Snobs and Conformists: Dynamic Adoption with Opposing Social Preferences

Tommaso Bondi

Cornell University

March 15, 2026

Abstract

We study dynamic adoption with private quality signals and social learning, introducing one departure from standard models: consumers disagree about whether popularity is desirable. *Snobs* value exclusivity; *conformists* value popularity. This preference conflict generates an endogenous boom-bust lifecycle and a quality-duration reversal: higher-quality products have shorter niche phases, because quality accelerates conformist entry and triggers earlier snob exit. An impossibility theorem shows that the reversal cannot occur when all consumers respond to popularity with the same sign, distinguishing it empirically from Bass diffusion, social learning, and network-effects models. Although quality enters primitive payoffs symmetrically, it polarizes equilibrium adoption: higher quality raises adoption, which lowers the conformist threshold but raises the snob threshold, so conformists are the more quality-responsive segment. On the supply side, a platform that broadcasts trending metrics boosts short-run adoption but compresses the niche phase. Welfare-optimal visibility is interior in the reversal regime. A profit-maximizing platform over-reveals when impatient but under-reveals when patient, and reveals less popularity information for higher-quality products – inverting the standard network-effects prescription. For platforms hosting identity-sensitive content, the results imply that suppressing popularity metrics can extend the discovery phase for high-quality products and improve both long-run profits and welfare.

1 Introduction

In early 2021, Clubhouse was the most exclusive app in tech. Access required an invitation, rooms were small, and the user base skewed toward Silicon Valley insiders. Within months, downloads surged past ten million, celebrity hosts replaced niche moderators, and the early users who had made the app interesting quietly left. By mid-2022, Clubhouse had lost roughly 90% of its peak traffic. The product had not changed; its audience had. The very popularity that validated the platform for new users destroyed what made it valuable to its original ones.

This tension arises in markets where adoption changes a product's positioning, increasing appeal for some consumers while eroding it for others. Fashion is the canonical case: an item first adopted by taste-makers becomes less attractive to them when adopted by the mass market, even as that same adoption makes it more attractive to consumers who value conformity. Similar dynamics arise for music, restaurants, venues, and online communities, where platforms broadcast popularity through rankings, trending modules, and algorithmic amplification. The central question is: how can a market scale when adoption itself changes who wants to adopt?

We formalize this tension in a dynamic adoption model. Consumers learn about product quality through private signals and public adoption outcomes – a standard learning environment (Banerjee, 1992; Bikhchandani et al., 1992). We introduce one key new ingredient: some consumers value popularity, while others value exclusivity. Borrowing the classic terms from Leibenstein (1950), one type (*snobs*) experiences disutility from broader adoption, while the other (*conformists*) experiences utility from it. This single departure suffices to alter both the equilibrium dynamics and the platform's information-design problem.

On the *demand side*, the equilibrium produces a boom-bust lifecycle with three phases. In a *discovery* phase, snobs adopt the product while it remains uncrowded, generating informative public signals about quality. Once beliefs rise far enough, conformists enter in force – a *surge* driven by the very popularity that snobs created. But the surge triggers its own unwinding: rising adoption pushes crowding costs above the quality benefit for snobs, and the product enters *decline*: its original audience exits. This pattern is not driven by exogenous satiation or taste decay; it emerges endogenously from the interaction of learning with preference conflict.

The model also delivers a *quality-duration reversal*. Higher quality raises adoption, which generates more favorable public signals and draws conformists in sooner. But each conformist who enters raises the crowding cost that pushes snobs out. The result is that higher-quality products can have *shorter* niche phases – quality attracts a mainstream audience so quickly that the exclusivity window collapses.

The pattern recurs in music: Arctic Monkeys built a devoted indie following through early demos shared on MySpace, then their debut became the fastest-selling in UK history and the original fans distanced themselves; Billie Eilish went from a SoundCloud discovery to the youngest artist to sweep the four major Grammys, and the underground audience that had prized her moved on.

This prediction is also consistent with evidence that cultural tastes spreading quickly also decline faster (Berger & Le Mens, 2009) – though the proposed mechanism there operates through memory decay and similarity-based interference rather than opposing social preferences – and that luxury products lose status as they diffuse (Bellezza, 2023). The two mechanisms generate a sharp empirical distinction: under memory-based interference, *all* adopters abandon fast-spreading tastes; under opposing preferences, only *snobs* abandon while conformists stay.

Conversely, in standard adoption models, popularity pushes all consumers in the same direction: it provides information, creates coordination value, or triggers bandwagon effects (Bass, 1969; Banerjee, 1992; Bikhchandani et al., 1992; Katz & Shapiro, 1985). We prove that the quality-duration reversal is impossible in any model where all consumers respond to popularity with the same sign, giving the reversal empirical content: observing it constitutes evidence of opposing social preferences.

Although quality enters primitive payoffs symmetrically, it polarizes equilibrium adoption by raising the adopter mass – which lowers the conformist threshold but raises the snob threshold. Conformists are therefore the more quality-responsive segment: for them, quality and adoption reinforce each other. Quality is self-defeating for snobs, because excellence attracts the crowd that destroys the exclusivity they value.

On the *supply side*, these demand patterns pose a concrete design problem for platforms that control the *visibility* of adoption information. Broadcasting trending metrics accelerates the conformist surge, boosting short-run engagement but compressing the niche phase – a *learning-composition tradeoff* between faster quality discovery and audience preservation.

The tradeoff is sharpened by an *endogenous informativeness* channel: the aggregate adoption signal becomes a noisier indicator of quality as conformists enter, because their adoption is driven partly by bandwagon effects orthogonal to quality. TikTok's For You page is the extreme case: the algorithm surfaces content almost entirely by engagement velocity, compressing the path from niche discovery to mass exposure into hours. Bandcamp, by contrast, displays no algorithmic feed, no trending charts, and no play counts – a design that has sustained independent music communities for over a decade.

The model shows that welfare-optimal visibility is interior: neither full transparency nor full opacity is efficient. Moreover, optimal visibility is *decreasing in product quality*, inverting the standard network-effects prescription that better products benefit most from transparency. Whether a profit-maximizing platform over- or under-reveals relative to a welfare benchmark depends on patience: impatient platforms over-reveal to capture the conformist surge; patient platforms in identity-heavy markets restrict visibility to preserve the discovery phase.

A natural objection is that snobs might dislike popularity *with conformists*, rather than popularity in general. The Online Appendix allows composition-sensitive preferences and shows that the core mechanism strengthens: each conformist entrant dilutes "coolness," amplifying the cascade. The extension also admits *cool equilibria* – steady states in which snobs adopt permanently and conformists are deterred – corresponding to markets like vinyl, literary fiction, and exclusive venues that sustain niche positioning indefinitely. The same appendix verifies robustness to convex social preferences.

Roadmap. Section 2 reviews the literature. Section 3 presents the model. Sections 4–5 characterize the consumer equilibrium and derive the quality-duration reversal. Section 6 analyzes platform strategy. Section 7 concludes. The Appendix contains proofs of the main results; the Online Appendix collects extensions and remaining proofs.

2 Related Literature

Fashion, status, and social preferences. Our opposing preferences for popularity capture distinction motives (Veblen, 1899; Bourdieu, 1984), identity economics (Akerlof & Kranton, 2000), signaling (Pesendorfer, 1995), network externalities (Katz & Shapiro, 1985), and social conformity (Bernheim, 1994). We take the coexistence of both types as primitive and study their *interaction* in a dynamic market.

Economic theories of fashion originate with Simmel (1957) and Leibenstein (1950). Corneo & Jeanne (1997a) formalize snob and bandwagon effects in a static framework; Amaldoss & Jain (2005) show how uniqueness-seeking generates upward-sloping demand; Amaldoss & Jain (2008) extend the analysis to reference group effects. We embed these opposing preferences in a dynamic environment with learning, generating lifecycle dynamics and comparative statics unavailable in static settings.

Several papers generate fashion cycles through alternative channels: Pesendorfer (1995) through signaling, Baumann & Olszewski (2021) through equilibrium multiplicity, and Ke et al. (2024) through cross-generation social product design. The mechanism here differs: cycles arise from heterogeneity in the *sign* of social payoffs, not from signaling, multiplicity, or generational turnover. To the best of our knowledge, the literature has not shown that opposing social preferences can reverse the quality-duration relationship.

Diffusion and social learning. The Bass diffusion model (Bass, 1969) is the workhorse framework for product adoption. Relative to Bass, the key departure is mixed-sign social payoffs, which matter in two places. First, Bass "innovators" adopt independently of the installed base; the snobs here respond *negatively* to aggregate adoption. Second, Bass predicts that higher quality extends the adoption lifecycle; we show quality can *shorten* it.

The "chasm" in technology adoption (Moore, 1991) is loosely analogous to the Phase I–II transition, though Moore's chasm reflects a demand-side friction between visionaries and pragmatists rather than a preference-based composition shift. In standard social-learning models (Banerjee, 1992; Bikhchandani et al., 1992), popularity matters because it is informative; here it is also payoff-relevant, with opposite signs across consumers.

Platforms and information design. Tucker & Zhang (2011) provide field-experimental evidence that toggling popularity displays shifts demand heterogeneously across product types – a direct analogue of the visibility parameter. Nistor et al. (2025) model authenticity-monetization tensions generating growth-then-decline patterns, and Cong & Li (2024) study influencer-seller matching with audience composition effects. On seeding, Godes & Mayzlin (2009) and Aral & Walker (2012) document positive spillovers under uniformly positive externalities; the model shows such strategies can backfire when spillovers have mixed signs. The visibility results connect to the broader platform information design literature (Ke et al., 2022; Berman & Oery, 2024; Berman et al., 2024) and to Bayesian persuasion (Kamenica & Gentzkow, 2011; Bergemann & Morris, 2019); Yao (2024) studies a related dynamic persuasion problem in which a firm endogenizes the consumer's information environment. To the best of our knowledge, the optimal signal has not previously been studied in a setting where information accelerates a *dynamic* composition shift.

Empirical evidence. Several empirical findings are consistent with the model's predictions. Berger & Heath (2007, 2008) show that consumers abandon products when outgroup members adopt – the micro-level mechanism underlying snob exit. Berger & Le Mens (2009) document that cultural tastes that spread quickly also decline faster, consistent with the quality-duration reversal. Bellezza (2023) finds that luxury products lose status as they diffuse.

Iyengar et al. (2011) document "opinion snobship" among early adopters, and Valsesia et al. (2020) show that social media users who follow fewer accounts are perceived as higher-status, consistent with low adoption signaling independence. Yoganarasimhan (2017) identifies fashion cycles in data, and Schoenmueller et al. (2021) find that influencer follower counts follow bell-shaped lifecycles consistent with the three-phase structure.

A related strand documents a "curse of recognition": Li et al. (2025) show that restaurants *losing* a Michelin star see improved consumer ratings, consistent with an expectations channel, while ruling out customer-mix changes. Rossi & Schleef (2026) find a related effect for Academy Award nominations. The present model offers a complementary channel through audience *composition* that may operate on a longer timescale. The two channels generate different timing predictions – expectations shift immediately, while composition shifts are lagged – and can be distinguished with consumer-level panel data.

3 Model

3.1 Environment

Time is discrete ($t = 0, 1, 2, \ldots$). A single product arrives at $t = 0$ with quality $\theta \in \{L, H\}$ drawn from the prior $\mathbb{P}(\theta = H) = p \in (0, 1)$.

We model adoption as a flow decision: agents choose each period whether to consume the product. This keeps the state space one-dimensional and is natural for markets where engagement is repeatedly reversible (streaming, following, repeat patronage).[1] Exit is endogenous: snobs leave only if crowding eventually outweighs improved beliefs.

There is a continuum of agents with mass normalized to 1. Agents are heterogeneous in their preferences over aggregate adoption. A fraction $\lambda \in (0, 1)$ are *snobs*, who value exclusivity; the remaining $1 - \lambda$ are *conformists*, who value popularity. Types are fixed and commonly known.

3.2 Payoffs

Let $n \in [0, 1]$ denote the mass of current adopters (flow, not stock), and let $v(\theta)$ denote the quality payoff, normalized to $v(H) = 1$ and $v(L) = 0$. Snobs derive utility

$$U^S(\text{adopt} \mid \theta, n) = v(\theta) - \alpha n \tag{1}$$

where $\alpha > 0$ measures originality preference.[2] Conformists derive utility

$$U^C(\text{adopt} \mid \theta, n) = v(\theta) + \beta n \tag{2}$$

where $\beta > 0$ measures conformity preference.

Each agent has access to an outside option yielding type-specific flow payoff c_τ, with $c_C > c_S \geq 0$. This ordering implies that snobs are more willing to experiment when adoption is low;

[1] Under stock adoption with an exit option, both entry and exit thresholds acquire an option-value premium $\omega(\hat{\theta}) > 0$ (Chamley, 2004), but the premium does not depend on the sign of social preferences.

[2] The Online Appendix shows robustness to convex specifications $U^S = v(\theta) - \alpha n^\rho$ ($\rho > 1$) and develops composition-dependent preferences.

it arises naturally from an exclusivity premium that effectively lowers the snob outside option. Absent, this, we would have $U^C(\cdot) \geq U^S(\cdot)$ irrespectively of both θ and n. This would, counterfactually, make conformists the early adopters, eliminating the snob-led discovery phase.

Assumption 1 (Interiority). $\alpha + \beta > c_C - c_S; \quad c_S < 1; \quad c_C < 1 + \beta.$

These ensure the two types' adoption thresholds cross in $(0, 1)$, that snobs can adopt at low adoption levels ($c_S < 1$), and that conformists participate for sufficiently high posteriors even at full adoption ($c_C < 1 + \beta$).[3]

3.3 Information

Every period, agent i receives a fresh private signal $s_{i,t} \in \mathbb{R}$ about the product's quality. Identifying the quality states with numerical values $L = 0$ and $H = 1$, signals are drawn from distributions with density $f(s \mid \theta)$ satisfying MLRP:

$$s_{i,t} \mid \theta \sim \mathcal{N}(\theta, \sigma^2) \tag{3}$$

where $\sigma > 0$ measures signal noise. Signals are conditionally independent across agents and time given θ. Since adoption is a flow decision, the payoff-relevant private information at time t is the agent's current-period posterior $\mu_{i,t}$, which combines the public belief $\hat{\theta}_t$ with the time-t signal $s_{i,t}$.

Assumption 2 (Signal Structure). $f(s \mid H)/f(s \mid L)$ is strictly increasing in s (MLRP). Densities $f(s \mid \theta)$ are continuous with $f(s \mid \theta) > 0$ for all $s \in \mathbb{R}$, $\lim_{s\to-\infty} f(s \mid H)/f(s \mid L) = 0$, and $\lim_{s\to\infty} f(s \mid H)/f(s \mid L) = \infty$.

These conditions are standard (Smith & Sørensen, 2000; Chamley, 2004) and ensure posteriors are strictly increasing in signals with full support on $(0, 1)$.

[3] The parameter values used in Figure 1 satisfy all conditions comfortably. When $\beta \geq c_C$, the conformist threshold can reach zero at moderate adoption levels; this does not affect the qualitative results.

Given prior belief $\hat{\theta}_t$ (the public belief at period t) and private signal s_i, agent i forms posterior:

$$\mu_i = \mathbb{P}(\theta = H \mid s_i, \hat{\theta}_t) = \frac{f(s_i \mid H)\hat{\theta}_t}{f(s_i \mid H)\hat{\theta}_t + f(s_i \mid L)(1 - \hat{\theta}_t)} \tag{4}$$

Under MLRP, μ_i is strictly increasing in s_i. Let $G(\mu; \hat{\theta}_t, \theta)$ denote the distribution of posteriors when public belief is $\hat{\theta}_t$ and true quality is θ, with continuous density $g > 0$ on $(0, 1)$.

Each period, agents simultaneously observe private signals, form posteriors via (4), and choose $a_i \in \{\text{adopt}, \text{pass}\}$ to maximize expected utility given rational expectations about aggregate adoption n^*.

4 Benchmark: Single-Period Equilibrium

4.1 Decision Rules

Consider agent i with posterior belief $\mu_i = \mathbb{P}(\theta = H \mid s_i, \hat{\theta}_t)$ and expected equilibrium adoption n^e.

Lemma 1 (Threshold Strategies). *Given posterior μ_i and expected adoption n^e, optimal strategies are:*

$$\textit{Snob adopts} \iff \mu_i \geq \alpha n^e + c_S =: \underline{\mu}^S(n^e) \tag{5}$$

$$\textit{Conformist adopts} \iff \mu_i \geq c_C - \beta n^e =: \underline{\mu}^C(n^e) \tag{6}$$

where c_S and c_C are type-specific reservation utilities.

The snob threshold $\underline{\mu}^S(n) = \alpha n + c_S$ is increasing in n: more adoption raises the quality bar for snobs. The conformist threshold $\underline{\mu}^C(n) = c_C - \beta n$ is decreasing: more adoption lowers the bar for conformists.[4]

At $n = 0$, $\underline{\mu}^S(0) = c_S < c_C = \underline{\mu}^C(0)$, so snobs face a lower adoption threshold and dominate early adoption – not because they are better informed, but because their preferences

[4] The opposing monotonicity has counterparts in Corneo & Jeanne (1997a)'s static analysis. The dynamic results are, to the best of our knowledge, new.

favor adoption when few others have adopted. The thresholds cross at $n^\dagger = (c_C - c_S)/(\alpha + \beta)$, which plays a key role in lifecycle dynamics.

4.2 Opposing Responses and Equilibrium Adoption

Snobs and conformists respond oppositely to aggregate adoption. The following corollary collects the key implications of the threshold structure for type-specific adoption.

Corollary 1 (Opposing Responses to Adoption)**.** *Under the threshold strategies in Lemma 1:*

(i) Snob adoption $n^S(n) = \lambda[1 - G(\alpha n + c_S)]$ *is strictly decreasing in* n *and reaches zero at* $\bar{n}^S \equiv (1 - c_S)/\alpha$*, where the threshold equals 1.*

(ii) Conformist adoption $n^C(n) = (1 - \lambda)[1 - G(c_C - \beta n)]$ *is strictly increasing in* n*.*

(iii) At any $n \in (0, \bar{n}^S)$*,* $\partial n^S/\partial n < 0$ *and* $\partial n^C/\partial n > 0$ *simultaneously.*

Snobs exit because rising adoption pushes their threshold $\alpha n + c_S$ toward 1, squeezing out all but the most optimistic agents. Conformists enter because rising adoption pulls their threshold $c_C - \beta n$ toward zero, admitting agents with increasingly pessimistic signals. The two groups' best responses thus move in opposite directions – a feature absent from any model where all agents respond to popularity with the same sign.

This opposing monotonicity is the primitive force behind both the lifecycle result and the impossibility result, consistent with evidence in Iyengar et al. (2011) and Berger & Heath (2008). In Pesendorfer (1995), abandonment reflects signal erosion, not crowding disutility; in Kuksov & Wang (2013), all consumers agree on trendiness. Pesendorfer's model predicts abandonment coinciding with a *new product introduction*, whereas the present model predicts abandonment in response to *audience composition shifts* without new entry.

4.3 Equilibrium Adoption Mass

Define the best-response mapping $\Phi : [0, 1] \to [0, 1]$ by:

$$\Phi(n; \hat{\theta}_t) = \lambda \cdot [1 - G(\underline{\mu}^S(n); \hat{\theta}_t)] + (1 - \lambda) \cdot [1 - G(\underline{\mu}^C(n); \hat{\theta}_t)] \tag{7}$$

where $G(\mu; \hat{\theta}_t)$ is the subjective CDF of posteriors given public belief $\hat{\theta}_t$. The term $1 - G(\underline{\mu}^\tau(n))$ is the mass of type-τ agents whose posteriors exceed the adoption threshold.[5]

Lemma 2 (Equilibrium Existence)**.** *Under Assumptions 2 and 1, the best-response mapping $\Phi : [0, 1] \to [0, 1]$ is continuous, with $\Phi(0) > 0$ and $\Phi(1) < 1$. Consequently there exists at least one equilibrium $n^* \in (0, 1)$ satisfying $\Phi(n^*) = n^*$.*

Uniqueness requires a contraction condition on Φ:

Assumption 3 (Uniqueness)**.** $\kappa \equiv \max_{n\in[0,1]} |\Phi'(n)| < 1$.

When $\kappa < 1$, Φ is a contraction, so equilibrium is unique and globally stable. When $\kappa > 1$, three equilibria can coexist: two stable (high and low adoption) and one unstable. All within-period results – opposing monotonicity, threshold strategies, the decomposition of adoption into snob and conformist components – carry through at any stable fixed point, because these depend on the local properties of Φ rather than on global uniqueness.

The dynamic results (lifecycle, reversal, impossibility) also extend: the FOSD-based duration ordering requires that $n^S(\hat{\theta})$ be decreasing in $\hat{\theta}$, which holds at any stable equilibrium where the crowding externality is locally strong enough (see Remark 2). The platform visibility results require $n^*(\hat{\theta})$ to be C^1 in $\hat{\theta}$, which holds at any stable fixed point by the implicit function theorem.[6]

5 Dynamic Equilibrium and Product Lifecycles

5.1 Dynamic Structure

Given public belief $\hat{\theta}_t$, period-t adoption is pinned down by the static fixed point, and dynamics arise only through belief updating. At the start of period t, agents observe a noisy public signal of previous-period adoption $\tilde{n}_{t-1} = n_{t-1} + \varepsilon_t$ where $\varepsilon_t \sim \mathcal{N}(0, \sigma_\varepsilon^2)$ represents

[5]The objective distribution $G(\mu; \hat{\theta}_t, \theta)$, which conditions additionally on true quality θ, governs belief updating and cross-quality comparisons (Section 5) but not within-period adoption decisions, which depend only on agents' subjective beliefs.

[6]Under multiplicity, the platform's visibility choice could additionally serve as an equilibrium-selection device. We maintain Assumption 3 throughout to focus on the novel mechanism.

measurement error in the platform's reported metrics, and update the public belief via Bayes' rule:

$$\hat{\theta}_t = \frac{f(\tilde{n}_{t-1} \mid \theta = H,\, \hat{\theta}_{t-1}) \cdot \hat{\theta}_{t-1}}{f(\tilde{n}_{t-1} \mid \theta = H,\, \hat{\theta}_{t-1}) \cdot \hat{\theta}_{t-1} + f(\tilde{n}_{t-1} \mid \theta = L,\, \hat{\theta}_{t-1}) \cdot (1 - \hat{\theta}_{t-1})} \tag{8}$$

where $f(\tilde{n} \mid \theta, \hat{\theta}_{t-1})$ is the density of observed adoption given true quality θ and prior belief $\hat{\theta}_{t-1}$. Let $n^*_\theta(\hat{\theta})$ denote realized adoption when agents hold belief $\hat{\theta}$ but quality is θ. Under MLRP, $n^*_H(\hat{\theta}) > n^*_L(\hat{\theta})$. The noise ensures Bayes' rule applies for all observations and that learning is gradual. In Section 6, the noise variance σ^2_ε becomes a platform choice variable.

Because adoption is a flow decision, continuation values cancel in the agent's Bellman equation: the adoption decision reduces to comparing current-period payoffs, and the myopic thresholds in Lemma 1 are dynamically optimal. A Markov Perfect Equilibrium exists under Assumption 1 and is unique under Assumption 3. A basic property of the equilibrium adoption function will be used throughout:

Lemma 3 (Monotonicity of Adoption in Beliefs)**.** *Under Assumptions 2 and 3, $n^*(\hat{\theta})$ is strictly increasing in $\hat{\theta}$.*

5.2 Three-Phase Lifecycle

The interaction of threshold strategies with social learning generates a distinctive lifecycle. When the product is high quality, beliefs drift upward, conformists eventually enter, and the resulting crowding drives snobs out. All dates and phase boundaries below refer to the expected path under a fixed quality state; on any given realization, the stochastic belief path may differ. Let $\epsilon > 0$ be a small constant.

Proposition 1 (Three-Phase Structure on the Expected Path)**.** *Suppose $\theta = H$, $\alpha + c_S \geq 1$, and $\alpha < \bar{\alpha}_{\text{growth}}(\hat{\theta}_0, \lambda)$.*[7] *Then along the expected path under $\theta = H$, aggregate participation exhibits three phases:*

(I) Growth *($t = 0, \ldots, t_1 - 1$), where $t_1 \equiv \min\{t \geq 1 : \mathbb{E}[n^C_t] \geq \mathbb{E}[n^S_t]\}$: $\mathbb{E}[n_{t+1}] > \mathbb{E}[n_t]$; $\mathbb{E}[n^S_t] > \mathbb{E}[n^C_t]$.*

[7] The growth bound $\bar{\alpha}_{\text{growth}}$ ensures that the snob threshold effect does not overwhelm belief-driven adoption growth in Phase I; see the Appendix for the formal expression.

(II) Snob Peak *($t = t_1, \ldots, t^*$): $\mathbb{E}[n_t^S] > 0$ and $\mathbb{E}[n_t^C] > 0$;* $\quad t^* \equiv \max\{t \geq 0 : \mathbb{E}[n_t^S] = \max_{s \geq 0} \mathbb{E}[n_s^S]\}$.

(III) Decline *($t > t^*$): $\mathbb{E}[n_t^S] \to 0$ as $t \to \infty$; $T = \min\{t > t^* : \mathbb{E}[n_t^S] < \epsilon\} < \infty$.*

The condition $\alpha + c_S \geq 1$ means that sufficiently widespread adoption can make even a known high-quality product unattractive to snobs. In the reversal regime ($\alpha > \bar{\alpha}(\lambda)$), this condition holds whenever $\bar{\alpha}(\lambda) \geq 1 - c_S$; under the closed-form approximation $\bar{\alpha} \approx \beta(1-\lambda)/\lambda$, this reduces to $c_S \geq 1 - \beta(1-\lambda)/\lambda$, which is satisfied by the simulation parameters.[8] When $\alpha + c_S < 1$, the snob threshold never reaches 1 and both types coexist permanently; the reversal requires $\alpha + c_S \geq 1$.

Niche-phase duration T measures how long the snob-supported phase persists. Conformists may continue buying long after snobs depart, but the product has lost its niche positioning. This three-phase structure echoes classical product lifecycle theory (Mahajan et al., 1990) but emerges from preference heterogeneity rather than technology diffusion.

For a given product, the model generates a single rise-and-decline episode. Fashion cycles in the sense of Simmel (1957) – recurring waves of adoption and abandonment – emerge from *sequential* product introductions: as product A enters decline, snobs migrate to product B.

The Phase I–II transition occurs through two reinforcing channels: rising beliefs push more conformist posteriors above their threshold, while rising n_t lowers the threshold itself ($\underline{\mu}^C = c_C - \beta n_t$). Phase III begins when conformist entry raises the snob threshold toward 1, squeezing out marginal snobs. Unlike Bass-model decline, snob exit is *endogenous*, driven by the same preference heterogeneity that generated growth (Schoenmueller et al., 2021).

Remark 1 (Low-Quality Products). When $\theta = L$, beliefs drift toward zero and adoption never surges. Because adoption stays low, the snob threshold remains near c_S and snob adoption declines only gradually as beliefs erode – without the conformist surge that drives rapid exit under H. The quality-duration reversal $T^H < T^L$ arises from this asymmetry: high quality triggers the conformist surge that compresses the snob-supported phase, while

[8] The growth condition ($\alpha < \bar{\alpha}_{\text{growth}}$) and the reversal condition ($\alpha > \bar{\alpha}(\lambda)$) are jointly satisfiable: the simulation parameters in Figure 1 ($\alpha = 1.5$, $\beta = 1.0$, $\lambda = 0.50$, $c_S = 0.05$, $c_C = 0.40$) satisfy both conditions comfortably, producing a clear three-phase lifecycle under H and the reversal $T^H < T^L$.

low quality never triggers it. In Figure 1(b), snobs exceed conformists at every point in time; the product lingers as a quiet niche because the mainstream never arrives.

5.3 The Quality-Duration Reversal

Standard diffusion models predict that higher quality unambiguously extends product lifecycles (Bass, 1969; Mahajan et al., 1990). We show this relationship can reverse.

Proposition 2 (Quality-Duration Reversal)**.** *There exists a unique* $\bar{\alpha}(\lambda) \in (0, \infty)$ *such that: for* $\alpha < \bar{\alpha}$*, higher quality extends the niche phase (*$T^H > T^L$*); for* $\alpha > \bar{\alpha}$*, higher quality shortens it (*$T^H < T^L$*).*[9]

The mechanism works through two reinforcing channels. Higher quality speeds belief improvement and conformist entry: better products generate more favorable public signals, beliefs rise faster, and conformists – whose threshold falls with adoption – enter sooner. But faster conformist entry raises the crowding burden on snobs. When snob aversion is strong enough ($\alpha > \bar{\alpha}$), the acceleration overwhelms the quality benefit: higher quality shortens the niche phase by attracting conformists so rapidly that snobs are crowded out.

The pattern is visible across identity-sensitive markets. HBO's *The Wire* spent five seasons as a cult favorite, championed by a small audience that treated the show as a mark of discerning taste. When critical consensus eventually crowned it the greatest television drama, mainstream viewership surged – and the early fans who had evangelized the show found that recommending it no longer signaled anything distinctive.

Remark 2 (Comparative Statics of the Reversal Threshold). The reversal threshold $\bar{\alpha}(\lambda)$ is defined exactly by $F(\bar{\alpha}) = 1$, where $F(\alpha) \equiv (1 - \lambda)\, g^C(n^*(\alpha))\, (\alpha + \beta)$ and g^C is the posterior density at the conformist threshold evaluated at the equilibrium adoption level. The function F is continuous, satisfies $F(0) < 1$ (by the contraction condition) and $F(\alpha) \to \infty$ as $\alpha \to \infty$ (by full support), so the intermediate value theorem gives existence and uniqueness. When the posterior density is approximately constant across the threshold

[9]The niche-phase duration T is well-defined (finite) under the maintained condition $\alpha + c_S \geq 1$ (Proposition 1).

region ($g(\underline{\mu}^S(n^*)) \approx g(\underline{\mu}^C(n^*))$), the threshold simplifies to

$$\bar{\alpha}(\lambda) \;=\; \frac{\beta(1-\lambda)}{\lambda}, \tag{9}$$

so the reversal occurs when $\alpha\lambda > \beta(1-\lambda)$ – i.e., when the aggregate crowding damage from conformist entry exceeds the aggregate conformity benefit.[10] Implicit differentiation gives $\partial\bar{\alpha}/\partial\lambda < 0$ (more snobs lower the bar for reversal) and $\partial\bar{\alpha}/\partial\beta > 0$ (stronger conformist preferences raise it).

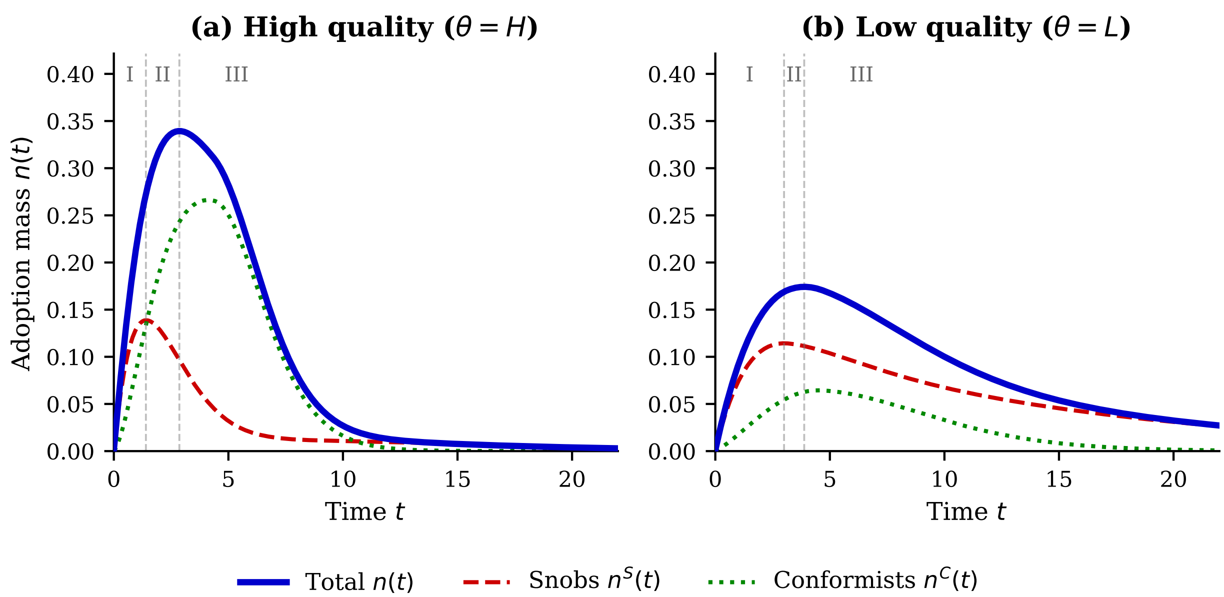

Figure 1: Equilibrium lifecycle dynamics. Panel (a): high-quality product; panel (b): low-quality product. High quality produces a larger but shorter cycle (the quality-duration reversal). Parameters: $\alpha = 1.5$, $\beta = 1.0$, $c_S = 0.05$, $c_C = 0.40$, $\lambda = 0.50$, $p_0 = 0.45$.[12]

Learning generates the timing; mixed-sign social payoffs generate the reversal. If agents observed θ directly – eliminating learning entirely – the mechanism would survive: higher quality would still raise adoption, which would still crowd snobs out. Learning produces the *dynamics* (gradual belief convergence triggers phase transitions), but the reversal itself is a

[10]The density approximation holds near the center of a symmetric posterior distribution and is exact when the prior is 1/2 and signal noise is moderate. The exact threshold, obtained by implicit differentiation of $F(\bar{\alpha}; \lambda, \beta) = 1$, inherits the same comparative statics without this restriction.

consequence of opposing preferences over n.

We have shown that the quality-duration reversal arises when consumers have opposing preferences over popularity. A natural question is whether the reversal could also arise in standard adoption models through other channels – for instance, through heterogeneity in the *magnitude* of social preferences rather than their sign. The answer is no.

Proposition 3 (Impossibility of Reversal Under Same-Sign Popularity Effects)**.** *Consider any threshold-based flow adoption model with utility*

$$u_i(\theta, n) = v(\theta) + \gamma_i \cdot h(n), \quad v(H) > v(L),\ h' > 0 \tag{10}$$

private posteriors satisfying MLRP, and Bayesian belief updating. If γ_i has the same sign for all agents, then $T^H \geq T^L$ along expected paths.

The "same sign" condition nests the leading adoption frameworks as special cases. In a flow reinterpretation of Bass diffusion (Bass, 1969), "innovators" adopt independently of the installed base ($\gamma_i = 0$) while "imitators" are drawn in by adoption ($\gamma_i > 0$); both signs are non-negative. Social learning (Banerjee, 1992; Bikhchandani et al., 1992) and network effects (Katz & Shapiro, 1985) similarly feature $\gamma_i \geq 0$ for every agent: adoption is (weakly) self-reinforcing. Congestion models have $\gamma_i \leq 0$ for all. Private-values models set $\gamma_i = 0$.

In every case, higher quality generates more favorable adoption signals at every t, so expected duration is weakly increasing in quality. The reversal $T^H < T^L$ therefore requires mixed-sign externalities: $\gamma_i > 0$ for some agents and $\gamma_j < 0$ for others. This gives the reversal empirical bite: observing that higher-quality products cycle faster is not merely evidence of fast diffusion, but evidence of sign heterogeneity in social preferences.

In the present model, T measures *niche-phase* duration; in same-sign models, the natural analogue is total adoption duration. The impossibility result applies to both: under same-sign externalities, Lemma 4 implies $\mathbb{E}[n_t \mid H] \geq \mathbb{E}[n_t \mid L]$ at every t, so *any* duration measure monotonically increasing in expected adoption satisfies $T^H \geq T^L$.

The Online Appendix shows that the reversal survives convex social payoffs and composition-sensitive crowding. The key requirement is that some agents are attracted to adoption while others are repelled by it.[13]

5.4 Quality Elasticities by Type

The reversal operates through aggregate adoption, but the underlying force is a divergence in how each type responds to quality. Quality enters both types' primitive utility identically (the $v(\theta)$ term is common), so neither type has a stronger primitive preference for quality. Yet quality is a polarizing factor in equilibrium, because it raises adoption, which shifts the two thresholds in opposite directions.

Corollary 2 (Asymmetric Quality Elasticity)**.** *At a stable within-period equilibrium where $dn^*/d\hat{\theta} > 0$:*

$$\frac{dn_t^C}{d\hat{\theta}} > 0 \quad \textit{always}; \qquad \frac{dn_t^S}{d\hat{\theta}} < 0 \iff \alpha > \frac{\partial_{\hat{\theta}}[1 - G(\underline{\mu}^S)]}{g(\underline{\mu}^S) \cdot dn^*/d\hat{\theta}} \tag{11}$$

Conformist adoption always rises with perceived quality; snob adoption falls with quality when crowding aversion is sufficiently strong.

Each type's equilibrium response to quality decomposes into a direct effect (better signals raise adoption above threshold) and an indirect effect (higher n shifts thresholds). For conformists, both effects reinforce entry. For snobs, the indirect channel works against adoption: higher n raises $\alpha n + c_S$, triggering exit. When α is large enough, the indirect channel dominates and better quality drives snobs out.[14]

The cumulative effect is even sharper than the within-period effect. Under high quality, snobs are crowded out early; under low quality, they persist because the conformist surge never materializes. In Figure 1, snobs account for 71% of cumulative adoption-periods under $\theta = L$ but only 29% under $\theta = H$ – a complete inversion of the adopter composition driven entirely by quality.

The same force operates at the micro level. In music, a band that achieves mainstream success through a hit single often watches its core following thin out as casual listeners flood in (Berger & Le Mens, 2009).

[13] Under multiplicative interaction ($v(\theta) \cdot (1 + \gamma_i h(n))$), the threshold's quality sensitivity depends on n, breaking the MCS argument. Under stock adoption, path-dependent option values break the static fixed-point structure.

[14] The Appendix shows that the quality-elasticity threshold coincides with the duration-reversal threshold $\bar{\alpha}(\lambda)$.

In terms of *preferences*, both types respond identically to quality. In terms of *equilibrium behavior*, conformists are more quality-responsive – because for them, quality and popularity are aligned. For snobs, quality and popularity work at cross purposes.

That snobs are the less quality-responsive segment may seem surprising, but it is a natural consequence of the equilibrium bundle they face: lower quality are often paired with greater exclusivity. In practice, many snob segments gravitate toward products whose apparent "deficiencies" serve as screening devices – vinyl's inconvenience, deliberately unglamorous – but expensive – fashion (the Balenciaga "ugly shoe" trend), obscure unissued records – because the product's limitations are not incidental but the barrier that preserves exclusivity.

For a platform monitoring creator health, this means that early-adopter churn following a popularity spike is not a sign of declining quality – it is a sign of *rising* quality that has attracted a mainstream audience. The appropriate response is to protect the creator's niche positioning (e.g., by moderating algorithmic amplification), not to infer that the product has deteriorated.

The snob share λ governs cycling speed: categories with larger early-adopter segments (e.g., streetwear, independent music) cycle faster, while mass consumer goods with small λ exhibit slow discovery and long lifecycles. Table 1 summarizes how the demand-side predictions differ from standard positive-spillover environments; the platform-design implications, developed in Section 6, are previewed in the bottom panel.

5.5 Welfare and Market Composition

The preceding results establish that snobs and conformists play complementary roles in the lifecycle. Snobs drive discovery: they adopt early, generating the informative public signals that allow the market to learn about quality. Conformists generate scale: their entry validates successful products and sustains adoption after the niche phase ends. Neither role alone produces a well-functioning market – without snobs, no one experiments; without conformists, no one scales – suggesting that welfare is maximized at an interior mix of the two types.

Proposition 4 (Optimal Market Composition)**.** *Let* $W(\lambda) = \lambda \cdot \mathbb{E}[U^S(\lambda)] + (1-\lambda) \cdot \mathbb{E}[U^C(\lambda)]$

Table 1: Predictions under positive spillovers vs. opposing preferences.

	Positive spillovers (network effects, Bass, social learning)	**Opposing preferences** (this paper)
Demand-side predictions		
Lifecycle shape	Monotone diffusion or S-curve	Endogenous boom-bust: growth → peak → decline (Prop. 1)
Quality and niche-phase duration	Higher quality ⇒ longer lifecycle	Higher quality ⇒ shorter niche phase (Prop. 2)
Reversal possible?	No: $T^H \geq T^L$ always (Prop. 3)	Yes: $T^H < T^L$ when $\alpha > \bar{\alpha}$ (Prop. 2)
Quality elasticity by type	Both types benefit from higher quality	Conformists benefit; snobs may be crowded out (Corollary 2)
Optimal audience composition	More consumers always weakly better	Interior snob share $\lambda^{**} \in (0,1)$ (Prop. 4)
Platform design		
Welfare-optimal visibility	Full transparency ($\varphi_W^* = 1$)	Interior: $\varphi_W^* \in (0,1)$ (Prop. 5)
Profit vs. welfare	Aligned at full transparency	Impatient platforms over-reveal; patient platforms under-reveal (Prop. 6)
Visibility and product quality	Show more for better products	Show *less* for better products (Prop. 7)
Dynamic visibility	Static: always reveal	Incubate then scale: $\varphi_E^* < \varphi_L^*$ (Prop. 8)
Hiding popularity metrics	Unambiguously harmful (slows positive feedback)	Can benefit niche content; weaker effects on mass content

Note: The positive-spillovers column represents a stylized benchmark with uniformly non-negative externalities ($\gamma_i \geq 0$ for all i), aggregating across network-effects, Bass, and social-learning models. Individual models within this class may differ in details, but all share the same-sign property that drives the contrasts shown here.

denote social welfare as a function of the snob share.

(i) The welfare-maximizing snob share is interior: $\lambda^{**} \in (0,1)$*.*

(ii) The welfare-maximizing snob share is decreasing in crowding aversion and increasing

in signal noise: $d\lambda^{**}/d\alpha < 0$ *and* $d\lambda^{**}/d\sigma > 0$.

The interiority result follows from a tension between what each type would prefer. Snobs prefer *fewer* snobs: each additional snob intensifies crowding, and while more snob adoption improves public signals, this benefit accrues primarily to conformists rather than to snobs (who already hold favorable private signals in Phase I). Hence $d\mathbb{E}[U^S]/d\lambda < 0$ at λ^{**}. Conformists prefer *more* snobs: each additional snob generates informative early-adoption signals and raises Phase I adoption, lowering the conformist threshold, so $d\mathbb{E}[U^C]/d\lambda > 0$ at λ^{**}.

Since welfare is a weighted average of the two types' surpluses, the optimum lies between the compositions each type would prefer: $\lambda^S < \lambda^{**} < \lambda^C$. Neither type internalizes the full social value of its participation: snobs ignore the informational externality they provide to conformists; conformists ignore the crowding externality they impose on snobs.

Part (ii) characterizes how the optimal composition shifts with market primitives. Noisier private signals (σ) raise the marginal informational value of snobs, tilting the optimum toward a snob-heavy mix. The platform cannot directly choose λ, but its visibility policy affects the active composition: broadcasting trending metrics accelerates conformist entry, shifting the mix away from the snob-heavy composition that maximizes welfare in identity-sensitive categories.

The analysis treats λ, α, and β as primitives; in practice, these may respond to platform design. The visibility results below are most applicable to short-run policy choices.[15]

6 Implications for Platform Strategy

The demand analysis shows that any policy accelerating conformist entry shortens the discovery period. Platforms typically do not sell the product directly; instead, they choose how much adoption information to reveal.

Popularity metrics – play counts, follower numbers, trending badges – are demand-

[15] Under equal per-capita welfare weights ($W = \frac{1}{2}\mathbb{E}[U^S] + \frac{1}{2}\mathbb{E}[U^C]$), the interiority of optimal visibility and the qualitative comparative statics are preserved. The platform results (Propositions 6–8) are independent of the welfare function.

composition instruments, not passive information displays.[16] Each platform's choice – Spotify and YouTube show real-time counts; Bandcamp shows almost nothing – trades off short-run engagement against the length of the discovery phase.

The setting connects to information design (Kamenica & Gentzkow, 2011): a more informative adoption signal helps conformists identify good products faster but also accelerates the crowding that drives snobs out. We formalize this tradeoff by parameterizing the precision of the public adoption signal.

Definition 1 (Platform Visibility)**.** The platform chooses visibility $\varphi \in [0, 1]$, which controls the precision of the public adoption signal. Agents observe a noisy lagged signal

$$\tilde{n}_{t-1} = n_{t-1} + \varepsilon_t, \qquad \varepsilon_t \sim \mathcal{N}(0,\, \sigma_\varepsilon^2(\varphi)) \tag{12}$$

where the noise variance is

$$\sigma_\varepsilon^2(\varphi) = \sigma_0^2 \cdot \frac{1-\varphi}{\varphi} \tag{13}$$

for a baseline noise level $\sigma_0^2 > 0$. At $\varphi = 1$, agents observe previous-period adoption perfectly ($\sigma_\varepsilon^2 = 0$); as $\varphi \to 0$, the signal becomes uninformative ($\sigma_\varepsilon^2 \to \infty$).

The lagged signal $\tilde{n}_{t-1}$ enters the belief-updating equation (8): higher φ makes the adoption signal more precise, so the market learns faster about product quality. In the reversal regime, this faster learning advances conformist entry and shortens niche-phase duration. Moreover, the informativeness of $\tilde{n}$ about θ is itself endogenous to the lifecycle phase: in Phase I, adoption is driven primarily by quality-sensitive snobs, so $\tilde{n}$ is a clean quality signal; in Phase II, conformist entry is driven partly by bandwagon effects (βn) that are orthogonal to θ, so $\tilde{n}$ becomes a noisier signal of quality precisely when adoption is highest.The composition shift from snobs to conformists thus degrades both the audience and the signal simultaneously – a reinforcing channel that strengthens the case for the "incubate then scale" policy derived below.

[16]The model's φ is an index of signal precision: $\varphi \approx 1$ corresponds to real-time play counts; φ near 0 to suppressed metrics with editorial curation only. Instagram's 2019 experiment hiding like counts (announced by Adam Mosseri at Facebook F8, April 30, 2019; expanded globally by November 2019) is a discrete shift in φ.

The platform earns per-period revenue $r(n_t)$ with $r' > 0$ and chooses φ before observing θ to maximize expected discounted revenue $\Pi(\varphi) = \mathbb{E}_\theta\big[\sum_t \delta_P^t r(n_t(\varphi,\theta))\big]$, where $\delta_P \in (0,1)$ is the platform's discount factor.[17] We first characterize a welfare benchmark, then compare profit-maximizing behavior against it. The planner maximizes

$$W(\varphi) = \lambda \cdot \mathbb{E}_\theta\big[\textstyle\sum_t \delta^t U_t^S(\varphi,\theta)\big] + (1-\lambda)\cdot \mathbb{E}_\theta\big[\textstyle\sum_t \delta^t U_t^C(\varphi,\theta)\big], \tag{14}$$

where $U_t^S \equiv \max\{v(\theta) - \alpha n_t,\, c_S\}$ and $U_t^C \equiv \max\{v(\theta) + \beta n_t,\, c_C\}$ are flow utilities, $\delta \in (0,1)$ is the common consumer discount factor, and expectations integrate over the quality state and the induced belief process.

Proposition 5 (Welfare-Optimal Visibility)**.** *Under the conditions of Proposition 2 with $\alpha > \bar{\alpha}(\lambda)$:*

(i) Welfare-optimal visibility is interior: $\varphi_W^* \in (0,1)$.

(ii) The welfare cost of full transparency is increasing in snob aversion and decreasing in snob share: $W(\varphi_W^*) - W(1)$ is increasing in α and decreasing in λ.

(iii) Welfare-optimal visibility is increasing in signal noise and decreasing in snob aversion: $\partial\varphi_W^*/\partial\sigma > 0$ and $\partial\varphi_W^*/\partial\alpha < 0$.

Neither extreme is optimal: full transparency destroys the discovery phase; full opacity prevents conformists from validating good products. The model predicts that experiments hiding popularity metrics (such as Instagram's 2019 like-count trial) generate the largest welfare gains in identity-sensitive categories (fashion, lifestyle) rather than informational ones (news, reviews).

A welfare planner would choose φ_W^*, but a profit-maximizing platform faces a different objective. Revenue depends on total adoption ($r' > 0$), so the platform values the conformist surge without accounting for the crowding cost it imposes on snobs.

[17] The ex ante commitment to φ rules out dynamic renegotiation. Without commitment, the platform faces a time-inconsistency problem: it would like to promise low early visibility to preserve the niche phase, but once snobs have adopted, the marginal revenue from raising visibility is positive and the platform has no incentive to maintain the restriction. Proposition 8 partially addresses this by allowing phase-specific visibility.

Proposition 6 (Profit-Maximizing Visibility)**.** *Under the conditions of Proposition 5:*

(i) There exists a unique patience threshold $\bar{\delta}_P \in (0,1)$*: impatient platforms over-reveal* ($\varphi^*_\Pi > \varphi^*_W$ *if* $\delta_P < \bar{\delta}_P$) *and patient platforms under-reveal* ($\varphi^*_\Pi < \varphi^*_W$ *if* $\delta_P > \bar{\delta}_P$)*.*

(ii) The patience threshold is lower in identity-heavy *markets (high* α*, high* λ*):* $\partial\bar{\delta}_P/\partial\alpha < 0$ *and* $\partial\bar{\delta}_P/\partial\lambda < 0$*.*

The intuition is that an impatient platform broadcasts trending badges and play counts to capture the conformist surge, trading niche-phase compression for short-run revenue. A patient platform suppresses popularity metrics to preserve the discovery phase that generates sustained engagement.

Part (ii) sharpens this: platforms in identity-heavy markets should under-reveal at almost any level of patience. Concretely: a music-streaming platform whose competitive advantage rests on curating independent and emerging artists should suppress or de-emphasize play counts and trending badges in its discovery interface. A platform oriented toward mainstream hits, where the audience is predominantly conformist, loses little from full transparency and gains from the faster quality sorting it provides.

At the extreme, platforms like the invite-only social network Raya build their entire value proposition on restricting access – keeping λ artificially high by screening out conformists at the door.[18]

The pattern is broadly consistent with observed practice, though cross-platform comparisons confound visibility choices with audience composition and business model. An interesting example is Spotify's interface: editorial playlists like New Music Friday and RADAR surface tracks *without* displaying play counts, while algorithmic playlists and artist pages show them prominently – lower visibility for curator-vetted content (identity-sensitive audience) and higher visibility for algorithmic recommendations (conformist-dominated).

How should visibility vary with product quality? A full-information benchmark establishes the direction.

[18] Clubhouse's trajectory illustrates the cost: after opening from invite-only to the general public in mid-2021, the original community dispersed rapidly – a discrete drop in λ that compressed the niche phase.

Remark 3 (Full-Information Benchmark). If the platform observes θ and can condition φ on it, the profit-maximizing visibility is lower for high-quality products: $\varphi^*_\Pi(H) < \varphi^*_\Pi(L)$. Better products attract conformists faster, making the niche-phase compression cost of visibility larger.

In practice, a platform does not observe θ directly but can condition visibility on observable quality indicators – category, critical reception, early engagement patterns. The above result holds more generally:

Proposition 7 (Optimal Visibility Decreasing in Perceived Quality)**.** *In the regime where the quality-duration reversal holds, suppose the platform observes a signal z correlated with θ and sets $\varphi(z)$. The profit-maximizing visibility is decreasing in perceived quality: $\varphi^*(z)$ is decreasing in* $\Pr(\theta = H \mid z)$.

A streaming platform should therefore display fewer popularity metrics for premium content than for mass content – consistent with Spotify's practice of suppressing play counts on editorial playlists while displaying them on algorithmic ones.

The monotonicity of φ^* in quality relies on the binary quality structure; with continuous quality, the relationship could be non-monotone if very-high-quality products attract conformists so rapidly that the niche phase is effectively zero regardless of visibility, eliminating the benefit of restricting it. More generally, the reversal regime expands as the quality gap $v(H) - v(L)$ increases, because a larger gap amplifies the differential learning speed across quality states and hence the conformist acceleration channel. The prediction "show less for better products" is sharpest in the interior range where the niche phase is positive but sensitive to visibility.

Tucker & Zhang (2011)'s field experiment provides a direct test: randomly toggling download-count displays for iPhone apps, they found that displaying popularity helped less popular products but hurt already popular ones. Their interpretation is informational (downloads signal quality for unknown products). The present model offers a complementary composition channel: displaying metrics helps conformist-oriented products but hurts snob-oriented ones. The informational channel predicts effects varying with *prior popularity*; the composition channel predicts effects varying with *audience type*. More broadly, hiding metrics benefits high-quality niche content while having little effect on mass-market products, distinguishing the model from "reduced social comparison" explanations that predict uniform effects.

6.1 Dynamic Visibility: Incubate Then Scale

The preceding analysis assumes the platform commits to a single φ for all periods. In practice, platforms often employ *dynamic* visibility policies: editorial curation during a product's early phase, followed by algorithmic amplification once the product has established a following. We now show that this pattern is optimal.

Suppose the platform can choose period-specific visibility $\varphi_t \in [0, 1]$. Partition the lifecycle into an *incubation phase* ($t < t_1$, when snobs dominate adoption) and a *scaling phase* ($t \geq t_1$, when conformists have entered). The platform chooses (φ_E, φ_L) – early and late visibility – to maximize discounted welfare.[19]

Proposition 8 (Optimal Dynamic Visibility)**.** *Under the conditions of Proposition 5:*

(i) The welfare-maximizing dynamic policy satisfies $\varphi_E^* < \varphi_L^*$: restrict visibility during incubation, increase it during scaling.

(ii) The welfare gain from dynamic visibility over the best static policy is increasing in α and decreasing in λ: $[W(\varphi_E^*, \varphi_L^*) - W(\varphi_W^*)]/W(\varphi_W^*)$ is increasing in α and decreasing in λ.

The intuition follows from the time-varying nature of the learning-composition tradeoff. During incubation, the composition cost dominates: low visibility preserves the discovery phase. During scaling, the learning cost dominates: snobs are already declining or have exited, so the marginal adopters are conformists who benefit from precise signals. Dynamic visibility exploits this asymmetry, and the value of doing so is largest precisely where the static reversal is strongest (high α, low λ).

[19] The two-phase formulation captures the key tradeoff while keeping the analysis tractable. Allowing fully time-varying φ_t yields a solution that is approximately step-shaped, with a discrete jump near t_1. The same ordering $\varphi_E^* < \varphi_L^*$ holds for a profit-maximizing platform with $\delta_P > \bar{\delta}_P$.

This "incubate then scale" pattern is consistent with observed platform practice. Apple's App Store implements it explicitly: the "Today" tab features editorially curated apps with no download counts visible, while the "Top Charts" tab ranks apps by downloads – a new app receives editorial incubation in the former before graduating to metric-driven discovery in the latter.

The prescription is concrete: delay aggregate engagement metrics until the early-adopter phase has run its course, then switch on algorithmic amplification. The optimal switching point is earlier for mass-market content and later for identity-sensitive content. Observable proxies for t_1 include a sustained increase in the share of new (vs. returning) users or a flattening of the early-adopter retention curve.

7 Conclusion

This paper develops a dynamic model of product adoption in markets where consumers have opposing preferences over popularity. The central result is the quality-duration reversal: better products have *shorter* niche phases, because excellence attracts a conformist audience that drives early adopters away. An impossibility theorem establishes that this reversal requires mixed-sign adoption externalities, giving it empirical bite as a diagnostic for opposing preferences. For platforms, the model suggests restricting popularity feedback in identity-heavy categories and disclosing *less* information for higher-quality products.

Four testable predictions emerge. First, within identity-sensitive categories, higher quality should predict shorter niche-phase durations. Second, a platform experiment hiding popularity metrics should extend the discovery phase for high-quality content but not mass-market products.

Two further predictions address the "curse of recognition" documented by Bondi (2025), Li et al. (2025) and Rossi & Schleef (2026): third, post-recognition satisfaction declines should partly reflect audience *composition* shifts, not only expectation inflation; and fourth, composition-driven declines should be *lagged* relative to the recognition event, whereas expectation-driven declines should be immediate. Both predictions are testable with consumer-level panel data linking adopter demographics to post-recognition outcomes.

Natural extensions include overlapping lifecycles, seller instruments such as artificial scarcity, and platform competition.[20] The model also abstracts from supply-side responses: creators who observe their own metrics may alter content to attract a broader audience, endogenously accelerating the composition shift. Finally, the binary quality structure delivers clean monotonicity in the visibility-quality relationship (Proposition 7); with continuous quality, the relationship could be non-monotone, and characterizing the optimal policy in this richer environment is an open problem.

The paper began with a question: how can a market scale when adoption itself changes who wants to adopt? The learning-composition tradeoff is the answer: in markets with strong originality preferences, scaling and niche preservation are in tension, and the platform's information policy determines where the balance falls. Every trending badge a platform displays is a bet that the conformist revenue the badge generates is worth the niche audience it displaces.

References

Akerlof, G. A., & Kranton, R. E. (2000). Economics and identity. *The Quarterly Journal of Economics*, 115(3), 715–753.

Amaldoss, W., & Jain, S. (2005). Conspicuous consumption and sophisticated thinking. *Management Science*, 51(10), 1449–1466.

Amaldoss, W., & Jain, S. (2008). Trading up: A strategic analysis of reference group effects. *Marketing Science*, 27(5), 932–942.

Aral, S., & Walker, D. (2012). Identifying influential and susceptible members of social networks. *Science*, 337(6092), 337–341.

Banerjee, A. V. (1992). A simple model of herd behavior. *The Quarterly Journal of Economics*, 107(3), 797–817.

Bass, F. M. (1969). A new product growth for model consumer durables. *Management Science*, 15(5), 215–227.

Baumann, L., & Olszewski, W. (2021). Demand cycles and heterogeneous conformity preferences. *Journal of Economic Theory*, 194, 105252.

[20] A competing platform that offers higher visibility may attract creators seeking faster audience growth, even if lower visibility would maximize long-run welfare. The monopolist visibility results here are a necessary first step; the competitive analysis is left for future work.

Bellezza, S. (2023). Distance and alternative signals of status: A unifying framework. *Journal of Consumer Research*, 50(2), 322–342.

Bergemann, D., & Morris, S. (2019). Information design: A unified perspective. *Journal of Economic Literature*, 57(1), 44–95.

Berger, J., & Heath, C. (2007). Where consumers diverge from others: Identity signaling and product domains. *Journal of Consumer Research*, 34(2), 121–134.

Berger, J., & Heath, C. (2008). Who drives divergence? Identity signaling, outgroup dissimilarity, and the abandonment of cultural tastes. *Journal of Personality and Social Psychology*, 95(3), 593–607.

Berger, J., & Le Mens, G. (2009). How adoption speed affects the abandonment of cultural tastes. *Proceedings of the National Academy of Sciences*, 106(20), 8146–8150.

Berman, R., Oery, A., & Zheng, X. (2024). Influence or advertise: The role of social learning in influencer marketing. Working paper.

Berman, R., Zhao, H., & Zhu, Y. (2024). Strategic design of recommendation algorithms. Working paper.

Bernheim, B. D. (1994). A theory of conformity. *Journal of Political Economy*, 102(5), 841–877.

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

Bondi, T. (2025). Alone, together: A model of social (mis)learning from consumer reviews. *Marketing Science*, 44(6), 1258–1277.

Bourdieu, P. (1984). *Distinction: A Social Critique of the Judgement of Taste*. Harvard University Press.

Chamley, C. (2004). *Rational Herds: Economic Models of Social Learning*. Cambridge University Press.

Cong, L. W., & Li, S. (2024). Influencer marketing and product competition. *Journal of Economic Theory*, 220, 105867.

Corneo, G., & Jeanne, O. (1997). Snobs, bandwagons, and the origin of social customs in consumer behavior. *Journal of Economic Behavior & Organization*, 32(3), 333–347.

Godes, D., & Mayzlin, D. (2009). Firm-created word-of-mouth communication: Evidence from a field test. *Marketing Science*, 28(4), 721–739.

Han, Y. J., Nunes, J. C., & Drèze, X. (2010). Signaling status with luxury goods: The role of brand prominence. *Journal of Marketing*, 74(4), 15–30.

Iyengar, R., Van den Bulte, C., & Valente, T. W. (2011). Opinion snobship and social contagion in new product diffusion. *Marketing Science*, 30(2), 195–212.

Kamenica, E., & Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, 101(6), 2590–2615.

Katz, M. L., & Shapiro, C. (1985). Network externalities, competition, and compatibility. *American Economic Review*, 75(3), 424–440.

Ke, T. T., Lin, S., & Lu, M. Y. (2022). Information design of online platforms. Working paper.

Ke, T. T., Liao, C., Long, F., & Lu, M. Y. (2024). Dynamic social product design and fashion classics.

Working Paper, available at SSRN 4720921.

Kuksov, D., & Wang, K. (2013). A model of the "it" products in fashion. *Marketing Science*, 32(1), 51–69.

Leibenstein, H. (1950). Bandwagon, snob, and Veblen effects in the theory of consumers' demand. *The Quarterly Journal of Economics*, 64(2), 183–207.

Li, X., Deng, Y., Manchanda, P., & De Reyck, B. (2025). Can lower(ed) expert opinions lead to better consumer ratings? The case of Michelin stars. *Management Science*, forthcoming.

Mahajan, V., Muller, E., & Bass, F. M. (1990). New product diffusion models in marketing: A review and directions for research. *Journal of Marketing*, 54(1), 1–26.

Moore, G. A. (1991). *Crossing the Chasm: Marketing and Selling High-Tech Products to Mainstream Customers*. HarperBusiness.

Morris, S., & Shin, H. S. (2003). Global games: Theory and applications. In M. Dewatripont, L. P. Hansen, & S. J. Turnovsky (Eds.), *Advances in Economics and Econometrics* (pp. 56–114). Cambridge University Press.

Nistor, C., Selove, M., & Villas-Boas, J. M. (2025). Influencer authenticity: To grow or to monetize. *Management Science*, 71(6), 5251–5267.

Pesendorfer, W. (1995). Design innovation and fashion cycles. *American Economic Review*, 85(4), 771–792.

Rossi, M., & Schleef, F. (2026). Quality disclosures and disappointment: Evidence from the Academy Awards. *Management Science*, forthcoming.

Schoenmueller, V., Libai, B., & Kogan, S. (2021). The rise and fall of influencers: Evidence from social trading. Working Paper, Bocconi University.

Simmel, G. (1957). Fashion. *American Journal of Sociology*, 62(6), 541–558.

Smith, L., & Sørensen, P. (2000). Pathological outcomes of observational learning. *Econometrica*, 68(2), 371–398.

Tucker, C., & Zhang, J. (2011). How does popularity information affect choices? A field experiment. *Management Science*, 57(5), 828–842.

Valsesia, F., Proserpio, D., & Nunes, J. C. (2020). The positive effect of not following others on social media. *Journal of Marketing Research*, 57(6), 1152–1168.

Veblen, T. (1899). *The Theory of the Leisure Class*. Macmillan.

Yao, Y. (2024). Dynamic persuasion and strategic search. *Management Science*, 70(10), 6024–6043.

Yoganarasimhan, H. (2017). Identifying the presence and cause of fashion cycles in data. *Journal of Marketing Research*, 54(1), 5–26.

Appendix: Proofs of Main Results

Notation. Product quality is $\theta \in \{L, H\}$ with $v(H) = 1$ and $v(L) = 0$. Signals $s \sim \mathcal{N}(\theta, \sigma^2)$ satisfy MLRP. The posterior given prior $\hat{\theta}_t$ and signal s is $\mu(s; \hat{\theta}_t) = \Pr(\theta = H \mid s, \hat{\theta}_t)$. We use two posterior distributions: the *subjective* distribution $G(\mu; \hat{\theta}_t)$ governs adoption decisions; the *objective* distribution $G(\mu; \hat{\theta}_t, \theta)$ governs belief updating. When the distinction is immaterial, we write $g(\mu)$ for brevity. The best-response mapping $\Phi(n; \hat{\theta})$ denotes the within-period subjective fixed-point mapping; $\Phi(n; \hat{\theta}, \theta)$ denotes the objective mapping under true state θ (used in cross-quality comparisons). Preference parameters: $\alpha, \beta > 0$. Reservation utilities: $c_C > c_S \geq 0$. Snob mass: $\lambda \in (0, 1)$.

Proofs of Lemma 1 and Lemma 2 are in the Online Appendix. Extensions (composition-dependent preferences, convex social preferences) also appear in the Online Appendix.

A.1 Proof of Proposition 1 (Three-Phase Structure)

Proof. We establish the three phases in sequence.

Step 1: Snob-dominated initial adoption. At $t = 0$, $\hat{\theta}_0 = p$. Since $c_S < c_C$ (Assumption 1), $\underline{\mu}^S(n) < \underline{\mu}^C(n)$ for all $n < n^\dagger$, so snobs face a lower threshold. At the equilibrium n_0^*, the snob component of Φ exceeds the conformist component whenever $\lambda/(1-\lambda) > [1 - G(c_C)]/[1 - G(c_S)]$, which holds for all $\lambda \geq 1/2$ and more generally when $c_C - c_S$ is not too small.[21]

Step 2: Belief dynamics. Under $\theta = H$, realized adoption satisfies $n_0^H > n_0^L$ (MLRP), so the adoption signal $\tilde{n}_0 \sim \mathcal{N}(n_0^\theta, \sigma_\varepsilon^2)$ is stochastically higher under H. Bayesian updating gives $\mathbb{E}[\hat{\theta}_1 \mid H] > \hat{\theta}_0$. By induction (Lemma 4), $\hat{\theta}_t$ is a bounded submartingale under H, so $\hat{\theta}_t \to 1$ a.s. by the martingale convergence theorem (the KL divergence between signal distributions is strictly positive at every interior $\hat{\theta}$, ruling out interior convergence).[22]

Step 3: Phase transitions. As $\hat{\theta}_t \to 1$, conformist adoption $\mathbb{E}[n_t^C] \to (1 - \lambda)$ while snob adoption $\mathbb{E}[n_t^S] \to 0$ (Step 4), so $t_1 < \infty$. Under the condition $\alpha < \bar{\alpha}_{\text{growth}}$, expected

[21] The condition fails only when λ is very small and $c_C \approx c_S$; the three-phase structure then degenerates. The quality-duration reversal (Proposition 2) does not require snob dominance at $t = 0$ – it requires only that $n^S(\hat{\theta})$ be globally decreasing in $\hat{\theta}$, which is a property of the static equilibrium independent of initial composition.

[22] Phase definitions characterize the expected path under $\theta = H$; on any realized path, stochastic beliefs may produce temporary deviations.

snob adoption is non-decreasing during Phase I (Step 5), ensuring $t^* \geq t_1$. The IFT gives $dn_t/d\hat{\theta}_t = (\partial\Phi/\partial\hat{\theta})/(1-\Phi') > 0$ at any stable equilibrium, confirming that adoption responds positively to beliefs within each period.

Step 4: Snob exit ($T < \infty$). Snob adoption vanishes when $\underline{\mu}^S(n_t) = \alpha n_t + c_S \geq 1$, i.e., $n_t \geq \bar{n}^S \equiv (1 - c_S)/\alpha$. Under $\alpha + c_S \geq 1$ (so $\bar{n}^S \leq 1$), we show the limiting equilibrium satisfies $n^S_\infty = 0$. As $\hat{\theta} \to 1$, posteriors concentrate near 1 and $1 - G(\mu; 1) \to \mathbf{1}[\mu < 1]$. If $\alpha n_\infty + c_S < 1$, the snob term approaches λ and the conformist term approaches $(1 - \lambda)$, giving $n_\infty \to 1$; but at $n = 1$ the threshold is $\alpha + c_S \geq 1$, contradicting $\alpha n_\infty + c_S < 1$. Hence $\alpha n_\infty + c_S \geq 1$ and $n^S_\infty = 0$. Dominated convergence gives $\mathbb{E}[n^S_t] \to 0$, so $T < \infty$. (When $\alpha + c_S < 1$, snobs persist permanently and the three-phase structure does not arise.)

Under $\theta = L$, beliefs drift downward, adoption stays low, and snobs decline only gradually – the conformist surge never materializes. The comparison $T^H < T^L$ is established in the proof of Proposition 2.

Step 5: Growth-phase monotonicity. Since $n^*(\hat{\theta})$ is C^1 and increasing, and $\hat{\theta}_t$ is a submartingale under H, growth requires the positive drift $(dn^*/d\hat{\theta}) \cdot \mathbb{E}[\Delta\hat{\theta}]$ to exceed the Jensen correction from the concavity of n^*. The correction is at most $\frac{1}{2}\|(n^*)''\|_\infty \cdot \mathrm{Var}(\hat{\theta}_{t+1})$, which is bounded. Define $\bar{\alpha}_{\text{growth}}(\hat{\theta}_0, \lambda) \equiv 1/[\lambda \cdot g(\underline{\mu}^S(0); \hat{\theta}_0)]$; for α below this bound, the snob threshold effect is weak enough that n^* is approximately linear in $\hat{\theta}$ and the drift dominates throughout Phase I. □

A.2 Proof of Proposition 3 (Impossibility Under Same-Sign Popularity Effects)

Proof. Agent i adopts iff $\mu_i \geq c_i - \gamma_i h(n)$. We show $\mathbb{E}[n_t \mid H] \geq \mathbb{E}[n_t \mid L]$ for all t, which implies $T^H \geq T^L$ under any duration measure monotonically increasing in expected adoption.

Key step: $n^*(\hat{\theta})$ increasing in $\hat{\theta}$ under either sign. The best-response mapping $\Phi(n; \hat{\theta}, \theta) = \sum_i \Pr(\mu_i \geq c_i - \gamma_i h(n) \mid \hat{\theta}, \theta)$ satisfies $\Phi(n; \hat{\theta}, H) \geq \Phi(n; \hat{\theta}, L)$ at every n by MLRP (higher θ shifts posteriors rightward, increasing the mass above any fixed threshold). When $\gamma_i \geq 0$ for all i (strategic complements), Φ is increasing in n; monotone comparative statics à la Milgrom and Roberts (1990) give $n_H^* \geq n_L^*$ at the largest (or smallest) equilibrium. When $\gamma_i \leq 0$ for all i (strategic substitutes), Φ is decreasing, the fixed point is unique, and $\Phi(\cdot; \hat{\theta}, H) \geq \Phi(\cdot; \hat{\theta}, L)$ directly gives $n_H^* \geq n_L^*$.[23]

Induction. At $t = 0$, $\hat{\theta}_0 = p$ is common, and $n_0^H \geq n_0^L$ by the key step. For the inductive step: $n_s^H \geq n_s^L$ for $s \leq t-1$ implies the observed signal $\tilde{n}_{t-1}$ is stochastically higher under H, so Bayesian updating yields $\hat{\theta}_t$ FOSD-higher under H (Lemma 4). The key step then gives $n_t^H \geq n_t^L$. Since $\mathbb{E}[n_t \mid H] \geq \mathbb{E}[n_t \mid L]$ at every t, $T^H \geq T^L$.

The reversal $T^H < T^L$ therefore requires mixed-sign γ_i: under H, agents with $\gamma_i > 0$ enter faster, raising n, which pushes agents with $\gamma_j < 0$ past their thresholds faster than improved beliefs can compensate. □

A.3 Belief Dominance Under Higher Quality

Lemma 4 (Stochastic Dominance of Beliefs)**.** *Starting from a common prior $\hat{\theta}_0 = p$, the distribution of public beliefs under $\theta = H$ first-order stochastically dominates the distribution under $\theta = L$ at every $t \geq 1$. That is, $\Pr(\hat{\theta}_t \geq x \mid H) \geq \Pr(\hat{\theta}_t \geq x \mid L)$ for all x and all $t \geq 1$, with strict inequality for interior x.*

Proof. By induction. At $t = 0$, $\hat{\theta}_0 = p$ is degenerate and identical under both states. Suppose $\hat{\theta}_{t-1} \mid H$ FOSD-dominates $\hat{\theta}_{t-1} \mid L$. The updated belief $\hat{\theta}_t = B(\hat{\theta}_{t-1}, \tilde{n}_{t-1})$ is increasing in both arguments. For any $\hat{\theta}_{t-1}$, the signal $\tilde{n}_{t-1} \sim \mathcal{N}(n_\theta^*(\hat{\theta}_{t-1}), \sigma_\varepsilon^2)$ has mean $n_H^*(\hat{\theta}_{t-1}) > n_L^*(\hat{\theta}_{t-1})$ (MLRP), so the conditional distribution of $\hat{\theta}_t \mid \hat{\theta}_{t-1}$ under H FOSD-dominates under L. Since B is also increasing in $\hat{\theta}_{t-1}$, higher realizations of $\hat{\theta}_{t-1}$ shift the conditional distribution rightward. Combining marginal FOSD of $\hat{\theta}_{t-1}$ (inductive hypothesis) with conditional FOSD of $\hat{\theta}_t \mid \hat{\theta}_{t-1}$ gives unconditional FOSD at t. □

This lemma underpins both the reversal proof (higher quality accelerates conformist entry) and the impossibility result (higher quality raises adoption at every t under uniform-sign preferences).

[23] In both cases, the result holds for any monotone equilibrium selection rule.

A.4 Proof of Proposition 2 (Quality-Duration Reversal)

Proof. Step 1. Let $T^\theta = \min\{t > t^*(\theta) : \mathbb{E}[n_t^S \mid \theta] < \epsilon\}$ denote niche-phase duration under quality state θ.

Step 2. Under threshold strategies, period-t adoption satisfies:

$$n_t(\hat{\theta}_t) = \lambda[1 - G(\alpha n_t + c_S; \hat{\theta}_t)] + (1 - \lambda)[1 - G(c_C - \beta n_t; \hat{\theta}_t)] \tag{15}$$

with $\partial n_t / \partial \hat{\theta}_t > 0$.

Step 3. By Lemma 4, the distribution of $\hat{\theta}_t$ under H FOSD-dominates under L for $t \geq 1$. In particular:

$$\mathbb{E}[\hat{\theta}_{t+1} - \hat{\theta}_t \mid \theta = H] > 0, \quad \mathbb{E}[\hat{\theta}_{t+1} - \hat{\theta}_t \mid \theta = L] < 0 \tag{16}$$

Step 4. Define $t^*(\theta) = \max\{t : \mathbb{E}[n_t^S] = \max_{s \geq 0} \mathbb{E}[n_s^S]\}$, the last period of peak expected snob adoption.

Step 5. Higher quality has two opposing effects on t^*. The *persistence effect*: under H, beliefs rise faster, keeping posteriors above $\underline{\mu}^S$ longer (t^* increases). The *acceleration effect*: under H, faster conformist entry raises n_t and hence $\underline{\mu}^S(n) = \alpha n + c_S$, triggering earlier snob exit (t^* decreases). Step 6 formalizes both effects as components of a single derivative $dn^S/d\hat{\theta}$: the persistence effect corresponds to the belief term and the acceleration effect to the threshold term.

Step 6. Snob adoption $n^S(\hat{\theta}) = \lambda[1 - G(\underline{\mu}^S(n^*(\hat{\theta})); \hat{\theta})]$. Differentiating:

$$\frac{dn^S}{d\hat{\theta}} = \lambda \left[\underbrace{-g^S \cdot \alpha \cdot \frac{dn^*}{d\hat{\theta}}}_{\text{threshold effect}} + \underbrace{\frac{\partial[1 - G(\underline{\mu}^S; \hat{\theta})]}{\partial \hat{\theta}}}_{\text{belief effect}} \right] \tag{17}$$

The sign of $dn^S/d\hat{\theta}$ depends on whether the threshold effect (crowding from conformist entry) dominates the belief effect (improved posteriors). Using the IFT expression for $dn^*/d\hat{\theta}$ from the equilibrium fixed point, the threshold effect dominates when $F(\alpha) \equiv (1 - \lambda)\, g^C(n^*(\alpha))\, (\alpha + \beta) > 1$.

Step 7: Existence of $\bar{\alpha}$. Define $F(\alpha; \hat{\theta}) \equiv (1 - \lambda)\, g(\underline{\mu}^C(n^*(\alpha, \hat{\theta})); \hat{\theta})\, (\alpha + \beta)$, where g^C

is the posterior density at the conformist threshold evaluated at the equilibrium. We need $F(\alpha;\hat{\theta}) > 1$ for *all* $\hat{\theta} \in (0,1)$ to ensure global monotonicity of n^S in $\hat{\theta}$ (Step 8). Define the uniform threshold:

$$\bar{\alpha} \equiv \inf\big\{\alpha > 0 : \inf_{\hat{\theta}\in(0,1)} F(\alpha;\hat{\theta}) \geq 1\big\}. \tag{18}$$

Under the Gaussian specification, $g(\mu;\hat{\theta})$ is continuous on $(0,1)\times(0,1)$ with $g > 0$ everywhere (Assumption 2), $n^*(\hat{\theta})$ is C^1 in $\hat{\theta}$ (IFT under Assumption 3), and $\underline{\mu}^C(n^*(\hat{\theta})) = c_C - \beta n^*(\hat{\theta}) \in (0,1)$ for all $\hat{\theta}$ at which the conformist margin is active.[24] Hence $F(\alpha;\hat{\theta})$ is continuous on $(0,\infty)\times[0,1]$, and $\inf_{\hat{\theta}} F(\alpha;\hat{\theta})$ is attained by compactness and is itself continuous in α. At $\alpha = 0$, $F(0;\hat{\theta}) = (1-\lambda)\, g^C\, \beta < 1$ for all $\hat{\theta}$ by Assumption 3. As $\alpha \to \infty$, the factor $(\alpha+\beta) \to \infty$ while g^C remains bounded below by a positive constant (full support and continuity on a compact domain), so $\inf_{\hat{\theta}} F(\alpha;\hat{\theta}) \to \infty$. By the intermediate value theorem, $\bar{\alpha} \in (0,\infty)$ and $\inf_{\hat{\theta}} F(\bar{\alpha};\hat{\theta}) = 1$.

Step 8: From within-period derivative to dynamic duration ordering. For any $\alpha > \bar{\alpha}$, we have $\inf_{\hat{\theta}} F(\alpha;\hat{\theta}) > 1$ by definition, so $F(\alpha;\hat{\theta}) > 1$ for every $\hat{\theta} \in (0,1)$. This means $dn^S/d\hat{\theta} < 0$ at every $\hat{\theta}$: snob adoption $n^S(\hat{\theta})$ is globally decreasing in the public belief.[25]

By Lemma 4, the distribution of $\hat{\theta}_t$ under H FOSD-dominates under L. Since $n^S(\hat{\theta})$ is decreasing, the integral characterization of FOSD gives $\mathbb{E}[n_t^S \mid H] \leq \mathbb{E}[n_t^S \mid L]$ for all $t \geq 1$. Under H, expected snob adoption converges to 0; under L, snob adoption declines only gradually since the conformist surge never materializes. The ϵ-stopping time therefore satisfies $T^H < T^L$.

Step 9: Comparative statics. Implicit differentiation of $F(\bar{\alpha};\lambda,\beta) = 1$ yields $\partial\bar{\alpha}/\partial\lambda < 0$ and $\partial\bar{\alpha}/\partial\beta > 0$.[26] □

[24]When $\beta \geq c_C$, the conformist threshold drops below zero at high enough $\hat{\theta}$ (all conformists adopt). At such beliefs, $g^C = 0$ and $F(\alpha;\hat{\theta}) = 0$, so the infimum of F is determined by beliefs where the conformist margin is active and $g^C > 0$. The intermediate value argument for $\bar{\alpha}$ is unaffected. As $\hat{\theta} \to 0$, adoption $n^*(\hat{\theta}) \to 0$ and the conformist threshold $c_C - \beta n^* \to c_C < 1$, so g^C remains bounded away from zero by full support.

[25]For non-Gaussian signals, the sufficient condition for global monotonicity is $\inf_{\hat{\theta}} g(\underline{\mu}^C(n^*(\hat{\theta}));\hat{\theta}) > 0$, which holds under full support. The argument extends to any signal structure satisfying Assumption 2.

[26]Under the approximation $g(\underline{\mu}^S(n^*)) \approx g(\underline{\mu}^C(n^*))$, the threshold has the closed form $\bar{\alpha}(\lambda) = \beta(1-\lambda)/\lambda$, with $\partial\bar{\alpha}/\partial\lambda = -\beta/\lambda^2$ and $\partial\bar{\alpha}/\partial\beta = (1-\lambda)/\lambda$. The exact threshold inherits the same signs.

A.5 Proof of Corollary 2 (Asymmetric Quality Elasticity)

Proof. Type-specific adoption satisfies $n_t^S = \lambda[1 - G(\alpha n^* + c_S; \hat{\theta}_t)]$ and $n_t^C = (1-\lambda)[1 - G(c_C - \beta n^*; \hat{\theta}_t)]$. For conformists:

$$\frac{dn_t^C}{d\hat{\theta}} = (1-\lambda)\left[\frac{\partial[1 - G(\underline{\mu}^C; \hat{\theta})]}{\partial\hat{\theta}} + \beta\, g(\underline{\mu}^C)\,\frac{dn^*}{d\hat{\theta}}\right] \tag{19}$$

Both terms are positive (higher beliefs shift posteriors rightward; more adoption lowers the conformist threshold), so $dn_t^C/d\hat{\theta} > 0$ unconditionally. For snobs, the indirect effect enters with a negative sign:

$$\frac{dn_t^S}{d\hat{\theta}} < 0 \iff \alpha > \frac{\partial_{\hat{\theta}}[1 - G(\underline{\mu}^S; \hat{\theta})]}{g(\underline{\mu}^S)\, dn^*/d\hat{\theta}} \tag{20}$$

The right-hand side is the ratio of the direct belief effect to the indirect crowding effect. Using the IFT expression for $dn^*/d\hat{\theta}$ from the equilibrium fixed point, the condition $dn^S/d\hat{\theta} < 0$ reduces to $F(\alpha) \equiv (1-\lambda)\, g(\underline{\mu}^C(n^*); \hat{\theta})\,(\alpha + \beta) > 1$ – the same condition that defines the reversal regime ($\alpha > \bar{\alpha}$). The quality-elasticity threshold thus coincides with $\bar{\alpha}(\lambda)$. □

A.6 Proof of Proposition 4 (Optimal Market Composition)

Proof. Part (i): Interior optimum. Social welfare is $W(\lambda) = \lambda \cdot \mathbb{E}[U^S(\lambda)] + (1-\lambda) \cdot \mathbb{E}[U^C(\lambda)]$. As $\lambda \to 0$, snob adoption $n_0^S = \lambda[1 - G(c_S)] \to 0$, the public signal is uninformative, beliefs stall near the prior, and the welfare contribution from the snob-conformist interaction vanishes. As $\lambda \to 1$, conformist surplus $(1-\lambda)\mathbb{E}[U^C] \to 0$ and the three-phase lifecycle degenerates. For intermediate λ, snobs generate informative public signals that improve conformist quality matching, so $W(\lambda) > \max\{W(0), W(1)\}$ at some interior point. Since W is continuous on $[0,1]$, the maximum occurs at $\lambda^{**} \in (0,1)$.

To establish uniqueness, we verify quasi-concavity. The welfare derivative is $W'(\lambda) = [\mathbb{E}[U^S] - \mathbb{E}[U^C]] + \lambda(d\mathbb{E}[U^S]/d\lambda) + (1-\lambda)(d\mathbb{E}[U^C]/d\lambda)$. At small λ, the third term dominates: $(1-\lambda)\, d\mathbb{E}[U^C]/d\lambda > 0$ because the marginal snob generates highly valuable informational externalities for the large conformist population. At large λ, the first bracket $\mathbb{E}[U^S] - \mathbb{E}[U^C]$ becomes negative (snobs are crowded and earn low surplus) and the second term $\lambda\, d\mathbb{E}[U^S]/d\lambda < 0$ dominates (additional crowding on the large snob population). Hence W'

transitions from positive to negative. Since the positive component (conformist informational benefit) is concave in λ and the negative component (snob crowding cost) is convex, $W'(\lambda)$ is strictly decreasing wherever $W'(\lambda) = 0$, implying the zero is unique and W is strictly quasi-concave on $[0, 1]$.

At the optimum, $d\mathbb{E}[U^S]/d\lambda < 0$ (additional snobs intensify crowding) and $d\mathbb{E}[U^C]/d\lambda > 0$ (additional snobs generate informative signals and raise Phase I adoption, lowering the conformist threshold).[27] The FOC $\mathbb{E}[U^S]-\mathbb{E}[U^C]+\lambda^{**}(d\mathbb{E}[U^S]/d\lambda)+(1-\lambda^{**})(d\mathbb{E}[U^C]/d\lambda) = 0$ balances these opposing forces, giving $\lambda^S < \lambda^{**} < \lambda^C$.

Part (ii): Comparative statics. By the implicit function theorem on the FOC $W'(\lambda^{**}) = 0$: $d\lambda^{**}/d\alpha = -(\partial^2 W/\partial\lambda\partial\alpha)/(d^2W/d\lambda^2)|_{\lambda^{**}}$. The denominator is negative (SOC). The cross-partial $\partial^2 W/\partial\lambda\partial\alpha < 0$: higher α makes the marginal snob's crowding cost larger. Hence $d\lambda^{**}/d\alpha < 0$. For σ: $\partial^2 W/\partial\lambda\partial\sigma > 0$ (noisier private signals raise the marginal snob's informational value), so $d\lambda^{**}/d\sigma > 0$. □

[27] Formally, $d\mathbb{E}[U^C]/d\lambda = \sum_t \delta^t[(\partial\mathbb{E}[U_t^C]/\partial n^*)(dn^*/d\lambda) + (\partial\mathbb{E}[U_t^C]/\partial\hat{\theta}_t)(d\hat{\theta}_t/d\lambda)]$. The within-period terms are positive; the lifecycle compression term is negative but weighted by late periods when the marginal snob's contribution is small.

Online Appendix

Snobs and Conformists: Dynamic Adoption with Opposing Social Preferences

Tommaso Bondi

Cornell University

This Online Appendix collects extensions to the baseline model and proofs of secondary results. Proofs of the main propositions appear in the paper's Appendix.

OA.1 Extensions

OA.1.1 Composition-Dependent Preferences

The baseline model assumes snobs dislike popularity measured by *how many* adopt. But in many markets – fashion, nightlife, social platforms – what matters is *who* adopts. Empirical evidence supports this distinction: Berger & Heath (2007, 2008) show that consumers abandon products specifically when *outgroup* members adopt, not merely when adoption increases; Han et al. (2010) document that consumers strategically manage brand prominence depending on the perceived audience; and Amaldoss & Jain (2005) show that uniqueness-seeking generates demand patterns that depend on the *type* of other adopters, not just their number. This section extends the model to allow snob utility to depend on adopter composition, showing that all core results survive with quantitative amplification.

Decompose total adoption into components $n_t = n_t^S + n_t^C$ and define coolness as $C_t \equiv n_t^S - \xi n_t^C$, where $\xi > 0$ captures how much conformist adoption dilutes cachet. Snob utility becomes $U^S = v(\theta) - \alpha n_t + \eta C_t$, where $\eta \geq 0$ captures how much snobs value coolness per se. Conformist utility is unchanged. We restrict $\eta < \alpha$ to ensure that snobs' net response to own-type adoption remains negative ($\partial U^S / \partial n^S = \eta - \alpha < 0$), preserving the contraction property that guarantees equilibrium uniqueness.[28]

Proposition 9 (Coolness-Amplified Dynamics). *With composition-dependent coolness ($\eta > 0$):*

(i) Each conformist reduces snob utility by $\alpha + \eta\xi$ rather than α alone.

(ii) $T_\eta < T_0$; to first order for small $\eta\xi$, $T_\eta / T_0 \approx \alpha/(\alpha + \eta\xi)$.

[28] When $\eta > \alpha$, snobs exhibit within-type strategic complementarities: more snob adoption makes the product cooler, encouraging further snob adoption. This can generate multiple equilibria and requires a separate uniqueness analysis.

(iii) C_t *peaks in Phase I and becomes negative in Phase II.*

(iv) A cool equilibrium ($n^S > 0$, $n^C = 0$ *permanently) exists when* c_C *is large,* β *is small, or* σ *is large.*

The proposition establishes that accounting for adopter composition *amplifies* the core mechanism rather than altering it. If $\eta\xi \approx \alpha$, conformist entry is twice as damaging as in the baseline, potentially halving lifecycle duration. The quality-duration reversal, three-phase structure, and impossibility result all carry through; only the quantitative magnitudes change.

Cool equilibria explain markets that sustain exclusivity despite demand pressure: vinyl records (inconvenience screens casual listeners), literary fiction (difficulty screens casual readers), and Berghain (door policy screens mainstream attendees). In each case, structural features raise conformist entry costs, sustaining positive coolness indefinitely.

Figure 2 illustrates with parameters $\eta = 2.0$, $\xi = 1.0$ (snob-on-snob crowding $\alpha - \eta = -0.5$; crowding from conformists $\alpha + \eta\xi = 3.5$). The three-phase lifecycle and reversal are preserved. Snob peaks are higher (negative self-crowding), snob exit is steeper (the "coolness cascade"), and conformist peaks are higher. Under low quality, the model predicts a "return to niche" as conformists decay faster than snobs.

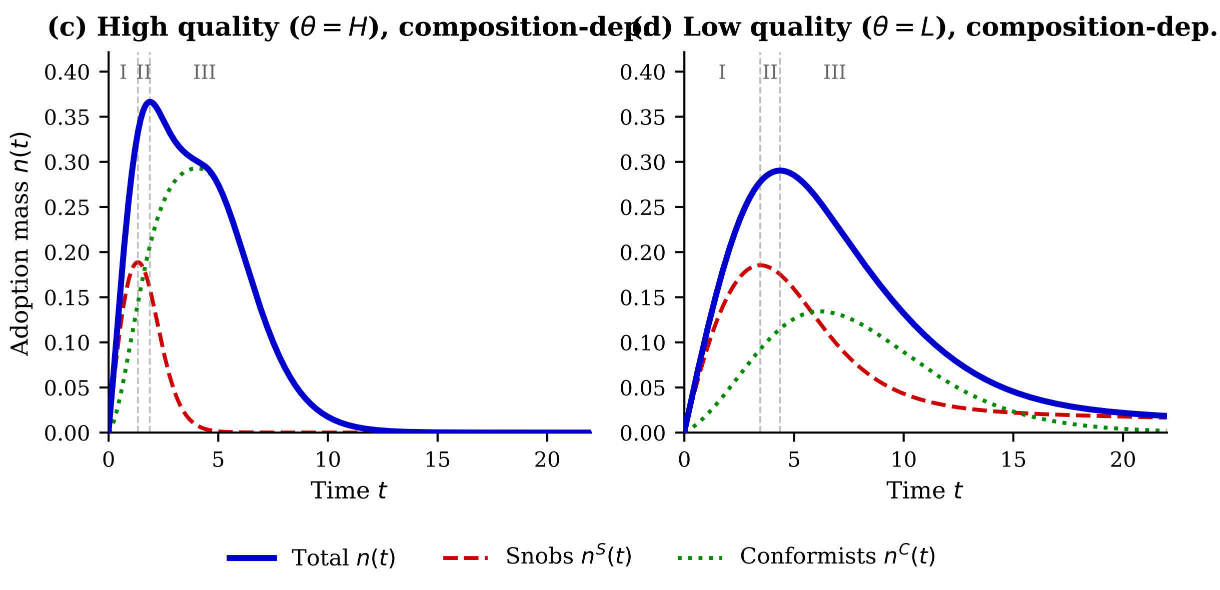

Figure 2: Lifecycle dynamics under composition-dependent preferences ($\eta = 2.0$, $\xi = 1.0$). Compared with Figure 1: snob peaks are 36–63% higher, snob half-life is 50–59% shorter (the coolness cascade), and conformist peaks are higher. All other parameters as in Figure 1.

OA.1.2 Robustness to Functional Form

The baseline assumes linear social preferences: $U^S = v(\theta) - \alpha n$ and $U^C = v(\theta) + \beta n$. This appendix shows that the main results are robust to convex specifications and documents how non-linearity affects cycle shape.

Setup. Consider convex crowding costs $U^S = v(\theta) - \alpha n^\rho$ and convex bandwagon benefits $U^C = v(\theta) + \beta n^\zeta$, where $\rho, \zeta > 1$. The baseline is the special case $\rho = \zeta = 1$. Threshold strategies still obtain, with modified cutoffs $\underline{\mu}^S(n) = \alpha n^\rho + c_S$ and $\underline{\mu}^C(n) = c_C - \beta n^\zeta$. At the peak, both marginal types are indifferent; for the symmetric quadratic case $\rho = \zeta = 2$ with $c_S = 0$, peak adoption is $n^* = \sqrt{c_C/(\alpha+\beta)}$.

Robustness of core results. The three-phase lifecycle persists because the opposing monotonicity of thresholds is unchanged: $\partial\underline{\mu}^S/\partial n = \alpha\rho n^{\rho-1} > 0$ while $\partial\underline{\mu}^C/\partial n = -\beta\zeta n^{\zeta-1} < 0$. The quality-duration reversal also survives: for the quadratic case with Beta-distributed posteriors[29] – the reversal condition has an exact closed form,

$$\alpha > \frac{\beta^2(1-\lambda)(1+\lambda)}{4c\lambda^2},$$

which is finite and well-defined without any density approximation (here $c \equiv c_C$, the conformist outside option). For $\lambda = 1/3$, $\beta = 1$, $c_C = 1$, this gives $\bar{\alpha}_Q = 2$. The impossibility result (Proposition 3) carries through unchanged because it depends only on the sign of γ_i, not on functional form.

Cycle shape. With $\rho > 1$, the marginal crowding cost $-\alpha\rho n^{\rho-1}$ is negligible at low n but accelerates sharply near the peak, producing right-skewed curves: gradual ascent and rapid descent. This matches the empirical profile of many fashion and cultural products, where trends build gradually then crash within weeks.

OA.2 Additional Proofs

OA.2.1 Proof of Lemma 1 (Threshold Strategies)

Proof. Under flow adoption, continuation values cancel (they are identical whether the agent adopts or waits), so the adoption decision reduces to comparing current payoffs. A snob adopts iff $\mu - \alpha n \geq c_S$, i.e., $\mu \geq \alpha n + c_S \equiv \underline{\mu}^S(n)$. A conformist adopts iff $\mu + \beta n \geq c_C$, i.e., $\mu \geq c_C - \beta n \equiv \underline{\mu}^C(n)$.[30] The thresholds satisfy $d\underline{\mu}^S/dn = \alpha > 0$ and $d\underline{\mu}^C/dn = -\beta < 0$, and cross at $n^\dagger = (c_C - c_S)/(\alpha+\beta) \in (0,1)$ by Assumption 1. □

[29] The Beta distribution on $(0,1)$, parameterized by shape parameters varying with $\hat{\theta}$ and σ, provides full support, MLRP, and closed-form densities at the thresholds – properties shared with the Gaussian specification but admitting exact integration.

[30] With irreversible adoption, the threshold includes an option-value premium $\omega(\hat{\theta}) > 0$, shifting the intercept without changing the slope; see Chamley (2004).

OA.2.2 Derivation: Opposing Responses to Adoption

Proof. Snob adoption $n^S(n) = \lambda[1-G(\alpha n+c_S)]$ satisfies $dn^S/dn = -\lambda\alpha g(\alpha n+c_S) < 0$ since $g > 0$ (Assumption 2). As $n \to \bar{n}^S \equiv (1-c_S)/\alpha$, the threshold $\alpha n + c_S \to 1$ and $n^S \to 0$. Conformist adoption $n^C(n) = (1-\lambda)[1-G(c_C-\beta n)]$ satisfies $dn^C/dn = (1-\lambda)\beta g(c_C-\beta n) > 0$. Both derivatives hold simultaneously for $n \in (0, \bar{n}^S)$, establishing part (iii). □

OA.2.3 Proof of Lemma 2 (Equilibrium Existence)

Proof. Fix $\hat{\theta} \in (0,1)$. The best-response mapping

$$\Phi(n;\hat{\theta}) = \lambda[1 - G(\alpha n + c_S;\hat{\theta})] + (1-\lambda)[1 - G(c_C - \beta n;\hat{\theta})] \tag{21}$$

is continuous on $[0,1]$ (thresholds are linear in n; G is C^1 by Assumption 2). At $n = 0$: $\Phi(0) \geq \lambda[1 - G(c_S)] > 0$ since $c_S < 1$. At $n = 1$: if $\alpha + c_S \geq 1$, the snob term is zero; if $\alpha + c_S < 1$, it is $\lambda[1 - G(\alpha + c_S)] < \lambda$. In either case $\Phi(1) < 1$.[31] By Brouwer's theorem, $\exists\, n^* \in (0,1)$ with $\Phi(n^*) = n^*$. □

Proof of MPE Existence (Section 5)

Proof. The public belief $\hat{\theta}_t$ is the sole aggregate state. Since continuation values cancel under flow adoption, optimal policies are myopic (Lemma 1). Belief consistency requires $\hat{\theta}_{t+1} = B(\hat{\theta}_t, \tilde{n}_t)$ where $\tilde{n}_t \sim \mathcal{N}(n^*_\theta(\hat{\theta}_t), \sigma^2_\varepsilon)$ and $n^*_\theta(\hat{\theta}_t)$ is realized adoption under true quality θ. Full-support Gaussian noise ensures Bayes' rule applies for all observations, pinning down off-path beliefs uniquely. The Bellman operator for a snob,

$$(\mathcal{T}V)(\mu,\hat{\theta}) = \max\Big\{\mu - \alpha n^*(\hat{\theta}) + \delta \int V(\mu',\hat{\theta}')\, dF(\mu',\hat{\theta}' \mid \hat{\theta}),\ c_S + \delta \int V(\mu',\hat{\theta}')\, dF(\mu',\hat{\theta}' \mid \hat{\theta})\Big\} \tag{22}$$

satisfies Blackwell's monotonicity and discounting conditions on the Banach space of bounded measurable functions with the sup-norm, hence is a contraction with modulus δ. The Banach Fixed Point Theorem gives a unique value function V^*. The MPE combines the static fixed point $n^*(\hat{\theta})$, myopic optimality, and consistent belief transitions. □

Proof of Uniqueness (Assumption 3)

Proof. From (21), $\Phi'(n) = -\lambda\alpha g(\underline{\mu}^S(n)) + (1-\lambda)\beta g(\underline{\mu}^C(n))$. If $\kappa \equiv \max_n |\Phi'(n)| < 1$, then Φ is a contraction and the fixed point is unique (Banach). When $\kappa > 1$, $h(n) \equiv \Phi(n) - n$

[31] The objective distribution $G(\mu;\hat{\theta},\theta)$ governs belief updating but not the within-period fixed-point argument, which uses the subjective $G(\mu;\hat{\theta})$.

satisfies $h(0) > 0$ and $h(1) < 0$; if $\Phi' > 1$ on a sufficiently wide interval, h crosses zero three times (IVT), giving two stable equilibria ($\Phi' < 1$) and one unstable ($\Phi' > 1$).[32] □

Derivation of Within-Period Comparative Statics

Proof. At a stable equilibrium ($\Phi'(n^*) < 1$), the IFT gives $\partial n^*/\partial\alpha = -\lambda g(\underline{\mu}^S)n^*/(1-\Phi') < 0$ and $\partial n^*/\partial\beta = (1-\lambda)g(\underline{\mu}^C)n^*/(1-\Phi') > 0$. Dynamically, higher β triggers earlier conformist entry, which accelerates snob exit, so the lifecycle peak satisfies $dn^*_{\text{peak}}/d\beta < 0$.[33] □

Proof of Proposition 9 (Coolness Dynamics)

Proof. We analyze the extended model with composition-dependent "coolness."

Step 1: Setup. Define coolness as $C_t = n_t^S - \xi n_t^C$, where $\xi > 0$ measures how much conformist presence dilutes coolness. Snob utility with coolness dependence is:

$$U^S = v(\theta) - \alpha n_t + \eta C_t = v(\theta) - \alpha(n_t^S + n_t^C) + \eta(n_t^S - \xi n_t^C) \tag{23}$$

Rearranging:

$$U^S = v(\theta) + (\eta - \alpha)n_t^S - (\alpha + \eta\xi)n_t^C \tag{24}$$

Step 2: Marginal effect of conformist entry. Differentiating:

$$\frac{\partial U^S}{\partial n^C} = -(\alpha + \eta\xi) < -\alpha \tag{25}$$

Each conformist reduces snob utility by $\alpha + \eta\xi$, exceeding the baseline crowding cost α by $\eta\xi$. This establishes part (i): snob exit accelerates.

Step 3: Lifecycle compression (part ii). The snob threshold becomes $\underline{\mu}^S(n^S, n^C) = (\alpha - \eta)n^S + (\alpha + \eta\xi)n^C + c_S$. Since conformist entry has amplified effect, the threshold is reached sooner, compressing the lifecycle.

Step 4: Coolness overshooting (part iii). The change in coolness is:

$$\Delta C_t = \Delta n_t^S - \xi \Delta n_t^C \tag{26}$$

In Phase I, $\Delta n_t^C \approx 0$, so $\Delta C_t \approx \Delta n_t^S > 0$ and coolness rises. When conformists enter (Phase II), $\Delta n_t^C > 0$. Coolness peaks when $\Delta n_t^S = \xi \Delta n_t^C$ and subsequently declines. If ξ is large, coolness becomes negative ($C_t < 0$) when $n_t^C > n_t^S/\xi$.

Step 5: Cool equilibria (part iv). A cool equilibrium requires snobs to adopt ($n^S > 0$) while conformists are deterred ($n^C \approx 0$). The snob-only fixed point n^{**} solves $\lambda[1 - G((\alpha -$

[32] Three crossings are standard in models with locally strong complementarities; see Morris & Shin (2003).

[33] Near the threshold crossing $n^\dagger = (c_C - c_S)/(\alpha + \beta)$, the comparative statics $\partial n^\dagger/\partial\alpha < 0$ and $\partial n^\dagger/\partial\beta < 0$ confirm the IFT results.

$\eta)n^{**} + c_S)] = n^{**}$, which exists and is unique by the intermediate value theorem. The conformist deterrence condition requires that the conformist threshold $c_C - \beta n^{**}$ remains high enough that $(1-\lambda)[1 - G(c_C - \beta n^{**}; \hat{\theta}_\infty)]$ is negligible, where $\hat{\theta}_\infty$ is the limiting belief under snob-only adoption. This holds when c_C is large (high conformist outside option), β is small (weak bandwagon preference), or σ is large (dispersed posteriors, so $\hat{\theta}_\infty$ stays interior and $c_C - \beta n^{**}$ exceeds most posteriors). □

Proof of Proposition 5 (Optimal Visibility)

Proof. We first establish the equilibrium response to visibility, then prove each part.

Equilibrium response. Under visibility φ, the noise variance is $\sigma_\varepsilon^2(\varphi) = \sigma_0^2(1-\varphi)/\varphi$, decreasing in φ. Lower noise makes the observed signal $\tilde{n}_{t-1}$ more informative about actual adoption n_{t-1}, increasing the responsiveness of the likelihood ratio $f(\tilde{n}_{t-1} \mid H)/f(\tilde{n}_{t-1} \mid L)$ to differences in adoption across quality states. This accelerates belief convergence toward θ, so higher φ raises $\hat{\theta}_t$ (on average, under H) at every $t \geq 1$. Since $n^*(\hat{\theta})$ is increasing in $\hat{\theta}$ (Lemma 3), the equilibrium adoption path $\{n_t(\varphi)\}$ is increasing in φ during the learning phase. Under Assumption 3, $n^*(\hat{\theta}_t(\varphi))$ is C^1 in φ through the belief channel: $\partial n_t^*/\partial\varphi = (dn^*/d\hat{\theta}) \cdot (\partial\hat{\theta}_t/\partial\varphi)$, where both factors are positive.

Part (i): Interior optimum. Social welfare is $W(\varphi) = \sum_{t=0}^{\infty} \delta^t[\lambda \cdot \mathbb{E}[U_t^S(\varphi)] + (1-\lambda) \cdot \mathbb{E}[U_t^C(\varphi)]]$. The welfare derivative decomposes as:

$$W'(\varphi) = \underbrace{\sum_{t=1}^{\infty} \delta^t \left[\lambda \frac{\partial \mathbb{E}[U_t^S]}{\partial \hat{\theta}_t} + (1-\lambda) \frac{\partial \mathbb{E}[U_t^C]}{\partial \hat{\theta}_t}\right] \cdot \frac{\partial \mathbb{E}[\hat{\theta}_t]}{\partial \varphi}}_{\text{quality assessment (positive)}} + \underbrace{\sum_{t=1}^{\infty} \delta^t \lambda\,(-\alpha)\,\frac{\partial \mathbb{E}[n_t^C]}{\partial \varphi}}_{\text{composition cost (negative)}} \tag{27}$$

Both terms operate through the belief channel ($\partial\mathbb{E}[\hat{\theta}_t]/\partial\varphi > 0$ under H). The assessment term is positive (both types benefit from beliefs closer to θ under H), weighted by the full population; the composition cost is negative but weighted by $\lambda < 1$.

At $\varphi \to 0^+$, both terms diverge (the marginal informational gain from reducing infinite noise is infinite), but the assessment term dominates by the factor $1/\lambda > 1$, so $W'(0^+) > 0$.[34] At $\varphi = 1$, marginally reducing φ delays conformist entry. The welfare gain is the snob pioneer-rent surplus over the additional exclusive-phase periods. The welfare cost is second-order: the marginal conformist earns $U^C = c_C$ by indifference, so the first-order welfare loss from shifting this agent out is zero. When $\alpha > \bar{\alpha}(\lambda)$, conformist entry triggers snob exit, so delaying it preserves substantial snob surplus. Hence $W'(1) < 0$.

By continuity and the intermediate value theorem, $\varphi_W^* \in (0,1)$.

[34]Under the Gaussian specification, the marginal welfare gain near $\varphi = 0$ is proportional to the Fisher information $(n_H^* - n_L^*)^2/\sigma_\varepsilon^4$. The assessment term scales with weight 1 and the composition cost with weight $\lambda\alpha$, bounded under the contraction condition.

Part (ii): Comparative statics of welfare gap. The welfare cost of full transparency $\Delta W \equiv W(\varphi^*_W) - W(1)$ satisfies $\partial \Delta W/\partial\alpha > 0$ because $\partial^2 W/\partial\alpha\partial\varphi < 0$ in the reversal regime (the marginal welfare cost of visibility increases with α), and $\partial\Delta W/\partial\lambda < 0$ because more snobs reduce the conformist mass that drives the acceleration externality.

Part (iii): First-order condition. The SOC $W''(\varphi^*_W) < 0$ holds because the assessment gain exhibits diminishing returns in signal precision while the acceleration cost is convex in φ.[35] By the IFT: $\partial\varphi^*_W/\partial\sigma = -(\partial^2 W/\partial\varphi\partial\sigma)/(\partial^2 W/\partial\varphi^2) > 0$ (noisier private signals raise the marginal value of public data) and $\partial\varphi^*_W/\partial\alpha < 0$ (higher α raises the marginal acceleration cost). □

Proof of Proposition 6 (Profit-Maximizing Visibility)

Proof. The platform maximizes $\Pi(\varphi;\delta_P) = \mathbb{E}_\theta\left[\sum_{t=0}^\infty \delta_P^t r(n_t(\varphi,\theta))\right]$ with $r' > 0$.

Part (i): Existence of $\bar{\delta}_P$. Define $M(\varphi,\delta_P) \equiv \partial\Pi/\partial\varphi = \mathbb{E}_\theta\left[\sum_{t=0}^\infty \delta_P^t r'(n_t)(\partial n_t/\partial\varphi)\right]$. This is well-defined because the sum converges (adoption is bounded and $\delta_P < 1$) and $\partial n_t/\partial\varphi$ exists by the IFT.

At $\delta_P \to 0$, only the first periods matter. Since n_0 does not depend on φ (no signal has arrived), $M \approx \delta_P r'(n_1)\mathbb{E}[\partial n_1/\partial\varphi]$. Under $\theta = H$, higher φ raises adoption (better signals reveal quality); under $\theta = L$, it lowers adoption. At moderate priors, the H-state effect dominates because adoption is higher under H and $r' > 0$ amplifies the contribution of higher-adoption states. Hence $M > 0$ for δ_P small: the platform benefits from increasing visibility.

At $\delta_P \to 1$, M sums $r'(n_t)(\partial n_t/\partial\varphi)$ without discounting. In the reversal regime, $\partial n_t/\partial\varphi > 0$ during the conformist surge (more visibility raises short-run adoption) but the lifecycle compresses: the number of positive-adoption periods shrinks. For δ_P close to 1 the lost late-period revenue dominates the early-period gain, so $M(\varphi^*_W,\delta_P) < 0$ and $\varphi^*_\Pi < \varphi^*_W$.

Since $M(\varphi^*_W,\cdot)$ is continuous in δ_P, positive near 0, and negative near 1, the IVT gives $\bar{\delta}_P \in (0,1)$. Uniqueness follows from $\partial M/\partial\delta_P < 0$: the derivative $\partial M/\partial\delta_P = \sum_{t=1}^\infty t\,\delta_P^{t-1}\,r'(n_t)(\partial n_t/\partial\varphi)$ reweights M by t, upweighting late (negative) terms, so at $\bar{\delta}_P$ where $M = 0$ the reweighted sum is strictly negative.

Part (ii): Comparative statics. By the IFT on $M(\varphi^*_W,\bar{\delta}_P;\alpha,\lambda) = 0$: $\partial\bar{\delta}_P/\partial\alpha = -(\partial M/\partial\alpha)/(\partial M/\partial\delta_P)$. Higher α makes the late-period terms in M more negative ($\partial M/\partial\alpha < 0$), so $\partial\bar{\delta}_P/\partial\alpha < 0$. Similarly, $\partial\bar{\delta}_P/\partial\lambda < 0$. □

Proof of Proposition 7 (Quality and Optimal Visibility)

Proof. We show $\varphi^*(H) < \varphi^*(L)$ by comparing FOCs. For each θ, the SOC gives $\partial^2 W/\partial\varphi^2 < 0$. Higher quality increases the acceleration cost of visibility (better products attract con-

[35]Diminishing returns follow from the concavity of mutual information in signal precision for Gaussian channels; the acceleration cost is convex because the belief-convergence rate is convex in signal precision.

formists faster) while the assessment benefit is bounded by $v(H) - v(L)$. In the reversal regime, $\partial W(\varphi, H)/\partial\varphi < \partial W(\varphi, L)/\partial\varphi$ for any fixed φ, so $\varphi_W^*(H) < \varphi_W^*(L)$. The same argument applies to profit: $r' > 0$ and Lemma 4 give $\partial\Pi(\varphi, H)/\partial\varphi < \partial\Pi(\varphi, L)/\partial\varphi$, so $\varphi_\Pi^*(H) < \varphi_\Pi^*(L)$. For the signal-contingent extension, $\Pi(\varphi; z) = q(z)\Pi(\varphi, H) + (1 - q(z))\Pi(\varphi, L)$ with $q(z) = \Pr(\theta = H \mid z)$; the weighted FOC implies $\varphi^*(z)$ is decreasing in $q(z)$. □

Proof of Proposition 8 (Optimal Dynamic Visibility)

Proof. The platform chooses $(\varphi_E, \varphi_L) \in [0, 1]^2$ to maximize

$$W(\varphi_E, \varphi_L) = \sum_{t=0}^{t_1-1} \delta^t \big[\lambda \mathbb{E}[U_t^S(\varphi_E)] + (1-\lambda)\mathbb{E}[U_t^C(\varphi_E)]\big] + \sum_{t=t_1}^{\infty} \delta^t \big[\lambda \mathbb{E}[U_t^S(\varphi_L)] + (1-\lambda)\mathbb{E}[U_t^C(\varphi_L)]\big] \tag{28}$$

where the transition date t_1 depends on φ_E (lower early visibility delays it).

Part (i). During incubation ($t < t_1$), higher φ_E improves quality assessment (positive) but advances t_1, compressing the niche phase (negative). In the reversal regime, the composition cost dominates, so $\varphi_E^* < \varphi_W^*$. During scaling ($t \geq t_1$), snobs are declining and the composition cost is second-order, so $\varphi_L^* > \varphi_W^*$. Combining: $\varphi_E^* < \varphi_W^* < \varphi_L^*$.

Part (ii). $\Delta W \equiv W(\varphi_E^*, \varphi_L^*) - W(\varphi_W^*, \varphi_W^*)$ is increasing in the heterogeneity of marginal welfare from visibility across phases. Higher α widens this gap ($\partial\Delta W/\partial\alpha > 0$); lower λ narrows it ($\partial\Delta W/\partial\lambda < 0$). □