee, J., and Luca, M. (2018). Aggregation of Consumer Ratings: An Application to Yelp.com. *Quantitative Marketing and Economics*, 16(3):289–339.

De Langhe, B., Fernbach, P. M., and Lichtenstein, D. R. (2016). Navigating by the Stars: Investigating the Actual and Perceived Validity of Online User Ratings. *Journal of Consumer Research*, 42(6):817–833.

Ellison, G. and Ellison, S. F. (2014). Match Quality, Search, and the Internet Market for Used Books. *American Economic Review*, 104(1):293–321.

Fainmesser, I. P., Lauga, D. O., and Ofek, E. (2021). Ratings, Reviews, and the Marketing of New Products. *Management Science*, 67(11):7023–7045.

Feldman, P., Papanastasiou, G., and Segev, E. (2019). Social Learning and the Design of New Experience Goods. *Management Science*, 65(9):4178–4198.

Fleder, D. and Hosanagar, K. (2009). Blockbuster Culture's Next Rise or Fall: The Impact of Recommender Systems on Sales Diversity. *Management Science*, 55(5):697–712.

Forman, C., Ghose, A., and Wiesenfeld, B. (2008). Examining the Relationship Between Reviews and Sales: The Role of Reviewer Identity Disclosure in Electronic Markets. *Information Systems Research*, 19(3):291–313.

Godes, D. and Mayzlin, D. (2004). Using Online Conversations to Study Word-of-Mouth Communication. *Marketing Science*, 23(4):545–560.

Guo, L. and Zhang, J. (2012). Consumer Deliberation and Product Line Design. *Marketing Science*, 31(6):995–1007.

Holtz, D., Carterette, B., Chandar, P., Nazari, Z., Cramer, H., and Aral, S. (2020). The Engagement-Diversity Connection: Evidence from a Field Experiment on Spotify. Working Paper, MIT Sloan School of Management.

Hong, Y. and Pavlou, P. A. (2014). Product Fit Uncertainty in Online Markets: Nature, Effects, and Antecedents. *Information Systems Research*, 25(2):328–344.

Hummel, P. and Morgan, J. (2014). A Model of Flops. *RAND Journal of Economics*, 45(1):75–96.

Jerath, K. and Ren, Q. (2021). Consumer Rational (In)Attention to Favorable and Unfavorable Product Information, and Firm Information Design. *Journal of Marketing Research*, 58(2):343–362.

Ke, T. T., Lin, S., and Lu, M. Y. (2022). Information Design of Online Platforms. Working Paper.

Lei, Z., Yin, D., Mitra, S., and Zhang, H. (2022). Swayed by the Reviews: Disentangling the Effects of Average Ratings and Individual Reviews in Online Word-of-Mouth. *Production and Operations Management*, 31(6):2697–2711.

Luca, M. (2016). Reviews, Reputation, and Revenue: The Case of Yelp.com. Harvard Business School Working Paper No. 12-016.

Mayzlin, D., Dover, Y., and Chevalier, J. A. (2014). Promotional Reviews: An Empirical Investigation of Online Review Manipulation. *American Economic Review*, 104(8):2421–2455.

Pocchiari, M., Proserpio, D., and Dover, Y. (2025). Online Reviews: A Literature Review and Roadmap for Future Research. *International Journal of Research in Marketing*, 42(2):275–297.

Rafieian, O. and Zuo, S. (2025). Personalization, Algorithmic Dependence, and Learning. Working paper, Cornell University.

Sayedi, A. (2024). Pricing in a Duopoly with Observational Learning. Working Paper, University of Washington.

Sun, M. (2012). How Does the Variance of Product Ratings Matter? *Management Science*, 58(4):696–707.

Vellodi, N. (2022). Ratings Design and Barriers to Entry. Working Paper.

Waldfogel, J. (2007). *The Tyranny of the Market: Why You Can't Always Get What You Want.* Harvard University Press.

Wang, X. (2021). Do Fit Opinions Matter? The Impact of Fit Context on Online Product Reviews. *Information Systems Research*, 32(3):839–857.

Yoganarasimhan, H. (2020). Search Personalization Using Machine Learning. *Management Science*, 66(3):1045–1070.

Zimmermann, S., Herrmann, P., Kundisch, D., and Nault, B. R. (2018). Decomposing the Variance of Consumer Ratings and the Impact on Price and Demand. *Information Systems Research*, 29(4):984–1002.

Appendix: Proofs

Monotone Comparative Statics of the Fixed Point. We repeatedly use the following observation. Let $F_\theta(\phi) = \psi(\gamma_\theta(\phi))$ be a family of maps indexed by a parameter θ, where ψ is strictly increasing and γ_θ is continuous and strictly decreasing on $[\lambda, 1]$ for each θ, so that F_θ is strictly decreasing and has a unique fixed point $\phi^*(\theta)$ by Proposition 1. If $\gamma_{\theta'}(\phi) \geq \gamma_\theta(\phi)$ for all $\phi \in [\lambda, 1]$, then $F_{\theta'}(\phi) \geq F_\theta(\phi)$ for all ϕ. At the old fixed point, $F_{\theta'}(\phi^*(\theta)) \geq F_\theta(\phi^*(\theta)) = \phi^*(\theta)$. Since $F_{\theta'}$ is strictly decreasing and id is strictly increasing, their unique crossing must satisfy $\phi^*(\theta') \geq \phi^*(\theta)$, with strict inequality if $\gamma_{\theta'}(\phi^*(\theta)) > \gamma_\theta(\phi^*(\theta))$. Since $\gamma^* = \psi^{-1}(\phi^*)$ and ψ^{-1} is strictly increasing, γ^* is increasing in θ as well. Propositions 4 and 5 apply this with $\theta \in \{\mu, L, \zeta\}$; Proposition 8 applies it with $\theta = \beta$. $\square$

Proof of Lemma 1. The expected coverage is $C_a(n_a, n_b) = \sum_{k \in \mathcal{F}_a} \mathbb{P}(k \text{ covered})$, by linearity of expectation. Partition $\mathcal{F}_a$ into exclusive features $\mathcal{F}_a \setminus \mathcal{F}_b$ (of which there are $|\mathcal{F}_a| - |\mathcal{F}_a \cap \mathcal{F}_b| = L - \mu$ by Assumption 3) and shared features $\mathcal{F}_a \cap \mathcal{F}_b$ (of which there are μ).

Exclusive features. Consider feature $k \in \mathcal{F}_a \setminus \mathcal{F}_b$. Since $k \notin \mathcal{F}_b$, no type-b reviewer mentions k. Each type-a reviewer mentions k with probability $1/L$ (Assumption 4); draws are independent across reviewers by assumption. Feature k is covered if at least one of the n_a type-a reviewers mentions it:

$$\mathbb{P}(k \text{ covered}) = 1 - \mathbb{P}(\text{no type-}a \text{ reviewer mentions } k) = 1 - (1 - 1/L)^{n_a}.$$

Summing over the $L - \mu$ exclusive features and using linearity of expectation gives the first term of (1).

Shared features. Consider feature $k \in \mathcal{F}_a \cap \mathcal{F}_b$. Both type-$a$ and type-b reviewers can mention k: each draws uniformly from her own $\mathcal{F}_\tau$ (both of size L), and $k \in \mathcal{F}_\tau$ for both types, so each mentions k with probability $1/L$. Feature k is covered if at least one of the

$n_a + n_b$ total reviewers mentions it:

$$\mathbb{P}(k \text{ covered}) = 1 - (1 - 1/L)^{n_a}(1 - 1/L)^{n_b} = 1 - (1 - 1/L)^{n_a+n_b}.$$

Summing over the μ shared features gives the second term.

The symmetry claim $C_b(n_a, n_b) = C_a(n_b, n_a)$ follows from swapping labels $a \leftrightarrow b$: this maps $(n_a, n_b) \mapsto (n_b, n_a)$ and $(\mathcal{F}_a, \mathcal{F}_b) \mapsto (\mathcal{F}_b, \mathcal{F}_a)$, leaving the formula invariant. □

Proof of Lemma 2. *Part (i): Quantity.* Fix ϕ and increase n. Both $n_a = \phi n$ and $n_b = (1 - \phi)n$ increase. In (1), each term is of the form $c[1 - (1 - 1/L)^m]$ where m increases, so coverage increases. Since $\gamma(\cdot)$ is increasing, Γ increases.

Part (i): Variety. Fix n and consider $d\Gamma/d\phi$ at $\phi > 1/2$. Since $n_a = \phi n$ and $n_b = (1-\phi)n$, the chain rule gives $dC_a/d\phi = n(\partial C_a/\partial n_a - \partial C_a/\partial n_b)$ and similarly for C_b. From (1), for shared features both $\partial/\partial n_a$ and $\partial/\partial n_b$ contribute the same term (both enter through $n_a + n_b = n$, which is fixed), so these cancel in $dC_a/d\phi$. Only exclusive features remain. Hence $d\Gamma/d\phi$ depends only on the marginal coverage of exclusive features, and

$$\frac{d\Gamma}{d\phi} \propto \lambda\gamma'(C_a)\big(1 - \tfrac{1}{L}\big)^{n_a} - (1 - \lambda)\gamma'(C_b)\big(1 - \tfrac{1}{L}\big)^{n_b},$$

where the proportionality constant is $n \ln(L/(L-1))(L - \mu) > 0$. At $\phi > 1/2$: $n_a > n_b$, so $(1 - 1/L)^{n_a} < (1 - 1/L)^{n_b}$ (the minority's marginal coverage is higher). Since $C_b < C_a$ when $n_b < n_a$, the concavity of γ gives $\gamma'(C_b) \geq \gamma'(C_a)$. Both effects make the second (negative) term dominate: the ratio of the second to the first is at least $\frac{(1-\lambda)\gamma'(C_b)(1-1/L)^{n_b}}{\lambda\gamma'(C_a)(1-1/L)^{n_a}} = \frac{1-\lambda}{\lambda} \cdot \frac{\gamma'(C_b)}{\gamma'(C_a)} \cdot (1-1/L)^{n_b-n_a}$. The first factor is less than 1, but the third factor is $(1-1/L)^{-n(2\phi-1)}$, which grows exponentially in n. For any fixed $\phi > 1/2$ and $\lambda < 1$, there exists $\underline{n}(\phi, \lambda)$ such that the ratio exceeds 1 for all $n \geq \underline{n}$, giving $d\Gamma/d\phi < 0$.

Part (ii): Diverse limit. As $n \to \infty$ with $\phi \in (0, 1)$ fixed, both $n_a = \phi n \to \infty$ and $n_b = (1 - \phi)n \to \infty$. Hence $(1 - 1/L)^{n_a} \to 0$ and $(1 - 1/L)^{n_b} \to 0$. From (1): $C_a \to (L-\mu)\cdot 1 + \mu \cdot 1 = L$. By symmetry, $C_b \to L$. Therefore $\Gamma \to \lambda\gamma(L) + (1-\lambda)\gamma(L) = \gamma(L) = \hat{\gamma}$.

Part (iii): Homogeneous limit. With $\phi = 1$: $n_b = 0$, so $(1 - 1/L)^{n_b} = 1$. From (1): $C_a \to (L - \mu) + \mu = L$. For C_b: by symmetry, $C_b(n_a, 0) = C_a(0, n_a)$. From (1) with $n_a = 0$ and $n_b = n_a$: $C_a(0, n_a) = (L - \mu)[1 - 1] + \mu[1 - (1 - 1/L)^{n_a}] \to \mu$. Hence $\Gamma \to \lambda\gamma(L) + (1 - \lambda)\gamma(\mu) = \bar{\gamma}$.

Part (iv). $\bar{\gamma} = \lambda\gamma(L) + (1 - \lambda)\gamma(\mu)$. Since $\gamma(\cdot)$ is strictly increasing, $\bar{\gamma}$ is strictly increasing in μ. Similarly, $\hat{\gamma} = \gamma(L)$ is strictly increasing in L (more features means more signals) and in ζ (each signal is more informative).

For the inequality $\hat{\gamma} > \bar{\gamma}$: $\hat{\gamma} - \bar{\gamma} = \gamma(L) - [\lambda\gamma(L) + (1 - \lambda)\gamma(\mu)] = (1 - \lambda)[\gamma(L) - \gamma(\mu)]$. Since $1 - \lambda > 0$ and $\gamma(L) > \gamma(\mu)$ whenever $L > \mu$ (which holds by assumption $\mu < L$), we have $\hat{\gamma} > \bar{\gamma}$. □

Proof of Proposition 1. The proof uses only the three abstract conditions on γ.

Existence. At $\phi = \lambda$: $\gamma(\lambda) > 1/2$ by condition (ii), so $F(\lambda) = \psi(\gamma(\lambda)) > \lambda$ (since $\psi(\gamma) > \lambda$ iff $\gamma > 1/2$). At $\phi = 1$: $\gamma(1) < 1$ by condition (iii), so $F(1) = \psi(\gamma(1)) < 1$. By continuity of F, the intermediate value theorem gives $\phi^* \in (\lambda, 1)$ with $F(\phi^*) = \phi^*$.

Uniqueness. γ is strictly decreasing on $(\lambda, 1)$ by condition (i). Since ψ is strictly increasing, $F = \psi \circ \gamma$ is strictly decreasing. A strictly decreasing function crosses the identity at most once.

Stability. The ODE $\dot{\phi} = F(\phi) - \phi$ has ϕ^* as its unique equilibrium on $[\lambda, 1]$. Since F is strictly decreasing and crosses the identity once, $F(\phi) > \phi$ for $\phi < \phi^*$ and $F(\phi) < \phi$ for $\phi > \phi^*$. Hence ϕ^* is globally asymptotically stable for the ODE. The stochastic approximation argument (Robbins–Monro) extends to the discrete-time stochastic system; the full verification is given in the quadratic-specification proof below.

Perfect matching inconsistent. $F(1) = \psi(\gamma(1)) < 1$ since $\gamma(1) < 1$ by condition (iii). □

Proof of Proposition 1 (quadratic specification). At $\phi = \lambda$: $\gamma(\lambda) = \bar{\gamma} + (\hat{\gamma} - \bar{\gamma}) \cdot 4\lambda(1 - \lambda)$. Since $\lambda \in (1/2, 1)$, $4\lambda(1 - \lambda) \in (0, 1)$, so $\gamma(\lambda) > \bar{\gamma}$. We need $\bar{\gamma} > 1/2$. From the definition: $\bar{\gamma} = \lambda\gamma(L) + (1 - \lambda)\gamma(\mu)$. For any $L \geq 1$ and $\zeta > 1/2$: $\gamma(L) > 1/2$.

Also $\gamma(\mu) \geq \gamma(0) = 1/2$. Hence $\bar{\gamma} \geq \lambda \cdot (\text{something} > 1/2) + (1-\lambda) \cdot 1/2 > 1/2$. (The hypothesis $\hat{\gamma} > 1/2$ ensures this formally.) Since $\gamma(\lambda) > 1/2$, we have $\psi(\gamma(\lambda)) > \lambda$. To see this: $\psi(\gamma) = \lambda\gamma/[\lambda\gamma + (1-\lambda)(1-\gamma)]$. $\psi(\gamma) > \lambda$ iff $\lambda\gamma > \lambda[\lambda\gamma + (1-\lambda)(1-\gamma)]$ iff $\gamma(1-\lambda) > (1-\lambda)(1-\gamma)$ iff $\gamma > 1-\gamma$ iff $\gamma > 1/2$. So $F(\lambda) = \psi(\gamma(\lambda)) > \lambda$.

At $\phi = 1$: $4\phi(1-\phi) = 0$, so $\gamma(1) = \bar{\gamma}$. Then $F(1) = \psi(\bar{\gamma})$. We need $\psi(\bar{\gamma}) < 1$. Since $\psi(\gamma) < 1$ iff $(1-\lambda)(1-\gamma) > 0$ iff $\gamma < 1$, and $\bar{\gamma} < 1$ by hypothesis, $F(1) < 1$.

Since F is continuous on $[\lambda, 1]$, $F(\lambda) > \lambda$, and $F(1) < 1$, by the intermediate value theorem there exists $\phi^* \in (\lambda, 1)$ with $F(\phi^*) = \phi^*$.

Uniqueness. We show F is strictly decreasing on $[\lambda, 1]$. The diversity index $d(\phi) = 4\phi(1-\phi)$ has derivative $d'(\phi) = 4(1-2\phi) < 0$ for $\phi > 1/2$. Since $\lambda > 1/2$, $d(\phi)$ is strictly decreasing on $[\lambda, 1]$. Since $\hat{\gamma} > \bar{\gamma}$ (Proposition 2(iv)), $\gamma(\phi) = \bar{\gamma} + (\hat{\gamma} - \bar{\gamma})d(\phi)$ is strictly decreasing. Since $\psi(\gamma) = \lambda\gamma/[\lambda\gamma + (1-\lambda)(1-\gamma)]$ is strictly increasing in γ (the numerator increases and the denominator decreases), the composition $F = \psi \circ \gamma$ is strictly decreasing.

A strictly decreasing function and the strictly increasing identity function $\mathrm{id}(\phi) = \phi$ cross at most once. Since existence gives at least one crossing, there is exactly one.

Stability. The update rule (5) takes the form

$$\phi_{t+1} = \phi_t + a_t[h(\phi_t) + \xi_t],$$

where $a_t = 1/(n_t + 1)$, $h(\phi) = F(\phi) - \phi$, and $\xi_t = \hat{\phi}_t - \psi(\gamma_t)$ is the deviation of the realized flow composition $\hat{\phi}_t$ from its conditional mean $\psi(\gamma_t)$, so that $\mathbb{E}[\xi_t \mid \mathcal{F}_t] = 0$. We verify the conditions of a standard stochastic approximation theorem (see, e.g., Borkar, 2008, Theorem 2.1; the compact state space $[\lambda, 1]$ simplifies the verification):

(i) Step sizes. Since $\gamma(\phi) \leq \gamma(\lambda) < 1$ for all $\phi \in [\lambda, 1]$, each product receives at least a fraction $(1-\lambda)(1-\gamma(\lambda)) > 0$ of the unit mass each period (from type-b consumers who buy A). Hence $n_t \geq c \cdot t$ for $c = (1-\lambda)(1-\gamma(\lambda)) > 0$, ensuring $a_t = O(1/t)$, $a_t \to 0$, $\sum a_t = \infty$, and $\sum a_t^2 < \infty$.

(ii) *Mean field.* $h(\phi) = F(\phi) - \phi$ is Lipschitz continuous on $[\lambda, 1]$ (since F is a smooth composition of ψ and γ), with $h(\phi^*) = 0$.

(iii) *Noise.* ξ_t is a martingale difference with bounded second moment: $\mathbb{E}[\xi_t^2 \mid \mathcal{F}_t] \leq 1$ since $\hat{\phi}_t \in [0, 1]$.

(iv) *Global attractivity.* The ODE $\dot{\phi} = h(\phi)$ has ϕ^* as its unique equilibrium on $[\lambda, 1]$, and $h(\phi) > 0$ for $\phi < \phi^*$, $h(\phi) < 0$ for $\phi > \phi^*$ (since F is strictly decreasing and crosses the identity once). Hence ϕ^* is globally asymptotically stable for the ODE.

These conditions guarantee $\phi_t \to \phi^*$ almost surely.

Remaining claims. Since $\phi^* \in (\lambda, 1)$ by construction, $\phi^* < 1$ (partial matching). Since $\gamma(\phi)$ is strictly decreasing on $[\lambda, 1]$ and $\phi^* \in (\lambda, 1)$, we have $\gamma^* = \gamma(\phi^*) \in (\gamma(1), \gamma(\lambda)) = (\bar{\gamma}, \bar{\gamma} + (\hat{\gamma} - \bar{\gamma}) \cdot 4\lambda(1 - \lambda)) \subset (\bar{\gamma}, \hat{\gamma})$. At the steady state, each consumer identifies the correct product with probability γ^* and buys it, so the fraction of all consumers correctly matched is γ^*. □

Proof of Propositions 4 and 5. *Part (i): ϕ^* and γ^* increasing in μ.* Increasing μ raises $\bar{\gamma} = \lambda\gamma(L) + (1-\lambda)\gamma(\mu)$ (since $\gamma(\mu)$ is increasing in μ). This shifts $\gamma(\phi) = \bar{\gamma} + (\hat{\gamma} - \bar{\gamma}) \cdot 4\phi(1-\phi)$ upward for every ϕ. The monotone comparative statics result (above) then gives ϕ^* and $\gamma^* = \psi^{-1}(\phi^*)$ both increasing in μ.

Part (ii). The argument is identical: increasing L or ζ raises both $\hat{\gamma}$ and $\bar{\gamma}$ (Proposition 2(iv)), shifting $\gamma(\phi)$ upward pointwise. The monotone comparative statics result gives ϕ^* and γ^* increasing.

Part (iii): Diminishing marginal effect of μ. $\partial\bar{\gamma}/\partial\mu = (1 - \lambda)\gamma'(\mu)$. From the formula $\gamma(m) = \zeta^m/[\zeta^m + (1 - \zeta)^m]$, write $\ell(m) = (\zeta/(1 - \zeta))^m$, so $\gamma(m) = \ell/(\ell + 1)$. Differentiating: $\gamma'(m) = \ell'/(\ell + 1)^2 = \ell \ln(\zeta/(1 - \zeta))/(\ell + 1)^2 = \gamma(m)(1 - \gamma(m)) \ln(\zeta/(1 - \zeta))$.

The factor $\gamma(m)(1 - \gamma(m))$ achieves its maximum of $1/4$ at $\gamma = 1/2$ (i.e., $m = 0$) and is strictly decreasing for $\gamma > 1/2$ (i.e., $m > 0$). Hence $\gamma'(\mu)$ is highest at $\mu = 0$, and $\partial\bar{\gamma}/\partial\mu$

is diminishing. By the monotone comparative statics result, the diminishing marginal effect on $\bar{\gamma}$ translates into a diminishing marginal effect on γ^*.

Part (iv): Baseline quality as the binding constraint. From (3): $\partial\gamma(\phi)/\partial\bar{\gamma} = 1-4\phi(1-\phi)$ and $\partial\gamma(\phi)/\partial\hat{\gamma} = 4\phi(1-\phi)$. We need $1-4\phi(1-\phi) > 4\phi(1-\phi)$, i.e., $8\phi^*(1-\phi^*) < 1$, which holds iff $\phi^* > \underline{\phi} \equiv (2+\sqrt{2})/4$. By the monotone comparative statics result, ϕ^* is increasing in $\bar{\gamma}$, L, and ζ, so there exists a threshold in each primitive above which the condition is satisfied. Below $\underline{\phi}$, the diversity gap $\hat{\gamma} - \bar{\gamma}$ matters more than the floor itself. □

Proof of Proposition 8. When a fraction β of consumers follow the recommender, recommender-routed buyers of product A have type-a share $\phi_{\text{rec}} = \psi(\gamma_R)$. The review base composition satisfies the fixed-point equation $\phi = T_\beta(\phi) \equiv \beta\phi_{\text{rec}} + (1-\beta)\psi(\gamma(\phi))$.

Existence and uniqueness. Since $\psi \circ \gamma$ is strictly decreasing on $[\lambda, 1]$ (Proposition 1), T_β is strictly decreasing in ϕ for each $\beta \in [0, 1)$. Hence T_β crosses the identity exactly once, giving a unique fixed point $\phi^*(\beta)$.

Monotonicity in β. When $\phi_{\text{rec}} > \psi(\gamma(\phi))$ – which holds at $\phi = \phi^*(0)$ whenever $\gamma_R > \gamma^*$ – raising β raises $T_\beta(\phi)$ pointwise. At the old fixed point $\phi^*(0)$: $T_\beta(\phi^*(0)) > T_0(\phi^*(0)) = \phi^*(0)$. Since T_β is strictly decreasing and id is strictly increasing, the new crossing point satisfies $\phi^*(\beta) > \phi^*(0)$. The same argument applies for incremental increases in β.

Part (i): When $\gamma_R > \gamma^*$: $\phi_{\text{rec}} = \psi(\gamma_R) > \psi(\gamma^*) = \phi^*$ (since ψ is increasing). By the monotonicity argument, $\phi^*(\beta)$ is increasing in β, so the review base becomes more homogeneous. Since γ is decreasing on $(\lambda, 1)$, $\gamma_H^*(\beta) = \gamma(\phi^*(\beta))$ decreases.

Part (ii): When $\gamma_R < \gamma^*$: $\phi_{\text{rec}} = \psi(\gamma_R) < \phi^*$. Raising β now lowers $T_\beta(\phi)$ at the old fixed point, so $\phi^*(\beta)$ decreases, diversifying the review base. Since γ is decreasing, $\gamma_H^*(\beta)$ increases. □

Proof of Proposition 2. The inverse selection function is derived by solving $\phi = \psi(\gamma)$ for γ:

$$\psi^{-1}(\phi) = \frac{\phi(1-\lambda)}{\lambda(1-\phi) + \phi(1-\lambda)}.$$

The fixed-point condition $\psi^{-1}(\phi^*) = \gamma(\phi^*)$ becomes, in terms of $m = 1 - \phi$:

$$\frac{(1-m)(1-\lambda)}{(1-\lambda) + m(2\lambda - 1)} = \bar{\gamma} + 4\Delta m(1-m).$$

Define $f(m) = \psi^{-1}(1-m) - \gamma(1-m)$. We need $f(m^*) = 0$. Note $f(0) = 1 - \bar{\gamma} > 0$ and $f'(0) = -(R + 4\Delta) < 0$, so f is positive and decreasing near $m = 0$.

Lower bound. Expanding around $m = 0$: $\psi^{-1}(1-m) \approx 1 - Rm$ (since $\frac{d}{dm}\psi^{-1}(1-m)\big|_{m=0} = -R$, where $R = \lambda/(1-\lambda)$) and $\gamma(1-m) \approx \bar{\gamma} + 4\Delta m$. The linearized equation $1 - Rm = \bar{\gamma} + 4\Delta m$ yields $m = (1-\bar{\gamma})/(R + 4\Delta)$. To establish this as a lower bound, note that $\psi^{-1}(\phi)$ is strictly convex in ϕ on $(0,1)$: differentiating twice gives $(\psi^{-1})''(\phi) = 2\lambda(1-\lambda)(2\lambda-1)/[\lambda(1-\phi)+\phi(1-\lambda)]^3 > 0$ for $\lambda > 1/2$. Meanwhile, $\gamma(\phi) = \bar{\gamma} + (\hat{\gamma} - \bar{\gamma}) \cdot 4\phi(1-\phi)$ is strictly concave: $\gamma''(\phi) = -8(\hat{\gamma} - \bar{\gamma}) < 0$. Hence $f(m) = \psi^{-1}(1-m) - \gamma(1-m)$ is convex in m (convex minus concave), and the linearized root undershoots: $m^* \geq (1-\bar{\gamma})/(R + 4\Delta)$.

Upper bound. At any steady state, $\gamma^* = \bar{\gamma} + 4\Delta m^*(1-m^*) > \bar{\gamma}$. Since ψ is increasing, $\phi^* = \psi(\gamma^*) > \psi(\bar{\gamma})$, so $m^* = 1 - \phi^* < 1 - \psi(\bar{\gamma})$.

Match efficiency. The approximation $\gamma^* \approx 1 - Rm^*$ follows from the linear expansion: $\psi^{-1}(1-m) \approx 1 - Rm$ and the equilibrium condition $\psi^{-1}(\phi^*) = \gamma^*$. The approximation error is quadratic in m^*, arising from the higher-order terms in the Taylor expansion of ψ^{-1}. □

Proof of Corollary 1. The efficiency gap closed is $(\gamma^* - \lambda)/(1-\lambda)$. From the comparative statics (Proposition 4), γ^* is increasing in μ and ζ. As $\bar{\gamma} \to 1$ (which requires both μ close to L and $\gamma(L)$ close to 1, the latter holding when L is large or ζ is close to 1), $\Delta \to 0$ and $m^* \to (1-\bar{\gamma})/(R + 4\Delta) \to 0$, so $\gamma^* \to 1$ and the fraction approaches 1. □

Proof of Proposition 6. Let α denote the fraction of score-followers. The composition of A's buyer pool satisfies

$$\phi = \alpha\lambda + (1-\alpha)\psi(\gamma(\phi)), \tag{9}$$

which implicitly defines $\phi^*(\alpha)$ and hence $\gamma^*(\alpha) = \gamma(\phi^*(\alpha))$. Score-followers match at rate λ;

text-readers match at rate $\gamma^*(\alpha)$. Hence $M(\alpha) = \alpha\lambda + (1-\alpha)\gamma^*(\alpha)$ with $M(0) = \gamma^*$ and $M(1) = \lambda$.

Implicit differentiation. Write $G(\phi, \alpha) = \alpha\lambda + (1-\alpha)\psi(\gamma(\phi)) - \phi$. At a fixed point, $G = 0$. Since $G_\phi = (1-\alpha)\psi'(\gamma)\gamma'(\phi) - 1 < -1$ (because $\psi' > 0$ and $\gamma' < 0$), the implicit function theorem gives

$$\frac{d\phi^*}{d\alpha} = -\frac{G_\alpha}{G_\phi} = \frac{\lambda - \psi(\gamma^*)}{1 - (1-\alpha)\psi'(\gamma^*)\gamma'(\phi^*)} = \frac{\lambda - \phi^*}{1 + (1-\alpha)|\psi'\gamma'|} < 0,$$

where $\lambda - \phi^* < 0$ and the denominator exceeds 1. Score-followers reduce ϕ^* (diversify the pool). Differentiating M:

$$\left.\frac{dM}{d\alpha}\right|_{\alpha=0} = (\lambda - \gamma^*) + \gamma'(\phi^*) \cdot \frac{\lambda - \phi^*}{1 - \psi'(\gamma^*)\gamma'(\phi^*)}. \tag{10}$$

The first term is the direct loss from replacing informed consumers with uninformed ones. The second is the indirect diversity gain from a less homogeneous pool.

Existence and uniqueness of $\bar{\lambda}$. Define $\Phi(\lambda) \equiv dM/d\alpha\big|_{\alpha=0}$; the threshold $\bar{\lambda}$ is defined as the largest zero of Φ. We show $\Phi < 0$ near $\lambda = 1/2$ and $\Phi > 0$ for λ sufficiently close to 1.

Write the indirect gain as $B(\lambda) = |\gamma'(\phi^*)| \cdot (\phi^* - \lambda)/(1 + |\psi'\gamma'|)$ and the direct loss as $D(\lambda) = \gamma^* - \lambda$, so $\Phi = B - D$.

(a) As $\lambda \to 1$: $\phi^* \to 1$, so $\gamma^* = \gamma(\phi^*) \to \gamma(1) = \bar{\gamma}$. Since $\bar{\gamma} < 1$ for all finite L and $\zeta < 1$, we have $D = \gamma^* - \lambda \to \bar{\gamma} - 1 < 0$, i.e., D becomes negative (score-followers actually match *better* than text-readers). Since $B \geq 0$ and $-D > 0$, $\Phi = B - D > 0$.

(b) As $\lambda \to 1/2^+$: at $\lambda = 1/2$ exactly, ψ is the identity ($\psi(\gamma) = \gamma$), so $\phi^* = \gamma^*$ and $\psi'(\gamma^*) = 1$. Hence $D = \gamma^* - 1/2 > 0$, $B = |\gamma'|(\phi^* - 1/2)/(1 + |\gamma'|)$, and $B/D = |\gamma'|/(1 + |\gamma'|) < 1$. So $\Phi = B - D < 0$ at $\lambda = 1/2$. By continuity, $\Phi < 0$ in a neighborhood of $\lambda = 1/2$.

(c) Existence and sign pattern. Since $\Phi = B - D$ with $B \geq 0$ and $D = \gamma^* - \lambda$, we show D is strictly decreasing in λ, which determines the sign pattern.

Total differentiation of the fixed-point system $\phi^* = \psi(\gamma^*, \lambda)$ and $\gamma^* = \gamma(\phi^*, \lambda)$ gives

$$\frac{d\gamma^*}{d\lambda} = \frac{\gamma_\lambda - |\gamma_\phi| \cdot \psi_\lambda}{1 + \psi_\gamma \cdot |\gamma_\phi|},$$

where $\gamma_\lambda = [\gamma(L) - \gamma(\mu)](1 - 4\phi^*(1 - \phi^*)) > 0$, $\psi_\lambda > 0$, $\psi_\gamma = \psi'(\gamma^*) > 0$, and $|\gamma_\phi| = |\gamma'(\phi^*)| > 0$. For $d\gamma^*/d\lambda < 1$, we need $\gamma_\lambda - |\gamma_\phi| \cdot \psi_\lambda < 1 + \psi_\gamma \cdot |\gamma_\phi|$. Since $\gamma_\lambda \leq \gamma(L) - \gamma(\mu) < 1$ (because $\gamma(L) < 1$ and $\gamma(\mu) \geq 1/2$) and $1 + \psi_\gamma \cdot |\gamma_\phi| > 1$: the inequality holds a fortiori. Hence $dD/d\lambda = d\gamma^*/d\lambda - 1 < 0$: D is strictly decreasing.

Since $D(\lambda)$ is strictly decreasing, continuous, positive near $\lambda = 1/2$ (from (b)), and negative near $\lambda = 1$ (from (a)), it has a unique zero λ_0. For $\lambda > \lambda_0$: $D < 0$ and $B \geq 0$, so $\Phi = B - D > 0$ (score-following improves match efficiency). For λ near $1/2$: $\Phi < 0$ (from (b)). Define $\bar{\lambda}$ as the largest value at which $\Phi(\bar{\lambda}) = 0$. Since $\Phi > 0$ for all $\lambda > \lambda_0$, we have $\bar{\lambda} \leq \lambda_0$. For $\lambda > \bar{\lambda}$, either $\lambda > \lambda_0$ (in which case $\Phi > 0$ directly) or $\lambda \in (\bar{\lambda}, \lambda_0]$ (in which case $\Phi > 0$ by the definition of $\bar{\lambda}$ as the largest zero).[32]

Comparative statics of $\bar{\lambda}$. We show $\bar{\lambda}$ is increasing in μ and ζ. Increasing μ raises $\bar{\gamma}$ and hence γ^* (Proposition 4(i)). At any fixed λ, the direct loss $D = \gamma^* - \lambda$ increases, shifting $\Phi = B - D$ downward. The region $\{\lambda : \Phi(\lambda) < 0\}$ expands, so $\bar{\lambda}$ shifts right.[33] The argument for ζ is analogous: increasing ζ raises both $\hat{\gamma}$ and $\bar{\gamma}$, strengthening the text signal at every composition. This increases D at each λ, shifting Φ downward and raising $\bar{\lambda}$.[34] □

Proof of Corollary 2. *Part (i).* At the steady state, product A's type-b reviewer stock is $n_b^A = (1 - \lambda)(1 - \gamma^*)n$, while product B's type-a reviewer stock is $n_a^B = \lambda(1 - \gamma^*)n$. Since $n_b^A / n_a^B = (1 - \lambda)/\lambda = 1/R < 1$, the minority's reviewer stock on its "off" product is smaller. Coverage exhibits diminishing returns (coupon-collector): the marginal coverage of a type-b

[32]Numerical computation confirms that Φ has a unique zero and is strictly increasing across the full parameter space, so $\bar{\lambda}$ is also the *smallest* zero. The analytical argument establishes the weaker claim that suffices for the proposition.

[33]More precisely: if $\Phi(\lambda; \mu) < 0$ for all $\lambda \leq \bar{\lambda}(\mu)$, then increasing μ makes Φ more negative at $\bar{\lambda}(\mu)$, so the largest zero must lie strictly to the right.

[34]For ζ, the indirect gain B also shifts; numerical verification confirms the net effect is monotone across the full parameter space.

review on A for exclusive features is $\frac{L-\mu}{L}(1-\frac{1}{L})^{n_b^A}$, which is larger when n_b^A is smaller. Hence each minority reviewer's marginal contribution to A's text accuracy exceeds each majority reviewer's marginal contribution to B's, by a factor that is increasing in R.

Part (ii). $R = \lambda/(1-\lambda)$ is strictly increasing in λ. □

Proof of Proposition 3. At the steady state, product A's text accuracy for type-b readers depends on coverage of $\mathcal{F}_b$. The exclusive features $\mathcal{F}_b \setminus \mathcal{F}_a$ (of which there are $L-\mu$) are covered only by type-b reviewers, each of whom mentions a given feature with probability $1/L$. The expected number of exclusive features covered is $(L-\mu)[1-(1-1/L)^{n_b^A}]$, giving total type-$b$ coverage of $m_b^* = \mu + (L-\mu)[1-(1-1/L)^{n_b^A}]$ (shared features are fully covered in the large-n regime). The marginal coverage from one additional type-b review is

$$\frac{\partial C_b^{\text{excl}}}{\partial n_b^A} = \frac{L-\mu}{L}\Big(1-\frac{1}{L}\Big)^{n_b^A},$$

which is the probability that the new review covers a previously uncovered exclusive feature. If it does, the type-b reader's coverage increases from m_b^* to m_b^*+1, improving her accuracy by $\gamma(m_b^*+1)-\gamma(m_b^*) \approx \gamma'(m_b^*)$. The marginal efficiency gain is the minority population weight $(1-\lambda)$ times the signal value $\gamma'(m_b^*)$ times the marginal coverage probability, giving (8). Part (i): increasing μ raises m_b^* (more shared features covered), and since γ' is decreasing for $m > 0$, the signal value falls; additionally, $(L-\mu)/L$ falls. Part (ii): increasing n_b^A raises m_b^* (reducing $\gamma'(m_b^*)$) and reduces the coupon-collector term. Part (iii): for fixed L, increasing $L-\mu$ (i.e., decreasing μ) raises $(L-\mu)/L$ directly and lowers m_b^* (fewer shared features means lower baseline coverage), increasing both the coupon-collector marginal and the signal value. □

Proof of Proposition 7. *Part (i).* Each type-a reviewer mentions a feature from $\mathcal{F}_a$ uniformly at random; each type-b reviewer mentions a feature from $\mathcal{F}_b$ uniformly at random. At the steady state, ϕ^* fraction of reviewers are type a and $1-\phi^*$ are type b. A feature f's mention probability per review is:

$$p(f) = \begin{cases} \phi^*/L + (1-\phi^*)/L = 1/L & \text{if } f \in \mathcal{F}_a \cap \mathcal{F}_b, \\ \phi^*/L & \text{if } f \in \mathcal{F}_a \setminus \mathcal{F}_b, \\ (1-\phi^*)/L & \text{if } f \in \mathcal{F}_b \setminus \mathcal{F}_a. \end{cases}$$

Since $\phi^* > 1/2$: $1/L > \phi^*/L > (1-\phi^*)/L$. By the law of large numbers, empirical mention frequencies converge to these values as $n \to \infty$, so the ranking by frequency places the μ shared features first, followed by the $L - \mu$ exclusive-a features, followed by the $L - \mu$ exclusive-b features, with probability approaching 1. For $k > \mu$, the top-k list includes all μ shared features and $k - \mu$ exclusive-a features. No exclusive-b features appear. A type-b reader attends to $\mathcal{F}_b = (\mathcal{F}_b \setminus \mathcal{F}_a) \cup (\mathcal{F}_a \cap \mathcal{F}_b)$; of the features in the summary, only the μ shared features are in her attention set. Her coverage is $\min(k, \mu) = \mu$, yielding accuracy $\gamma(\mu)$.

Part (ii). A type-a reader attends to $\mathcal{F}_a$; of the top-k features, both shared and exclusive-a features are in her set, giving coverage $\min(k, L)$. A type-b reader has coverage $\min(k, \mu)$. Hence $\Gamma_{\text{sum}}(k) = \lambda\gamma(\min(k, L)) + (1-\lambda)\gamma(\min(k, \mu))$. For $\mu \le k \le L$: the type-b reader's coverage is capped at μ (no $\mathcal{F}_b$-exclusive features appear in the top-k list). At $k = L$: $\Gamma_{\text{sum}} = \lambda\gamma(L) + (1-\lambda)\gamma(\mu) = \bar{\gamma}$. For $\mu \le k < L$: $\gamma(k) < \gamma(L)$ while $\gamma(\min(k, \mu)) = \gamma(\mu)$, so $\Gamma_{\text{sum}}(k) < \bar{\gamma}$. (For $k > L$, the summary begins to include $\mathcal{F}_b$-exclusive features, and Γ_{sum} increases above $\bar{\gamma}$; the policy-relevant regime is $k \le L$.)

Part (iii). At the steady state with full text, $\gamma^* > \bar{\gamma}$ (Proposition 1). For $k \le L$: $\Gamma_{\text{sum}}(k) \le \bar{\gamma} < \gamma^*$. We show $dM/d\alpha\big|_{\alpha=0} < 0$ and $M(1) < \gamma^*$, establishing that introducing summary-readers strictly reduces match efficiency at the margin and in the limit.

Endpoint. $M(1) = \Gamma_{\text{sum}}(k) \le \bar{\gamma} < \gamma^*$.

Marginal effect. Summary-readers diversify the buyer pool: their composition $\psi(\Gamma_{\text{sum}})$ is strictly below $\phi^* = \psi(\gamma^*)$ (since $\Gamma_{\text{sum}} < \gamma^*$ and ψ is increasing). The equilibrium with summary fraction α satisfies $\phi = \alpha\psi(\Gamma_{\text{sum}}) + (1-\alpha)\psi(\gamma(\phi))$. Implicit differentiation (as in the proof of Proposition 6) gives

$$\left.\frac{dM}{d\alpha}\right|_{\alpha=0} = \big(\Gamma_{\text{sum}} - \gamma^*\big) + |\gamma'(\phi^*)| \cdot \frac{\phi^* - \psi(\Gamma_{\text{sum}})}{1 + |\psi'(\gamma^*)\gamma'(\phi^*)|}.$$

The first term (direct loss) equals $-(\gamma^* - \Gamma_{\text{sum}})$. The second (indirect diversity gain) equals $|\gamma'| \cdot (\phi^* - \psi_S)/(1 + |F'(\phi^*)|)$, where $\psi_S = \psi(\Gamma_{\text{sum}})$ and $|F'(\phi^*)| = \psi'(\gamma^*) \cdot |\gamma'(\phi^*)|$. By the mean value theorem, $\phi^* - \psi_S = \psi(\gamma^*) - \psi(\Gamma_{\text{sum}}) = \psi'(\xi)(\gamma^* - \Gamma_{\text{sum}})$ for some $\xi \in (\Gamma_{\text{sum}}, \gamma^*)$. Substituting:

$$\left.\frac{dM}{d\alpha}\right|_{\alpha=0} = (\gamma^* - \Gamma_{\text{sum}})\left[-1 + \frac{|\gamma'| \cdot \psi'(\xi)}{1 + |F'(\phi^*)|}\right].$$

Since $\gamma^* - \Gamma_{\text{sum}} > 0$, the sign equals the sign of the bracketed term. We need $|\gamma'| \cdot \psi'(\xi) < 1 + |F'|$, i.e., $|\gamma'| \cdot [\psi'(\xi) - \psi'(\gamma^*)] < 1$. Since ψ is strictly concave for $\lambda > 1/2$ ($\psi''(\gamma) = -2\lambda(1-\lambda)(2\lambda-1)/[\lambda\gamma + (1-\lambda)(1-\gamma)]^3 < 0$) and $\xi < \gamma^*$: $\psi'(\xi) > \psi'(\gamma^*)$, so the difference is positive. We bound it via a second application of the mean value theorem: $\psi'(\xi) - \psi'(\gamma^*) = |\psi''(\eta)|(\gamma^* - \xi)$ for some $\eta \in (\xi, \gamma^*)$.

We bound each factor. First, $|\gamma'(\phi^*)| = 4\Delta(2\phi^* - 1) < 4\Delta$. Second, since $\eta > \xi > \Gamma_{\text{sum}} \geq 1/2$ and $|\psi''|$ is decreasing in γ (the denominator $[\lambda\gamma + (1-\lambda)(1-\gamma)]^3$ is increasing): $|\psi''(\eta)| \leq |\psi''(1/2)| = 16\lambda(1-\lambda)(2\lambda-1)$. Third, $\gamma^* - \xi \leq \gamma^* - \frac{1}{2} < \frac{1}{2}$ (since $\gamma^* < 1$ by Proposition 1). Hence

$$|\gamma'| \cdot |\psi''(\eta)| \cdot (\gamma^* - \xi) < 4\Delta \cdot 16\lambda(1-\lambda)(2\lambda-1) \cdot \tfrac{1}{2} = 32\Delta\lambda(1-\lambda)(2\lambda-1).$$

Since $\Delta = (1-\lambda)[\gamma(L) - \gamma(\mu)] < (1-\lambda)/2$ (because $\gamma(L) < 1$ and $\gamma(\mu) \geq 1/2$):

$$32\Delta\lambda(1-\lambda)(2\lambda-1) < 16\lambda(1-\lambda)^2(2\lambda-1).$$

The right-hand side is the product of four positive terms λ, $(1-\lambda)$, $(1-\lambda)$, and $(2\lambda-1)$ times 16. Their sum is $1 + \lambda \leq 2$, so by the AM–GM inequality their product is at most

$((1+\lambda)/4)^4 \leq (1/2)^4 = 1/16$, and the inequality is strict because equality in AM–GM requires all factors to be equal, which would entail $\lambda = 1-\lambda$ and $\lambda = 2\lambda - 1$ simultaneously – a contradiction. Hence $16\lambda(1-\lambda)^2(2\lambda-1) < 16 \cdot (1/16) = 1$, and $dM/d\alpha\big|_{\alpha=0} < 0$.[35] $\square$

A.10 Numerical Validation of the Quadratic Reduced Form

The reduced-form text accuracy $\gamma(\phi) = \bar{\gamma} + (\hat{\gamma} - \bar{\gamma}) \cdot 4\phi(1-\phi)$ is adopted for tractability. This subsection verifies that it captures the shape of the accuracy derived from the coupon-collector microfoundation.

For a given composition ϕ and total reviews n, the microfounded accuracy is $\Gamma_{\text{micro}}(\phi) = \lambda\gamma(C_a(\phi n, (1-\phi)n)) + (1-\lambda)\gamma(C_b(\phi n, (1-\phi)n))$, where C_τ is the expected coverage from Lemma 1 and $\gamma(m) = \zeta^m/(\zeta^m + (1-\zeta)^m)$. The reduced form matches the microfounded function at $\phi = 1/2$ (maximum diversity) and $\phi = 1$ (full homogeneity) by construction; the question is how well it approximates the interior.

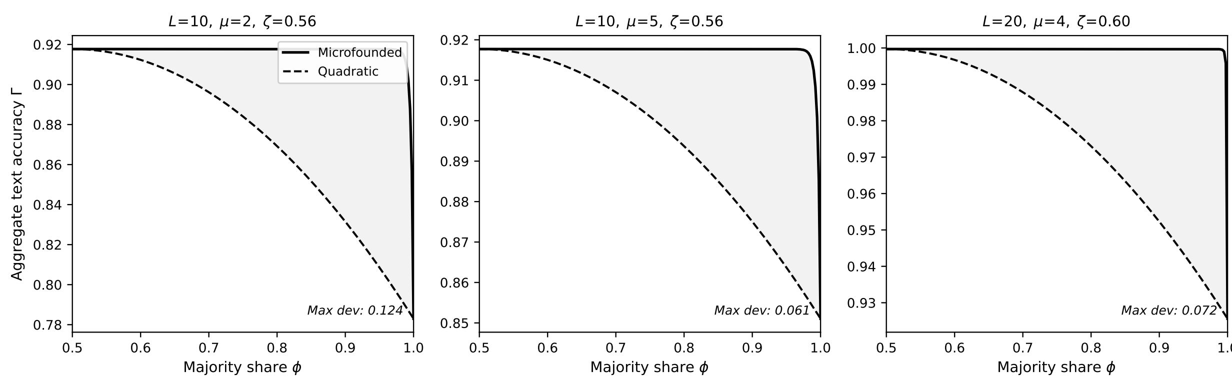

Figure 1: Microfounded text accuracy (solid) versus the quadratic reduced form (dashed) for three parameterizations in the large-n regime (coverage saturated). The two curves match at the endpoints by construction; the interior deviation is largest in low-overlap markets (panel a), where the microfounded function curves more sharply near $\phi = 1$.

Figure 1 shows that the quadratic captures the monotone, concave shape of the microfounded function across all three parameterizations. The deviation is largest near $\phi = 1$,

[35] For $k > L$, the summary includes $\mathcal{F}_b$-exclusive features, and for $k = 2L - \mu$, $\Gamma_{\text{sum}} = \hat{\gamma} > \gamma^*$: a "summary" that reproduces all features strictly improves on the decentralized equilibrium. But such a summary is the full review text, and the interesting regime is $k \ll L$.

where the coupon-collector's geometric coverage decay produces a sharper decline than the quadratic allows; the maximum pointwise deviation grows with n (at $n = 200$: roughly 0.07 for moderate overlap and 0.12 for low overlap; at $n = 500$: roughly 0.09 and 0.15). In low-overlap markets, the microfounded function is slightly more S-shaped than quadratic.

The key qualitative properties that drive our results – monotone decrease on $[\lambda, 1]$, concavity, and boundary values $\gamma(1/2) = \hat{\gamma}$ and $\gamma(1) = \bar{\gamma}$ – hold for both functions. The impossibility result (Proposition 1) and global stability require only these qualitative properties, not the quadratic form. The closed-form diversity equation and comparative statics use the quadratic for tractability; the figure confirms that the quadratic faithfully represents the microfounded tradeoff between pool homogeneity and text accuracy.