CHOICE OVERLAP AND REPUTATION FEEDBACK IN CREATOR–BRAND MATCHING

Anonymous authors
Paper under double-blind review

ABSTRACT

Content platforms create reputation feedback when past bookings improve creators' rankings for future paid promotions. We study how the alternatives available to different brands shape the stability of this feedback and creators' access to bookings. At any positive fixed point, a normalized choice-overlap matrix determines the local stability threshold of the mean booking dynamics. Markets with identical creator positions and effective choice can lie on opposite sides of this threshold. A separate reinforcement threshold governs attraction toward monopoly allocations, allowing stable interior and monopoly states to coexist. We derive a sharp constraint linking effective choice to admissible reinforcement, so choice expansion and reputation weighting must be evaluated jointly. Finite-sample spectral bounds support policy selection under explicit calibration assumptions. Uniform exploration removes boundary equilibria and provides a finite-window access bound, although small exploration shares can still sustain concentrated interior states. Controlled simulations test the stability thresholds, spectral bounds, and event-level access. Among booked creators, offline estimates associate a one-standard-deviation increase in booking evenness with a 7.13% higher audience-tier advancement probability relative to the sample mean; a reported 60-day online comparison shows a 3.4% increase in gross promotion transaction value (GMV). The framework distinguishes creator variety, brand-facing choice, and sustained access as separate design objectives.

1 INTRODUCTION

Content platforms intermediate paid promotion between brands seeking relevant audiences and creators seeking commercial opportunities. A booking is both a transaction and a reputation signal: when rankings reward past commercial success, today's allocation changes tomorrow's visibility. This feedback connects matching design to participation in a two-sided marketplace (Rochet & Tirole, 2003; Armstrong, 2006).

A broad creator roster need not offer each brand close substitutes: discrete-choice models link substitution to option attributes and heterogeneous preferences (Train, 2009, Chapter 6). Yet additional alternatives change the channels through which reputation reallocates future bookings. Our central question is: *How does overlap among brand-facing choices determine reputation feedback, and how should a platform coordinate choice expansion with reinforcement control?* The relevant distinction is three-way: creator coverage, effective choice, and the direction in which feedback is strongest. Neither coverage nor an effective-choice count alone identifies that direction. For example, a roster of 100 creators can offer each brand only two credible alternatives; different brands may consider the same pair or different pairs.

We characterize local stability through a normalized choice-overlap matrix. Two markets with identical creator positions, effective choice, and baseline shares can have opposite stability at the same reinforcement strength. A trace bound links choice expansion to the admissible feedback strength: a higher effective-choice floor raises the minimum spectral exposure at a balanced allocation. These controls must therefore be tuned jointly.

Two thresholds determine the mean-map geometry. Choice overlap governs local stability at a positive allocation, while the reinforcement exponent governs attraction toward monopoly corners.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

Above unit reinforcement, stable interior and monopoly allocations can coexist. A statistical certificate can admit this regime: it controls local response while leaving the risk of escape unresolved. An explicit probability mixture removes boundary equilibria and guarantees persistent access. Its own covariance spectrum controls local stability, but small mixtures can leave near-monopoly interior equilibria. Market size, memory, and booking noise determine whether the local approximation is useful.

Our contributions are:

- We formulate creator–brand matching as a feedback-driven allocation problem: heterogeneous brand needs and past bookings jointly shape future opportunities. Our model separates creator coverage, effective choice, and sustained access (Section 2).
- We establish the overlap-sensitive interior threshold, boundary classification, and sharp choice–reinforcement constraint, together with finite-sample policy bounds that distinguish local stability from access (Sections 3–4).
- We test the mechanism in controlled simulations and examine creator advancement, transaction value, and user response through offline diagnostics and a reported 60-day online comparison (Section 5). Platform analyses provide descriptive context; without candidate slates they cannot test the overlap-based theory directly.

Related work. Increasing-returns and nonlinear urn models explain path-dependent concentration and the unit-exponent transition (Arthur, 1989; Pemantle, 2007). Recommendation-feedback and performative-prediction studies examine how allocation changes later observations or preferences (Chaney et al., 2018; Jiang et al., 2019; Perdomo et al., 2020). Our model fixes creator positions and updates commercial reputation under heterogeneous brand needs, separating the interior spectral threshold from boundary attraction. Stochastic approximation connects mean dynamics to noisy updates (Benaïm, 1999; Kushner & Yin, 2003); the overlap spectrum characterizes our matching rule's local response.

Creator-equilibrium research studies strategic content, specialization, production, and entry (Ben-Porat & Tennenholtz, 2018; Hron et al., 2023; Jagadeesan et al., 2023; Yao et al., 2023; 2024; Wei et al., 2026). We study reputation stability conditional on fixed content. Provider viability (Mladenov et al., 2020), exposure fairness (Singh & Joachims, 2018; Morik et al., 2020), and reciprocal matching (Su et al., 2022) impose different allocation requirements. Individual and aggregate diversity can diverge (Fleder & Hosanagar, 2009); our distinction concerns coverage, effective choice, and spectral response. Matrix concentration (Tropp, 2012) supports local policy evaluation, while explicit exploration bounds access. These guarantees address feedback and participation rather than welfare optimization or causal effects. Appendix I extends the comparison to discrete-choice economics and the mathematical foundations.

2 Matching Model and Choice Diagnostics

2.1 Reputation-dependent booking

A market has $N \geq 2$ creators at fixed positions $x_i \in \mathbb{R}^d$. Independent brand needs $\theta \sim F$ induce match utilities $q_i(\theta)$, decreasing with distance from x_i. Creator i has reputation $a_i^t > 0$. With $\alpha \geq 0$, $\varepsilon > 0$, and $\beta = \alpha/\varepsilon$, define reinforced scores (Pemantle, 2007) and multinomial-logit choice (McFadden, 1974):

$$z_i(\theta; a) = q_i(\theta) + \alpha \log a_i, \tag{1}$$

$$p_i(\theta; a) = \frac{\exp(z_i(\theta; a)/\varepsilon)}{\sum_j \exp(z_j(\theta; a)/\varepsilon)}, \qquad e_i(a) = \mathbb{E}_\theta p_i(\theta; a). \tag{2}$$

One creator is booked at each event, $Y^t \mid \theta_t, a^t \sim \mathrm{Categorical}(p(\theta_t; a^t))$, and

$$a^{t+1} = \rho a^t + Y^t, \qquad 0 < \rho < 1. \tag{3}$$

The conditional mean map is $\mathcal{T}(a) = \rho a + e(a)$; its iterates need not equal the nonlinear process's unconditional mean trajectory. The derivative holds retrieval, eligibility, and match utilities fixed; reputation-dependent retrieval would add a derivative term. The logit temperature ε also denotes the utility-shock scale in the separate pricing model; match utility is not measured campaign performance.

2.2 COVERAGE, EFFECTIVE CHOICE, AND OVERLAP

Definition 1 (Coverage and effective choice)**.** Let $\Sigma = N^{-1}\sum_i (x_i - \bar{x})(x_i - \bar{x})^\top$ and $\delta_D > 0$. Aggregate diversity is $\mathcal{D} = \log\det(\Sigma + \delta_D I)$. At a specified reputation state, define the probability Gram matrix (Train, 2009, cf. Section 6.4) and effective counts:

$$P = \mathbb{E}[p(\theta)p(\theta)^\top], \qquad n(\theta) = \left(\sum_i p_i(\theta)^2\right)^{-1}, \qquad \mathcal{D}_{\mathrm{loc}} = \frac{1}{\mathrm{tr}(P)} \in [1, N].$$

Thus $\mathcal{D}_{\mathrm{loc}}$ is the harmonic mean of $n(\theta)$, not $\mathbb{E}[n(\theta)]$.

For m equally weighted alternatives, $n(\theta) = m$; a retrieved set of m candidates need not satisfy this equality. The worst-case effective count is $\mathcal{D}_{\mathrm{loc}}^{\min} = (\operatorname{ess\,sup}_\theta \sum_i p_i(\theta)^2)^{-1}$. The Herfindahl–Hirschman index (HHI), $\mathrm{HHI}(e) = \sum_i e_i^2$, measures concentration of aggregate booking shares e_i: $1/N$ indicates equal shares, while 1 indicates that one creator receives all bookings. Larger HHI means greater concentration. Equal-demand disjoint niches have $\mathrm{HHI} = 1/N$ but $\mathcal{D}_{\mathrm{loc}} = 1$.

Assumption 1 (Balanced baseline)**.** For the balanced special case, $a^* = \mathbf{1}/[N(1-\rho)]$ satisfies $e(a^*) = \mathbf{1}/N$. The general stability result allows unequal shares.

Definition 2 (Spectral feedback exposure)**.** At a strictly positive fixed point, let $s = e(a^*) = (1-\rho)a^*$, $S = \mathrm{diag}(s)$, $P^* = \mathbb{E}[p(\theta; a^*)p(\theta; a^*)^\top]$, and $B = S^{-1/2}P^*S^{-1/2}$. Define $\sigma_* = 1 - \lambda_{\min}(B) \in [0, 1]$. At a balanced baseline, $B = NP$ and $\sigma_* = \sigma = 1 - N\mu_{\min}(P)$.

Effective choice depends on $\mathrm{tr}(P)$, whereas local stability depends on the smallest eigenvalue of B. Markets can therefore agree on one and differ on the other. Disconnected interchangeable groups can also have $\sigma = 1$.

3 CHOICE OVERLAP AND REPUTATION STABILITY

3.1 COVERAGE AND EFFECTIVE CHOICE ARE INSUFFICIENT

Proposition 3 (Coverage does not identify local choice)**.**

(a) At any state, $\mathcal{D}_{\mathrm{loc}} = (\sum_k \mu_k(P))^{-1}$; at a balanced baseline, $\sigma = 1 - N\mu_{\min}(P)$. Both depend on choice probabilities, whereas $\mathcal{D}$ depends only on creator positions.

(b) In the closure of the choice model, disjoint equal-mass choice has $(\mathcal{D}_{\mathrm{loc}}, \sigma) = (1, 0)$; uniform full overlap has $(N, 1)$.

(c) Fix creator positions $-c, c$. Equal demand at these positions gives $\mathcal{D}_{\mathrm{loc}} \to 1$ as $\varepsilon \downarrow 0$; demand at their midpoint gives $\mathcal{D}_{\mathrm{loc}} = 2$. Aggregate diversity is identical.

(d) For $N = 4$, $P_t = (1-t)I/N + t\mathbf{1}\mathbf{1}^\top/N^2$ gives $(\sigma, \mathcal{D}_{\mathrm{loc}}) = (t, [1 - 3t/4]^{-1})$. At $t = 0.9$ this is $(0.9, 3.08)$; two disconnected uniform pairs instead give $(1, 2)$. Small positive perturbations preserve this increase in σ with a decrease in $\mathcal{D}_{\mathrm{loc}}$. Together with (b), this rules out both global monotone orderings.

Appendix A.1 constructs finite-score approximations. The controlled comparison below keeps both coverage and effective choice fixed.

3.2 THE NORMALIZED-OVERLAP THRESHOLD

Theorem 4 (Normalized-overlap stability)**.** *At any strictly positive fixed point of the mean map,*

$$\lambda_k(J) = \rho + \beta(1-\rho)[1 - \lambda_k(B)].$$

The linearized map is asymptotically stable iff $\beta\sigma_* < 1$. The nonlinear mean map is locally asymptotically stable under this strict inequality and locally unstable when $\beta\sigma_* > 1$. Equality requires higher-order analysis.

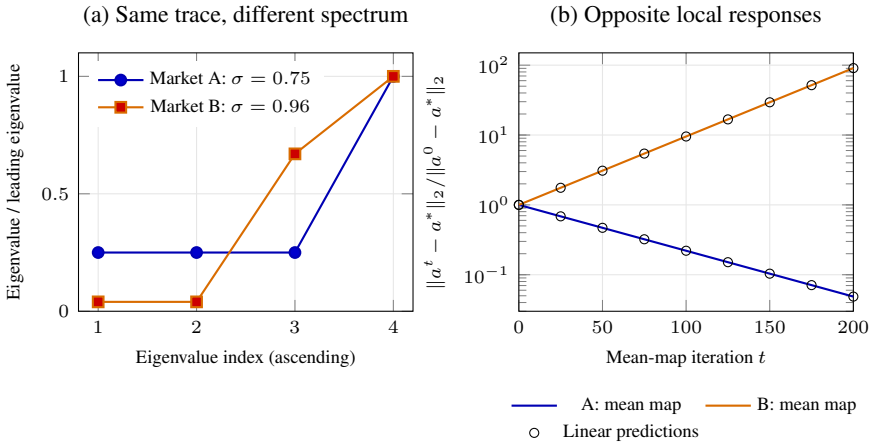

Figure 1: Controlled model comparison. Both markets have the same four creator positions, four candidates per need, balanced booking shares, and $\mathcal{D}_{\text{loc}} = 16/7$. With $\beta = 1.2$ and $\rho = 0.85$, the dominant linear factors are 0.985 and 1.0228. Panel (b) starts from the same 10^{-6} zero-sum perturbation along a common dominant direction. Equal effective choice does not imply equal feedback stability.

Proof sketch. Differentiation gives $J = \rho I + \beta(1-\rho)(I - P^* S^{-1})$ and $S^{-1/2} J S^{1/2} = \rho I + \beta(1-\rho)(I - B)$. The matrix B is positive semidefinite, satisfies $B\sqrt{s} = \sqrt{s}$, and is similar to the row-stochastic matrix $S^{-1}P^*$, so its eigenvalues lie in $[0, 1]$. The spectral radius is therefore $\rho + \beta(1-\rho)\sigma_*$. See Appendix B.1. □

Homogeneity $e(ca) = e(a)$ implies $De(a)a = 0$, explaining the aggregate eigenvalue ρ. Under Assumption 1, $B = NP$.

Proposition 5 (A positive sublinear fixed point)**.** *Suppose $w_i(\theta) = \exp(q_i(\theta)/\varepsilon)$ satisfy $0 < \underline{w} \le w_i(\theta) \le \overline{w} < \infty$ almost surely, and every creator is eligible. For $0 \le \beta < 1$, the mean map has exactly one strictly positive fixed point.*

The strictly concave potential $\mathbb{E} \log \sum_i w_i(\theta) s_i^\beta$ has a unique interior maximizer, whose first-order conditions are $e_i(s) = s_i$ (Appendix B.3).

Proposition 6 (Boundary threshold)**.** *Under the bounded-weight and eligibility assumptions of Proposition 5, write $\nu = (1-\rho)a$ on the unit simplex. For $\beta > 1$, every vertex of $T(\nu) = \rho\nu + (1-\rho)e(\nu)$ is locally attracting, with tangent derivative ρI. For $0 < \beta < 1$, every proper face repels perturbations that introduce a positive missing coordinate. At $\beta = 1$, the invasion factor from vertex i toward j is $\rho + (1-\rho)\mathbb{E}[w_j/w_i]$, so stability depends on match utilities.*

Near a vertex, missing mass b_∂ satisfies $b'_\partial = \rho b_\partial + O(b_\partial^\beta)$ for $\beta > 1$ (Appendix B.4). Thus $1 < \beta < 1/\sigma_*$ can support a stable interior allocation and attracting monopoly corners at the same time. The interior certificate alone does not quantify escape probabilities.

Corollary 7 (Effective choice constrains exposure)**.** *At a balanced fixed point,*

$$\sigma \ge L_N(\mathcal{D}_{\text{loc}}), \qquad L_N(x) = \frac{N}{N-1}\left(1 - \frac{1}{x}\right).$$

Equality holds when the nonleading eigenvalues of NP coincide. Equal effective choice nevertheless need not imply equal exposure or stability.

The smallest nonleading eigenvalue is at most their average, proving the bound (Appendix B.5). Figure 1 gives a finite-score counterexample to the converse.

Corollary 8 (Choice–reinforcement feasibility)**.** *At a balanced fixed point, achieved effective choice $\mathcal{D}_{\text{loc}} \ge m > 1$ and a margin $\beta\sigma \le 1 - \eta$ require*

$$\beta \le \frac{1-\eta}{L_N(m)}.$$

This necessary restriction is sharp; it is not sufficient without a bound on the actual σ.

A higher achieved-choice floor lowers the best-case admissible reinforcement ceiling. This is the floor–cap tension. It does not say that adding a candidate always increases realized exposure: irrelevant candidates can leave probabilities almost unchanged, and different overlap directions can respond differently.

3.3 Local fluctuations and persistent access

Corollary 9 (Local relaxation)**.** *The dominant linear factor is $\lambda_{\max} = \rho + \beta(1-\rho)\sigma_*$; its gap is $(1-\rho)(1-\beta\sigma_*)$.*

Proposition 10 (Frozen-noise amplification)**.** *Let $\zeta = 1-\rho$ and $\nu = \zeta a$, initialized with total mass one. At a stable fixed point, the frozen-coefficient linearization of $u = S^{-1/2}(\nu - s)$ has innovation covariance $\zeta^2(I - \sqrt{s}\sqrt{s}^\top)$. Its stationary variance along a unit contrast eigenvector with factor λ_k is*

$$\mathrm{Var}(v_k^\top u) = \frac{\zeta^2}{1-\lambda_k^2}.$$

For an average of b independent bookings per update, divide this by b.

This standard linear-noise calculation (van Kampen, 2007) shows how the local margin amplifies booking shocks. Appendix C.1 derives the formula; its use for the nonlinear process requires small fluctuations.

At a balanced isotropic configuration, all contrast factors equal λ. The surrogate then predicts the relative per-creator fluctuation scale

$$\sqrt{\mathrm{Var}(N\nu_i)} = \sqrt{\frac{(N-1)(1-\rho)^2}{b(1-\lambda^2)}}.$$

The predicted scale grows with market size at fixed memory and batch size: it is about 0.80 for the 128-creator event baseline and 0.11 for the highest-reinforcement batched check. At fixed interior margin, either $1-\rho \propto 1/N$ or $b \propto N$ controls this growth.

Proposition 11 (Permanent exclusion despite mean stability)**.** *For two identical fully overlapping creators, $0 < \rho < 1$, $0 < \beta < 1$, and positive initial stocks, either creator has strictly positive probability of receiving every future booking, although the balanced mean map is stable.*

Along such a history the losing stock decays as ρ^t, so its booking probabilities are bounded by a summable geometric sequence. The complementary infinite product is positive (Appendix C.2). A separate mixture $p_i^\gamma = (1-\gamma)p_i + \gamma/N$, $\gamma > 0$, ensures infinitely many bookings for every always-eligible creator almost surely (Appendix C.3); Proposition 14 gives the modified stability condition and finite-window bound. Interior stability alone does not guarantee access.

As a consistency check, two fully overlapping creators at $\rho = 0$ have $\sigma_* = 1$ and spectral radius β, matching the exact log-odds recursion (Proposition 16, Appendix C.4).

4 Spectral Certification and Policy Design

4.1 A selection-valid reinforcement cap

Corollary 12 (Policy-specific cap)**.** *Suppose $U \in (0,1]$ bounds σ_* uniformly over the final policies and positive fixed points considered during parameter selection. The choice $\alpha \le (1-\eta)\varepsilon/U$, $\eta \in (0,1)$, ensures $\beta\sigma_* \le 1-\eta$ and spectral gap at least $(1-\rho)\eta$.*

An accepted $\beta > 1$ also has attracting corners under Proposition 6; $\beta = 1$ requires its utility-dependent boundary test. Algorithm 1 separates this unmixed warning from the mixed-policy check below. A certificate for a single pre-treatment policy cannot justify changing its retrieval rule or clipping α to an uncertified value. Use either a uniform certificate over the contemplated parameter range or simultaneous certificates for a prespecified finite family. The worst-case $U = 1$ is valid without calibration. An adjustment to ε changes both β and the choice matrix, so an exploration-temperature adjustment also requires recertification.

Theorem 13 (Finite-family spectral certificates)**.** *Consider K fixed policy–fixed-point pairs with known positive shares s (the pair index is suppressed). Observe probabilities $p_k = p(\theta_k; a^*)$ on n iid needs, independent of policy construction. Suppose $\|S^{-1/2}p_k\|^2 \le L$ almost surely. Define*

$$\widehat{P} = \frac{1}{n}\sum_{k=1}^{n} p_k p_k^\top, \quad \widehat{B} = S^{-1/2}\widehat{P}S^{-1/2}, \quad R_n = \sqrt{\frac{2L\ell}{n}} + \frac{2L\ell}{3n}, \quad \ell = \log\frac{2NK}{\delta_c}.$$

With probability at least $1-\delta_c$, every pair satisfies

$$\sigma_* \le U = \min\{1, \max\{0, 1 - \lambda_{\min}(\widehat{B}) + R_n\}\}.$$

Any selected pair satisfying $\beta U \le 1-\eta$ preserves the margin on this simultaneous-coverage event.

Matrix Bernstein and eigenvalue perturbation give the result (Appendix D.2). One may take $L = 1/\min_i s_i$; a smaller valid leverage bound is sharper. The appendix gives both an estimated-share variant and the improvement $L \mapsto L-1$ in the variance term for $L \ge 2$. A local certificate above unit reinforcement permits stronger reputation weighting while preserving interior response, but also permits coexisting corner attractors. Explicit mixing changes that boundary geometry.

A fitted-model workflow. First fit choice scores on logged eligible sets and bookings. Freeze each candidate policy, solve its fitted mean map for a positive fixed point, and evaluate its probabilities on independent calibration needs. Sampling error, model error, and fixed-point error require separate bounds. If $\|B_{\text{true}} - B_{\text{fit}}\| \le \Delta_B$ at their respective fixed points, the raw upper bound gains Δ_B before clipping to $[0,1]$ (Appendix D.3). A small solver residual alone is insufficient. The fitted-score studies test this workflow within specified synthetic families. The available platform logs contain only booked creators, which cannot identify $\mathbb{E}[pp^\top]$: winner outer products estimate $\mathrm{diag}(s)$ instead.

4.2 A MIXED POLICY WITH AN ACCESS GUARANTEE

Proposition 14 (Mixed-policy stability and access)**.** *Fix an always-eligible pool and $p_i^\gamma = (1-\gamma)p_i + \gamma/N$, $0 < \gamma \le 1$. Under Proposition 5's weight assumptions, the mean map has a fixed point and all fixed points are interior. At such a point, let $e = \mathbb{E}p$, $E = \mathrm{diag}(e)$, $s = (1-\gamma)e + \gamma\mathbf{1}/N$, and $S = \mathrm{diag}(s)$. Define*

$$H_\gamma = (1-\gamma)S^{-1/2}(E - P^*)S^{-1/2}, \qquad \sigma_\gamma = \lambda_{\max}(H_\gamma) < 1.$$

The linearized map is stable iff $\beta\sigma_\gamma < 1$. For any next T events, the expected fraction of creators booked is at least $1-(1-\gamma/N)^T$; every creator is booked infinitely often almost surely.

The Jacobian is similar to $\rho I + \beta(1-\rho)H_\gamma$. At $\gamma = 0$, $H_0 = I - B$, recovering Theorem 4. Appendix D.5 bounds σ_γ from the largest eigenvalue of the empirical conditional-choice covariance, including estimated-share and model-error allowances. Small γ can still support concentrated interior equilibria. A conservative contraction bound and the numerical branch analysis in Appendices D.6 and F.9 address that separate question.

For a target expected access fraction $A_{\text{tar}} \in (0,1)$ within T events, Proposition 14 gives the sufficient rule $\gamma \ge N[1-(1-A_{\text{tar}})^{1/T}]$, provided the right-hand side is at most one. Larger pools and shorter horizons require more exploration under this bound; values above one leave the target uncertified. The guarantee concerns expected coverage of an always-eligible pool. The platform then checks stability at the resulting mixed-policy fixed point, since exploration changes both the fixed-point shares and feedback response.

4.3 CHOICE EXPANSION AND ECONOMIC TRADE-OFFS

The retrieval policy seeks at least m eligible candidates above a relevance threshold, records infeasibility, and measures achieved $\mathcal{D}_{\text{loc}}$. The spectral certificate applies to the entire final retrieval/scoring policy. Algorithm 1 and Figure 11 detail selection, fallback, and logging.

A separate logit pricing model explains the economic trade-off. Unlike the booking-conditioned dynamics, it includes an outside option. At fixed demand, more comparable creators reduce price and

individual profit while increasing transaction probability and brand inclusive value (Propositions 17–18). Aggregate creator profit can rise or fall: Appendix Table 2's example falls by approximately 30.8% as the pool grows from two to five. Appendix E.3 derives the sign condition. These static results motivate separate measurements of brand value and creator earnings; they are not an integrated pricing–reputation equilibrium.

5 Experiments and Observational Diagnostics

Synthetic studies test spectral response, estimation, and event-level access with explicit score matrices and no platform calibration. Platform diagnostics D1–D2 assess the roster proxy; H1 examines creator advancement; H2–H3 examine reported transaction value and user response. This sequence connects the model's allocation mechanism to distinct marketplace outcomes.

5.1 Controlled overlap and finite-sample certification

Matched-market check. Four fixed creators and four equally likely brand types yield strictly positive choice matrices with $\mathcal{D}_{\mathrm{loc}} = 16/7$, baseline shares $1/4$, and identical coverage. Their normalized spectra are $(1, 0.25, 0.25, 0.25)$ and $(1, 0.67, 0.04, 0.04)$. At $\beta = 1.2$, $\rho = 0.85$, the dominant factors are 0.985 and 1.0228. After 200 nonlinear mean-map steps, the same 10^{-6} perturbation shrinks to 0.0487 of its initial size in A and grows to 90.8 in B (Figure 1). Across the parameter grid, finite differences agree with the analytic Jacobian; every tested perturbation contracts in A and expands in B. Appendix F.1 gives the finite distance-score construction.

Calibration protocol. For $N \in \{16, 64, 128\}$, use $C_\kappa = (1-\kappa)I + \kappa \mathbf{1}\mathbf{1}^\top/N$ with uniform brand types, $\kappa \in \{0.2, 0.6\}$, and known balanced shares. Hence $B = C_\kappa^\top C_\kappa$ and $\sigma = 1 - (1-\kappa)^2$. We sample $n \in \{10^3, 10^4, 10^5\}$ needs in 100 independent replicates, using $\delta_c = 0.05$, $K = 8$, and $L = N\|C_{\kappa,1\cdot}\|^2$. At $N = 128, n = 10^5$, the mean certified ceilings are 1.797 and 1.083 times the universal ceiling 0.9. Smaller samples often yield no improvement over that reference. Every reported calibration replicate covers its true exposure; this finite check complements the coverage theorem. Selection stays inside the evaluated family or a valid uniform range. The accepted $\beta = 1.4$, $\kappa = 0.2$ setting also appears in the concentrated event-level baseline below: its valid local certificate leaves boundary attraction intact.

5.2 Local noise and event-level access

Local approximation. With $N = 32$, $\kappa = 0.4$, $\rho = 0.98$, and batches of 256 independent bookings, we simulate 64 independent paths for each $\beta \in \{0.4, 1, 1.4\}$ over 6,000 updates, discarding the first 2,000. Observed-to-predicted dominant-coordinate second-moment ratios are 0.997, 1.023, and 1.050. Appendix Figure 3 compares the local prediction to nonlinear simulations. The single-booking experiment below uses a larger-noise regime.

Event-level protocol. We simulate 64 independent paths of 20,000 single-booking events for each $N \in \{16, 64, 128\}$, starting balanced, with $\rho = 0.995$. Choice expansion changes κ from 0.2 to 0.6; the cap changes β from 1.4 to 0.9. A fifth arm adds a 0.05 uniform-access mixture. All creators remain eligible: this is a probability-level expansion within a fixed pool. Concentration is booking HHI divided by its equal-share reference $1/N$. Access shortfall is $100(1 - A/A_{\mathrm{unif}})$, where A is the fraction booked in the final 2,000 events and $A_{\mathrm{unif}} = 1 - (1 - 1/N)^{2000}$ is expected access under independent uniform booking. Lower values are better for both. For $N = 128$, expansion with the cap has concentration 6.73 times the equal-share reference and an 82.13% access shortfall; the explicit mixture reduces these to 2.16 times and 22.55%. Even the stable baseline concentrates completely over the final windows at $N = 64, 128$. The cap-only arm has a 60.90% shortfall at $N = 128$, compared with effectively none at $N = 16$. These outcomes support separate measurements of concentration and access. Appendix F.5 reports all market sizes and intervals.

5.3 Boundary behavior, mixing, and fitted scores

We hold $\beta\sigma = 0.50$ exactly by setting $\kappa = 1 - \sqrt{1 - 0.50/\beta}$, with $\rho = 0.995$ and $\beta \in \{0.6, 0.9, 1.0, 1.2, 1.4\}$. Each setting has 64 replicates of 20,000 events. At $N = 128$, the access shortfall rises from 56.7% at $\beta = 0.6$ to 71.9% at 0.9 and 99.2% at 1.4; all runs in the last setting

Table 1: Synthetic outcomes at $N = 128$. Concentration is HHI divided by $1/N$; access shortfall is $100(1 - A/A_{\text{unif}})$, where $A_{\text{unif}} = 1 - (1 - 1/N)^{2000}$. Lower is better for both. Values are means $\pm$ 95% Monte Carlo half-widths across 64 paths.

Policy	HHI / equal shares	Access shortfall (%)
Baseline	128.00 ± 0.00	99.22 ± 0.00
Choice expansion	128.00 ± 0.00	99.22 ± 0.00
Cap only	2.97 ± 0.05	60.90 ± 0.59
Expansion + cap	6.73 ± 0.13	82.13 ± 0.31
Expansion + cap + access	2.16 ± 0.03	22.55 ± 0.79

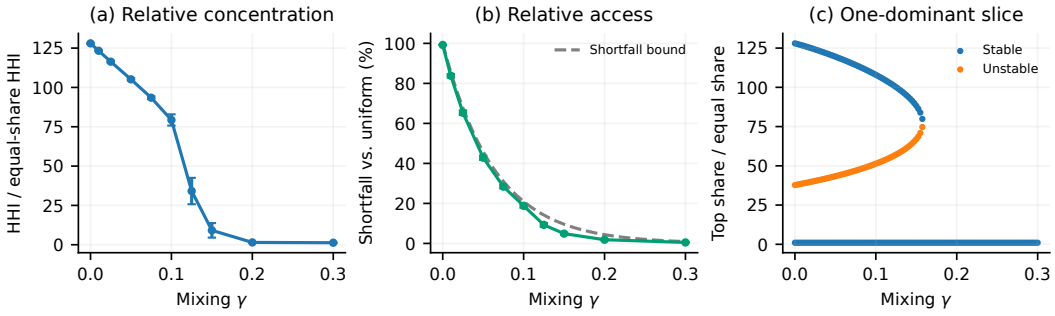

Figure 2: Synthetic mixed-policy evaluation at $N = 128$, $\beta = 1.4$, $\kappa = 0.2$, $\rho = 0.995$. (a–b) Final-window means and 95% Monte Carlo intervals over 64 paths. HHI is relative to equal shares; access shortfall uses the uniform-booking reference in Table 1. The dashed curve bounds that shortfall. (c) Equilibrium top shares are relative to $1/N$, with stability evaluated in the full state space. Small positive mixing preserves a stable high-share branch; its fold is specific to this construction.

have complete final-window concentration. At $\beta = 1.2$, the monopoly count is zero at $N = 64$ but 96.9% of runs at $N = 128$. Equal interior margins therefore leave substantial differences in finite-event outcomes (Appendix Figure 5 and Section F.6).

At $\beta = 1.4$, $\kappa = 0.2$, we compare memory scaling $1 - \rho = 0.08/N$ and batch scaling $b = N$ at $1 - \rho = 0.005$, with $N \in \{16, 64, 128\}$. All runs span 100 reputation-time units, with 50 discarded. Memory scaling gives observed/predicted second-moment ratios $1.512, 1.532, 1.531$; batches give $1.002, 1.015, 1.008$. Thus size control alone need not make fluctuations small enough for an accurate linear approximation (Appendix F.7).

Mixed-policy check. At $N = 128$, $\beta = 1.4$, $\kappa = 0.2$, and $\rho = 0.995$, we run 64 paths for each of ten mixing rates. The increase from $\gamma = 0.05$ to 0.2 reduces concentration by 98.6% and increases access by 71.8% relative to the lower-mixing setting. A stable one-dominant mean-map branch remains at small mixing rates and terminates near 79.0% of the higher mixing rate in this symmetric family (Figure 2). This branch calculation does not enumerate all possible equilibria. Appendix F.9 gives the protocols.

Asymmetric score fit. Generated bookings with observed need types fit a score model with unequal fixed-point shares. At $N = 8$, $\beta = 1.4$, $\gamma = 0.9$, the shares range from 0.935 to 1.066 times equal allocation. A contraction bound controls solver error; independent calibration estimates shares and the covariance matrix. Mean exposure bounds are $1.000, 1.000, 0.402$ times the analytic mixing bound at training sizes $10^3, 10^4, 10^5$ (100 fits each). Only the largest sample tightens it. Appendices D.7 and F.10 state the parametric and demand assumptions.

5.4 Observational marketplace diagnostics

The platform analyses combine promotion outcomes, creator-week attributes, and content exposures. Records describe completed matches without candidate slates. The diagnostics therefore complement the theory through observed allocation and advancement, without estimating its over-

lap matrix. Associations are percentages of the corresponding estimation-sample outcome mean per predictor SD; intervals use the same fixed reference. Appendix G.1 gives the measurement scope.

D1. Is audience balance associated with booking concentration? The audience-balance proxy is the inverse concentration of follower mass among the top 50 creators. Its relative concentration association is -17.23% (Holm–wild $p = 0.343$). **D2.** Is the same proxy associated with average booked fees? The relative association is -6.28% (Holm–wild $p = 0.141$). Both diagnostics are inconclusive. The nonidentification claim rests on Proposition 3; these estimates document uncertainty about the roster proxy, rather than empirical superiority of local choice. This distinction motivates evaluation of realized allocation in H1. Appendix G.2 reports sensitivity benchmarks.

H1. Is more even booking allocation associated with subsequent creator advancement? Among already-booked peers, a one-SD increase in booking evenness is associated with a 7.13% higher audience-tier advancement probability relative to its sample mean (95% interval $[3.92, 10.35]\%$; Holm–wild $p = 0.003$). The one-week association is 12.87% ($[6.74, 19.00]\%$; unadjusted wild $p < 0.001$). Category-clustered intervals and agreement across horizons provide observational support for the opportunity-allocation motivation of the thesis: more even bookings are associated with creator advancement. Appendix G.2 gives the outcome and horizon comparisons.

5.5 Online evaluation over 60 days

The authors report a 60-day A/B comparison. Relative changes are $100(Y_T/Y_C - 1)$ for treatment and control aggregates Y_T and Y_C.

H2 (marketplace value). Does the online treatment increase gross promotion transaction value (GMV) relative to control? The reported GMV increase is 3.4%, directionally consistent with the aim of improving marketplace value alongside opportunity allocation. The comparison concerns the test period, not a long-run revenue optimum.

H3 (user response). Does the same comparison improve user response? The reported changes are $+12.0628\%$ in click-through rate, $+2.1687\%$ in marketplace-wide engagement, and -4.8515% in negative comments. Together with H2, these directions are consistent with an alignment between marketplace value and user response.

These author-reported point estimates provide directional context. Assignment records, uncertainty estimates, and implementation of Algorithm 1 remain undocumented, so statistical significance and the policy mechanism are not established (Appendix G.4, Table 11).

6 Conclusion

Creator–brand matching requires separate attention to creator coverage, brand-facing choice, and sustained access. We show that choice overlap determines the local response to reputation feedback, while reinforcement strength governs attraction toward monopoly allocations. Markets with identical coverage and effective choice can therefore have different stability. Finite-sample bounds allow candidate policies to be assessed from choice probabilities, and explicit exploration provides a finite-window access bound under the stated eligibility conditions. Numerical studies show that local stability and small exploration shares can coexist with substantial concentration. Observational platform analyses associate booking evenness with creator advancement; the reported 60-day comparison shows favorable directions for GMV and user response, while audience-balance diagnostics remain inconclusive. Together, these findings distinguish the allocation of commercial opportunities from the variety of the creator roster.

The practical implication is to evaluate retrieval, reputation weights, and exploration jointly against effective choice, concentration, and access. The formal guarantees depend on the stated model assumptions; causal validation of the policy mechanism requires treatment records and uncertainty estimates beyond the available online summary.

THE USE OF LARGE LANGUAGE MODEL

We used large language models (LLMs) only to aid or polish the writing; no content generation, ideation, or data analysis was performed by LLMs.

ETHICS STATEMENT

The data owner has authorized research use of the platform exports and publication of aggregate results. Raw user, creator, and brand identifiers are not included in the manuscript or submission package. The proposed policy uses eligibility, relevance, and competition signals rather than sensitive attributes. Absolute platform levels, sample counts, proprietary labels, and source identifiers are omitted from the release.

REPRODUCIBILITY STATEMENT

Appendices A–E state the formal assumptions and arguments. Numerical constructions and protocols appear in Appendices F.1, F.3, and F.5. The source package includes model scripts, fixed seeds, and relative outputs for the reported numerical scenarios. Data processing and observational specifications appear in Appendices G.1, G.2, and G.3. Raw platform logs are not distributed; confidential reference levels and record counts are withheld. Relative outputs reproduce the publication figures, but do not permit an independent refit of the platform models. Online estimates alone cannot reproduce randomized inference. Boundary, mixing, scaling, and fitted-model protocols appear in Appendices F.6–F.10.

REFERENCES

Simon P. Anderson, André de Palma, and Jacques-François Thisse. *Discrete Choice Theory of Product Differentiation*. MIT Press, 1992.

Mark Armstrong. Competition in two-sided markets. *The RAND Journal of Economics*, 37(3): 668–691, 2006. doi: 10.1111/j.1756-2171.2006.tb00037.x.

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

Omer Ben-Porat and Moshe Tennenholtz. A game-theoretic approach to recommendation systems with strategic content providers. In *Advances in Neural Information Processing Systems*, volume 31, 2018. URL https://proceedings.neurips.cc/paper/2018/hash/a9a1d5317a33ae8cef33961c34144f84-Abstract.html.

Michel Benaïm. Dynamics of stochastic approximation algorithms. *Séminaire de Probabilités*, 33: 1–68, 1999. URL https://numdam.org/item/SPS_1999__33__1_0/.

Steven Berry, James Levinsohn, and Ariel Pakes. Automobile prices in market equilibrium. *Econometrica*, 63(4):841–890, 1995. doi: 10.2307/2171802.

Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, 100(5):992–1026, 1992. doi: 10.1086/261849.

A. Colin Cameron and Douglas L. Miller. A practitioner's guide to cluster-robust inference. *Journal of Human Resources*, 50(2):317–372, 2015. doi: 10.3368/jhr.50.2.317.

Allison J. B. Chaney, Brandon M. Stewart, and Barbara E. Engelhardt. How algorithmic confounding in recommendation systems increases homogeneity and decreases utility. In *Proceedings of the 12th ACM Conference on Recommender Systems*, pp. 224–232, 2018.

Daniel Fleder and Kartik Hosanagar. Blockbuster culture's next rise or fall: The impact of recommender systems on sales diversity. *Management Science*, 55(5):697–712, 2009. doi: 10.1287/mnsc.1080.0974.

Sture Holm. A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics*, 6(2):65–70, 1979.

Roger A. Horn and Charles R. Johnson. *Matrix Analysis*. Cambridge University Press, 2 edition, 2012. doi: 10.1017/CBO9781139020411.

Jiri Hron, Karl Krauth, Michael I. Jordan, Niki Kilbertus, and Sarah Dean. Modeling content creator incentives on algorithm-curated platforms. In *International Conference on Learning Representations*, 2023. URL https://arxiv.org/abs/2206.13102.

Meena Jagadeesan, Nikhil Garg, and Jacob Steinhardt. Supply-side equilibria in recommender systems. In *Advances in Neural Information Processing Systems*, volume 36, 2023.

Ray Jiang, Silvia Chiappa, Tor Lattimore, András György, and Pushmeet Kohli. Degenerate feedback loops in recommender systems. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 2019. doi: 10.1145/3306618.3314288. URL https://arxiv.org/abs/1902.10730.

Ramesh Johari, Hannah Li, Inessa Liskovich, and Gabriel Y. Weintraub. Experimental design in two-sided platforms: An analysis of bias. *Management Science*, 68(10):7069–7089, 2022. doi: 10.1287/mnsc.2021.4247.

Harold J. Kushner and G. George Yin. *Stochastic Approximation and Recursive Algorithms and Applications*. Springer, 2 edition, 2003. doi: 10.1007/b97441.

Lydia T. Liu, Horia Mania, and Michael Jordan. Competing bandits in matching markets. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pp. 1618–1628, 2020. URL https://proceedings.mlr.press/v108/liu20c.html.

James G. MacKinnon and Matthew D. Webb. Wild bootstrap inference for wildly different cluster sizes. *Journal of Applied Econometrics*, 32(2):233–254, 2017. doi: 10.1002/jae.2508.

Daniel McFadden. Conditional logit analysis of qualitative choice behavior. In Paul Zarembka (ed.), *Frontiers in Econometrics*, pp. 105–142. Academic Press, 1974. URL https://eml.berkeley.edu/reprints/mcfadden/zarembka.pdf.

Martin Mladenov, Elliot Creager, Omer Ben-Porat, Kevin Swersky, Richard Zemel, and Craig Boutilier. Optimizing long-term social welfare in recommender systems: A constrained matching approach. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 6987–6998, 2020. URL https://proceedings.mlr.press/v119/mladenov20a.html.

Marco Morik, Ashudeep Singh, Jessica Hong, and Thorsten Joachims. Controlling fairness and bias in dynamic learning-to-rank. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2020. doi: 10.1145/3397271.3401100. URL https://arxiv.org/abs/2005.14713.

Robin Pemantle. A survey of random processes with reinforcement. *Probability Surveys*, 4:1–79, 2007.

Juan Perdomo, Tijana Zrnic, Celestine Mendler-Dünner, and Moritz Hardt. Performative prediction. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 7599–7609, 2020. URL https://proceedings.mlr.press/v119/perdomo20a.html.

Jean-Charles Rochet and Jean Tirole. Platform competition in two-sided markets. *Journal of the European Economic Association*, 1(4):990–1029, 2003. doi: 10.1162/154247603322493212.

Tobias Schnabel, Adith Swaminathan, Ashudeep Singh, Navin Chandak, and Thorsten Joachims. Recommendations as treatments: Debiasing learning and evaluation. In *Proceedings of the 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pp. 1670–1679, 2016. URL https://proceedings.mlr.press/v48/schnabel16.html.

Ashudeep Singh and Thorsten Joachims. Fairness of exposure in rankings. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 2219–2228, 2018. doi: 10.1145/3219819.3220088.

Yi Su, Magd Bayoumi, and Thorsten Joachims. Optimizing rankings for recommendation in matching markets. In *Proceedings of the ACM Web Conference 2022*, 2022. doi: 10.1145/3485447.3511961.

Kenneth E. Train. *Discrete Choice Methods with Simulation*. Cambridge University Press, 2 edition, 2009. URL https://eml.berkeley.edu/books/choice2.html.

Joel A. Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of Computational Mathematics*, 12(4):389–434, 2012. doi: 10.1007/s10208-011-9099-z. URL https://arxiv.org/abs/1004.4389.

N. G. van Kampen. *Stochastic Processes in Physics and Chemistry*. Elsevier, 3 edition, 2007. ISBN 9780444529657.

Roman Vershynin. *High-Dimensional Probability: An Introduction with Applications in Data Science*, volume 47 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, 2018. doi: 10.1017/9781108231596.

Lequn Wang, Yiwei Bai, Wen Sun, and Thorsten Joachims. Fairness of exposure in stochastic bandits. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 10686–10696, 2021. URL https://proceedings.mlr.press/v139/wang21b.html.

Xiukun Wei, Min Shi, and Xueru Zhang. Market games for generative models: Equilibria, welfare, and strategic entry. In *International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=T7c5LXuebY.

Fan Yao, Chuanhao Li, Denis Nekipelov, Hongning Wang, and Haifeng Xu. How bad is top-k recommendation under competing content creators? In *Proceedings of the 40th International Conference on Machine Learning*, pp. 39674–39701, 2023.

Fan Yao, Yiming Liao, Jingzhou Liu, Shaoliang Nie, Qifan Wang, Haifeng Xu, and Hongning Wang. Unveiling user satisfaction and creator productivity trade-offs in recommendation platforms. In *Advances in Neural Information Processing Systems*, volume 37, 2024. doi: 10.52202/079017-2759.

Cai-Nicolas Ziegler, Sean M. McNee, Joseph A. Konstan, and Georg Lausen. Improving recommendation lists through topic diversification. In *Proceedings of the 14th International Conference on World Wide Web*, pp. 22–32, 2005. doi: 10.1145/1060745.1060754.

A Choice Diagnostics and Nonidentification

These proofs support the choice diagnostics in Section 2 and the nonidentification result in Section 3.

A.1 Aggregate and Local Diversity (Prop. 3)

Lemma 15 (Structure of the overlap matrix)**.** *At the balanced interior baseline, $P = \mathbb{E}_\theta[p(\theta)p(\theta)^\top]$ is symmetric positive semidefinite, satisfies $P\mathbf{1} = \mathbf{1}/N$, and has eigenvalues in $[0, 1/N]$. Moreover,*

$$\operatorname{tr}(P) = \mathbb{E}_\theta\|p(\theta)\|_2^2, \qquad \frac{1}{N} \le \operatorname{tr}(P) \le 1.$$

Proof. Symmetry is immediate. For every $v \in \mathbb{R}^N$,

$$v^\top P v = \mathbb{E}_\theta\left[v^\top p(\theta)p(\theta)^\top v\right] = \mathbb{E}_\theta[(v^\top p(\theta))^2] \ge 0,$$

so P is positive semidefinite. Since $\mathbf{1}^\top p(\theta) = 1$,

$$P\mathbf{1} = \mathbb{E}_\theta[p(\theta)p(\theta)^\top \mathbf{1}] = \mathbb{E}_\theta[p(\theta)] = \frac{1}{N}\mathbf{1}$$

under Assumption 1. Hence $1/N$ is an eigenvalue. For a nonnegative matrix, the spectral radius is bounded by the maximum row sum; every row sum here equals $1/N$. Positive semidefiniteness then places every eigenvalue in $[0, 1/N]$. Finally,

$$\operatorname{tr}(P) = \mathbb{E}_\theta[\operatorname{tr}\{p(\theta)p(\theta)^\top\}] = \mathbb{E}_\theta\|p(\theta)\|_2^2.$$

Each probability vector obeys $1/N \le \|p(\theta)\|_2^2 \le 1$, by Cauchy–Schwarz and the simplex constraint. Expectation gives the trace bounds. □

Proof of Prop. 3. Part (a). The definition and the spectral theorem give, at any state,

$$\mathcal{D}_{\mathrm{loc}}^{-1} = \operatorname{tr}(P) = \sum_{k=1}^{N} \mu_k.$$

At a balanced baseline, Lemma 15 shows $\mu_{\min}(P) \in [0, 1/N]$, hence $\sigma = 1 - N\mu_{\min}(P) \in [0, 1]$. These quantities depend on F, match utility, and the scoring rule through $p(\theta)$. By contrast, $\mathcal{D} = \log\det(\Sigma + \delta_D I)$ depends only on the empirical covariance of the fixed creator positions. Thus there is no algebraic map from Σ to P without additional restrictions on demand and matching.

Part (b). At the one-hot endpoint of the closure—obtained, for example, as $\varepsilon \downarrow 0$ with unique best matches—every need selects exactly one creator and creators are selected with equal probability. Then $p(\theta)$ is a uniformly distributed coordinate vector. Therefore $P = I/N$, $\operatorname{tr}(P) = 1$, and $\mu_{\min}(P) = 1/N$, yielding $(\mathcal{D}_{\mathrm{loc}}, \sigma) = (1, 0)$. If every need sees the same uniform choice vector, $p(\theta) = \mathbf{1}/N$ almost surely. Therefore $P = \mathbf{1}\mathbf{1}^\top/N^2$, whose eigenvalues are $1/N$ in the aggregate direction and zero in all $N - 1$ contrast directions. Hence $(\mathcal{D}_{\mathrm{loc}}, \sigma) = (N, 1)$.

Part (c). Fix $x_1 = -c$ and $x_2 = c$, so the creator mean is zero and the position covariance is c^2 regardless of brand demand. Take a distance-based match-utility function strictly decreasing in distance. In Market A, let demand place equal mass on $-c$ and c. As $\varepsilon \downarrow 0$, the creator located at the realized need is chosen with probability approaching one. Thus $p(\theta)$ approaches a coordinate vector, $P \to I/2$, and $\mathcal{D}_{\mathrm{loc}} \to 1$. In Market B, put all demand at zero. The two match utilities are identical, and equal reputation at the baseline gives $p(\theta) = (1/2, 1/2)$ exactly. Hence $P = \mathbf{1}\mathbf{1}^\top/4$ and $\mathcal{D}_{\mathrm{loc}} = 2$. The creator positions, covariance, and aggregate diversity are identical across the two markets, while local diversity differs.

Part (d). The matrix $P_t = (1-t)I/N + t\mathbf{1}\mathbf{1}^\top/N^2$ is a feasible choice Gram matrix in the closure of the finite-score softmax model: with probability $1 - t$, let p be a uniformly drawn coordinate vector, and with probability t, let $p = \mathbf{1}/N$. For $0 \le t < 1$, its aggregate eigenvalue is $1/N$ and each of the $N - 1$ contrast eigenvalues equals $(1 - t)/N$. Consequently,

$$\sigma = t, \qquad \operatorname{tr}(P_t) = \frac{1}{N} + (N-1)\frac{1-t}{N} = 1 - t\left(1 - \frac{1}{N}\right).$$

At $N = 4$ and $t = 0.9$, this yields $(\sigma, \mathcal{D}_{\mathrm{loc}}) = (0.9, 1/0.325) \approx (0.9, 3.08)$. Now divide four creators into two equally likely demand blocks and choose uniformly between the two creators inside the realized block. Then P contains two rank-one 2×2 diagonal blocks with entries $1/8$. Its eigenvalues are $1/4, 1/4, 0, 0$, so $\sigma = 1$, while $\mathrm{tr}(P) = 1/2$ and $\mathcal{D}_{\mathrm{loc}} = 2$. The move from the first construction to the second raises σ but lowers $\mathcal{D}_{\mathrm{loc}}$, ruling out a globally increasing relation. Part (b) rules out a globally decreasing relation. To obtain finite-score interior constructions, replace each boundary probability vector by $p^{(\xi)} = (1 - \xi)p + \xi \mathbf{1}/N$, $\xi > 0$. This preserves its balanced expectation and gives strictly positive probabilities, with $P^{(\xi)} = (1 - \xi)^2 P + (2\xi - \xi^2)\mathbf{1}\mathbf{1}^\top/N^2$. All eigenvalues and traces vary continuously with ξ, so sufficiently small ξ preserves the strict reversal between the two constructions. Such vectors can be realized by distance scores: take creator positions x_i to be the coordinate vectors in $\mathbb{R}^N$ and each demand coordinate $\theta_i = (\varepsilon/2) \log p_i^{(\xi)}$. Then $-\|x_i - \theta\|^2/\varepsilon$ differs from $\log p_i^{(\xi)}$ by a creator-independent constant, giving exactly $p^{(\xi)}$ at equal reputation. □

B Reputation Stability and Choice–Reinforcement Bounds

This appendix proves the interior and boundary stability results and the choice–reinforcement bounds in Section 3.2.

B.1 Unequal-Share Stability

Proof of Theorem 4 at a general positive fixed point. Let $a^* > 0$ satisfy $a^* = \rho a^* + e(a^*)$, and set $s = e(a^*) = (1-\rho)a^*$ and $S = \mathrm{diag}(s)$. Then $s_i > 0$ and $\mathbf{1}^\top s = 1$. In a neighborhood bounded away from zero, the softmax derivatives are uniformly bounded, allowing differentiation under the expectation. With $P^* = \mathbb{E}[p(\theta; a^*)p(\theta; a^*)^\top]$,

$$\left.\frac{\partial e_i}{\partial a_j}\right|_{a^*} = \frac{\beta(1-\rho)}{s_j}(s_i\delta_{ij} - P^*_{ij}).$$

Consequently,

$$J = \rho I + \beta(1-\rho)(S - P^*)S^{-1} = \rho I + \beta(1-\rho)(I - P^*S^{-1}),$$

and the similarity transform is

$$S^{-1/2}JS^{1/2} = \rho I + \beta(1-\rho)(I - B), \qquad B = S^{-1/2}P^*S^{-1/2}.$$

The Gram matrix P^* is positive semidefinite and $P^*\mathbf{1} = s$. Thus B is positive semidefinite and $B\sqrt{s} = \sqrt{s}$. Also $S^{-1}P^*$ is nonnegative with row sums one and is similar to B. Its spectral radius is one, so the real eigenvalues of B lie in $[0, 1]$. It follows that all eigenvalues of J are nonnegative and

$$\mathrm{spr}(J) = \rho + \beta(1-\rho)[1 - \lambda_{\min}(B)].$$

Since $1 - \rho > 0$, this is below one exactly when $\beta\sigma_* < 1$. Continuous differentiability implies local asymptotic stability for the nonlinear mean map under the strict inequality and instability when the spectral radius exceeds one. Equality does not determine nonlinear stability. The matrix need not be symmetric in the original coordinates: the symmetric representation uses the weighted norm $\|h\|_{S^{-1}} = (h^\top S^{-1}h)^{1/2}$. The substitution $s_i = 1/N$ recovers $B = NP$ and the balanced calculation in Appendix B.2. □

B.2 Jacobian Spectrum and the $\beta\sigma = 1$ Threshold (Thm. 4)

Proof. Write $p(\theta, a)$ to make the state dependence explicit. Since reputation enters the normalized score as $\beta \log a_i$, the standard softmax derivative and the chain rule give

$$\frac{\partial p_i(\theta, a)}{\partial a_j} = \frac{\beta}{a_j}p_i(\theta, a)\{\delta_{ij} - p_j(\theta, a)\}.$$

The probabilities and their derivatives are bounded on a neighborhood of the strictly positive baseline, permitting differentiation under the integral that defines $e(a)$. Under Assumption 1, $e^* = \mathbf{1}/N$ and the fixed-point equation $a^* = \rho a^* + e^*$ yields $a^*_i = 1/\{N(1-\rho)\}$. Therefore $\beta/a^*_j = \beta N(1-\rho)$ for every j, and

$$\begin{aligned} G_{ij} &:= \left.\frac{\partial e_i(a)}{\partial a_j}\right|_{a=a^*} \\ &= \beta N(1-\rho)\mathbb{E}_\theta[p_i(\theta)\{\delta_{ij} - p_j(\theta)\}] \\ &= \beta N(1-\rho)\left(\frac{1}{N}\delta_{ij} - P_{ij}\right). \end{aligned}$$

The Jacobian of the conditional mean map is consequently

$$J := D\mathcal{T}(a^*) = \rho I + \beta(1-\rho)(I - NP).$$

Lemma 15 makes P symmetric and supplies an orthonormal eigenbasis. If $Pv_k = \mu_k v_k$, the same vector is an eigenvector of J with

$$\lambda_k = \rho + \beta(1-\rho)(1 - N\mu_k).$$

The aggregate direction $v_1 = \mathbf{1}/\sqrt{N}$ has $\mu_1 = 1/N$ and eigenvalue ρ. It changes the common reputation level but not relative creator shares. Every orthogonal direction is a competitive contrast. Since $0 \le \mu_k \le 1/N$, all $\lambda_k \ge \rho > 0$, and the largest is attained at $\mu_{\min}(P)$:

$$\lambda_{\max}(J) = \rho + \beta(1-\rho)\{1 - N\mu_{\min}(P)\} = \rho + \beta(1-\rho)\sigma.$$

All eigenvalues are nonnegative; therefore, asymptotic stability of the linearized discrete-time map is equivalent to $\lambda_{\max}(J) < 1$. With $\rho \in (0,1)$,

$$\lambda_{\max}(J) < 1 \iff \beta(1-\rho)\sigma < 1-\rho \iff \beta\sigma < 1.$$

Continuous differentiability of $\mathcal{T}$ on a neighborhood of the strictly positive fixed point ensures that a Jacobian spectral radius below one implies local asymptotic stability of that fixed point for the nonlinear discrete-time map. If $\beta\sigma > 1$, an eigenvector associated with $\mu_{\min}(P)$ expands to first order; the corresponding linearization instability criterion establishes local instability. At equality the linearization is non-hyperbolic and higher-order terms determine stability. □

B.3 Existence and Uniqueness Below Unit Reinforcement

Proof of Proposition 5. Homogeneity gives $e(ca) = e(a)$ for $c > 0$. Every fixed point has total stock $(1-\rho)^{-1}$, so its normalized shares $s = (1-\rho)a$ satisfy $s = e(s)$. For $0 < \beta < 1$, consider on the probability simplex

$$\mathcal{L}(s) = \mathbb{E}_\theta \log\left(\sum_i w_i(\theta) s_i^\beta\right).$$

The positive bounded weights make this continuous, including on the boundary. Each weighted sum is strictly concave in s; composing with the increasing concave logarithm preserves strict concavity. Expectation therefore gives a strictly concave function with a unique maximizer. If $s_i = 0$, replace s by $(1-u)s + ue_i$. For each need, the sum inside the logarithm becomes $(1-u)^\beta A + w_i u^\beta$, where $\underline{w} \le A \le \overline{w} N^{1-\beta}$. Uniformly over needs, the positive term of order u^β dominates the loss of order u. For sufficiently small $u > 0$ the objective strictly increases. Thus the maximizer is interior. There, $\partial_i \mathcal{L} = \beta e_i(s)/s_i$. The simplex first-order condition makes these derivatives equal. Their s-weighted sum identifies the multiplier as β, so $e_i(s) = s_i$. Conversely, every positive fixed point satisfies these conditions and is the unique maximizer. For $\beta = 0$, e is independent of reputation and strictly positive, so $a^* = e/(1-\rho)$ is unique directly. □

For $\beta > 0$ the continuous extension to the simplex fixes each vertex, because zero shares receive zero weight. At $\beta = 0$ this is false: the positive reputation-independent choice vector is selected. With $\beta = 1$ and identical fully overlapping creators, every simplex point is fixed, so uniqueness cannot be extended to that boundary without further assumptions.

B.4 Attraction and Invasion at the Boundary

Proof of Proposition 6. Let $R_w = \overline{w}/\underline{w}$ and $T(\nu) = \rho\nu + (1-\rho)e(\nu)$ on the simplex. Every vertex is fixed for $\beta > 0$. Near vertex i, set $b_\partial = \sum_{j \ne i} \nu_j$. For $\beta > 1$, positivity of the denominator and $\sum_{j \ne i} \nu_j^\beta \le b_\partial^\beta$ imply

$$\sum_{j \ne i} e_j(\nu) \le R_w \left(\frac{b_\partial}{1-b_\partial}\right)^\beta, \qquad b'_\partial \le \rho b_\partial + (1-\rho) R_w \left(\frac{b_\partial}{1-b_\partial}\right)^\beta.$$

The second term is $o(b_\partial)$. A sufficiently small neighborhood therefore has $b'_\partial \le c b_\partial$ for some $\rho < c < 1$, proving local attraction. Each missing coordinate of e is $O(b_\partial^\beta)$, so the derivative on the simplex tangent space is ρI.

For $0 < \beta < 1$, concavity gives $\sum_i \nu_i^\beta \le N^{1-\beta}$. For any positive coordinate ν_j introduced at a face,

$$e_j(\nu) \ge \frac{\nu_j^\beta}{R_w N^{1-\beta}}, \qquad \frac{T_j(\nu)}{\nu_j} \ge \rho + \frac{1-\rho}{R_w N^{1-\beta}} \nu_j^{\beta-1} \longrightarrow \infty.$$

Such a coordinate grows until it exits a sufficiently small neighborhood of zero. Exactly zero coordinates remain zero under the extended map. For $\beta = 1$, a perturbation $\nu_j = u$, $\nu_i = 1 - u$ has $e_j(\nu) = u\mathbb{E}[w_j/w_i] + O(u^2)$ by boundedness of weight ratios. The invasion multiplier is consequently $\rho + (1 - \rho)\mathbb{E}[w_j/w_i]$. Multipliers equal to one require higher-order analysis. For $\beta = 0$, the choice vector is strictly positive and independent of reputation, so boundary fixed points disappear. □

B.5 Effective Choice and a Lower Bound on Feedback Exposure

Proof of Corollary 7. At the balanced baseline, $B = NP$ has leading eigenvalue one and all eigenvalues in $[0, 1]$. The smallest of its $N - 1$ nonleading eigenvalues is at most their average, so

$$\lambda_{\min}(B) \leq \frac{\mathrm{tr}(B) - 1}{N - 1} = \frac{N/\mathcal{D}_{\mathrm{loc}} - 1}{N - 1}.$$

The definition $\sigma = 1 - \lambda_{\min}(B)$ gives the bound. Equality holds when all nonleading eigenvalues coincide. For $\beta > 1$, strict linearized stability requires $\mathcal{D}_{\mathrm{loc}} < \beta N/(\beta N - N + 1)$; the converse need not hold.

□

B.6 Choice–Reinforcement Feasibility

Proof of Corollary 8. At a balanced fixed point, Corollary 7 gives $\sigma \geq L_N(\mathcal{D}_{\mathrm{loc}})$ with $L_N(x) = N(1 - 1/x)/(N - 1)$, which is increasing for $x \geq 1$. If $\mathcal{D}_{\mathrm{loc}} \geq m > 1$, then $\beta L_N(m) \leq \beta\sigma$. Therefore the margin condition $\beta\sigma \leq 1 - \eta$ requires $\beta \leq (1 - \eta)/L_N(m)$. Strict local stability instead requires $\beta < 1/L_N(m)$. Equality in the spectral bound is attainable for isotropic nonleading eigenvalues, including $C = (1 - \kappa)I + \kappa\mathbf{1}\mathbf{1}^\top/N$ under uniform demand, for which $B = C^\top C$. Thus the bound is a sharp necessary feasibility restriction within the balanced model. □

B.7 Local Relaxation Rate

For a general fixed point the diagonalization in Appendix B.1 gives the same calculation below with σ_* replacing σ, in the S^{-1}-weighted coordinates. The balanced specialization makes the relaxation factor explicit.

Proof. The proof of Theorem 4 gives the largest contrast eigenvalue

$$\lambda_{\max} = \rho + \beta(1 - \rho)\sigma.$$

Below the threshold it lies in $(0, 1)$, so first-order perturbations along its eigenvector decay geometrically by factor $\lambda_{\max}$ per mean-map iteration. More explicitly, if $v_{\max}$ is a unit eigenvector associated with $\mu_{\min}(P)$ and the initial displacement is $h_0 = cv_{\max}$, differentiability of $\mathcal{T}$ gives

$$h_t = c\lambda_{\max}^t v_{\max} + o(\|h_0\|)$$

for every fixed linearization horizon t. Direct subtraction gives

$$1 - \lambda_{\max} = 1 - \rho - \beta(1 - \rho)\sigma = (1 - \rho)(1 - \beta\sigma).$$

For fixed $\rho < 1$, the reciprocal of this gap diverges as $\beta\sigma \uparrow 1$. □

C Stochastic Fluctuations, Access, and Lock-In

This appendix develops the stochastic results in Section 3.3 and the two-creator lock-in special case used in Section 3.2.

C.1 Local Noise Amplification

Proof of Proposition 10. Initialize total reputation at $(1-\rho)^{-1}$ and set $\zeta = 1-\rho$, $\nu = \zeta a$. Then $\mathbf{1}^\top \nu = 1$ at all times and $\nu^{t+1} = \rho\nu^t + \zeta Y^t$. At the fixed point $\nu = s$, the marginal one-hot booking has mean s and covariance $S - ss^\top$, even when the conditional probabilities vary across brand needs. In the coordinates $u = S^{-1/2}(\nu - s)$, the linearized drift is $A = \rho I + \beta\zeta(I - B)$ and the frozen innovation covariance is

$$Q = \zeta^2 S^{-1/2}(S - ss^\top)S^{-1/2} = \zeta^2(I - \sqrt{s}\sqrt{s}^\top).$$

For independent batches of b bookings at a fixed state, this covariance is divided by b. The contrast subspace is orthogonal to $\sqrt{s}$; both A and Q preserve it. The diagonalization of symmetric B gives, for a unit contrast eigenvector v_k, a scalar stable autoregression with factor $\lambda_k = \rho + \beta\zeta(1 - \lambda_k(B))$ and innovation variance ζ^2/b. Its stationary variance is the geometric sum

$$\frac{\zeta^2}{b}\sum_{j=0}^{\infty}\lambda_k^{2j} = \frac{\zeta^2}{b(1-\lambda_k^2)}.$$

At fixed $\beta\sigma_* < 1$, the leading expression as $\zeta \downarrow 0$ is $\zeta/[2b(1-\beta\sigma_*)]$. At fixed ζ, this variance diverges as the interior stability boundary is approached. In the balanced isotropic case all contrast eigenvalues coincide, so the surrogate covariance is $\zeta^2(I - \mathbf{1}\mathbf{1}^\top/N)/[b(1-\lambda^2)]$. Since $N\nu_i - 1 = \sqrt{N}u_i$, this gives $\mathrm{Var}(N\nu_i) = (N-1)\zeta^2/[b(1-\lambda^2)]$. □

C.2 Stochastic Permanent Exclusion

Proof of Proposition 11. For identical fully overlapping creators, choice probabilities are $p_i^t = (a_i^t)^\beta/[(a_1^t)^\beta + (a_2^t)^\beta]$. The balanced fixed point has $\sigma = 1$, so $0 < \beta < 1$ makes the mean map locally stable. Let E_T be the event that creator 1 receives the first T bookings. Conditional on the preceding all-winner history, its reputation path and its competitor's path are deterministic:

$$a_1^t = \rho^t a_1^0 + \frac{1-\rho^t}{1-\rho}, \qquad a_2^t = \rho^t a_2^0.$$

With $L = \min\{a_1^0, (1-\rho)^{-1}\} > 0$, we obtain

$$0 < p_2^t \le \left(\frac{a_2^0}{L}\right)^\beta \rho^{\beta t}, \qquad \sum_{t=0}^{\infty} p_2^t < \infty.$$

Each $p_2^t < 1$. For all sufficiently large t, $p_2^t \le 1/2$ and $\log(1-p_2^t) \ge -2p_2^t$, so $\sum_t \log(1-p_2^t) > -\infty$. The finite prefix factors are also positive. By the chain rule for conditional probabilities and continuity from above,

$$\Pr\left(\bigcap_{T\ge 1} E_T\right) = \lim_{T\to\infty}\prod_{t=0}^{T-1}(1-p_2^t) > 0.$$

The same argument gives a positive probability for the opposite all-winner history. The summability argument holds for every $\beta > 0$; choosing $\beta < 1$ demonstrates coexistence with mean-map stability. □

C.3 A Separate Persistent-Exploration Safeguard

Fix a pool of N creators that remains eligible at every event and a constant $\gamma \in (0, 1]$. Under the modified conditional rule $p_i^\gamma = (1-\gamma)p_i + \gamma/N$, any history gives creator i a booking probability at least γ/N. Repeated conditioning bounds the probability of no booking for i over the next L events by $(1-\gamma/N)^L$. The limit $L \to \infty$ shows that, for each deterministic T, the probability of

never booking i after T is zero. A countable union over T shows that finitely many bookings for i has probability zero. A finite union over creators gives infinitely many bookings for every creator almost surely. Proposition 14 supplies the matching stability analysis and a finite-window version of this bound.

C.4 Corner Stability in the Two-Creator Special Case

Proposition 16 (Restricted deterministic lock-in)**.** *For two fully overlapping creators with constant utility gap $\Delta q > 0$, the $\rho = 0$ mean-map extension is $T(\chi) = r_q\chi^\beta/[r_q\chi^\beta + (1-\chi)^\beta]$, $r_q = \exp(\Delta q/\varepsilon)$. For $\beta > 1$, both corners are locally superstable and the unique interior fixed point is unstable.*

Mean-map corner stability (Prop. 16). Label creator 1 as the creator with higher match utility, so $\Delta q = q_1 - q_2 > 0$ and $r_q = \exp(\Delta q/\varepsilon) > 1$. In the fully overlapping case, brand type does not affect the utility difference. This proof studies only the deterministic mean iteration. With $\rho = 0$, mean-map reputations equal the preceding mean shares. For creator 1's share $\chi \in (0,1)$, the map is

$$T(\chi) = \frac{\exp(q_1/\varepsilon)\chi^\beta}{\exp(q_1/\varepsilon)\chi^\beta + \exp(q_2/\varepsilon)(1-\chi)^\beta} = \frac{r_q\chi^\beta}{r_q\chi^\beta + (1-\chi)^\beta}.$$

First consider interior fixed points. The odds transformation is exact:

$$\frac{T(\chi)}{1-T(\chi)} = r_q\left(\frac{\chi}{1-\chi}\right)^\beta.$$

Let $u = \log\{\chi/(1-\chi)\}$. An interior fixed point must satisfy

$$u = \log r_q + \beta u, \qquad u^* = -\frac{\log r_q}{\beta - 1}.$$

Thus there is exactly one interior fixed point whenever $\beta \neq 1$. In log-odds coordinates the entire map is the affine recursion $u_{t+1} = \log r_q + \beta u_t$. Its solution is

$$u_t - u^* = \beta^t(u_0 - u^*).$$

For $\beta > 1$, any initial condition above u^* diverges to $+\infty$ and any initial condition below u^* diverges to $-\infty$. Under the logistic transformation, the former trajectories converge to $\chi = 1$ and the latter to $\chi = 0$. Both basins are nonempty open intervals, separated by the unstable interior fixed point.

It remains to verify stability at the boundary, where log-odds are not finite. For $b_\partial \downarrow 0$,

$$T(b_\partial) = \frac{r_q b_\partial^\beta}{r_q b_\partial^\beta + (1-b_\partial)^\beta} = r_q b_\partial^\beta\{1 + o(1)\},$$

and

$$1 - T(1-b_\partial) = \frac{b_\partial^\beta}{r_q(1-b_\partial)^\beta + b_\partial^\beta} = \frac{b_\partial^\beta}{r_q}\{1 + o(1)\}.$$

When $\beta > 1$, each boundary error is of smaller order than its predecessor; equivalently, the one-sided derivatives at both corners are zero. Hence both corners are locally superstable. Since $r_q > 1$, $\chi = 0$ assigns all demand to the creator with lower match utility. Its nonempty basin proves path-dependent, non-self-correcting lock-in in this special case. □

D Spectral Certification and Mixed Policies

These proofs and error bounds support the policy certificates and explicit mixing rule in Section 4.

D.1 Reinforcement Guardrail (Cor. 12)

Proof. Let σ_* refer to the positive fixed point of the final policy. The upper bound must hold uniformly over the parameter range used to choose that policy, or simultaneously for every member of a prespecified finite family. A bound for one policy cannot certify an unevaluated change to it. Since $\beta = \alpha/\varepsilon$ and $\sigma_* \le U$,

$$\beta\sigma_* = \frac{\alpha}{\varepsilon}\sigma_* \le \frac{1-\eta}{U}\sigma_* \le 1-\eta < 1.$$

Theorem 4 therefore gives local asymptotic stability under the linearized mean map and, by the discrete-time linearization theorem, local asymptotic stability of the nonlinear mean map at that fixed point. Substitution into the spectral-gap expression further yields the quantitative guarantee

$$1-\lambda_{\max} = (1-\rho)(1-\beta\sigma_*) \ge (1-\rho)\eta > 0.$$

Thus the policy does more than remain on the stable side of the boundary: it reserves a chosen minimum fraction η of the unreinforced spectral margin. Theorem 4 establishes $\sigma_* \le 1$, so setting $U = 1$ remains a valid worst-case bound when no tighter validated bound is available. □

D.2 Simultaneous Spectral Certificates

Proof of Theorem 13. For one fixed policy, set $X_k = v_k v_k^\top$ with $v_k = S^{-1/2}p(\theta_k; a^*)$. The vectors are computed at the same fixed policy and state. Then $\mathbb{E}X_k = B$, $0 \preceq X_k \preceq LI$, and $\|B\| = 1$, so $L \ge 1$. For $Z_k = X_k - B$, $\|Z_k\| \le L$ and

$$\mathbb{E}Z_k^2 = \mathbb{E}X_k^2 - B^2 \preceq LB - B^2 \preceq LI.$$

The self-adjoint matrix Bernstein inequality (Tropp, 2012) therefore gives

$$\Pr\{\|\widehat{B} - B\| > R\} \le 2N \exp\left(-\frac{nR^2}{2(L + LR/3)}\right).$$

The choice $R = R_n = \sqrt{2L\ell/n} + 2L\ell/(3n)$ with $\ell = \log(2NK/\delta_c)$ bounds this probability by δ_c/K. Weyl's inequality gives $\lambda_{\min}(B) \ge \lambda_{\min}(\widehat{B}) - R_n$. The resulting upper bound on σ_* retains coverage after clipping to $[0, 1]$. A union bound covers all K prespecified policy–fixed-point pairs, whether or not their evaluations share the same sampled needs. Thus any subsequently selected pair satisfying its own certificate retains the stated margin with probability at least $1 - \delta_c$. □

The bound $L = 1/\min_i s_i$ is always valid, since $\sum_i p_i^2/s_i \le 1/\min_i s_i$. A smaller certified uniform bound on this leverage can be substituted. Logs of realized bookings alone do not provide the probability vectors used here. If probabilities are estimated, model error must be bounded separately. Data used to fit the score model or choose the policy family must be separated from calibration, or an appropriate uniform bound must replace the finite-family argument. A certificate at an arbitrary observed state is not a fixed-point certificate.

A tighter variance bound. The preceding argument also gives $\|\mathbb{E}Z_k^2\| \le v_L := \max_{0\le x\le 1}(Lx - x^2)$, where

$$v_L = \begin{cases} L^2/4, & 1 \le L \le 2, \\ L - 1, & L \ge 2. \end{cases}$$

Thus R_n may be replaced by $\sqrt{2v_L\ell/n} + 2L\ell/(3n)$, with the same coverage. Figure 3 uses the radius in Theorem 13; Appendix F.6 compares both radii on the same draws.

Unknown fixed-point shares. An explicitly more conservative variant estimates s from the same iid needs. Let $\widehat{s} = n^{-1}\sum_k p_k$, $\widehat{S} = \operatorname{diag}(\widehat{s})$, and $\widehat{P} = n^{-1}\sum_k p_k p_k^\top$. If any estimated share is zero, return $U = 1$ without inverting $\widehat{S}$. Otherwise define $\widehat{b} = \lambda_{\min}(\widehat{S}^{-1/2}\widehat{P}\widehat{S}^{-1/2})$. For a family of size K, define

$$u = \sqrt{\frac{\log(4NK/\delta_c)}{2n}}, \quad t = \sqrt{\frac{2\log(4NK/\delta_c)}{n}} + \frac{2\log(4NK/\delta_c)}{3n}, \quad l = \min_i \widehat{s}_i - u.$$

Hoeffding's inequality and unnormalized matrix Bernstein, followed by union bounds, give $\|\widehat{s} - s\|_\infty \le u$ and $\|\widehat{P} - P\| \le t$ simultaneously with probability at least $1 - \delta_c$. If $l \le 0$, use $U = 1$. Otherwise every vector x satisfies

$$x^\top P x \ge \widehat{b}\, x^\top S x - (t + \widehat{b}u)\|x\|^2 \ge \left[\widehat{b} - \frac{t + \widehat{b}u}{l}\right] x^\top S x,$$

because $s_i \ge l$. Therefore $U = \min\{1, \max\{0, 1 - \widehat{b} + (t + \widehat{b}u)/l\}\}$ is valid. This variant can be uninformative for rare creators; the tighter known-share result must not be used with estimated shares without accounting for their uncertainty.

D.3 Model and Fixed-Point Error

Let B_{fit} and B_{true} be normalized matrices for the fitted and target models at their respective positive fixed points, with the same creator labels. Suppose an independent or selection-valid argument gives $\|B_{\mathrm{true}} - B_{\mathrm{fit}}\| \le \Delta_B$ with probability $1 - \delta_m$. Conditional on the fitted model, the calibration theorem gives its sampling bound with probability $1 - \delta_c$. Weyl's inequality and a union bound yield, with probability at least $1 - \delta_m - \delta_c$,

$$\sigma_{\mathrm{true}} \le \min\{1, \max\{0, 1 - \lambda_{\min}(\widehat{B}_{\mathrm{fit}}) + R_n + \Delta_B\}\}.$$

The model-error term must include changes in choice probabilities, fixed-point location, and share normalization. An implementation also needs a bound on the error of its numerical fixed point. A residual bounds location error only with an additional local conditioning or contraction argument. The synthetic fit below has exactly known balanced shares by symmetry, which avoids this extra estimation step. Its model-error bound depends on correct specification.

D.4 Stability and Access under Explicit Mixing

Proof of Proposition 14. Let $M_\gamma(\nu) = (1-\gamma)e(\nu) + \gamma\mathbf{1}/N$ and $T_\gamma(\nu) = \rho\nu + (1-\rho)M_\gamma(\nu)$. For $0 < \gamma < 1$, the compact simplex subset $\mathcal{K}_\gamma = \{\nu : \nu_i \ge \gamma/N,\ \mathbf{1}^\top\nu = 1\}$ maps into itself. Positive bounded weights imply continuity there, so Brouwer's theorem gives a fixed point. Every fixed point satisfies $s_i = (1-\gamma)e_i + \gamma/N \ge \gamma/N$; hence none is on the boundary. For $\gamma = 1$, the unique fixed point is uniform.

The score derivative gives $De(a^*) = \beta(E - P^*)\operatorname{diag}(a^*)^{-1}$, with $a^* = s/(1-\rho)$. Thus

$$J_\gamma = \rho I + \beta(1-\rho)(1-\gamma)(E - P^*)S^{-1}, \qquad S^{-1/2}J_\gamma S^{1/2} = \rho I + \beta(1-\rho)H_\gamma.$$

Each $\operatorname{diag}(p) - pp^\top$ is a categorical covariance matrix, so $H_\gamma \succeq 0$ and $H_\gamma\sqrt{s} = 0$. Also

$$\begin{aligned} H_\gamma &\preceq (1-\gamma)S^{-1/2}ES^{-1/2}, \\ \sigma_\gamma &\le \max_i \frac{(1-\gamma)e_i}{s_i} = \max_i \left(1 - \frac{\gamma}{Ns_i}\right) \\ &\le c_\gamma := 1 - \frac{\gamma}{N - (N-1)\gamma} < 1. \end{aligned}$$

All Jacobian eigenvalues are nonnegative; its spectral radius is $\rho + \beta(1-\rho)\sigma_\gamma$. This proves the strict linearized condition and the usual nonlinear local implications. Equality requires higher-order analysis. At $\gamma = 0$, $E = S$ and $H_0 = I - B$.

At every event, an always-eligible creator's conditional booking probability is at least γ/N. Repeated conditional expectation bounds the probability of no booking in any specified next T events

by $(1-\gamma/N)^T$. The expected fraction booked is at least $1-(1-\gamma/N)^T$. The no-booking probability after any deterministic time tends to zero as $T \to \infty$; a countable union over such times proves infinitely many bookings almost surely. These probability bounds apply to the event process; the deterministic stock lower bound above applies to its mean map. □

D.5 A Certificate for the Mixed Policy

For known positive fixed-point shares, define

$$X_k = (1-\gamma)S^{-1/2}[\mathrm{diag}(p_k) - p_k p_k^\top]S^{-1/2}, \qquad \widehat{H}_\gamma = n^{-1}\sum_k X_k.$$

Here p_k is the base choice vector, before uniform mixing, evaluated at the mixed policy's fixed point. For K prespecified policy–fixed-point pairs, suppose $0 \preceq X_k \preceq LI$. One may use $L = (1-\gamma)/\min_i s_i$. The centered matrices satisfy $\|X_k - H_\gamma\| \le L$ and $\mathbb{E}[(X_k - H_\gamma)^2] \preceq LH_\gamma \preceq LI$. The same matrix Bernstein argument as Theorem 13 gives

$$U_\gamma = \min\{c_\gamma, \lambda_{\max}(\widehat{H}_\gamma) + R_n\}, \qquad \sigma_\gamma \le U_\gamma$$

simultaneously with probability at least $1-\delta_c$, using that theorem's R_n and ℓ. At $\gamma = 1$ take $U_\gamma = 0$ directly. Policy selection requires $\beta U_\gamma \le 1-\eta$. An operator-norm bound Δ_H on model and fixed-point error is added before clipping, with its failure probability included in the union bound. Unlike the unmixed estimator, this bound uses a maximum covariance eigenvalue.

Estimated shares. Set $\widehat{s} = (1-\gamma)n^{-1}\sum_k p_k + \gamma\mathbf{1}/N$ and $\widehat{A} = (1-\gamma)n^{-1}\sum_k[\mathrm{diag}(p_k) - p_k p_k^\top]$. Let $\widehat{h} = \lambda_{\max}(\widehat{S}^{-1/2}\widehat{A}\widehat{S}^{-1/2})$, $\ell' = \log(4NK/\delta_c)$, and

$$u = (1-\gamma)\sqrt{\frac{\ell'}{2n}}, \quad t = (1-\gamma)\left(\sqrt{\frac{2\ell'}{n}} + \frac{2\ell'}{3n}\right), \quad l = \min_i \widehat{s}_i - u.$$

Hoeffding and matrix Bernstein give $\|\widehat{s} - s\|_\infty \le u$ and $\|\widehat{A} - A\| \le t$ simultaneously. For $l > 0$,

$$x^\top A x \le \widehat{h}\, x^\top S x + (t + \widehat{h}u)\|x\|^2 \le \left(\widehat{h} + \frac{t + \widehat{h}u}{l}\right) x^\top S x.$$

Thus $U_\gamma = \min\{c_\gamma, \widehat{h} + (t + \widehat{h}u)/l\}$ is valid; use c_γ if $l \le 0$. The same argument estimates the covariance normalized by $M_\gamma(\nu)$ at a numerical state ν; its difference from the target fixed-point matrix must then enter Δ_H separately.

D.6 A Conservative Contraction and Solver Bound

On $\mathcal{K}_\gamma$, the absolute column sum of $Dp(\nu)$ is $2\beta p_j(1-p_j)/\nu_j \le \beta N/(2\gamma)$. Hence M_γ has ℓ_1 Lipschitz constant

$$c = \frac{\beta(1-\gamma)N}{2\gamma}.$$

If $c < 1$, it has a unique fixed point and iterates initialized in $\mathcal{K}_\gamma$ converge to it. For any numerical solution $\widetilde{\nu} \in \mathcal{K}_\gamma$,

$$\|\widetilde{\nu} - s\|_1 \le \frac{\|M_\gamma(\widetilde{\nu}) - \widetilde{\nu}\|_1}{1-c}.$$

All fixed points belong to this set, so uniqueness also excludes other fixed points elsewhere on the simplex. In a permutation-symmetric market the unique point is uniform. The sufficient condition $\gamma > \beta N/(\beta N + 2)$ is conservative; the numerical branch fold below can occur at much smaller mixing rates.

D.7 Error Bound for the Asymmetric Fitted Example

The example in Appendix F.10 uses the correctly specified family C_κ and observed need types at equal training reputations. The diagonal-match frequency gives the same Hoeffding parameter

interval as Appendix F.8. Let its radius be b_κ, and its intersection with $(0,1)$ be I_κ. On this interval, a valid choice sensitivity is

$$L_\kappa = \frac{1}{2} \sup_{k \in I_\kappa} \frac{1}{k(1-k+k/N)}, \qquad L_s = \frac{\beta N}{2\gamma}.$$

The first expression follows from the two log-weight derivatives and $2p_j(1-p_j) \le 1/2$; the second is the preceding column-sum bound. If the numerical residual is r_{sol}, contraction gives

$$\sup_\theta \|p_{\mathrm{true}}(\theta; s) - p_{\mathrm{fit}}(\theta; \widetilde{\nu})\|_1 \le D_p := \frac{L_\kappa b_\kappa + L_s r_{\mathrm{sol}}}{1-c}.$$

Normalize the fitted matrix by $\bar{s} = M_{\gamma,\mathrm{fit}}(\widetilde{\nu})$. Both share vectors have entries at least $m_0 = \gamma/N$, and $\|s - \bar{s}\|_\infty \le (1-\gamma)D_p$. The covariance difference is at most $3(1-\gamma)D_p$ in operator norm. The categorical covariance norm is at most $1/2$, while $\|S^{-1/2} - \bar{S}^{-1/2}\| \le \|s - \bar{s}\|_\infty / (2m_0^{3/2})$. Therefore a joint score/fixed-point/normalization allowance is

$$\Delta_H = \left[\frac{3(1-\gamma)}{m_0} + \frac{(1-\gamma)^2}{2m_0^2} \right] D_p.$$

The estimated-share sampling bound plus Δ_H has coverage at least $1 - \delta_m - \delta_c$. This derivation requires the stated parametric family, known need distribution, and the contraction condition; it does not bound misspecification of a platform choice model.

E PRICING AND ECONOMIC TRADE-OFFS

This appendix gives the separate pricing model and economic comparative statics summarized in Section 4.

E.1 LOGIT–BERTRAND MARKUP (PROP. 17)

The single-type pricing model includes an outside option normalized to one: $d_i = \exp((q_i - \phi_i)/\varepsilon)/[1 + \sum_j \exp((q_j - \phi_j)/\varepsilon)]$, $\pi_i = ((1-\tau)\phi_i - c_i)d_i$, $c_i \geq 0$, $0 < \tau < 1$.

Proposition 17 (Markup relation)**.** *An interior single-type price optimum satisfies* $\phi_i - c_i/(1-\tau) = \varepsilon/(1-d_i)$.

Proof. Fix the prices of competing creators and a brand type θ. Let $V_i = (q_i(\theta) - \phi_i)/\varepsilon$ and write the logit denominator as $Z_{\text{price}} = 1 + \sum_j \exp(V_j)$. Direct differentiation yields

$$\frac{\partial d_i}{\partial \phi_i} = -\frac{1}{\varepsilon}\frac{\exp(V_i)}{Z_{\text{price}}} + \frac{1}{\varepsilon}\frac{\exp(2V_i)}{Z_{\text{price}}^2} = -\frac{1}{\varepsilon}d_i(1-d_i) < 0.$$

Creator i's expected profit from this type is

$$\pi_i(\phi_i) = \{(1-\tau)\phi_i - c_i\}d_i(\phi_i).$$

At an interior optimum, the product rule gives the necessary first-order condition

$$0 = \frac{\partial \pi_i}{\partial \phi_i} = (1-\tau)d_i - \{(1-\tau)\phi_i - c_i\}\frac{d_i(1-d_i)}{\varepsilon}.$$

Since an interior logit share satisfies $d_i \in (0,1)$, division by $(1-\tau)d_i$ is valid and gives

$$1 = \left(\phi_i - \frac{c_i}{1-\tau}\right)\frac{1-d_i}{\varepsilon},$$

or equivalently

$$\phi_i - \frac{c_i}{1-\tau} = \frac{\varepsilon}{1-d_i}.$$

The right-hand side is strictly increasing in the equilibrium share: $\partial\{\varepsilon/(1-d)\}/\partial d = \varepsilon/(1-d)^2 > 0$. This is an implicit equilibrium relation, not a demand-independent pricing formula, because d_i itself depends on ϕ_i, competitors' prices and match utilities, and the outside option. □

E.2 COMPARABLE SUBSTITUTES AND BRAND VALUE (PROP. 18)

Proposition 18 (Comparable substitutes)**.** *For $m \geq 2$ interchangeable creators with utility q, cost c, and type-I extreme-value shocks of scale ε, a unique symmetric interior pricing equilibrium exists. An increase in m lowers price ϕ_m, individual demand d_m, and individual profit $\pi_m = (1-\tau)\varepsilon d_m/(1-d_m)$, while raising transaction probability $\Lambda_m = md_m$ and brand inclusive value $W_m = \varepsilon \log[1 + m\exp((q-\phi_m)/\varepsilon)]$. Conditional effective choice equals m.*

Proof. Let $\bar{c} = c/(1-\tau)$ denote creator cost adjusted for the platform take. At a symmetric interior equilibrium, Proposition 17 implies

$$\phi_m = \bar{c} + \frac{\varepsilon}{1-d_m}. \tag{E.1}$$

The normalized outside-option share is $1 - md_m$. The logit odds of selecting any particular creator rather than the outside option therefore satisfy

$$\frac{d_m}{1-md_m} = \exp\left(\frac{q-\phi_m}{\varepsilon}\right). \tag{E.2}$$

We substitute (E.1) into (E.2) and define $\omega = (q-\bar{c})/\varepsilon$ to obtain the scalar equilibrium equation

$$\Psi(d,m) := \log d - \log(1-md) + \frac{1}{1-d} - \omega = 0, \qquad 0 < d < \frac{1}{m}. \tag{E.3}$$

For fixed m,

$$\frac{\partial \Psi}{\partial d} = \frac{1}{d} + \frac{m}{1-md} + \frac{1}{(1-d)^2} > 0. \tag{E.4}$$

Moreover,

$$\lim_{d \downarrow 0} \Psi(d, m) = -\infty, \qquad \lim_{d \uparrow 1/m} \Psi(d, m) = +\infty.$$

Continuity and strict monotonicity therefore give exactly one root $d_m \in (0, 1/m)$ for every integer $m \geq 2$. This root is a pricing equilibrium, not merely a solution to a necessary condition. With rivals' prices fixed, creator i's profit derivative has the sign of

$$H_i(\phi_i) = 1 - \frac{(\phi_i - \bar{c})(1 - d_i(\phi_i))}{\varepsilon}.$$

On $\phi_i \geq \bar{c}$, $H_i(\bar{c}) = 1$, $H_i(\phi_i) \to -\infty$ as $\phi_i \to \infty$, and

$$H_i'(\phi_i) = -\frac{1 - d_i}{\varepsilon} + \frac{(\phi_i - \bar{c}) d_i'}{\varepsilon} < 0,$$

because $d_i' < 0$. Prices below $\bar{c}$ yield negative profit, while the unique zero of H_i gives positive profit. Hence the root supplies each creator's unique best response at the symmetric profile.

For fixed admissible d,

$$\frac{\partial \Psi}{\partial m} = \frac{d}{1-md} > 0. \tag{E.5}$$

Take two integers $m_2 > m_1$ and write their unique equilibrium shares as d_2 and d_1. If $d_1 \geq 1/m_2$, feasibility of the second equilibrium immediately gives $d_2 < 1/m_2 \leq d_1$. If $d_1 < 1/m_2$, then d_1 lies in the domain of $\Psi(\cdot, m_2)$, and Equation (E.5) implies $\Psi(d_1, m_2) > \Psi(d_1, m_1) = 0$. Since $\Psi(\cdot, m_2)$ is strictly increasing by (E.4) and $\Psi(d_2, m_2) = 0$, it again follows that $d_2 < d_1$. Equation (E.1), whose right-hand side is strictly increasing in d_m, then yields

$$\phi_{m_2} < \phi_{m_1}.$$

Set $A_m = \exp\{(q - \phi_m)/\varepsilon\}$. The strict price decline implies $A_{m_2} > A_{m_1} > 0$. Consequently, $m_2 A_{m_2} > m_1 A_{m_1}$. Symmetric logit demand gives

$$\Lambda_m = \frac{m A_m}{1 + m A_m}.$$

The function $x \mapsto x/(1+x)$ is strictly increasing on $\mathbb{R}_+$; therefore, $\Lambda_{m_2} > \Lambda_{m_1}$. Likewise,

$$W_m = \varepsilon \log(1 + m A_m)$$

is strictly increasing in mA_m, so $W_{m_2} > W_{m_1}$.

Finally, conditional on choosing a creator, symmetry gives probability $1/m$ for each eligible creator. Hence the conditional inverse Herfindahl is

$$\left\{ \sum_{i=1}^{m} (1/m)^2 \right\}^{-1} = m.$$

Thus the comparative static is exactly an increase in local effective choice within this symmetric comparable-substitute environment. The markup relation also yields individual expected profit per arriving brand,

$$\pi_m = ((1-\tau)\phi_m - c) d_m = (1-\tau)\varepsilon \frac{d_m}{1 - d_m}.$$

Since $d/(1-d)$ is strictly increasing and d_m strictly decreases with m, π_m strictly decreases. Multiplication by any fixed positive arrival mass preserves that ordering. Appendix E.3 derives the separate aggregate-profit comparative static. □

E.3 Aggregate Creator Profit

For the symmetric pricing model write $A = (1-\tau)\varepsilon$ and $G(m) = m\pi_m/A = md_m/(1-d_m)$. For continuous $m \geq 2$, implicit differentiation of $\Psi(d, m) = 0$ yields

$$d'_m = -\frac{d_m^2(1-d_m)^2}{(1-d_m)^2 + d_m(1-md_m)}, \qquad G'(m) = \frac{d_m[1-(m+1)d_m + d_m^2]}{(1-d_m)[(1-d_m)^2 + d_m(1-md_m)]}.$$

The denominator is positive. Unlike individual profit, aggregate profit has no universal monotonicity: the sign depends on the bracketed expression. For the parameters in Table 2, total expected creator profit per arriving brand falls by approximately 30.8% relative to $m = 2$ when m increases to five. The included script checks the derivative against finite differences and also evaluates lower-value settings where aggregate profit increases with entry. Thus increased participation probability need not offset reduced markups, and a positive-sum claim requires demand expansion or additional stakeholder evidence.

E.4 Compatibility of Static Competition and Mean-Map Stability (Cor. 19)

Corollary 19 (Compatibility of separate modules)**.** *For fully overlapping equal creators, $\sigma = 1$ and $\beta < 1$ stabilizes the balanced mean map for any m. An increase in m separately improves the static brand outcomes above, while lowering individual creator profit.*

Proof. Conditional creator probabilities are $p(\theta) = \mathbf{1}/m$ in the symmetric fully overlapping market. Therefore $P = \mathbf{1}\mathbf{1}^\top/m^2$, whose smallest eigenvalue is zero for $m \geq 2$. Definition 2 gives $\sigma = 1$. The condition $\beta < 1$ is then equivalent to $\beta\sigma < 1$, so Theorem 4 establishes asymptotic stability of the linearized mean map and local asymptotic stability of the nonlinear mean map at the balanced allocation. Proposition 18 separately establishes the price, transaction-probability, and inclusive-value comparative statics as m increases. The pricing module does not feed prices or the outside-option decision into the reputation transition, so combining these results proves compatibility, not persistence in an integrated dynamic equilibrium. □

F Numerical Constructions and Simulation Protocols

F.1 Paired Finite-Score Markets with Equal Effective Choice

This construction evaluates the specified choice model with synthetic needs. There are four fixed creator positions $x_i = e_i \in \mathbb{R}^4$, where e_i is the ith coordinate vector. Thus both markets share $\Sigma = I/4 - \mathbf{1}\mathbf{1}^\top/16$ and the same ridge log-determinant for any common $\delta_D > 0$. Each market has four equally likely brand types.

Let

$$H = \frac{1}{2}\begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & -1 & -1 \\ 1 & -1 & 1 & -1 \\ 1 & -1 & -1 & 1 \end{pmatrix}, \qquad C_A = H\,\mathrm{diag}(1, 1/2, 1/2, 1/2)H^\top,$$

$$C_B = H\,\mathrm{diag}(1, \sqrt{0.67}, 0.2, 0.2)H^\top.$$

Each C is strictly positive and doubly stochastic. The diagonal/off-diagonal entries of C_A are $0.625/0.125$; the smallest entry of C_B is $(1 - \sqrt{0.67})/4 > 0.045$. With $\varepsilon = 1$, define brand type $\theta_k = \frac{1}{2}(\log C_{k1}, \ldots, \log C_{k4})$ and $q_i(\theta) = -\|x_i - \theta\|^2$. Because $q_i(\theta_k) = -1 - \|\theta_k\|^2 + \log C_{ki}$, the softmax at equal reputations is exactly row k of C.

Consequently $s = \mathbf{1}/4$, $a^* = \mathbf{1}/[4(1-\rho)]$, $P = C^\top C/4$, and $B = C^\top C$. Their spectra are

$$\mathrm{spec}(B_A) = (1, 0.25, 0.25, 0.25), \qquad \mathrm{spec}(B_B) = (1, 0.67, 0.04, 0.04).$$

Both have $\mathrm{tr}(P) = 7/16$, hence $\mathcal{D}_{\mathrm{loc}} = 16/7$, and baseline booking concentration $1/4$. In fact every row has squared norm $7/16$, so every brand has the same effective count in either market. Nevertheless $\sigma_A = 0.75$ and $\sigma_B = 0.96$. At $\beta = 1.2$, $\rho = 0.85$, Theorem 4 gives dominant factors 0.985 and 1.0228: one balanced mean map is locally stable and the other unstable. The comparison fixes creator positions and effective choice while varying brand-type distributions and score profiles.

Dynamics and checks. Figure 1 directly iterates the nonlinear mean map for 200 steps from $a^0 = a^* + 10^{-6}v$, where $v = (1, -1, 1, -1)^\top/2$ is a unit dominant direction in both markets. The final deviation divided by its initial value is 0.0486683 in A and 90.8121 in B. The plotted linear predictions are 0.985^t and 1.0228^t.

For a parameter check, replace $1/2$ in C_A by $a \in \{0.2, 0.35, 0.5, 0.6\}$, and use nonleading eigenvalues (b, c, c) for C_B, where $c = 0.4a$ and $b = \sqrt{3a^2 - 2c^2}$. All resulting choice matrices remain positive and each pair retains equal effective choice. For both markets, cross $\beta \in \{0.7, 0.9, 1.05, 1.2, 1.4\}$ with $\rho \in \{0.2, 0.6, 0.85, 0.95\}$, giving 160 configurations. Central-difference Jacobians with coordinate-relative step 10^{-5} agree with the analytic derivative throughout the grid.

The same primary comparison uses 20 random zero-sum unit directions at each initial magnitude 10^{-7}, 10^{-6}, and 10^{-5} (seed 20260924), paired across A and B. At 200 steps all 60 tested perturbations contract in A and expand in B; B's final/initial norm ratios range from 8.44 to 90.80. B also has contracting directions. The script `scripts/check_overlap_mechanism.py` writes the choice matrices, brand types, trajectories, parameter grid, and random-start results in `model_checks/`.

F.2 Parameterized Policy Comparison

The numerical illustration uses the fully overlapping symmetric model with $\sigma = 1$, $q = 0.60$, $c = 0.10$, $\tau = 0.20$, $\varepsilon = 0.12$, and $\rho = 0.85$. Static prices solve Equation (E.3). The reported concentration is that of normalized reputation, $\sum_i (a_i / \sum_j a_j)^2$, after 4,000 deterministic mean-map iterations from a fixed perturbed seed. It equals booking HHI at a fixed point, but need not do so at a finite nonstationary iterate. The static pricing and dynamic allocation columns are separate modules; the table is not a joint pricing–reputation equilibrium or an online ablation. Figure 8 also shows the $N = 50$ numerical trajectory and the analytic markup curve.

The scripts reproduce Table 2 and Figure 8 from these parameters.

Table 2: Synthetic policy comparison. Static payment, transaction, and value indices use the baseline policy as 100. Dynamic concentration is divided by the equal-share HHI $1/m$. The static and dynamic columns describe separate model modules.

Policy	m	$\beta\sigma$	Payment	Transactions	Value	HHI / equal
Baseline	2	1.10	100.0	100.0	100.0	2.00
Substitute floor only	5	1.10	78.1	104.9	153.5	5.00
Reinforcement guardrail only	2	0.90	100.0	100.0	100.0	1.00
Joint floor + guardrail	5	0.90	78.1	104.9	153.5	1.00

Numerical verification. Additional numerical checks construct nine positive unequal-share fixed points for $N \in \{2, 4, 9\}$ and $\beta \in \{0.4, 0.9, 1.4\}$. Finite-difference Jacobians agree with the analytic derivative. For $\rho = 0.8$, $\beta = 0.5$, and $a_1^0 = a_2^0 = 1$, a 300-factor product with a geometric remainder bound confirms a strictly positive probability of the all-creator-1 history. Pricing-root checks reproduce the table's static columns and the decreasing individual-profit comparative static.

F.3 Spectral Calibration Protocol

For each $N \in \{16, 64, 128\}$, creator positions are the coordinate vectors in $\mathbb{R}^N$. Uniformly distributed brand types correspond to rows of $C_\kappa = (1-\kappa)I + \kappa \mathbf{1}\mathbf{1}^\top/N$, with $\kappa \in \{0.2, 0.6\}$. Set $\varepsilon = 1$ and $\theta_{ki} = \frac{1}{2}\log C_{ki}$. Squared-distance utilities produce exactly $C_{k\cdot}$ at equal reputation, as in the paired-market construction. Every entry is positive, all creators are eligible, and the balanced state is a fixed point at every reinforcement strength. Its shares are known exactly: $s_i = 1/N$. Consequently $B = C_\kappa^\top C_\kappa$, $\sigma = 1 - (1-\kappa)^2$, and $\mathcal{D}_{\mathrm{loc}} = [(1-\kappa)^2 + (2\kappa - \kappa^2)/N]^{-1}$.

The finite family comprises the two overlap settings crossed with $\beta \in \{0.9, 1.2, 1.4, 2.0\}$, so $K = 8$ at each N. At the balanced state the probability matrices do not vary with β; each sampled matrix can therefore evaluate all four associated candidates. We use $\delta_c = 0.05$ and the exact leverage bound $L = N\|C_{\kappa,1\cdot}\|^2$ in Theorem 13. For each $n \in \{10^3, 10^4, 10^5\}$, 100 independent multinomial draws of type counts generate $\widehat{B} = NC^\top \operatorname{diag}(n_1/n, \ldots, n_N/n)C$. This is exactly equivalent to sampling n iid needs and accumulating their probability outer products. The seed is 2027.

Every calibration replicate covers the true exposure. For $N = 128, n = 10^5$ and margin $\eta = 0.1$, mean numerical ceilings are 1.797 and 1.083 times the universal reference 0.9, at $\kappa = 0.2$ and 0.6, respectively. Only prespecified candidates satisfying their own certificates may be selected; these ceilings do not certify arbitrary interpolated policies. For example, $\beta = 1.4$ can pass in the low-overlap setting, whereas $\beta = 1.2$ cannot pass in the high-overlap setting. The simulations evaluate known-score, balanced fixed points, not learned scores or uncertain shares. The released indices summarize every setting, including small-sample cases that do not improve on the universal cap.

F.4 Batched Local-Noise Protocol

We use $N = 32$, $\kappa = 0.4$, $\rho = 0.98$, and batch size $b = 256$. Each update averages 256 independent categorical bookings at the current state, then applies the normalized reputation transition. We initialize $\nu_i = 1/N$ and simulate 6,000 updates, omitting the first 2,000 when computing moments. The three strengths $\beta \in \{0.4, 1, 1.4\}$ have stable balanced mean maps. The statistic is the time-averaged squared projection of $S^{-1/2}(\nu - s)$ onto $(e_1 - e_2)/\sqrt{2}$. Replicate means estimate its second moment about the balanced state, not a within-replicate demeaned variance. Symmetry makes the population contrast mean zero from this start. There are 64 independent replicates per setting, with seed 2029.

Observed/predicted second-moment ratios are 0.9972, 1.0229, and 1.0497. Figure 3 divides both the estimates and their Monte Carlo uncertainty by the predicted moment, and uses t_{63} intervals across independent replicates; time points within a trajectory are not treated as independent. Maximum relative coordinate deviations across the recorded runs are 0.214, 0.296, and 0.596. No trajectory is censored for departing from the local neighborhood. These checks test the linear approximation in specified finite runs, not global convergence to a stationary distribution.

Table 3: Synthetic outcomes for $N = 16, 64$. Concentration is HHI divided by $1/N$; access shortfall is $100(1 - A/A_{\text{unif}})$, where $A_{\text{unif}} = 1 - (1 - 1/N)^{2000}$. Lower is better for both. Values are means $\pm$ 95% Monte Carlo half-widths across 64 paths.

N	Policy	HHI / equal shares	Access shortfall (%)
16	Baseline	1.07 ± 0.01	2.93 ± 1.08
16	Choice expansion	16.00 ± 0.00	93.75 ± 0.00
16	Cap only	1.02 ± 0.00	0.00 ± 0.00
16	Expansion + cap	1.16 ± 0.02	6.25 ± 1.30
16	Expansion + cap + access	1.07 ± 0.01	0.00 ± 0.00
64	Baseline	64.00 ± 0.00	98.44 ± 0.00
64	Choice expansion	64.00 ± 0.00	98.44 ± 0.00
64	Cap only	1.37 ± 0.02	21.00 ± 1.04
64	Expansion + cap	3.29 ± 0.06	64.62 ± 0.68
64	Expansion + cap + access	1.49 ± 0.02	5.25 ± 0.68

F.5 Event-Level Stochastic Policy Protocol

At each $N \in \{16, 64, 128\}$, initialize normalized reputation uniformly, set $\rho = 0.995$, and simulate 20,000 one-booking events in 64 independent replicates. The five settings (κ, β, γ) are baseline $(0.2, 1.4, 0)$, choice expansion $(0.6, 1.4, 0)$, cap only $(0.2, 0.9, 0)$, expansion plus cap $(0.6, 0.9, 0)$, and expansion plus cap plus access $(0.6, 0.9, 0.05)$. Choice expansion changes the probabilities among an unchanged eligible pool, not the number of retrieved candidates. The last setting mixes choice with uniform access at weight γ; the unmodified spectral formula is not claimed as its derivative.

Let $w_i = \nu_i^\beta$ and $z_i = (1-\kappa)w_i + \kappa \sum_j w_j/N$. The integral of the finite choice matrix over uniform needs gives the exact marginal booking probability

$$e_i(\nu) = w_i \left[\frac{1-\kappa}{N z_i} + \frac{\kappa}{N^2} \sum_j \frac{1}{z_j} \right].$$

We sample a one-hot booking from $(1-\gamma)e(\nu) + \gamma \mathbf{1}/N$ and update $\nu' = \rho\nu + (1-\rho)Y$. This has the same booking law as first sampling the independent brand type, then sampling its conditional choice. Uniform variates are shared across policy arms to pair replicates; replicates are independent, and the master seed is 2028.

Window booking HHI is computed from empirical shares during the final 2,000 events. Access is the fraction of eligible creators booked at least once in that window. Table 3 gives every market size and arm; intervals are t_{63} Monte Carlo intervals for replicate means, not platform-effect confidence intervals. Each arm's trajectory is also recorded every 500 events. At $N = 64, 128$, the baseline and expansion arms assign all recorded final-window bookings to one creator in every replicate. The window statistic records concentration over 2,000 events. At $N = 128$ the joint cap/expansion arm has lower access than the cap-only arm, while uniform mixing improves access in the joint setting. These outcomes show that a locally stable mean map can coexist with poor event-level access.

Reproduction and additional checks. `scripts/spectral_policy_experiments.py` writes the full calibration, noise, stochastic, and algebra outputs to `spectral_checks/`. The algebra stage (seed 2030) checks the aggregate-profit derivative against finite differences at 15 parameter combinations and solves the strictly concave fixed-point problem at three starts for each of nine configurations.

F.6 Boundary and Matched-Margin Experiments

We compare $\beta \in \{0.6, 0.9, 1.0, 1.2, 1.4\}$ at $\kappa = 1 - \sqrt{1 - 0.50/\beta}$. All settings therefore have $\beta\sigma = 0.50$ and, with $\rho = 0.995$, the same interior linear factor 0.9975. Each of $N \in \{64, 128\}$ uses 64 independent replicates of 20,000 events initialized at equal shares. Replicates share uniform draws across arms within each market size. The last 2,000 events define booking HHI, access, and the indicator that every booking went to one creator. The seed is 41027. The score matrix changes

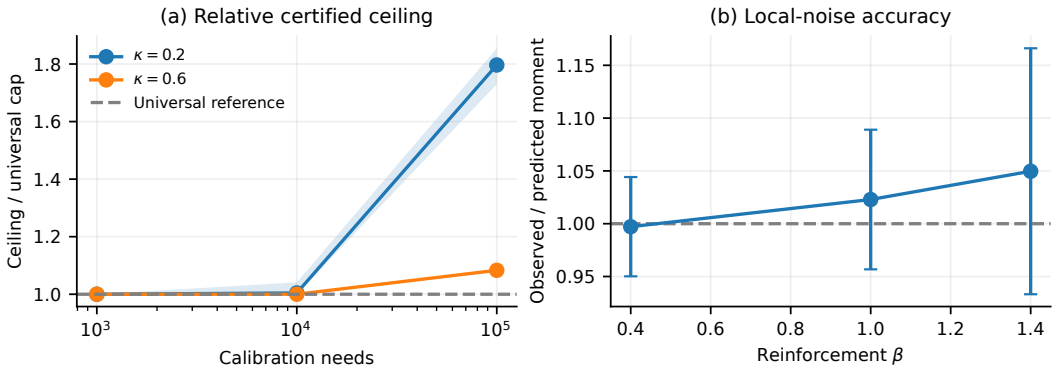

Figure 3: Synthetic spectral calibration and local-noise checks. (a) Certified ceilings relative to the universal cap at $N = 128$; shading shows the range across 100 calibration replicates. (b) Observed-to-predicted second-moment ratios across 64 batched-process replicates, with 95% Monte Carlo intervals on the same relative scale.

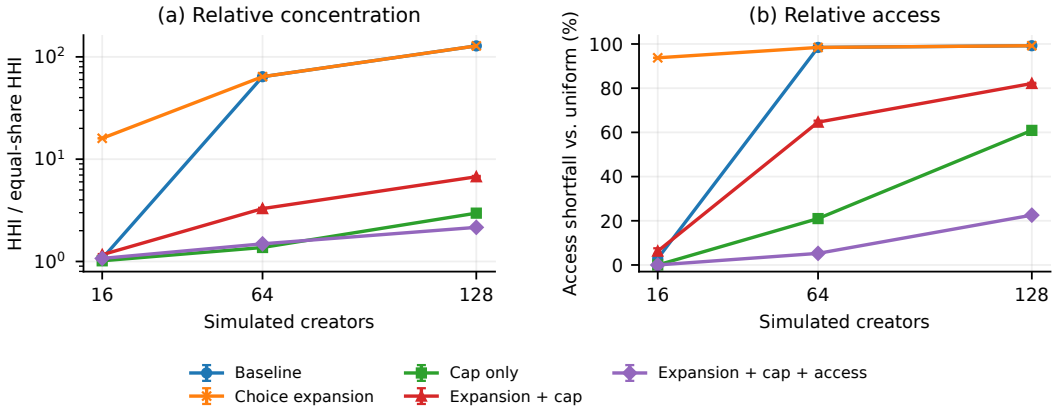

Figure 4: Synthetic market-size sensitivity under a fixed update rule. Points are 64-replicate means with 95% Monte Carlo intervals; all five arms are shown, including coincident curves. Concentration is relative to equal-share HHI on a logarithmic scale. Access shortfall uses the uniform reference in Table 1.

with β to match the margin; higher-order dynamics therefore differ even though the local contraction factor agrees.

At $\beta = 1.2$, no run has a complete final-window monopoly at $N = 64$, whereas 96.9% of runs do at $N = 128$. At $\beta = 1.4$, every run does at both sizes. Thus attracting corners can coexist with appreciable finite-window access; the boundary proposition does not specify a hitting time. Separate finite perturbations at distances 10^{-3}, 10^{-6}, and 10^{-9} check the missing-mass update near a vertex for $N \in \{4, 16, 128\}$.

The tighter Bernstein variance bound is evaluated using the 1,800 calibration draws. Its radius never exceeds the conservative radius, and every observed upper bound covers the true exposure. The aggregate outputs report both radii; Figure 3 uses the conservative radius in Theorem 13.

F.7 Memory and Batch Scaling

At $\beta = 1.4$, $\kappa = 0.2$, and $N \in \{16, 64, 128\}$, we compare fixed $1 - \rho = 0.005$, scaled $1 - \rho = 0.08/N$, and batches of $b = N$ at $1 - \rho = 0.005$. Each setting uses 64 independent replicates, seed 41028, and a horizon of 100 reputation-time units, where one update advances time by $1 - \rho$. We discard the first 50 units. Thus scaled-memory horizons are 20,000, 80,000, and 160,000 events; the other settings use 20,000 updates. The two $N = 16$ single-booking settings have iden-

Table 4: Matched $\beta\sigma = 0.50$. Relative outcomes use the uniform references in Table 1; monopoly frequency is the percentage of replicates. Intervals are 95% Monte Carlo half-widths.

N	β	κ	HHI / equal	Access shortfall (%)	Monopoly (%)
64	0.6	0.5918	1.46 ± 0.02	24.34 ± 0.90	0.0
64	0.9	0.3333	1.93 ± 0.03	42.11 ± 0.93	0.0
64	1.0	0.2929	2.25 ± 0.04	49.15 ± 0.81	0.0
64	1.2	0.2362	4.69 ± 0.97	65.41 ± 1.00	0.0
64	1.4	0.1982	64.00 ± 0.00	98.44 ± 0.00	100.0
128	0.6	0.5918	2.60 ± 0.03	56.70 ± 0.50	0.0
128	0.9	0.3333	4.20 ± 0.09	71.92 ± 0.43	0.0
128	1.0	0.2929	5.55 ± 0.13	77.37 ± 0.42	0.0
128	1.2	0.2362	127.99 ± 0.02	99.19 ± 0.03	96.9
128	1.4	0.1982	128.00 ± 0.00	99.22 ± 0.00	100.0

Table 5: Observed-to-predicted second-moment ratios with 95% Monte Carlo half-widths. Concentrated settings fall outside the local approximation regime. Update settings are simulation inputs.

N	Rule	$1-\rho$	b	Updates	Obs./pred.
16	Fixed	0.005000	1	20,000	1.405 ± 0.104
16	Memory	0.005000	1	20,000	1.512 ± 0.150
16	Batch	0.005000	16	20,000	1.002 ± 0.016
64	Fixed	0.005000	1	20,000	Concentrated
64	Memory	0.001250	1	80,000	1.532 ± 0.067
64	Batch	0.005000	64	20,000	1.015 ± 0.010
128	Fixed	0.005000	1	20,000	Concentrated
128	Memory	0.000625	1	160,000	1.531 ± 0.047
128	Batch	0.005000	128	20,000	1.008 ± 0.006

tical parameters and independent replicates. Each batch samples independent needs/bookings at the current state.

The measured moment is the time average of $N^{-1}\sum_i (N\nu_i - 1)^2$, with prediction $(N-1)(1-\rho)^2/[b(1-\lambda^2)]$. Memory scaling gives ratios $1.512, 1.532, 1.531$: it controls size dependence, but fluctuations remain large enough to exceed the linear prediction. Batch scaling gives $1.002, 1.015, 1.008$. Every fixed-memory run at the two larger sizes has a single-creator final window. At a vertex the relative second moment is $N-1$, so its ratio to the local prediction is about 198. These concentrated regimes lie outside the approximation's neighborhood; the table labels them separately. Intervals use independent replicate means and t_{63} critical values. Batch access fractions are not compared to single-booking access because equal numbers of updates contain different numbers of bookings.

F.8 A Fitted Synthetic Score Model

This experiment fits the symmetric choice family to generated booking data. With $N = 32$, $\kappa = 0.2$, and equal reputations, the probability that a booking matches its need's designated creator is $f_{\mathrm{diag}} = 1 - \kappa + \kappa/N$. The need type is observed for every generated booking. For $n_{\mathrm{fit}} \in \{1000, 10000, 100000\}$ independent need–booking pairs, we fit $\widetilde{\kappa} = N(1 - \widehat{f}_{\mathrm{diag}})/(N-1)$, projected to $[10^{-8}, 1 - 10^{-8}]$, which contains the true parameter. The fitted utility scores are $q_i(\theta_j) = \log C_{\widetilde{\kappa},ji}$ with $\varepsilon = 1$. The diagonal count is the sufficient statistic for this one-parameter model.

We solve the fitted mean map at $\beta = 1.4$, $\rho = 0.85$ from a small random perturbation of equal shares, to residual below 10^{-13}; symmetry identifies the exact fixed point as $s_i = 1/N$. Calibration uses 100,000 independent needs per fit. Hoeffding's inequality gives

$$|\widetilde{\kappa} - \kappa| \le \frac{N}{N-1}\sqrt{\frac{\log(2/\delta_m)}{2n_{\mathrm{fit}}}}$$

Table 6: Synthetic fitted-score results. Bounds and model-error allowances are divided by the true exposure; reinforcement ceilings are divided by the universal cap. Coverage and acceptance are percentages of fits.

Training	U/σ	Ceiling / universal	Δ_B/σ	Covers (%)	Accepts (%)
1,000	1.487	1.870	0.268	100	100
10,000	1.314	2.114	0.085	100	100
100,000	1.255	2.214	0.027	100	100

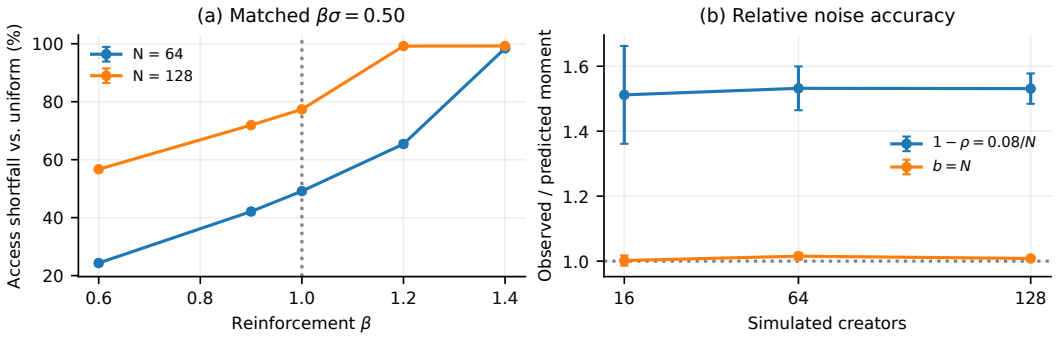

Figure 5: Synthetic boundary and noise checks, with 64 replicates and 95% Monte Carlo intervals. (a) All settings share $\beta\sigma = 0.50$ and $\rho = 0.995$; access shortfall is relative to uniform booking over the final 2,000 of 20,000 events, as in Table 1. (b) The moment ratio compares observed and predicted second moments. Memory scaling controls their growth but leaves approximation error; batches of size N closely match the linear-noise prediction.

with probability at least $1-\delta_m$. Since the normalized matrices are $(1-\kappa)^2(I-\mathbf{1}\mathbf{1}^\top/N)+\mathbf{1}\mathbf{1}^\top/N$, twice this bound is a valid Δ_B. We set $\delta_m = \delta_c = 0.025$ and apply Appendix D.3, using the conservative sampling radius of Theorem 13. There are 100 independent fits per training size, seed 41029. All runs cover the true exposure and accept the evaluated $\beta = 1.4$ candidate. Every accepted candidate nevertheless carries the boundary-attraction flag. This calculation assumes the specified score family; it does not bound model misspecification for platform data.

F.9 Mixing, Access, and Concentrated Interior States

For $N = 128$, $\beta = 1.4$, $\kappa = 0.2$, and $\rho = 0.995$, we simulate 64 independent paths at each $\gamma \in \{0, 0.01, 0.025, 0.05, 0.075, 0.1, 0.125, 0.15, 0.2, 0.3\}$. Each path starts balanced, lasts 20,000 events, and reports booking HHI and access over its final 2,000 events. Replicates share uniform draws across mixing rates; uncertainty uses independent replicate means and t_{63} critical values. The seed is 51027. All 640 paths are included. The access bound in Proposition 14 holds for any starting history, whereas the observed means depend on this initialization and horizon.

One-dominant branch. We restrict mean-map shares to $s = (x, (1-x)/(N-1), \ldots, (1-x)/(N-1))$ and solve $(1-\gamma)e_1(s) + \gamma/N - x = 0$. A grid of 121 mixing rates spans $[0, 0.3]$; root bracketing locates both stable and unstable solutions. Stability uses the full N-dimensional Jacobian, including directions outside this symmetric slice. The nonuniform branches meet at a mixing rate approximately 79.0% of the $\gamma = 0.2$ reference, from a simultaneous root and zero-slope calculation. The top share at this fold is 77.3 times equal allocation. At $\gamma = 0.05$, the stable high-share solution is approximately 119.6 times equal allocation. The calculation identifies this branch's fold; it does not exclude other multi-dominant equilibria. The stronger sufficient contraction condition of Appendix D.6 gives a global uniqueness statement, but requires a mixing rate above 4.945 times the 0.2 reference here.

Fluctuation scale and finite-window concentration. Figure 6 plots the matched-margin runs against their surrogate fluctuation scale, about 0.562 at $N = 64$ and 0.797 at $N = 128$. Dif-

Table 7: Synthetic mixing outcomes, expressed relative to uniform allocation. Concentration is HHI divided by $1/N$; access shortfall uses the reference in Table 1. Intervals are 95% Monte Carlo half-widths across 64 paths.

γ	HHI / equal shares	Access shortfall (%)
0.000	128.00 ± 0.00	99.22 ± 0.00
0.010	123.31 ± 0.35	83.81 ± 0.88
0.025	116.36 ± 0.57	65.33 ± 1.18
0.050	105.14 ± 0.75	42.86 ± 1.09
0.075	93.45 ± 0.89	28.39 ± 0.86
0.100	79.32 ± 3.57	18.80 ± 0.94
0.125	34.14 ± 8.38	9.28 ± 0.93
0.150	9.14 ± 4.62	4.90 ± 0.52
0.200	1.48 ± 0.02	1.86 ± 0.26
0.300	1.28 ± 0.01	0.50 ± 0.15

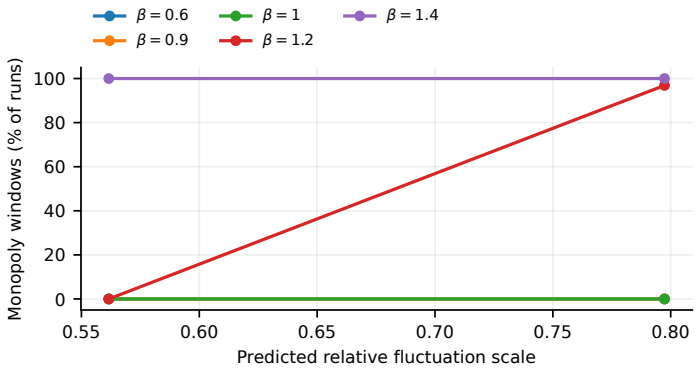

Figure 6: Synthetic matched-margin diagnostic. Points show the percentage of runs with complete final-window concentration. The two relative fluctuation scales come from the local surrogate; different reinforcement exponents can yield different concentration frequencies at the same scale.

ferent reinforcement exponents share each scale but have different monopoly fractions. Thus this two-size diagnostic illustrates sensitivity; it is not a fitted escape-time model or an independently validated prediction rule.

F.10 An Asymmetric Fitted-Score Pipeline

This synthetic study uses $N = 8$, $\beta = 1.4$, $\gamma = 0.9$, and true $\kappa = 0.2$. The known need probabilities are $F_j = j/36$ for $j = 1, \ldots, 8$. Training data consist of independent observed need types and bookings at equal reputations under the base choice rule. A booking's diagonal-match indicator is sufficient for κ; this indicator cannot be constructed without the need type. The fitted family is correctly specified.

Each training size $10^3, 10^4, 10^5$ has 100 fits, seed 51028. For each fit, we solve the mixed mean map numerically to ℓ_1 residual below 10^{-12}. The uniform contraction bound is $c = 0.622222$, so the residual yields the location bound in Appendix D.6. Independent calibration uses 100,000 need draws; both the mean shares and the conditional-choice covariance are estimated from those draws. The model/fixed-point allowance comes from Appendix D.7, with $\delta_m = \delta_c = 0.025$. The upper bound applies to the evaluated $\beta = 1.4$ candidate, not to an uncalibrated change of reinforcement.

The true fixed-point shares range from 0.935 to 1.066 times equal allocation. Every fitted upper bound covers the true exposure. At the two smaller training sizes, the parameter-error allowance is loose enough that the final bound equals the analytic c_γ. At 100,000 training pairs the mean bound is 40.2% of that reference. All candidates pass the chosen local margin, including those accepted by the analytic mixing bound alone. This high-mixing example tests estimated shares and numerical fixed-point error; it does not demonstrate platform model fit.

Table 8: Asymmetric synthetic fit. Exposure bounds and error allowances are divided by the analytic mixing bound c_γ. Coverage and acceptance are percentages of fits.

Training	U_γ / c_γ	Δ_H / c_γ	Covers (%)	Accepts (%)
1,000	1.000	3.654	100	100
10,000	1.000	0.948	100	100
100,000	0.402	0.285	100	100

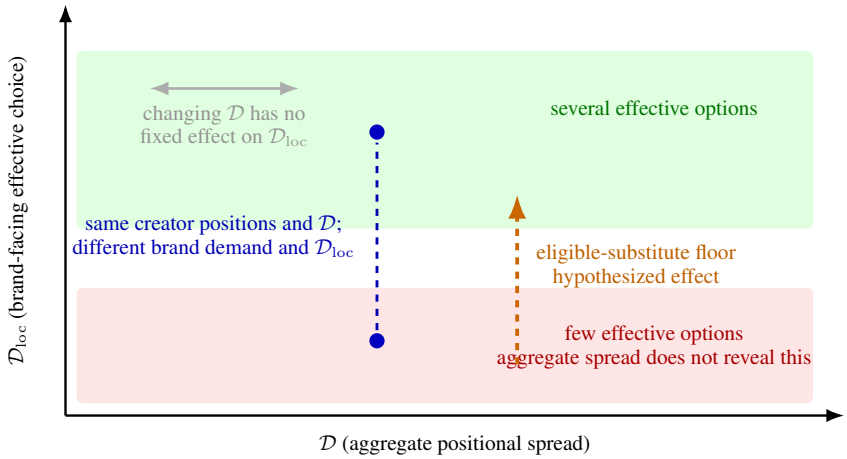

Figure 7: Conceptual separation of aggregate coverage and effective choice. The eligible-candidate floor's effect on achieved choice requires measurement.

F.11 Coverage and Perfect-Overlap Illustrations

Figure 7 separates coverage from effective choice; Figure 8 displays the perfect-overlap mean map and the separate pricing relation.

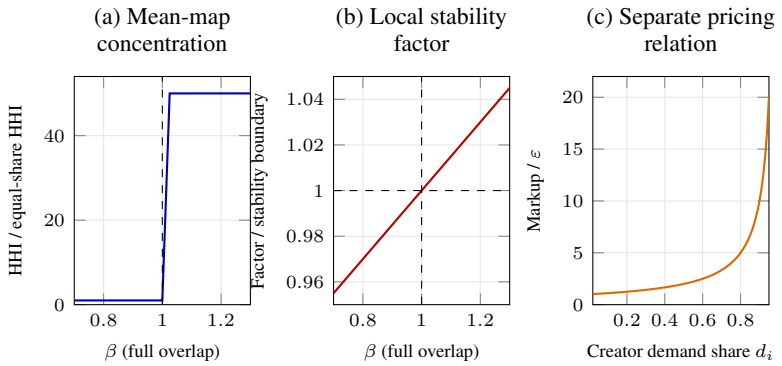

Figure 8: Numerical model illustration. (a) Concentration relative to equal shares after 4,000 mean-map steps, averaged over 12 perturbed starts with $N = 50$ and $\rho = 0.85$; this equals booking HHI at a fixed point before rescaling. (b) Factor relative to its unit stability boundary. (c) Markup relative to ε in the single-type model with an outside option; share and price are jointly determined.

G Observational Diagnostics and Relative Online Outcomes

G.1 Measurement Scope

Commercial records describe completed creator–brand matches rather than the candidate alternatives considered by each brand. They therefore do not identify the probability matrix required by the theory. Performance fields without usable variation and monetary fields with incompatible units are excluded from the reported outcome set. These restrictions prevent estimates of match quality or realized pricing power from the available records. The manuscript uses generic descriptions of audience size, content exposure, engagement, and promotion payments in place of production field names.

G.2 Observational Results

The analyses address D1–D2 and H1 in Section 5.4. D1–D2 examine the empirical informativeness of a roster-level proxy; H1 examines whether the distribution of booking opportunities is associated with creator advancement. These roles follow the paper's distinction between roster composition and the allocation creators experience. Every reported effect is a relative association per one-standard-deviation increase in the predictor. Specifically, we report $100\widehat{b}/\bar{y}$, where $\bar{y}$ is the outcome mean in that model's estimation sample, with the original regression weights. Standard errors and interval endpoints use the same positive divisor. This fixed-reference rescaling preserves test statistics and p-values; the intervals describe the rescaled coefficient conditional on the sample reference, rather than a population ratio estimand. Absolute outcome means, monetary amounts, sample counts, and source identifiers are withheld. The released indices support comparison of the reported estimates but do not permit independent reconstruction of the underlying regressions.

G.2.1 Audience-Balance Diagnostics: D1 and D2

For category g, week w, and its top-K creators by follower count T_K, define $f_{iw} = 1 + \text{followers}_{iw}$ and

$$\widehat{D}^{\text{fol}}_{g,w,K} = \frac{(\sum_{i \in T_K} f_{iw})^2}{\sum_{i \in T_K} f_{iw}^2}.$$

We refer to this follower effective count as the *audience-balance proxy*. It measures the spread of audience mass within the roster; brand preferences do not enter its definition. Raw followers with $K = 50$ form the primary specification. Square-root followers, follower–activity scores, and $K \in \{20, 100\}$ provide sensitivity checks. Post embeddings have not been validated as creator-position embeddings for the commercial marketplace.

Models control for log campaign volume, category fixed effects, and calendar-week fixed effects. Inference uses category-clustered standard errors (Cameron & Miller, 2015) and a restricted wild-cluster bootstrap-t with 9,999 draws and Webb six-point weights (MacKinnon & Webb, 2017). The two planned diagnostics form one Holm-adjusted family (Holm, 1979). The initial-booking association is exploratory. For a two-sided 5% test with 80% power, normal-approximation sensitivity benchmarks are 44.7% of the outcome mean for concentration and 9.5% for payment per predictor SD. These are reference-scale benchmarks, rather than power calculations for the multiplicity-adjusted bootstrap tests.

Table 9: Relative associations with the audience-balance proxy. Effects are percentages of each outcome's estimation-sample mean per predictor SD. Both planned diagnostics remain inconclusive.

Outcome	Relative effect (%)	Cluster t	Holm–wild p
Booking concentration (D1)	-17.23	-1.08	0.343
Mean promotion payment (D2)	-6.28	-1.85	0.141

D1 has a negative relative estimate, but its interval spans zero ($[-49.57, 15.10]\%$). Figure 9 indexes descriptive concentration to the lowest audience-balance decile. The highest decile is lower, while intermediate deciles need not be monotone. This descriptive pattern does not resolve the uncertainty in the controlled estimate. D2 likewise has a negative relative estimate with an interval that spans

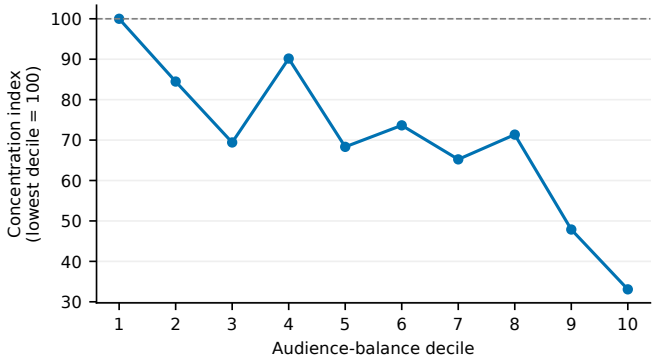

Figure 9: Relative booking concentration by audience-balance decile. The lowest decile is indexed to 100. Points are descriptive means; the controlled D1 estimate appears in Table 9.

zero ($[-13.19, 0.62]\%$). Average payment measures a transaction price, not a cost-adjusted markup. Neither diagnostic measures brand-facing substitution. The exploratory initial-booking association is $+0.16\%$ of its outcome mean (unadjusted wild $p = 0.819$). The analytical nonidentification result in Proposition 3 establishes the distinction between coverage and choice. D1–D2 add an empirical assessment of the roster proxy, with uncertainty that leaves both questions open; their non-rejection is not evidence that roster measures are irrelevant or that a local-choice measure would perform better.

G.2.2 Audience-Tier Advancement: H1

For a category–week–audience-tier peer cell C with weekly bookings b_i, the opportunity-evenness predictor is

$$\mathrm{Even}(C) = \frac{(\sum_{i \in C} b_i)^2}{|C| \sum_{i \in C} b_i^2}.$$

Peer count and evenness are therefore related by construction. Ordered audience-size tiers define peer cells; proprietary tier boundaries and category names are not disclosed. The primary outcome is advancement to a higher tier within 12 weeks. Weighted cell regressions control for mean observation gap, log peer count, log average audience size, posting activity, average bookings, tier, category, and calendar week. All retained creators have a weekly booking, so the estimates concern allocation among already-booked peers. Category-clustered intervals use the relevant cluster degrees of freedom. The restricted wild-cluster bootstrap uses the same draw count and weights as D1–D2. H1's three primary outcomes form a separate Holm family; the one-week check reports its unadjusted bootstrap p-value.

Table 10: Relative associations between opportunity evenness and audience advancement. Each coefficient is a percentage of its own outcome mean per predictor SD. Primary rows use Holm-adjusted wild-bootstrap p-values; the one-week check uses its unadjusted value.

Outcome / window	Relative effect (%)	Cluster t	Reported p
Tier advancement, ≤ 12 weeks	7.13	4.50	0.003
Any audience growth, ≤ 12 weeks	1.66	4.29	0.004
Audience growth rate, ≤ 12 weeks	4.69	2.66	0.008
Tier advancement, one week	12.87	4.26	< 0.001

A one-SD increase in evenness is associated with a 7.13% higher tier-advancement probability relative to its sample mean (95% interval $[3.92, 10.35]\%$). The corresponding associations are 1.66% for any audience growth ($[0.88, 2.45]\%$) and 4.69% for the continuous growth rate ($[1.11, 8.26]\%$). The one-week advancement association is 12.87% ($[6.74, 19.00]\%$). Each horizon uses its own sample mean; the relative magnitudes therefore cannot be read as differences in absolute advancement.

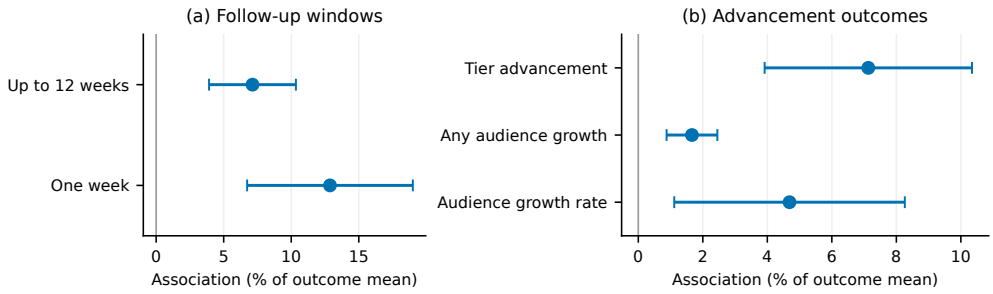

Figure 10: Relative opportunity-evenness associations. Points and 95% category-clustered intervals are divided by the corresponding outcome mean. Panel (a) compares observation windows; panel (b) compares outcome definitions. Reference means differ across outcomes and windows.

The positive associations agree across the reported horizons and outcome definitions. They provide observational support for the paper's focus on how commercial opportunities are distributed: more even allocation is associated with subsequent creator advancement among already-booked peers. Selection, prior momentum, and common demand remain possible explanations. H1 complements the reputation analysis as a downstream association; the model itself does not include a follower-growth equation. H2 separately evaluates the reported online GMV change in Appendix G.4.

G.3 Auxiliary User-Side Embedding Validation

A separate content-exposure analysis assesses preference signals. A deterministic 1% viewer sample is screened for inconsistent timestamps and invalid embeddings before preference-history construction. Sensitivity checks preserve the results' signs and adjusted significance decisions without the timestamp screen. Embedding-training provenance remains undocumented. Preference vectors use strictly earlier clicks. Models control for position, prior activity, predicted relevance, content type, taxonomy, ranking model, and time; viewer fixed effects and viewer-clustered standard errors account for repeated observations. Holm adjustment covers eight outcomes.

All coefficients use the relative reporting convention of Appendix G.2, with means computed on each actual estimation sample. A one-SD increase in text-preference similarity is associated with 10.50% more clicks relative to the click-rate mean ($p < 0.001$), 20.01% more active engagement relative to its mean ($p = 0.023$), and a 9.81% increase relative to mean log-transformed dwell time ($p < 0.001$). Individual non-click outcomes remain inconclusive after adjustment. In five-exposure sessions, a one-SD increase in semantic breadth is associated with -12.14% relative clicks ($p = 0.0029$), -44.06% relative active engagement ($p = 0.0375$), and -14.39% relative log-transformed dwell time ($p = 0.0017$). Continuation remains inconclusive ($p = 0.074$). Percentages for log-transformed time refer to that transformed outcome's mean; they are not percentage changes in dwell time itself. These associations concern immediate user feedback and do not measure brand-facing competition or causal effects of diversity.

G.4 Relative Outcomes from the Reported 60-Day A/B Test

The authors report a 60-day comparison between treatment and control. Table 11 uses generic outcome labels and expresses each change as $100(Y_T/Y_C - 1)$. A negative value for negative comments denotes a reduction relative to control. H2 concerns GMV; H3 concerns click-through rate, engagement, and negative feedback. The remaining rows provide the additional online outcomes recorded in the reported summary.

GMV describes realized transaction value during the comparison. CTR and marketplace-wide engagement describe interaction, while negative comments capture a reported negative-feedback outcome. Engagement covers the overall marketplace. The negative-comment figure retains its reported metric label; no exposure denominator is specified, so we do not interpret it as a change in the comment rate. The 60-day duration is reported by the authors; the experiment records needed to verify assignment, metric construction, uncertainty, and treatment implementation are not available in the

Table 11: Author-reported relative changes over a 60-day A/B test. Values are point estimates against control. The first four rows address H2–H3; the remaining rows are additional reported outcomes.

Metric	Relative change
Gross promotion transaction value (H2)	$+3.4\%$
Recommendation click-through rate (H3)	$+12.0628\%$
Marketplace-wide engagement (H3)	$+2.1687\%$
Negative comments (H3)	-4.8515%
Qualified visual-content views per viewer	$+0.1468\%$
Composite platform measure	$+0.0216\%$
Clicks on emerging creators' content	$+0.1758\%$
Visual-content completion	$+0.1942\%$
Recommendation-feed dwell time	$+0.1525\%$

analysis package. The table supports descriptive comparisons, not a significance test or a causal attribution to a particular policy component.

H Policy Specification and Implementation Requirements

H.1 Proposed Matching Policy

Algorithm 1 specifies the proposed candidate and scoring controls; actual deployment settings remain unverified. The policy records floor feasibility, the local spectral margin, and realized access separately. Its mixed mode uses Proposition 14 and the certificate in Appendix D.5. Access claims require a fixed pool that stays eligible at every event; a changing retrieved pool needs its own eligibility analysis. Small mixtures can preserve concentrated interior states, so the policy also records concentration and multi-start mean-map checks.

Algorithm 1 Relevance-filtered eligible-candidate floor matching policy

Require: Brand need θ; candidate creators with reputations $a_i > 0$; integer floor $m \geq 2$; exploration scale $\varepsilon > 0$; exposure upper bound $U \in (0, 1]$ uniform over the final retrieval rule and the contemplated reputation-weight range, if available; base reputation weight $\alpha_{\text{base}} \geq 0$; persistence $\rho \in (0, 1)$; safety margin $\eta \in (0, 1)$; fit score $q(x_i, \theta)$; eligibility rules and minimum credible-fit threshold $q_{\min}$

Require: Access window $T \geq 1$; mixing rate $\gamma \in [0, 1)$; for $\gamma > 0$, a fixed always-eligible pool of N creators with finite scores, covered by the selected policy's certificate

Ensure: Selected creator and updated reputation state, or a logged safe fallback

Retrieve candidates ranked by $q(x_i, \theta)$
Filter by eligibility and keep candidates with $q(x_i, \theta) \geq q_{\min}$
Backfill eligible candidates meeting $q_{\min}$ until at least m candidates are present, or log that the floor is infeasible
if no eligible candidate remains **then**
    Invoke the logged production-safe fallback and stop
end if
For $\gamma > 0$, invoke the logged fallback and stop if retrieval or filtering changes the certified fixed pool
For $\gamma = 0$ use an unmixed exposure bound; for $\gamma > 0$ use the mixed covariance bound U_γ as U
Without a statistical bound, use $U \leftarrow 1$ for $\gamma = 0$, or $U \leftarrow 1 - \gamma/[N - (N-1)\gamma]$ for the fixed-pool mixed policy
Set $\alpha \leftarrow \min\{\alpha_{\text{base}}, (1-\eta)\varepsilon/U\}$
Alternatively, select a prespecified policy whose own simultaneous certificate satisfies $\beta U \leq 1 - \eta$; do not clip to an uncertified setting
if $\gamma = 0$ and $\beta = \alpha/\varepsilon \geq 1$ **then**
    Flag boundary risk: under Proposition 6's assumptions, $\beta > 1$ attracts monopoly corners and $\beta = 1$ needs a utility-dependent test
end if
For $\gamma > 0$, check concentrated interior states separately; record the window-access bound $1 - (1 - \gamma/N)^T$
Compute $z_i \leftarrow q(x_i, \theta) + \alpha \log a_i$ for each eligible creator i
Set $p_i \leftarrow \exp(z_i/\varepsilon)/\sum_j \exp(z_j/\varepsilon)$; sample creator i from $(1-\gamma)p + \gamma \mathbf{1}/N$
Observe booking indicators $Y_j = \mathbf{1}\{j = i\}$
Update each reputation stock $a_j \leftarrow \rho a_j + Y_j$
Log eligibility, floor feasibility, candidate set, randomization assignment, score version, full probability vector, pre-booking reputation state, achieved effective choice, fit, and realized access
Record the local stability margin separately from the boundary flag and realized creator access

H.2 Feedback Architecture

Figure 11 locates retrieval and scoring controls within the reputation loop. The cap concerns local response at a positive fixed point; the boundary flag and persistent-access safeguard address different risks.

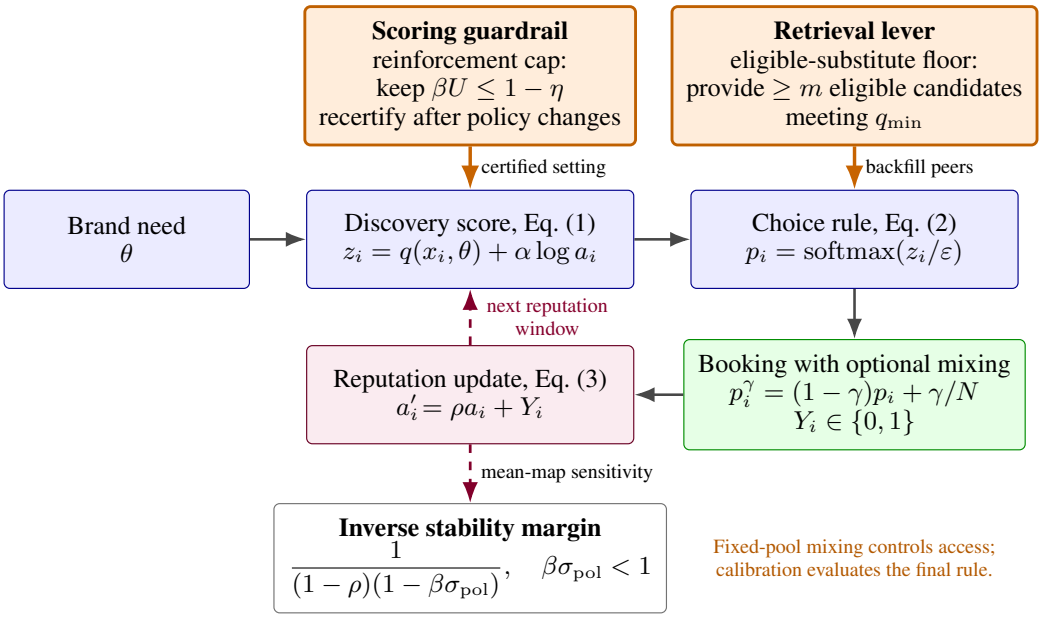

Figure 11: Proposed event-level reputation loop and intervention points. The exposure σ_{pol} is σ_* for the unmixed policy and σ_γ for the mixed policy. Fixed-pool mixing adds an access bound; local stability can still coexist with concentrated interior equilibria.

I Related Work

This section places the analysis within research on recommendation ecosystems, reinforced stochastic processes, and differentiated-product markets. The central comparison concerns what evolves, which allocation property is controlled, and what information is needed to evaluate a policy.

Reinforcement, social influence, and algorithmic feedback. Arthur (1989) explains how increasing returns can turn early advantages into persistent market dominance. Classical nonlinear urns provide a probabilistic foundation for reinforcement; the distinction between sublinear equalization and superlinear monopoly is established in this literature (Pemantle, 2007, Section 4.1). Informational cascades provide a different mechanism: agents infer information from previous actions and may disregard their private signals (Bikhchandani et al., 1992). Our reputation update represents an allocation rule without such Bayesian inference. In recommendation, Chaney et al. (2018) demonstrate how feedback-contaminated training data can homogenize behavior, while Jiang et al. (2019) analyze degeneracy in coupled recommendation and user-interest dynamics. Performative prediction formalizes the distribution changes induced by deployed predictors and studies the stability of retraining (Perdomo et al., 2020). Our model holds creator positions and match utilities fixed and updates commercial reputation. This restriction permits an explicit separation between the overlap-dependent interior threshold and the unit-exponent boundary threshold. The contribution is this separation under heterogeneous demand and depreciation; the classical reinforcement transition itself is an antecedent.

Strategic creators and endogenous supply. Ben-Porat & Tennenholtz (2018) study recommendation mechanisms with strategic content providers and introduce a mediator with fairness and equilibrium-stability properties. Hron et al. (2023) formulate exposure games in which creators adapt their content to a ranking algorithm, and show that modeling choices change the existence and character of equilibria. Jagadeesan et al. (2023) characterize specialization in multidimensional content markets and show that specialization can sustain positive producer profits. Under top-K recommendation and random-utility choice, Yao et al. (2023) bound user-welfare losses from creator competition with no-regret adaptation. Yao et al. (2024) examine the trade-off between immediate relevance and creator production volume as exploration changes. In generative-model markets, Wei et al. (2026) analyze platform model selection, welfare, and strategic entry. These studies make content, production, or entry endogenous. Our fixed-position model instead asks whether a reputation allocation returns after a small perturbation and whether creators continue to receive bookings. A Nash equilibrium of content choices and a locally stable reputation fixed point answer different questions; neither condition implies the other without a model of their joint dynamics.

Diversity and exposure fairness. Topic diversification balances relevance with the breadth of a recommendation list (Ziegler et al., 2005). Fleder & Hosanagar (2009) show that recommendations can increase individual consumption diversity while reducing aggregate sales diversity. Their result is an important antecedent to the need to distinguish levels of aggregation. Our positional coverage and conditional effective-choice count describe different objects from consumption diversity: the latter counts plausible alternatives for a specified brand need. Moreover, two markets can agree on both measures and differ in their normalized overlap spectrum. Fair-ranking methods impose constraints on the distribution of exposure (Singh & Joachims, 2018); dynamic learning-to-rank can control amortized group fairness while estimating relevance from biased feedback (Morik et al., 2020). Our stability diagnostic complements such allocation objectives. It measures the response to reputation perturbations, whereas a fairness constraint specifies how exposure should be distributed across items or groups.

Provider viability, matching, and persistent access. Mladenov et al. (2020) model provider viability thresholds and formulate long-run welfare optimization as constrained matching and equilibrium selection. Su et al. (2022) jointly optimize rankings in markets with reciprocal preferences and capacity constraints. Bandit matching studies competition when agents must learn unknown preferences (Liu et al., 2020), and merit-based exposure bandits control both reward regret and fairness regret (Wang et al., 2021). Our probability mixture addresses a narrower requirement: every creator in a fixed eligible pool retains a positive chance of a booking at each event. This yields persistent access and a finite-window access bound. It leaves merit-proportional exposure, sufficient earnings

for participation, and regret minimization as separate objectives. The mixed-policy examples further show that repeated access can coexist with a highly concentrated allocation.

Two-sided markets and differentiated-product demand. Platform economics emphasizes that participation on one side affects value on the other, and that the structure of charges matters alongside their total level (Rochet & Tirole, 2003; Armstrong, 2006). Discrete-choice theory provides the demand system used here: conditional logit links utilities to choice probabilities (McFadden, 1974), and heterogeneous tastes produce richer aggregate substitution patterns (Train, 2009). Anderson et al. (1992) develop the connection between discrete choice and price competition; Berry et al. (1995) combine differentiated demand and supply in an estimable market-equilibrium framework. Our static pricing module uses this established machinery to explain why comparable substitutes can discipline markups. The reputation dynamics condition on a booking, while the pricing module includes an outside option. They therefore provide separate comparative statics, rather than a jointly estimated equilibrium of prices, participation, and reputation. This distinction also prevents an access or stability guarantee from being interpreted as a welfare guarantee for every market participant.

Mean dynamics and stochastic fluctuations. The relation between a stochastic recursion and its mean field is central to stochastic-approximation theory (Benaïm, 1999; Kushner & Yin, 2003). With cumulative reinforcement, normalization commonly produces diminishing steps; reputation depreciation instead sustains a nonvanishing update scale. Consequently, local attraction of the mean map leaves room for noise-driven departures from its neighborhood. Linear-noise expansions describe small fluctuations around deterministic trajectories or equilibria (van Kampen, 2007). Our frozen-noise calculation applies this standard idea to identify the dependence on persistence, batch size, market size, and the overlap eigenvalues. The numerical scaling study examines the approximation in this setting. The matrix arguments use classical similarity, positive-semidefinite ordering, and eigenvalue perturbation theory (Horn & Johnson, 2012). The model-specific result is the normalized probability matrix that governs reputation response, together with the trace restriction linking effective choice to its extreme eigenvalue.

Finite-sample spectral evaluation. Matrix concentration provides nonasymptotic bounds for empirical averages of random matrices (Tropp, 2012; Vershynin, 2018). Our certificates apply these tools to probability Gram matrices for the base policy and conditional choice-covariance matrices for the mixed policy. Share normalization connects their spectra to the corresponding Jacobians. The statistical target is a local-response bound at a policy's own positive fixed point. This differs from the prediction-error, utility, and regret objectives above. A bound for the final policy must account for changes in retrieval, utilities, and reinforcement strength; a pre-intervention estimate need not remain valid. The fitted-model examples separate probability-model error, share error, numerical fixed-point error, and sampling error. Thus the concentration inequality is established machinery, while its policy-specific target and its integration with the feedback model constitute the application.

Measurement and empirical identification. Recommendation logs reflect earlier exposure decisions and user selection. Schnabel et al. (2016) develop propensity-based methods for debiased learning and evaluation under this selection process. In two-sided platforms, competition creates interference between treatment and control; Johari et al. (2022) analyze the resulting experimental bias and the role of randomization design. These concerns motivate distinct evidence requirements in our setting. Booked-creator records support descriptive concentration and mobility diagnostics, but the overlap matrix requires conditional probabilities over eligible alternatives. Policy-effect estimation additionally requires a documented assignment mechanism and appropriate inference. The controlled numerical studies test the proposed mechanism within the model, while the platform analyses provide descriptive context. This division keeps mathematical identification, statistical estimation, and causal validation as separate steps in the evaluation of reputation-aware matching.