Snobs and Conformists: Platform Design and Product Lifecycles

Tommaso Bondi

Cornell University

February 2026

Abstract

We develop a dynamic model of product adoption with two consumer types: *snobs*, who value exclusivity, and *conformists*, who value popularity. Their interaction generates an endogenous boom-bust lifecycle and a quality-duration reversal: higher-quality products can have *shorter* lifecycles, because superior quality accelerates conformist entry and triggers earlier snob exit. We then study two supply-side design problems. A welfare-maximizing platform should restrict the visibility of popularity information to delay conformist entry, with optimal visibility decreasing in product quality; whether a profit-maximizing platform over- or under-reveals depends on its patience. A seller optimally manages launch through scarcity, snob-targeted advertising, and type-based pricing that discounts early adopters and charges latecomers a premium; broad reach can backfire by compressing the lifecycle. The model's testable predictions – that rapid growth shortens lifecycles, that early-adopter exit precedes peak adoption, and that the direction of visibility distortion depends on platform patience – distinguish the framework from standard diffusion, social learning, and network-effects models.

1 Introduction

A niche creator builds a following on a streaming platform. For months the audience is small, engagement is high, and part of the appeal is precisely that the creator feels undiscovered. Then a burst of attention arrives: an algorithmic recommendation, a repost by a larger account, a surge in followers. New followers interpret the growing audience as proof of quality and subscribe. The original fans react in the opposite way: the creator is no longer "theirs," the comment section feels different, and the very fact that the account went mainstream changes what it delivers. Popularity validates for some and dilutes for others.

This tension appears whenever consumption has an *identity* or *positional* component, so that adoption by others changes what the product *means*. 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 cultural products, venues, and online communities, where "mainstreaming" simultaneously creates *validation* for some and *dilution* for others. Interestingly, these are settings where intermediaries actively broadcast popularity through design choices – rankings, trending modules, "most popular" badges, follower counts, and algorithmic amplification.

Motivated by these examples, we develop a dynamic adoption model with a central assumption: there are two types of consumers with opposing preferences over popularity. Borrowing the classic terms from Leibenstein (1950), one type (*snobs*) experiences disutility from broader adoption; the other (*conformists*) experiences utility from it. Consumers also learn about product quality over time, so adoption is informative in addition to any direct taste effect. This single ingredient – opposing responses to popularity – is enough to generate a rich set of results on both the demand and the supply side.

On the *demand side*, the equilibrium adoption path naturally exhibits a lifecycle with three phases. In a *discovery* phase, adoption is driven by snobs: they are willing to try the product before it is widely consumed, and their early purchases generate information about quality. As learning progresses and adoption becomes sufficiently validating, conformists enter, producing a *surge* in demand. Finally, the same surge triggers snob exit: as adoption rises, crowding costs dominate, and the product enters a *decline* phase in which adoption falls. This boom-bust pattern is not imposed through 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 accelerates learning, which speeds up the conformist takeoff. But faster conformist entry accelerates the crowding externality that makes the product unattractive to snobs. As a result, higher-quality products can have *shorter* culturally active lifecycles: the period in which the product is attractive to early adopters ends sooner precisely because the product is good enough to be validated quickly. This prediction matches the intuition behind many of the motivating examples – "going viral" can end the niche phase, and a rave review can compress the window in which the product feels distinctive – and finds direct empirical support in Berger & Le Mens (2009), who show that cultural tastes adopted more quickly die out faster, and in Bellezza (2023), who documents that luxury products lose status as they diffuse.

These demand-side patterns would be impossible in standard adoption models. In the Bass model (Bass, 1969), social learning frameworks (Banerjee, 1992; Bikhchandani et al., 1992), and network-effects models (Katz & Shapiro, 1985), popularity affects all consumers in the same direction: adoption is self-reinforcing, higher quality extends lifecycles, and markets tip toward winner-take-all outcomes. In our setting, popularity is a *two-sided force*: it increases demand among conformists while decreasing demand among snobs. We prove formally that the quality-duration reversal is impossible in any adoption model where all consumers respond to popularity with the same sign, making it a diagnostic for the presence of opposing preferences.

We then study the *supply side*. First, we analyze a *platform* that chooses how visible popularity is. We show that a welfare-maximizing platform should *restrict* visibility: adding noise to popularity signals delays conformist entry and extends the discovery phase, with optimal visibility decreasing in product quality. Under social learning or network effects, the opposite holds: transparency is unambiguously welfare-improving. Whether a profit-maximizing platform over- or under-reveals depends on its patience.

Second, we analyze a *seller* who can respond through scarcity, targeting, and pricing. Artificial scarcity sustains the snob-dominated phase indefinitely; under network effects, such restriction is unambiguously harmful. Optimal targeting favors snobs even when conformists are the larger segment. Type-based pricing discounts early adopters and charges latecomers a premium, reversing network-effects pricing logic. Broad-reach advertising can even *reduce* profits by accelerating the conformist transition. In both cases, policy that accelerates adoption can be privately attractive while shortening the period during which the product appeals to its early adopters. Seller instruments and the platform's visibility choice are complements.

A natural objection is that snobs might not dislike popularity per se; some might dislike popularity with conformists only. The Online Appendix addresses this concern by allowing richer, composition-sensitive social preferences and showing that the core mechanism remains: whenever some consumers are attracted to broader adoption while others are repelled, popularity management becomes a first-order determinant of product lifecycles and design incentives.

Roadmap. Section 2 reviews related literature. Section 3 presents the model. Sections 4–6 characterize the consumer equilibrium, derive the quality-duration reversal, and develop comparative statics. Section 7 analyzes firm strategy. Section 8 studies market efficiency and platform design. Section 9 organizes the model's testable predictions and discusses existing evidence. Section 10 concludes.

2 Related Literature

Fashion and status. 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 opposing preferences in a dynamic environment with learning, generating lifecycle dynamics that static models cannot produce. 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. Our mechanism is distinct: cycles arise from preference conflict along a unique equilibrium path, and our focus on the quality-duration relationship is new. On exclusivity, Kuksov & Xie (2012) show firms may restrict supply to preserve status; we add the intertemporal dimension.

Microfoundations. Our reduced-form utilities $-\alpha n$ and $+\beta n$ abstract from deeper mechanisms – distinction (Veblen, 1899; Bourdieu, 1984), conformity (Bernheim, 1994), network externalities (Katz & Shapiro, 1985), identity (Akerlof & Kranton, 2000), signaling (Pesendorfer, 1995) – taking the coexistence of both types as primitive.

Diffusion and social learning. The Bass model (Bass, 1969) and social learning frameworks (Banerjee, 1992; Bikhchandani et al., 1992) assume uniformly positive adoption externalities. Our snobs are not Bass innovators: they respond *negatively* to adoption, so higher quality can shorten lifecycles rather than extending them. The "chasm" in technology adoption (Moore, 1991) corresponds to our Phase I–II transition.[1]

Influencers and platforms. Tucker & Zhang (2011) provide field-experimental evidence that toggling popularity displays shifts demand heterogeneously across product types – a direct analogue of our visibility parameter. Valsesia et al. (2020) offer a complementary microfoundation for the snob response: following fewer others on social media signals autonomy and increases perceived influence. Nistor et al. (2024) model authenticity-monetization tensions generating growth-then-decline patterns parallel to our lifecycle, and Schoenmueller et al. (2021) document that influencer follower counts follow bell-shaped lifecycles. Berman et al. (2024) analyze influencer-driven social learning, complementing our preference-based mechanism, and Cong & Li (2024) study influencer-seller matching with audience composition effects. On seeding and targeting, Godes & Mayzlin (2009) and Aral & Walker (2012) document positive spillovers under uniformly positive externalities; our model shows such strategies can backfire when spillovers have mixed signs. Our visibility results connect to the Bayesian persuasion framework (Kamenica & Gentzkow, 2011; Bergemann & Morris,

[1] On the information side, Bondi (2025) shows how the changing composition of adopters distorts the informativeness of aggregate signals; we focus on how preferences over adoption shape diffusion itself.

2019). A growing body of work applies information design to platform settings;[2] our contribution is the lifecycle dimension: the optimal signal depends on how information accelerates a *dynamic* composition shift, not just on a static match or attention problem.

Empirical evidence. Berger & Heath (2007, 2008) show that consumers abandon products when outgroup members adopt. Yoganarasimhan (2017) identifies fashion cycles and proposes methods to detect them, Iyengar et al. (2011) document "opinion snobship" among early adopters, and Han et al. (2010) show consumers strategically manage brand prominence.

3 Model

Two types of consumers disagree about popularity: snobs lose utility when others adopt, conformists gain it. Because snobs have lower outside options, they try new products first. Their early adoption generates public information that eventually draws conformists in – but the resulting crowding drives snobs out. The question is when this composition shift becomes self-reinforcing, and how it depends on product quality. Answering that requires a formal dynamic framework.

3.1 Environment

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

Adoption is a flow decision: agents choose each period whether to consume the product. An agent who adopts in period t receives utility in that period; they may continue or exit in subsequent periods. There are no switching costs or durable commitments.[3]

There is a continuum of agents with mass normalized to 1. Agents are heterogeneous in their preferences over aggregate adoption. Snobs (mass $\lambda \in (0, 1)$) derive utility from distinctiveness, while conformists (mass $1 - \lambda$) derive utility from conformity. Agent types are permanent and publicly known in equilibrium. The binary type space is a tractable special case of a continuous parameter $\gamma_i \in \mathbb{R}$.[4]

[2] See Iyer & Zhong (2022) on dynamic push-notification design, Shi et al. (2023) on platform reputation systems, Hagiu & Wright (2024) on optimal product discoverability, Yao (2024) on dynamic persuasion in consumer search, Iyer & Katona (2016) on competition for attention with status-based utility, Gardete & Hunter (2024) on multiattribute search information, Ke et al. (2023) on platform recommendations, Jerath & Ren (2021) on firm information design under rational inattention, and Bergemann & Bonatti (2019) for a survey.

[3] The flow specification fits fashion, cultural consumption, subscriptions, and platforms with low switching costs.

[4] With continuous types $\gamma_i \sim F$, the sharp three-phase structure becomes smooth but the lifecycle remains hump-shaped. The binary specification captures the essential economics while delivering closed-form

Each period, each agent i receives a fresh private signal $s_{i,t} \in \mathbb{R}$ about the product's quality. Conditional on quality θ, signals are drawn from distributions with density $f(s \mid \theta)$ satisfying MLRP (see Assumption 1). For concreteness, we assume:

$$s_{i,t} \mid \theta \sim \mathcal{N}(\theta, \sigma^2) \tag{1}$$

where $\sigma > 0$ measures signal noise. Signals are conditionally independent across agents and time given θ.

Assumption 1 (Signal Structure)**.** Private signals satisfy three conditions (standard in the social learning literature). First, the monotone likelihood ratio property (MLRP): $f(s \mid H)/f(s \mid L)$ is strictly increasing in s. Second, bounded densities: $f(s \mid \theta)$ is continuous with $f(s \mid \theta) > 0$ for all $s \in \mathbb{R}$. Third, symmetric tails: $\lim_{s \to -\infty} f(s \mid H)/f(s \mid L) = 0$ and $\lim_{s \to \infty} f(s \mid H)/f(s \mid L) = \infty$. These conditions ensure posteriors are well-behaved: strictly increasing in signals, with full support on $(0, 1)$.

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

$$\mu_i = \mathbb{P}(\theta = H \mid s_i, \hat{\theta}_t) = \frac{f(s_i \mid H)\hat{\theta}_t}{f(s_i \mid H)\hat{\theta}_t + f(s_i \mid L)(1 - \hat{\theta}_t)} \tag{2}$$

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 $(0, 1)$.[5]

3.2 Utility Functions

Let $n \in [0, 1]$ denote the mass of current adopters (flow, not stock). Product quality is $\theta \in \{L, H\}$ with $v(H) = 1$ and $v(L) = 0$.

Snobs derive utility:

$$U^S(\text{adopt} \mid \theta, n) = v(\theta) - \alpha n \tag{3}$$

where $\alpha > 0$ measures originality preference.[6] Conformists derive utility:

$$U^C(\text{adopt} \mid \theta, n) = v(\theta) + \beta n \tag{4}$$

where $\beta > 0$ measures conformity preference.[7]

comparative statics.

[5] The fresh-signal assumption fits re-sampling (a diner revisits) or cohort turnover (each period's agents are new).

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

[7] The linear specification $+\beta n$ is the reduced form of several models: Bernheim's (1994) conformity model,

Each agent has access to an outside option yielding type-specific flow payoff c_τ, with $c_C > c_S \geq 0$. The ordering has a natural microfoundation. As written, the social component is purely negative for snobs ($-\alpha n$), but the more primitive specification is $U^S = v(\theta)+\kappa-\alpha n$, where $\kappa > 0$ captures the exclusivity premium snobs derive from being among a select few. Defining $c_S \equiv c - \kappa$ for a common outside option c absorbs the premium into the reservation utility, recovering our specification exactly. The ordering $c_C > c_S$ is therefore the reduced form of snobs deriving positive utility from exclusivity, not an ad hoc assumption about outside options.[8]

The ordering ensures that snobs adopt first while conformists wait; if $c_C = c_S$, both types face the same threshold at $n = 0$, eliminating the sequential entry that produces the three-phase structure. The key qualitative results – opposing threshold monotonicity, the reversal, and the impossibility – require only $\alpha, \beta > 0$, not this ordering.

Assumption 2 (Parameter Restrictions)**.** We impose: (i) $\alpha + \beta > c_C - c_S$, ensuring the threshold crossing point $n^\dagger = (c_C - c_S)/(\alpha + \beta) \in (0, 1)$; (ii) $c_C > c_S \geq 0$; and (iii) $\sigma > 0$.

Each period, agents simultaneously observe private signals, form posteriors via (2), and choose $a_i \in \{\text{adopt}, \text{pass}\}$ to maximize expected utility given rational expectations about aggregate adoption n^*. Because agents are atomistic, no individual affects n_t or the public belief $\hat{\theta}_t$; each agent takes the aggregate state as given. We focus on symmetric Bayesian Nash Equilibrium where expectations are correct.[9]

4 Benchmark: Single-Period Equilibrium

We first characterize the single-period equilibrium to isolate the preference mechanism from learning dynamics.

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*

network effects (Katz & Shapiro, 1985) with constant per-adopter benefits, and social identity models (Akerlof & Kranton, 2000) with linear group-size utility.

[8] An equivalent microfoundation works through the outside option directly: if the established alternative has adoption $\bar{n} > 0$, conformists value it more, yielding $c_C - c_S = (\alpha + \beta)\bar{n} > 0$.

[9] Asymmetric equilibria exist but are not robust to small perturbations in signals; see Morris & Shin (2003).

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 thresholds have intuitive properties. The snob threshold $\underline{\mu}^S(n) = \alpha n + c_S$ is increasing in n: as more agents adopt, snobs require higher expected quality to justify the crowding cost. The conformist threshold $\underline{\mu}^C(n) = c_C - \beta n$ is decreasing in n: as more agents adopt, conformists require lower expected quality because bandwagon benefits compensate.[10] 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 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 to Adoption

Our model's defining feature is that snobs and conformists respond oppositely to aggregate adoption.

Proposition 1 (Type-Specific Responses)**.** *In equilibrium, snobs and conformists exhibit opposite responses to adoption:*

(i) Snob exit region. *For $n > \bar{n}^S \equiv (1 - c_S)/\alpha$, the snob threshold exceeds one ($\underline{\mu}^S(n) > 1$), so no snob adopts: $\lambda[1 - G(\underline{\mu}^S(n))] = 0$.*

(ii) Conformist entry region. *For $n > \underline{n}^C(\hat{\theta}_t) \equiv (c_C - \bar{\mu}(\hat{\theta}_t))/\beta$ where $\bar{\mu}(\hat{\theta}_t) = \mathbb{E}[\mu_i \mid \hat{\theta}_t]$, a positive mass of conformists adopts.*

(iii) Simultaneous opposition. *When $\underline{n}^C < n < \bar{n}^S$, the snob adoption mass is decreasing in n while the conformist adoption mass is increasing in n.*

As n rises, $\partial \underline{\mu}^S / \partial n = \alpha > 0$ while $\partial \underline{\mu}^C / \partial n = -\beta < 0$. This opposing monotonicity is the central structural feature of the model, directly supported by evidence in Iyengar et al. (2011) and Berger & Heath (2008).[11]

[10] This opposing monotonicity appears in Corneo & Jeanne (1997a)'s static analysis. Our contribution is embedding it in a dynamic environment with learning, producing non-monotone lifecycle paths and an impossibility result.

[11] In Pesendorfer (1995), abandonment reflects signal erosion, not crowding disutility. In Kuksov & Wang (2013), all consumers agree on trendiness; here the same popularity level attracts some and repels others.

4.3 Equilibrium Adoption Mass

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

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

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

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

Since $\Phi : [0,1] \to [0,1]$ is continuous, existence follows from Brouwer. For uniqueness:

Assumption 3 (Sufficient Condition for Uniqueness)**.** Social preferences are moderate relative to signal precision: $\kappa \equiv \max_{n \in [0,1]} |\Phi'(n)| < 1$.

When $\kappa < 1$, Φ is a contraction and equilibrium is unique – a standard condition analogous to global games (Morris & Shin, 2003). We impose Assumption 3 throughout. The *within-period* comparative statics – opposing threshold monotonicity (Proposition 1), the direction of the reversal (Proposition 4), and the impossibility result (Proposition 5) – hold at any stable fixed point of Φ, regardless of uniqueness. The *dynamic path* additionally requires uniqueness, since the belief update (9) uses $n^*_\theta(\hat{\theta}_t)$ as a single-valued function.[12]

5 Dynamic Equilibrium and Product Lifecycles

We now embed these opposing threshold responses in a dynamic framework.

5.1 Dynamic Structure

A single product with quality $\theta \in \{L, H\}$ is available. Each period $t = 0, 1, 2, \ldots$, agents first observe their private signals and (for $t \geq 1$) a noisy signal of previous adoption, then simultaneously decide whether to adopt, after which payoffs realize and adoption mass n_t is determined.

The payoff-relevant state is $(n_t, \hat{\theta}_t)$ where n_t is current adoption and $\hat{\theta}_t$ is the public belief about quality. We focus on Markov Perfect Equilibrium where strategies depend only on the current state.

[12] If $\kappa > 1$, three equilibria can exist (see Appendix A). Under multiplicity, one could select the Pareto-dominant stable equilibrium; all within-period results carry through unchanged.

At the start of period t, agents observe $\tilde{n}_{t-1}$ and update the public belief from $\hat{\theta}_{t-1}$ to $\hat{\theta}_t$ via Bayes' rule (equation (9) below). They then receive private signals and form posteriors μ_i. With threshold strategies, period-t adoption is:

$$n_t = \lambda \cdot \Pr(\mu_i \geq \alpha n_t^e + c_S) + (1-\lambda) \cdot \Pr(\mu_i \geq c_C - \beta n_t^e) \tag{8}$$

where n_t^e is agents' rational expectation of period-t adoption and probabilities are over the distribution of posteriors μ_i induced by private signals given public belief $\hat{\theta}_t$. In equilibrium, $n_t^e = n_t$, so (8) is a fixed-point condition that determines adoption each period. This fixed point is computed conditional on $(\hat{\theta}_t, \theta)$; the learning mechanism driven by noisy lagged signals $\tilde{n}_{t-1}$ operates across periods.

Agents observe a noisy public signal of previous-period adoption: $\tilde{n}_{t-1} = n_{t-1} + \varepsilon_t$ where $\varepsilon_t \sim \mathcal{N}(0, \sigma_\varepsilon^2)$. The noise ensures Bayes' rule applies for all observations (no off-path belief issues) and that learning is gradual.[13]

Public beliefs update via Bayes' rule:

$$\hat{\theta}_t = \frac{f(\tilde{n}_{t-1} \mid \theta = H) \cdot \hat{\theta}_{t-1}}{f(\tilde{n}_{t-1} \mid \theta = H) \cdot \hat{\theta}_{t-1} + f(\tilde{n}_{t-1} \mid \theta = L) \cdot (1 - \hat{\theta}_{t-1})} \tag{9}$$

where $f(\tilde{n} \mid \theta)$ is the density of observed adoption given true quality. This density depends on the *equilibrium* adoption function:

$$f(\tilde{n} \mid \theta) = \frac{1}{\sigma_\varepsilon} \phi\left(\frac{\tilde{n} - n_\theta^*(\hat{\theta}_{t-1})}{\sigma_\varepsilon}\right) \tag{10}$$

where ϕ is the standard normal density and $n_\theta^*(\hat{\theta})$ is the equilibrium adoption when public belief is $\hat{\theta}$ and true quality is θ. Under MLRP, $n_H^*(\hat{\theta}) > n_L^*(\hat{\theta})$: high quality generates more adoption, so observing high $\tilde{n}_{t-1}$ raises beliefs. We treat aggregate adoption as the sole public signal for quality learning. In practice, adopter composition can distort the information content of aggregate signals; Bondi (2025) formalizes this channel. Our model complements that analysis by holding the learning technology fixed.

Agents also receive fresh private signals each period, so individual posteriors μ_i combine the public belief $\hat{\theta}_t$ with private signal s_i via Bayes' rule (2).

5.2 Value Functions and Equilibrium

Let $\delta \in (0, 1)$ denote the common discount factor. Given public belief $\hat{\theta}_t$ and private signal s_i, agent i forms posterior μ_i via (2). The value function for type $\tau \in \{S, C\}$ with posterior

[13] The noise can be microfounded as sampling variation in a large finite population or as heterogeneous observation timing. What matters is $\sigma_\varepsilon^2 > 0$; all qualitative results hold for any positive value.

μ_i in state $(n_t, \hat{\theta}_t)$ satisfies the Bellman equation:

$$V^\tau(\mu_i; \hat{\theta}_t) = \max \left\{ \underbrace{\mu_i + u^\tau(n^*(\hat{\theta}_t)) + \delta \mathbb{E}[V^\tau(\mu_i'; \hat{\theta}_{t+1})]}_{\text{adopt}}, \underbrace{c_\tau + \delta \mathbb{E}[V^\tau(\mu_i'; \hat{\theta}_{t+1})]}_{\text{outside option}} \right\} \tag{11}$$

where $u^S(n) = -\alpha n$, $u^C(n) = \beta n$, $n^*(\hat{\theta}_t)$ is equilibrium adoption given belief $\hat{\theta}_t$, c_τ is the flow payoff from the outside option, and expectations are over next-period posteriors and beliefs. Because adoption is a flow decision with no commitment, the continuation value $\delta \mathbb{E}[V^\tau(\mu_i'; \hat{\theta}_{t+1})]$ appears on both sides of the max operator and cancels.[14] The adoption decision therefore reduces to comparing current-period payoffs, and the myopic thresholds in Lemma 1 are dynamically optimal.

Definition 1 (Markov Perfect Equilibrium)**.** A Markov Perfect Equilibrium consists of value functions V^S, V^C and policy functions $\sigma^S, \sigma^C : (0,1) \times (0,1) \to \{\text{adopt}, \text{pass}\}$ satisfying three conditions: V^τ satisfies (11) given the transition induced by (σ^S, σ^C); $\sigma^\tau(\hat{\theta}, s)$ attains the maximum in (11) for each belief and signal; and adoption is consistent in that $n_t = \Phi(n_t; \hat{\theta}_t, \sigma^S, \sigma^C)$ per (8) at each t, with beliefs evolving via (9).

Proposition 2 (MPE Existence)**.** *Under Assumption 2, a Markov Perfect Equilibrium exists. It is unique when $\kappa \equiv \max_n |\Phi'(n)| < 1$ (Assumption 3).*[15]

5.3 Three-Phase Lifecycle

The dynamic equilibrium exhibits a predictable lifecycle when the product is high quality ($\theta = H$):

Definition 2 (Product Lifecycle)**.** A product lifecycle consists of three phases: growth ($t < t_1$, snobs dominate, n_t increasing), peak ($t_1 \leq t \leq t^*$, conformists enter, n_t reaches maximum n^*), and decline ($t > t^*$, snobs exit, n_t^S decreasing). Duration $T \equiv \min\{t > t^* : n_t^S < \epsilon\}$ measures the length of the culturally active lifecycle – the period during which snob participation remains above a threshold $\epsilon > 0$.[16]

This three-phase structure echoes classical product lifecycle theory (Mahajan et al., 1990) but emerges from preference heterogeneity rather than technology diffusion or market saturation. The model generates a single boom-bust for each product; once beliefs decline, snobs do not return because the belief path is absorbing. Fashion cycles in the sense of Simmel

[14] With switching costs, option-value considerations modify thresholds (Chamley, 2004), but the phase structure is robust since the option-value term shifts the intercept without changing the slope in n.

[15] When multiple equilibria exist, we select the Pareto-dominant (highest-adoption) equilibrium, which is also the unique stable one under best-response dynamics. See Appendix A for the full construction.

[16] All comparative statics on T hold uniformly across $\epsilon \in (0, n_0^S)$.

(1957) – recurring waves of adoption and abandonment – emerge from *sequential* product introductions: as product A enters decline, snobs migrate to product B.

Proposition 3 (Lifecycle Characterization). *Suppose $c_C > c_S$ and $\theta = H$. Define*

$$\bar{\alpha}_{growth} \equiv \frac{\sqrt{2\pi \operatorname{Var}(\mu \mid \hat{\theta}_0)}}{\lambda} \tag{12}$$

If $\alpha < \bar{\alpha}_{growth}$, then in any MPE with threshold strategies, the lifecycle proceeds in three phases:

(I) Growth *($t = 0, \ldots, t_1 - 1$): snobs dominate; n_t is increasing; $n_t^S > n_t^C$.*

(II) Peak *($t = t_1, \ldots, t^*$): conformists enter; both types adopt; $n_{t^*} = n^*$.*

(III) Decline *($t > t^*$): snob participation falls as $\underline{\mu}^S(n_t) = \alpha n_t + c_S$ rises. The lifecycle ends at $T = \min\{t > t^* : n_t^S < \epsilon\}$ for $\epsilon > 0$.*[17]

When $\alpha \geq \bar{\alpha}_{growth}$, the three-phase structure obtains qualitatively but Phase I may not exhibit monotone increases.

In Phase I, snobs adopt based on private quality signals, generating informative public signals that raise beliefs.[18] The Phase I–II transition occurs when rising beliefs and rising n_t jointly lower $\underline{\mu}^C = c_C - \beta n_t$ enough for a critical mass of conformists to enter. This transition can be sharp when social preferences are strong enough for the equilibrium to "tip" between states.[19]

Phase III begins when snob adoption starts to decline. Rising n_t pushes 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. After snobs depart, conformists may continue buying, but the product has lost its cultural cachet.[20]

Remark 1 (The Role of Posterior Variance). Higher posterior variance disperses beliefs, ensuring sufficient mass above the snob threshold even as it rises. The lifecycle is therefore most robust when quality uncertainty is moderate: enough signal to learn, but enough noise to prevent the snob threshold from binding too tightly.

[17] When $\alpha + c_S \geq 1$, the snob threshold eventually exceeds 1 and snob adoption reaches zero exactly. When $\alpha + c_S < 1$, snob adoption declines toward zero but remains positive.

[18] $\bar{\alpha}_{\text{growth}}$ is large when the prior is diffuse, the snob population is small, or signals are precise enough to spread the posterior distribution.

[19] The phase boundaries t_1 and t^* are stopping times defined formally in Appendix A. When multiple equilibria exist (Lemma 2), the transitions can be discontinuous.

[20] This distinguishes our model from Ke et al. (2024), who generate fashion classics via cross-generation signaling, and from Baumann & Olszewski (2021), whose cycles arise from equilibrium multiplicity. Our scarcity result provides the within-generation mechanism.

Remark 2 (Low-Quality Products). When $\theta = L$, beliefs drift toward zero and adoption never surges. Because adoption stays low, the snob threshold remains well below 1 and snobs *persist* – never crowded out, only progressively less convinced. The quality-duration reversal $T^H < T^L$ arises from this asymmetry: high quality triggers the conformist surge that drives snobs out rapidly, while low quality never triggers it.

5.4 Peak Adoption Mass

Assumption 4 (Density Regularity)**.** $g(\underline{\mu}^S(n^*)) \approx g(\underline{\mu}^C(n^*))$: the posterior density is approximately constant across the threshold region at peak adoption.[21]

Assumption 4 delivers closed-form expressions. The *signs* of all comparative statics hold exactly at any stable equilibrium by the implicit function theorem; only the magnitudes require the approximation.

Corollary 1 (Peak Adoption)**.** *Peak adoption satisfies* $n^* = \Phi(n^*; \hat{\theta}_{t^*})$*. Under Assumption 4:*

$$n^* = \frac{c_C - c_S}{\alpha + \beta} = n^\dagger, \quad \mu^* = \frac{\alpha c_C + \beta c_S}{\alpha + \beta} \tag{13}$$

with $\partial n^*/\partial \alpha < 0$ *and* $\partial n^*/\partial \beta < 0$.[22]

Peak adoption coincides with the threshold crossing point $n^\dagger$ under the density approximation. Stronger social preferences of either type shrink the peak: markets with intense snob aversion or strong bandwagon effects sustain smaller audiences at their height.

5.5 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 4 (Quality-Duration Reversal)**.** *There exists a threshold* $\bar{\alpha}(\lambda) > 0$*, with closed-form approximation* $\bar{\alpha}(\lambda) = \beta(1-\lambda)/\lambda$ *under Assumption 4, such that:*

$$T^H > T^L \iff \alpha < \bar{\alpha}(\lambda), \qquad T^H < T^L \iff \alpha > \bar{\alpha}(\lambda) \tag{14}$$

The threshold is decreasing in λ *and increasing in* β*. When snob aversion exceeds the threshold, higher-quality products have shorter lifecycles.*

[21] This holds when both thresholds fall in the bulk of the posterior distribution. The Online Appendix provides a worked example where exact closed forms obtain without this approximation.

[22] We use n^* for both the within-period equilibrium and the peak adoption level; the two coincide at t^*. Where the within-period equilibrium at arbitrary t is needed, we write n_t.

The mechanism is intuitive: higher quality generates stronger signals, which attract conformists faster. If snob aversion is strong enough, this acceleration triggers snob exit before the product reaches its natural peak. A restaurant that receives a rave review may see its regulars depart within weeks as reservations become impossible to get; a niche musician who goes viral may lose the core audience that valued the intimacy. The threshold $\bar{\alpha}(\lambda)$ pins down when this pattern dominates: fewer snobs or weaker conformist preferences make the reversal easier to trigger.

The quality-duration reversal is not merely an unusual parameter region of a general model. It is *impossible* in a broad class of adoption models that lack the preference heterogeneity we study. The following result establishes this formally.

Proposition 5 (Impossibility of Reversal Under Homogeneous Preferences)**.** *Consider any adoption model with: (a) flow utility* $u_i(\theta, n) = v(\theta) + \gamma_i \cdot h(n)$*,* $v' > 0$*,* $h' > 0$*; (b) threshold strategies based on private posteriors; (c) duration* $T(\theta) \equiv \sup\{t : n_t(\theta) > \varepsilon\}$ *for any* $\varepsilon \geq 0$*, where* n_t *denotes total adoption.*[23] *If* γ_i *has the same sign for all agents –* $\gamma_i \geq 0 \; \forall i$ *(bandwagon),* $\gamma_i \leq 0 \; \forall i$ *(congestion), or* $\gamma_i = 0 \; \forall i$ *(private values) – then:*

$$\theta^H > \theta^L \implies T(\theta^H) \geq T(\theta^L) \tag{15}$$

The reversal $T^H < T^L$ *requires* $\gamma_i > 0$ *for some agents and* $\gamma_j < 0$ *for others.*

The result rules out the reversal under Bass diffusion (Bass, 1969), social learning (Banerjee, 1992; Bikhchandani et al., 1992), network effects, and congestion – regardless of heterogeneity in $|\gamma_i|$, signal precision, or outside options.[24] The impossibility makes the reversal a diagnostic: observing that higher-quality products cycle faster in a market points specifically to the coexistence of consumers who value popularity and consumers who flee from it.

When originality preferences are weak ($\alpha < \bar{\alpha}$), the standard intuition holds: quality extends duration. When they are strong ($\alpha > \bar{\alpha}$), acceleration dominates: excellence speeds decline by attracting conformists too rapidly. The reversal is robust to non-linear social preferences (Online Appendix).[25]

Remark 3 (Compatibility of Uniqueness and the Reversal)*.* The reversal requires $\alpha > \bar{\alpha}(\lambda)$; uniqueness requires $\kappa < 1$. These are jointly satisfiable: the reversal binds on the *ratio* α/β, while uniqueness binds on *levels* relative to the posterior density. Increasing σ flattens the density and reduces κ, so the reversal obtains as a unique equilibrium outcome for a wide range of parameters.

[23] When a negative-externality subgroup exists ($\gamma_j < 0$), the result also holds with duration defined as $T(\theta) \equiv \sup\{t : n_t^-(\theta) > 0\}$.

[24] The additive structure accommodates multiplicative utilities via logarithms and extends to any model where $\partial \underline{\mu}_i / \partial n$ has the same sign for all agents.

[25] The key requirement is that snob disutility grows at least proportionally with conformist utility as adoption rises.

6 Comparative Statics

The three preference parameters – snob aversion α, conformist attraction β, and the snob share λ – vary across industries and platforms. This section derives how each shapes the lifecycle.

Corollary 2 (Preference Parameter Effects)**.** *Stronger preferences of either type reduce peak adoption ($\partial n^*/\partial\alpha < 0$, $\partial n^*/\partial\beta < 0$). Stronger snob aversion shortens lifecycles ($\partial T/\partial\alpha < 0$). Stronger conformist preferences raise the bar for the quality-duration reversal ($\partial\bar{\alpha}/\partial\beta > 0$): bandwagon effects must be overcome by commensurately stronger crowding aversion.*

Higher α shortens cycles – fashion-forward categories like streetwear turn over more rapidly than categories where originality matters less. Both parameters reduce peak adoption because stronger social preferences shrink the region where both types coexist. (Proofs in Online Appendix.)

Proposition 6 (Non-Monotonic Composition Effects)**.** *The relationship between snob share λ and cycle duration is non-monotonic:*

$$\frac{\partial T}{\partial \lambda} \begin{cases} > 0 & \text{if } \lambda < \lambda^* \\ < 0 & \text{if } \lambda > \lambda^* \end{cases} \tag{16}$$

where $\lambda^* = \frac{\beta}{\alpha+\beta}$. Duration is maximized when the snob share balances the relative strength of social preferences.

Too few snobs and products are never discovered; too many and products cycle too rapidly. The longest lifecycles occur at an interior composition, which varies across markets.

Proposition 7 (Information Precision Can Hurt)**.** *More precise signals (lower σ) accelerate conformist entry ($\partial t_1/\partial\sigma > 0$). When $\alpha > \bar{\alpha}(\lambda)$, the net effect is to shorten lifecycles ($\partial T/\partial\sigma > 0$) and reduce expected snob welfare.*

Better information helps conformists but can hurt snobs: improved quality assessment accelerates conformist entry, compressing pioneer rents.[26] Review aggregators and recommendation algorithms thus inadvertently accelerate product cycling in identity-relevant markets (Berger & Le Mens, 2009).

[26] This connects to Bayesian persuasion (Kamenica & Gentzkow, 2011): revealing more about the state can make some receivers worse off by triggering action that imposes negative externalities.

7 Implications for Firm Strategy

We begin with the seller's problem: a monopolist who can choose pricing, supply, and marketing instruments.[27]

7.1 Dynamic Pricing

Consider a monopolist setting a uniform price p_t each period.

Corollary 3 (Optimal Pricing Path). *The optimal uniform price path extracts the marginal consumer's surplus each period. When snobs are marginal (Phase I), the firm prices to the snob threshold; when conformists are marginal (Phases II–III), to the conformist threshold:*

$$p_t = \begin{cases} \hat{\theta}_t - \alpha n_t - c_S & t < t_1 \text{ (Phase I: snobs marginal)} \\ \hat{\theta}_t + \beta n_t - c_C & t \geq t_1 \text{ (Phases II–III: conformists marginal)} \end{cases} \tag{17}$$

The path is non-monotonic: typically decreasing in Phase I as rising crowding costs squeeze snob margins, jumping upward at the Phase I–II transition as conformists arrive with higher willingness to pay, then decreasing in Phase III as n_t falls.

The price path contrasts sharply with network-effects pricing (Katz & Shapiro, 1985), where prices typically increase as the installed base grows. Here, the Phase I–II price jump is a compositional effect: the marginal consumer shifts from snob to conformist, and the Phase I discount reflects the snob's willingness-to-pay schedule, not penetration pricing (Bass, 1969).[28]

7.2 Type-Specific Dynamic Pricing

A monopolist who can identify types faces a richer problem, because *who* adopts affects future adoption. Acquiring a snob extends the growth phase; acquiring a conformist accelerates crowding.[29]

Corollary 4 (Type-Specific Pricing). *With observable types, the myopically optimal per-*

[27] We treat the firm's problem as a sequence of per-period optimizations. The supply-side prescriptions rely on the sequential entry structure ensured by $c_C > c_S$.

[28] For moderate α, beliefs may rise faster than crowding costs early in Phase I, causing price to initially increase.

[29] Type identification is feasible via purchase history, access mechanisms, or self-selection (invite-only launches, pre-order windows). Adverse selection moderates the optimal price gap but does not reverse its sign.

period prices satisfy:

$$p_t^S = \hat{\theta}_t - \alpha n_t - c_S \tag{18}$$

$$p_t^C = \hat{\theta}_t + \beta n_t - c_C \tag{19}$$

yielding a price gap:

$$p_t^C - p_t^S = (\alpha + \beta)n_t - (c_C - c_S) \tag{20}$$

The gap is positive for $n_t > n^\dagger$ and maximal at n^*.

The prescription – discount snobs, charge conformists a premium – reverses network-effects pricing (Katz & Shapiro, 1985), where early adopters are subsidized to *build* the network rather than *preserve exclusivity.*[30] The profit gain has a static component (surplus extraction) and a dynamic one: subsidizing snobs extends Phase I, generating better quality signals that raise conformist willingness to pay.

Remark 4 (Incentive Compatibility with Unobservable Types). When types are unobservable, adoption timing is a natural screening instrument: an early-access price $p^{\text{early}} < p^{\text{late}}$ is incentive-compatible because snobs prefer the uncrowded early period and conformists prefer the validated late period.

7.3 Artificial Scarcity

A patient firm may prefer to cap total adoption below the natural peak, keeping the product in the snob-dominated phase indefinitely.

Proposition 8 (Artificial Scarcity and Firm Profits)**.** *A monopolist that caps adoption at $\bar{n} < n^*$ in every period strictly prefers scarcity over market clearing if and only if $\delta > \bar{\delta}(\bar{n})$, where $\bar{\delta} \in (0,1)$ is a unique threshold discount factor.*

The scarcity strategy sacrifices the conformist-surge revenue for an indefinite stream of snob-dominated revenue. The tradeoff favors scarcity when the firm is patient (δ high), snob margins are large, or the conformist transition is rapid.[31]

The result connects to a broader pattern in luxury and cultural markets. Hermès wait-lists, Supreme's limited drops, and Berghain's capacity caps all implement the quantity cap $\bar{n}$. The standard explanation is supply-side signaling or Veblen pricing; our model provides a demand-side rationale grounded in lifecycle dynamics. Under network effects or social learning, artificial scarcity is unambiguously harmful: restricting adoption weakens the externality that makes the product valuable (Katz & Shapiro, 1985). Here, restricting adoption *preserves* value by keeping the wrong type out.

[30] Cong & Li (2024) show that more powerful influencers command higher wages because they attract a more desirable audience composition – the influencer-level analogue of the snob discount.

[31] Formally, scarcity is preferred iff $\pi^{\text{scarce}}(\bar{n})/(1-\delta) > \sum_{t=0}^{T} \delta^t \pi_t$; the LHS diverges as $\delta \to 1$, guaranteeing a unique crossing. See Kuksov & Xie (2012) for the static analogue.

7.4 Advertising and Targeting

Lemma 3 (Advertising and Lifecycle Duration). *Advertising with effectiveness $a > 0$ that increases adoption satisfies $\partial T/\partial a < 0$ when $n > n^\dagger$, where $n^\dagger = (c_C - c_S)/(\alpha + \beta)$ is the threshold crossing point: successful advertising shortens product lifecycles by accelerating snob exit.*

Any force that increases n faster pushes the product through the reversal regime more quickly, so the firm's problem is to *target* adoption toward the right type. In diffusion and social learning models, broad advertising is unambiguously profit-increasing: more adopters generate more positive externalities, speeding diffusion. Here, broad advertising can be self-defeating because it attracts the type whose entry triggers exit of the type that sustains the lifecycle.

Proposition 9 (Optimal Targeting). *Let $\tau \in [0, 1]$ denote the share of advertising directed at snobs. The optimal targeting strategy satisfies:*

$$\tau^* = \min \left\{ 1, \underbrace{\frac{\alpha}{\alpha + \beta}}_{\text{static term}} + \underbrace{\frac{\delta}{1-\delta} \cdot \frac{\lambda}{1-\lambda} \cdot \frac{\partial T/\partial n^S}{|\partial T/\partial n^C|}}_{\text{dynamic term}} \right\} \tag{21}$$

where $\partial T/\partial n^\tau$ denotes the marginal effect on lifecycle duration of a unit increase in type-τ advertising.[32] *When $\delta > \bar{\delta}$ or $\alpha > \bar{\alpha}$, the optimum is $\tau^* = 1$.*

Firms should disproportionately target snobs even when conformists are the majority, contrasting with Godes & Mayzlin (2009) and Aral & Walker (2012). The result gives a demand-side foundation for why firms should prefer influencers with snob-like audiences (Nistor et al., 2024; Cong & Li, 2024) – not because those audiences are larger, but because they sustain the lifecycle that eventually attracts conformists. Valsesia et al. (2020) provide complementary evidence: influencers who follow fewer others are perceived as more autonomous and influential, consistent with our model's prediction that snob-like selectivity confers disproportionate market value.[33]

8 Market Efficiency and Platform Design

We now turn to the platform's problem. The firm controls quantity; the platform controls information. Both slow conformist entry, but through different channels.

[32] Formally, $\partial T/\partial n^S \equiv (\partial T/\partial a^S)|_{a^C}$ where a^τ is advertising intensity directed at type τ, well-defined by the implicit function theorem.

[33] Berman et al. (2024) show that social learning through mega-influencers is most valuable when product quality is uncertain – precisely the Phase I regime where our targeting result bites hardest.

8.1 How the Lifecycle Distributes Surplus

The lifecycle creates a fundamental surplus asymmetry between the two types. The expected lifetime utility of a representative agent of each type is:

$$EU^S = \sum_{t=0}^{t^*} \mathbb{E}\left[\theta - \alpha n_t - c_S \mid \text{adopt}\right] \cdot \Pr(\text{adopt}_t^S) \cdot \delta^t \tag{22}$$

$$EU^C = \sum_{t=t_1}^{T} \mathbb{E}\left[\theta + \beta n_t - c_C \mid \text{adopt}\right] \cdot \Pr(\text{adopt}_t^C) \cdot \delta^t \tag{23}$$

where $\Pr(\text{adopt}_t^\tau) = 1 - G(\underline{\mu}^\tau(n_t); \hat{\theta}_t)$ and expectations are over θ conditional on adoption.

Corollary 5 (Welfare Comparison)**.** *Let* $\bar{v}_t^\tau \equiv \mathbb{E}[\theta \mid adopt_t^\tau]$ *and* $\pi_t^\tau = \Pr(adopt_t^\tau)$*. Then:*

$$EU^S - EU^C = \underbrace{\sum_{t=0}^{t_1-1} [\bar{v}_t^S - \alpha n_t - c_S]\pi_t^S \delta^t}_{\text{Snob pioneer rents (Phase I)}} + R - \underbrace{\sum_{t=t^*+1}^{T} [\bar{v}_t^C + \beta n_t - c_C]\pi_t^C \delta^t}_{\text{Conformist bandwagon benefits (Phase III)}} \tag{24}$$

where $R = \sum_{t=t_1}^{t^*} \{[\bar{v}_t^S - \alpha n_t - c_S]\pi_t^S - [\bar{v}_t^C + \beta n_t - c_C]\pi_t^C\}\delta^t$.

Scarcity extends Phase I at the expense of conformist surplus; broad advertising does the reverse. The lifecycle is therefore a distributional contest: every instrument that benefits snobs (scarcity, opacity, targeted seeding) comes at a cost to conformists, and vice versa.

8.2 Market Composition and Platform Design

Definition 3 (Social Welfare)**.** Social welfare is the population-weighted sum of expected utilities:

$$W(\lambda) = \lambda \cdot EU^S(\lambda) + (1 - \lambda) \cdot EU^C(\lambda) \tag{25}$$

where the dependence on λ reflects how market composition affects equilibrium dynamics.

A natural conjecture is that opposing preferences are purely wasteful, and the optimal market would be homogeneous. This is wrong.

Proposition 10 (Welfare and Market Composition)**.** *Social welfare* $W(\lambda)$ *is maximized at an interior snob share* $\lambda^{**} \in (0, 1)$*. The welfare-maximizing share is lower when snob aversion is strong (*$\partial\lambda^{**}/\partial\alpha < 0$*) and higher when private signals are noisy (*$\partial\lambda^{**}/\partial\sigma > 0$*).*

With only conformists ($\lambda = 0$), no agent adopts unproven products – the coordination failure of Chamley (2004). With too many snobs, bandwagon effects are too weak to sustain adoption. The welfare-maximizing λ^{**} balances discovery against stability.

The population parameter λ is not a choice variable for any single agent, but platforms can shift the *effective* composition through design choices: community guidelines, invitation mechanisms, algorithmic curation, and onboarding friction all affect whether a platform attracts more snob-like or conformist-like users.[34] This makes Proposition 10 a design prescription.

The two externalities are asymmetric. Conformists impose a *crowding externality*: entry raises n_t, pushing marginal snobs out. Snobs impose an *information externality*: exit reduces the quality-sensitivity of the adopter pool, slowing belief updating.[35] Beyond composition, a platform has a second instrument: controlling the *visibility* of adoption information.

8.3 Optimal Visibility Design

Platforms routinely choose how much adoption information to display: Spotify shows real-time play counts; Bandcamp shows almost nothing; Instagram experimented with hiding likes entirely. We formalize this as a visibility parameter $\varphi \in [0, 1]$. When $\varphi = 1$, agents observe n_t exactly. When $\varphi < 1$, agents observe $\tilde{n}_t = n_t + \varepsilon_t$ where $\varepsilon_t \sim \mathcal{N}(0, \sigma_0^2(1 - \varphi)/\varphi)$.

Under noisy visibility, conformists cannot easily verify that adoption has crossed their entry threshold, while snobs – whose decisions at low n are driven primarily by private quality signals – are less affected.[36]

Lemma 4 (Equilibrium Response to Visibility)**.** *The equilibrium adoption path* $\{n_t(\varphi)\}$ *satisfies:*

(i) $\partial t_1/\partial\varphi < 0$*: higher visibility accelerates conformist entry.*

(ii) When $\alpha > \bar{\alpha}(\lambda)$*:* $\partial T/\partial\varphi < 0$ *(faster conformist entry compresses the lifecycle).*

(iii) Under Assumption 3, $n^*(\hat{\theta}, \varphi)$ *and* $W(\varphi)$ *are* C^1 *in* φ*.*

The proof is in the Online Appendix. Lower visibility mutes the popularity signal, selectively slowing conformist entry without comparably affecting snob decisions, which at low n depend primarily on private quality signals.

The platform chooses φ to maximize $W(\varphi) = \lambda \cdot EU^S(\varphi) + (1-\lambda) \cdot EU^C(\varphi)$. Lower φ slows conformist entry and extends lifecycle duration but reduces quality assessment precision – a three-way tradeoff.

Proposition 11 (Welfare-Maximizing Visibility)**.** *Under the conditions of Proposition 4 with* $\alpha > \bar{\alpha}(\lambda)$*:*

[34] The underground music scene operates with effective λ near 1; mass consumer goods near $\lambda = 0$.

[35] Bondi (2025) formalizes a related channel.

[36] In Berman et al. (2024), restricting visibility harms consumers because it reduces social learning. Here the mechanism is preference-based: reducing visibility curtails a preference externality, not an informative one.

(i) Neither full transparency ($\varphi = 1$) nor full opacity ($\varphi = 0$) maximizes welfare; the optimum $\varphi^* \in (0, 1)$ is interior.

(ii) The welfare cost of full transparency, $W(\varphi^*) - W(1)$, is larger when snobs care more intensely about exclusivity (increasing in α) and when snobs are scarce (decreasing in λ).

(iii) The platform should reveal more when private signals are noisy ($\partial\varphi^*/\partial\sigma > 0$) and less when snob aversion is strong ($\partial\varphi^*/\partial\alpha < 0$).

The optimal φ^* balances two marginal effects: transparency helps agents assess quality, but it also accelerates conformist entry and compresses the lifecycle. Full transparency is suboptimal because it ignores this lifecycle cost. Part (ii) says the welfare loss from transparency is most severe in markets with few, intensely distinction-seeking snobs – precisely the identity-driven categories (fashion, music, dining) where platforms most aggressively broadcast popularity.[37] Under social learning or network effects, this result reverses entirely: transparency is unambiguously welfare-improving because adoption is a positive signal, and more information helps agents coordinate on high-quality products (Banerjee, 1992; Bikhchandani et al., 1992). The prescription to *hide* popularity information arises only when some consumers respond negatively to it.

A platform that maximizes profit rather than welfare faces a different tradeoff. The source of misalignment is not patience per se, but the fact that the platform's objective – aggregate adoption n_t – ignores the *composition* of adopters. The welfare planner weights snob and conformist surplus separately, penalizing conformist entry that destroys snob surplus even when total adoption rises. The platform does not: it cares about conformist flooding only insofar as it triggers snob *exit*, reducing future n_t. A conformist who enters and crowds a snob who nonetheless stays is welfare-reducing (the snob's utility falls) but profit-neutral (total adoption is unchanged). This composition blindness means the platform and planner would disagree even at identical discount factors.

What patience determines is the *direction* of the resulting distortion. A profit-maximizing platform earns per-period revenue $r(n_t)$ with $r' > 0$, so it benefits from the conformist surge that higher visibility triggers. But it also loses revenue when the lifecycle ends prematurely. Which force dominates depends on how much the platform values future periods relative to the present.

Proposition 12 (Platform Patience and Visibility Distortion)**.** *Let φ^*_W denote welfare-maximizing visibility and let the platform maximize $\Pi(\varphi) = \sum_t \delta_P^t r(n_t(\varphi))$ with $r' > 0$, where $\delta_P \in (0, 1)$ is the platform's discount factor. Under the conditions of Proposition 11:*

[37] Kamenica & Gentzkow (2011) show that a sender can benefit from partial revelation; our visibility result identifies conformist acceleration as the specific mechanism making partial opacity welfare-improving.

(i) There exists a unique threshold $\bar{\delta}_P \in (0,1)$ such that the platform over-reveals ($\varphi^*_\Pi > \varphi^*_W$) if $\delta_P < \bar{\delta}_P$ and under-reveals ($\varphi^*_\Pi < \varphi^*_W$) if $\delta_P > \bar{\delta}_P$.

(ii) The threshold is decreasing in snob aversion ($\partial\bar{\delta}_P/\partial\alpha < 0$) and decreasing in the snob share ($\partial\bar{\delta}_P/\partial\lambda < 0$).

Whether a platform shows too much or too little popularity information depends on its patience. An impatient platform over-reveals: it broadcasts trending badges and play counts to trigger the conformist surge, capturing short-run adoption revenue at the cost of compressing the lifecycle. A patient platform under-reveals: it suppresses popularity metrics to preserve the discovery phase. At $\delta_P = \bar{\delta}_P$, the platform's optimum coincides with the welfare planner's. The comparative statics in part (ii) pin down when each regime applies. Strong snob aversion makes the conformist surge self-defeating: snobs flee rapidly, the adoption spike is brief, and the lifecycle collapse severe, so even moderately impatient platforms prefer restraint. Abundant snobs shrink the conformist mass, reducing the short-run gain from over-revealing relative to the lifecycle extension benefit.

These comparative statics generate a cross-platform prediction. Platforms in identity-heavy markets (high α, high λ) should under-reveal at almost any level of patience – and indeed, Bandcamp and Letterboxd display minimal popularity information. Platforms in mass markets (low α, low λ) should over-reveal unless unusually patient – and indeed, TikTok and Instagram aggressively broadcast trending metrics.[38] The model predicts faster product cycling on over-revealing platforms, holding quality constant.

Remark 5 (Visibility Outside the Reversal Regime). When $\alpha < \bar{\alpha}(\lambda)$, optimal φ^* is closer to full transparency but generically remains interior. Full transparency is optimal only when $\lambda = 0$ or $\alpha = 0$. In all cases, the platform obscures *popularity* information, not quality information, since private signals are unaffected.

Corollary 6 (Optimal Visibility Decreasing in Product Quality)**.** *In the regime where the quality-duration reversal holds, the optimal visibility satisfies $d\varphi^*/d\theta < 0$.*

Better products attract conformists faster, making the lifecycle compression cost of visibility larger. The result yields a cross-category design prescription: a streaming platform should display fewer popularity metrics for premium content (curated playlists, acclaimed releases) than for mass content (viral hits, trending tracks), because the acceleration externality is strongest where product quality is highest.[39]

The firm's scarcity strategy (Proposition 8) and the platform's opacity both slow conformist entry, but through different channels: the firm restricts *quantity*; the platform restricts *information*. The two instruments are complements: when the firm's cap $\bar{n}$ is below

[38] The platform's effective δ_P reflects its revenue model: subscription platforms (Substack, Patreon) are effectively patient; impression-based platforms (Instagram, TikTok) are effectively impatient.

[39] If a platform could condition φ on realized θ, agents might infer quality from the platform's visibility choice, creating a signaling problem outside the current model.

the unconstrained peak but still high enough that conformists would enter under full transparency, both lowering $\bar{n}$ and lowering φ independently delay conformist entry.[40] Supreme's limited drops implement the quantity cap; Berghain's combination of restricted entry and a photography ban illustrates the complementarity.

9 Testable Predictions and Empirical Evidence

We organize the model's predictions into consumer-side and supply-side, noting where they run *opposite* to standard models.

9.1 Consumer Equilibrium Predictions

Prediction 1: Early-adopter exit precedes the adoption peak. Snob participation should begin declining *before* total adoption peaks (Proposition 3), driven by crowding rather than product deterioration. Standard diffusion models predict that early adopters persist through the growth phase; here, they leave because of it. The same pattern appears in influencer markets: early followers of a niche creator disengage as the audience goes mainstream (Schoenmueller et al., 2021).[41] A test requires tracking cohort-level churn, with identifying variation from shocks to mainstream visibility.

Prediction 2: Rapid growth shortens lifecycles. Products with rapid initial growth should have *shorter* lifecycles at equal quality (Proposition 4). Berger & Le Mens (2009) show that cultural tastes adopted more quickly die out faster; Bellezza (2023) documents that luxury products lose status as they diffuse. The test regresses lifecycle duration on initial adoption speed, controlling for quality proxies; the model predicts a negative coefficient stronger in identity-relevant categories.

Prediction 3: The quality-duration correlation flips sign across markets. The correlation should be positive where snob preferences are weak and negative where they are strong (Proposition 4). The discriminating test is cross-category: the quality-duration slope should be positive for utilitarian goods and negative for identity goods.[42]

9.2 Supply-Side Predictions

Prediction 4: Visibility accelerates cycling, and the distortion depends on platform patience. Platforms displaying trending indicators should observe faster product cycling than platforms suppressing such information (Proposition 11). Whether a profit-maximizing platform over-

[40] When the cap fully deters conformists even under full transparency, the two instruments are substitutes.

[41] Products that *persisted* – Pabst Blue Ribbon, vinyl records – are those whose low perceived quality or format inconvenience screened out conformists.

[42] The acceleration of chart turnover coincides with the rise of platform features – algorithmic playlists, trending indicators – that Proposition 11 predicts should compress lifecycles.

or under-reveals depends on its effective patience (Proposition 12). Field experiments on popularity displays (Tucker & Zhang, 2011) confirm that toggling visibility shifts demand heterogeneously across product types; Instagram's experiment hiding like counts (Mosseri, 2019) provides a natural φ-variation design.[43]

Prediction 5: Broad advertising can backfire. Broad-reach advertising should shorten lifecycles in markets with strong originality preferences (Lemma 3), even when it increases short-run adoption. The identifying variation could come from quasi-random exposure shocks, with duration as the outcome and category identity-relevance as the moderator.

Prediction 6: Information precision can hurt. More precise quality information can *reduce* snob welfare and shorten lifecycles (Proposition 7), consistent with Berger & Le Mens (2009). In social learning models, more precise information unambiguously improves welfare; here, better quality assessment accelerates the composition shift that destroys discovery value.

10 Conclusion

This paper develops a dynamic model of markets where consumers disagree about whether popularity is a feature or a bug – fashion, restaurants, music, social platforms, and other taste-driven categories. Snobs seek exclusivity; conformists seek validation; their interaction generates endogenous boom-bust lifecycles and inverts the quality-duration relationship. An impossibility result establishes that this reversal requires mixed-sign adoption externalities, making it a diagnostic for the presence of opposing preferences in a market.

Two supply-side actors can respond to these dynamics. A seller benefits from slowing the conformist transition through quantity restriction – artificial scarcity, snob-targeted advertising, and type-based pricing. A platform benefits from slowing it through information restriction – hiding popularity metrics to delay conformist entry. These instruments are *complements*: a firm's scarcity strategy is more profitable when the platform also restricts visibility, and vice versa. The joint optimization problem – where firm and platform simultaneously choose quantity and information policies – remains open.

The framework leaves open several directions worth pursuing. A continuous type distribution $\gamma_i \sim F$ would smooth phase transitions; the core mechanism survives as long as some agents have $\gamma_i < 0$. Overlapping lifecycles, where snobs fleeing product A pioneer product B, would endogenize reservation utilities and connect the single-product analysis to portfolio dynamics. The full mechanism design problem with unobservable types remains open: timing-based screening appears feasible (Remark 4), but a complete treatment would require jointly optimizing visibility, pricing, and information release across types. Finally, the visibility parameter φ is a special case of a richer information design problem: a platform

[43] The lifecycle effect of visibility should be larger for identity-relevant categories (high α) than for utilitarian categories (low α).

could condition visibility on product age, adopter count, or category, targeting opacity where the acceleration externality is strongest.

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.

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

Bergemann, D., & Bonatti, A. (2019). Markets for information: An introduction. *Annual Review of Economics*, 11, 85-107.

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). The role of social learning in influencer marketing. Working Paper, The Wharton School.

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.

Gardete, P. M., & Hunter, M. (2024). Multiattribute search: Empirical evidence and information design. *Marketing Science*, 43(5), 1062-1080.

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

Hagiu, A., & Wright, J. (2024). Optimal discoverability on platforms. *Management Science*, 70(11), 7770-7790.

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.

Iyer, G., & Katona, Z. (2016). Competing for attention in social communication markets. *Management Science*, 62(8), 2304-2320.

Iyer, G., & Zhong, Z. (2022). Pushing notifications as dynamic information design. *Marketing Science*, 41(1), 51-72.

Jerath, K., & Ren, Q. (2021). Consumer rational (in)attention to favorable and unfavorable product information, and firm information design. *Journal of Marketing Research*, 58(2), 343-362.

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. (2023). Information design of online platforms. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC '23)*, 912–931.

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.

Kuksov, D., & Xie, Y. (2012). Competition in a status goods market. *Journal of Marketing Research*, 49(5), 609-623.

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

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.

Mosseri, A. (2019). Announcement: Testing private like counts. *Instagram Blog*, April 30. Expanded to seven countries July 2019; to the U.S. November 2019.

Nistor, C., Selove, M., & Villas-Boas, J. M. (2024). Influencer authenticity: To grow or to monetize. *Management Science*, forthcoming.

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

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

Shi, Z., Srinivasan, K., & Zhang, K. (2023). Design of platform reputation systems: Optimal information disclosure. *Marketing Science*, 42(3), 500-520.

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.

Warren, C., Batra, R., Loureiro, S. M. C., & Bagozzi, R. P. (2019). Brand coolness. *Journal of Marketing*, 83(5), 36-56.

Yao, Y. (2024). Dynamic persuasion and strategic search. *Management Science*, 70(12), 8541-8561.

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

A Main Proofs

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. Preference parameters: $\alpha, \beta > 0$. Reservation utilities: $c_C > c_S \geq 0$. Snob mass: $\lambda \in (0, 1)$.

A.1 Proof of Lemma 1 (Threshold Strategies)

Proof. Fix adoption mass $n \in [0, 1]$ and posterior $\mu \in (0, 1)$.

Snob's problem. From the Bellman equation (11), a snob adopts in period t if the adoption payoff exceeds the waiting payoff:

$$\underbrace{\mu - \alpha n}_{\text{adopt}} + \delta \mathbb{E}[V^S \mid \text{adopt}] \geq \underbrace{c_S}_{\text{outside option}} + \delta \mathbb{E}[V^S \mid \text{wait}] \tag{26}$$

Continuation value cancellation. Under per-period adoption (agents re-optimize each period independently), today's adoption does not constrain future choices, so the continuation value is identical whether the agent adopts or waits. The terms cancel, yielding:

$$\mu - \alpha n \geq c_S \quad \Rightarrow \quad \mu \geq \alpha n + c_S \equiv \underline{\mu}^S(n) \tag{27}$$

[44]

Conformist's problem. By analogous reasoning, a conformist adopts iff:

$$\mu + \beta n \geq c_C \quad \Rightarrow \quad \mu \geq c_C - \beta n \equiv \underline{\mu}^C(n) \tag{28}$$

Properties. The thresholds have derivatives $\frac{d\underline{\mu}^S}{dn} = \alpha > 0$ and $\frac{d\underline{\mu}^C}{dn} = -\beta < 0$. The thresholds cross at $n^\times = (c_C - c_S)/(\alpha + \beta)$, which lies in $(0, 1)$ when $\alpha + \beta > c_C - c_S$ and $c_C > c_S$ (both hold by Assumption 2). □

A.2 Proof of Proposition 1 (Type-Specific Responses)

Proof. We establish each part using the continuous signal framework.

Part (a): Snob exit region. With continuous signals satisfying Assumption 1, posteriors have full support on $(0, 1)$, so $1 - G(\underline{\mu}^S(n)) > 0$ for any threshold below 1. The relevant

[44] With irreversible adoption, the threshold includes an option-value premium $\omega(\hat{\theta}) > 0$, but this shifts the intercept without changing the slope $\partial \underline{\mu}^S / \partial n = \alpha$. All qualitative results survive; see Chamley (2004).

concept is therefore *mass*: snob adoption becomes negligible. Define snob participation as $\pi^S(n) = 1 - G(\alpha n + c_S)$. Since G is strictly increasing and $\alpha n + c_S$ is increasing in n:

$$\frac{d\pi^S}{dn} = -\alpha g(\alpha n + c_S) < 0 \tag{29}$$

As $n \to \bar{n}^S \equiv (1 - c_S)/\alpha$, the threshold $\alpha n + c_S \to 1$. Since posteriors have support on $(0, 1)$, $G(\mu) \to 1$ as $\mu \to 1^-$, so $\pi^S(n) \to 0$: snob adoption mass vanishes. For $n > \bar{n}^S$, the threshold exceeds 1 and snob adoption is exactly zero (no posterior can exceed 1).

Part (b): Conformist entry region. Define conformist participation as $\pi^C(n) = 1-G(c_C-\beta n)$. Since $c_C - \beta n$ is decreasing in n:

$$\frac{d\pi^C}{dn} = \beta g(c_C - \beta n) > 0 \tag{30}$$

For $n > c_C/\beta$, the threshold $c_C - \beta n < 0$, so $\pi^C(n) = 1$ (all conformists adopt since all posteriors exceed a negative threshold). More generally, π^C is increasing in n: higher adoption encourages more conformists.

Part (c): Overlap. In the region $\underline{n}^C < n < \bar{n}^S$ (where $\underline{n}^C(\hat{\theta}_t) = (c_C - \bar{\mu}(\hat{\theta}_t))/\beta$ and $\bar{n}^S = (1-c_S)/\alpha$ are defined in the proposition), snob participation is declining ($d\pi^S/dn < 0$, Part (a)) while conformist participation is increasing ($d\pi^C/dn > 0$, Part (b)). This region is non-empty when $\underline{n}^C < \bar{n}^S$, i.e., $(c_C - \bar{\mu})/\beta < (1 - c_S)/\alpha$, which holds when the prior is not too pessimistic (so $\bar{\mu}$ is not too small) and Assumption 2 holds. □

A.3 Proof of Lemma 2 (Equilibrium Existence)

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

$$\Phi(n) = \lambda \cdot [1 - G(\underline{\mu}^S(n))] + (1 - \lambda) \cdot [1 - G(\underline{\mu}^C(n))] \tag{31}$$

where $G(\mu) \equiv G(\mu; \hat{\theta}_t, \theta)$ is the CDF of posteriors.

Step 1: Setup. Fix public belief $\hat{\theta} \in (0, 1)$. Let $G(\mu; \hat{\theta}) = \Pr(\mu_i \leq \mu \mid \hat{\theta})$ denote the CDF of posteriors. The best-response mapping is:

$$\Phi(n; \hat{\theta}) = \lambda[1 - G(\alpha n + c_S; \hat{\theta})] + (1 - \lambda)[1 - G(c_C - \beta n; \hat{\theta})] \tag{32}$$

Step 2: Continuity and range. Under Assumption 1, $G(\cdot; \hat{\theta})$ is continuously differentiable with density $g > 0$ on $(0, 1)$. Since thresholds are linear in n, Φ is continuous. As a weighted sum of probabilities, $\Phi : [0, 1] \to [0, 1]$.

Step 3: Boundary conditions. At $n = 0$: since $c_S < 1$ and posteriors have full support, $\Phi(0) \geq \lambda[1-G(c_S)] > 0$. At $n = 1$: if $\alpha+c_S \geq 1$, the snob term vanishes and $\Phi(1) \leq 1-\lambda < 1$; if $\alpha + c_S < 1$, then $G(\alpha + c_S) \in (0, 1)$, giving $\Phi(1) < 1$.

Step 4: Existence. $\Phi : [0,1] \to [0,1]$ is continuous with $\Phi(0) > 0$ and $\Phi(1) < 1$. By Brouwer's fixed point theorem, $\exists n^* \in (0,1)$ with $\Phi(n^*) = n^*$. □

A.4 Proof of Proposition 2 (MPE Existence)

Proof. We construct the Markov Perfect Equilibrium following Smith & Sørensen (2000).

Step 1: Strategy and state. The public belief $\hat{\theta}_t$ is the sole payoff-relevant aggregate state: $n_t = n^*(\hat{\theta}_t)$ and $\hat{\theta}_{t+1}$ depends only on $(\hat{\theta}_t, \tilde{n}_t)$. The value function satisfies the Bellman equation (11). Since continuation values cancel (they appear identically in both branches), the optimal policy is to adopt iff $\mu_i + u^\tau(n^*(\hat{\theta}_t)) \geq c_\tau$, recovering the myopic thresholds from Lemma 1.

Step 2: Belief consistency. The belief updating equation involves the equilibrium adoption function $n^*(\hat{\theta})$. Specifically:

$$f(\tilde{n} \mid \theta, \hat{\theta}_t) = \frac{1}{\sigma_\varepsilon} \phi \left(\frac{\tilde{n} - n^*_\theta(\hat{\theta}_t)}{\sigma_\varepsilon} \right) \tag{33}$$

where $n^*_\theta(\hat{\theta}_t) = \lambda[1 - G(\alpha n^*(\hat{\theta}_t) + c_S; \hat{\theta}_t, \theta)] + (1-\lambda)[1 - G(c_C - \beta n^*(\hat{\theta}_t); \hat{\theta}_t, \theta)]$ is the realized adoption when agents expect $n^*(\hat{\theta}_t)$ but true quality is θ. The distribution $G(\cdot; \hat{\theta}_t, \theta)$ reflects the actual posterior distribution, which depends on the true θ through signal realizations. Note that $n^*_H(\hat{\theta}_t) > n^*_L(\hat{\theta}_t)$ by MLRP: high quality generates better signals, hence higher posteriors and more adoption. The equilibrium requires both:

(a) Within-period: $n^* = \Phi(n^*; \hat{\theta}_t)$ (Lemma 2)

(b) Across-period: $\hat{\theta}_{t+1} = B(\hat{\theta}_t, \tilde{n}_t; n^*_H, n^*_L)$ where B is Bayesian updating

Both conditions are satisfied simultaneously: (a) holds by Brouwer's theorem for each $\hat{\theta}_t$; (b) is then well-defined given the solution to (a).

Step 3: Off-path beliefs. Since $\varepsilon_t \sim N(0, \sigma^2_\varepsilon)$ has full support, any observed $\tilde{n}_t \in \mathbb{R}$ has positive probability under both $\theta = H$ and $\theta = L$. Thus Bayes' rule applies for all observations and off-path beliefs are uniquely pinned down.

Step 4: Contraction and existence. Let $\mathcal{B}$ denote the Banach space of bounded, measurable functions $V : (0,1) \times (0,1) \to \mathbb{R}$ (mapping $(\mu, \hat{\theta})$ pairs to values), equipped with the sup-norm $\|V\|_\infty = \sup_{(\mu,\hat{\theta})} |V(\mu, \hat{\theta})|$. This is a complete metric space.

Define the Bellman operator $\mathcal{T}$ by:

$$(\mathcal{T}V)(\mu, \hat{\theta}) = \max\Big\{ \mu + u^\tau(n^*(\hat{\theta})) + \delta \int V(\mu', \hat{\theta}')\, dF(\mu', \hat{\theta}' \mid \hat{\theta}),$$
$$c_\tau + \delta \int V(\mu', \hat{\theta}')\, dF(\mu', \hat{\theta}' \mid \hat{\theta}) \Big\} \tag{34}$$

where $F(\cdot \mid \hat{\theta})$ is the joint distribution of next-period posteriors and beliefs given current state $\hat{\theta}$.

Monotonicity: If $V \geq W$ pointwise, then $\mathcal{T}V \geq \mathcal{T}W$ (the max of larger quantities is larger).

Discounting: For any constant $c \geq 0$: $\mathcal{T}(V + c) \leq \mathcal{T}V + \delta c$, since the continuation value term is multiplied by $\delta < 1$.

By Blackwell's sufficient conditions, $\mathcal{T}$ is a contraction with modulus δ. By the Banach Fixed Point Theorem, $\mathcal{T}$ has a unique fixed point $V^* \in \mathcal{B}$. The MPE combines the static fixed point $n^*(\hat{\theta})$ (Lemma 2), myopic threshold optimality (Step 1), consistent belief transitions (Steps 2–3), and the unique value function (Step 4). □

A.5 Proof of Uniqueness (Assumption 3)

Proof. We characterize when the best-response mapping Φ has multiple fixed points.

Step 1: Derivative of Φ. From (31), differentiating with respect to n:

$$\Phi'(n) = -\lambda\alpha \cdot g(\underline{\mu}^S(n)) + (1-\lambda)\beta \cdot g(\underline{\mu}^C(n)) \tag{35}$$

where $g(\mu) > 0$ is the posterior density at μ. The first term is negative (snobs respond negatively to higher n), the second is positive (conformists respond positively). Note that Φ' can be positive, negative, or zero depending on which effect dominates at each n.

Step 2: Sufficient condition for uniqueness. If $|\Phi'(n)| < 1$ for all $n \in [0, 1]$, then Φ is a contraction on $[0, 1]$. By the Banach fixed point theorem, the fixed point is unique. The condition $\kappa = \max_n |\Phi'(n)| < 1$ suffices.

Step 3: Sufficient condition for multiplicity. We show that under certain conditions, $h(n) = \Phi(n) - n$ has at least three zeros. Recall $h(0) > 0$ and $h(1) < 0$ (Steps 4–5 of the existence proof). By the intermediate value theorem, h has at least one zero. To show the existence of three zeros, we construct conditions under which h has a local minimum below zero followed by a local maximum above zero (or vice versa) in the interior.

Suppose there exist $n_1, n_2 \in (0, 1)$ with $n_1 < n_2$ such that: (a) $h(n_1) < 0$ (i.e., $\Phi(n_1) < n_1$), and (b) $h(n_2) > 0$ (i.e., $\Phi(n_2) > n_2$). Then: since $h(0) > 0$ and $h(n_1) < 0$, IVT gives a zero in $(0, n_1)$. Since $h(n_1) < 0$ and $h(n_2) > 0$, IVT gives a zero in (n_1, n_2). Since $h(n_2) > 0$ and $h(1) < 0$, IVT gives a zero in $(n_2, 1)$. This gives at least three distinct zeros.

Such n_1, n_2 exist when $\Phi'(n) > 1$ over a sufficiently wide interval. Specifically, if $\Phi' > 1$ on an interval $[a, b] \subset (0, 1)$, and the first zero n_1 satisfies $n_1 < a$, then h increases on $[a, b]$ faster than the identity function, pushing h positive on $[a, b]$ provided $b - a$ is large enough relative to $|h(a)|$. Conversely, when $\kappa < 1$, $h' = \Phi' - 1 < 0$ everywhere, so h is strictly decreasing, yielding exactly one zero. This motivates Assumption 3.[45]

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

Step 4: Stability. At a fixed point n^* with $\Phi'(n^*) < 1$, tatonnement dynamics $n_{k+1} = \Phi(n_k)$ converge locally to n^* (stable). At a fixed point with $\Phi'(n^*) > 1$, the dynamics diverge (unstable). When three fixed points exist, the outer two have $\Phi' < 1$ (stable) and the middle one has $\Phi' > 1$ (unstable). □

A.6 Proof of Proposition 3 (Lifecycle Characterization)

Proof. We construct the equilibrium lifecycle in four parts: (1) belief monotonicity under $\theta = H$ (Step 3); (2) monotonicity of snob adoption in Phase I (Step 7); (3) existence and characterization of a unique peak (Step 5); and (4) eventual snob exit (Step 6). Steps 1–2 and 4 establish the phase definitions.

Step 1: Initial conditions. At $t = 0$, public belief is $\hat{\theta}_0 = p$ (the prior). No adoption history exists, so agents base decisions purely on private signals.

Step 2: Phase I (Growth). Consider adoption at $t = 0$. The snob threshold is $\underline{\mu}^S(n) = \alpha n + c_S$. At $n = 0$, $\underline{\mu}^S(0) = c_S$. Snobs with posterior $\mu > c_S$ adopt. Since $g(\mu) > 0$ on $(0, 1)$ by Assumption 1, and $c_S < 1$, we have $\Pr(\mu > c_S) > 0$. Thus snob adoption is positive: $n_0^S = \lambda \cdot [1 - G(c_S)] > 0$.

For conformists, the threshold is $\underline{\mu}^C(0) = c_C > c_S$. If c_C is large enough that $\Pr(\mu > c_C)$ is small at $t = 0$, conformist adoption is negligible. Formally, if $c_C > \mathbb{E}[\mu \mid \hat{\theta}_0 = p]$, most conformists wait.

Step 3: Belief updating. Observing adoption $n_0 > 0$, the public belief updates via Bayes' rule. Let $n_0^H = \mathbb{E}[n_0 \mid \theta = H]$ and $n_0^L = \mathbb{E}[n_0 \mid \theta = L]$. Under MLRP, $n_0^H > n_0^L$ (high quality generates more favorable signals, hence more adoption). The likelihood ratio update is:

$$\frac{\hat{\theta}_1}{1 - \hat{\theta}_1} = \frac{\hat{\theta}_0}{1 - \hat{\theta}_0} \cdot \frac{f(n_0 \mid H)}{f(n_0 \mid L)} \tag{36}$$

where $f(\cdot \mid \theta)$ is the density of observed adoption. If $\theta = H$, beliefs drift upward on average.

Step 4: Formal phase definitions. We define phases without arbitrary ϵ:

- *Phase I (Growth):* Periods $t \in \{0, 1, \dots, t_1 - 1\}$ where $t_1 = \min\{t \geq 1 : n_t^C > n_t^S\}$ is the first period where conformist adoption exceeds snob adoption. In Phase I, $n_t^S > n_t^C$ and $n_{t+1} > n_t$.

- *Phase II (Peak):* Periods $t \in \{t_1, \dots, t^*\}$ where $t^* = \arg\max_{t \geq 0} n_t$ is the period of maximum adoption. In Phase II, both types have positive adoption: $n_t^S > 0$ and $n_t^C > 0$.

- *Phase III (Decline):* Periods $t \in \{t^* + 1, \dots, T\}$ where $T = \min\{t > t^* : n_t^S < \epsilon\}$ is the first period after peak where snob adoption falls below ϵ. In Phase III, snob

An explicit Gaussian example appears in the Online Appendix.

participation is declining: $n^S_{t+1} < n^S_t$.

These definitions are sharp: t_1, t^*, and T are well-defined stopping times depending only on the adoption path $\{n_t\}$ and its decomposition $\{n^S_t, n^C_t\}$.

Step 5: Phase II (Peak) characterization. Both types are active. At t^*, $n_{t^*} \geq n_{t^*-1}$ and $n_{t^*+1} \leq n_{t^*}$.

The equilibrium adoption satisfies $n_t = \Phi(n_t; \hat{\theta}_t)$. By the implicit function theorem:

$$\frac{dn_t}{d\hat{\theta}_t} = \frac{\partial\Phi/\partial\hat{\theta}}{1 - \partial\Phi/\partial n} \tag{37}$$

The denominator $1 - \Phi'(n_t) > 0$ at any stable equilibrium. The numerator $\partial\Phi/\partial\hat{\theta} > 0$ under MLRP: higher $\hat{\theta}$ shifts posteriors rightward, increasing adoption. Thus $dn_t/d\hat{\theta}_t > 0$. Note this is a within-period comparative static; the dynamic belief path $\{\hat{\theta}_t\}$ is stochastic and need not be monotone. The peak of *snob* adoption occurs at t^*; total adoption n_t may continue to rise after t^* as conformists sustain the installed base.

Step 6: Phase III (Snob exit). We prove that snob adoption vanishes in finite time when $\theta = H$ and α is in the reversal region, establishing that $T < \infty$.

Snob adoption at time t is $n^S_t = \lambda[1 - G(\alpha n_t + c_S; \hat{\theta}_t)]$. The snob threshold $\underline{\mu}^S(n_t) = \alpha n_t + c_S$ is increasing in total adoption n_t. Snob adoption reaches zero when $\underline{\mu}^S(n_t) \geq 1$, i.e., $n_t \geq \bar{n}^S \equiv (1 - c_S)/\alpha$.

Under $\theta = H$, beliefs satisfy $\mathbb{E}[\hat{\theta}_{t+1}|\hat{\theta}_t, \theta = H] > \hat{\theta}_t$ (the martingale property of posterior means under the true state). As $\hat{\theta}_t \to 1$, the posterior distribution concentrates near 1, and adoption by both types increases: both thresholds are exceeded by nearly all agents. The equilibrium $n^*(\hat{\theta}_t) \to 1$ as $\hat{\theta}_t \to 1$, since $\Phi(n; \hat{\theta}) \to \lambda + (1-\lambda) = 1$ when all posteriors exceed all thresholds.

Case 1: $\alpha + c_S \geq 1$ (equivalently, $\bar{n}^S \leq 1$). Since $n^*(\hat{\theta}_t) \to 1 > \bar{n}^S$, there exists t' such that $n_{t'} > \bar{n}^S$, at which point $\underline{\mu}^S(n_{t'}) > 1$ and snob adoption is exactly zero. Thus $T \leq t' < \infty$.

Case 2: $\alpha + c_S < 1$ (equivalently, $\bar{n}^S > 1$). The snob threshold never reaches 1, so snob adoption is always positive. However, snob adoption still declines toward zero as n_t rises, and the reversal manifests: $n^{S,H}_t$ falls faster than $n^{S,L}_t$.[46]

Low-quality comparison. Under $\theta = L$, beliefs drift downward ($\mathbb{E}[\hat{\theta}_{t+1}|\hat{\theta}_t, L] < \hat{\theta}_t$). Adoption stays low: the conformist threshold $c_C - \beta n_t$ remains high, limiting total adoption. With low n_t, the snob threshold $\alpha n_t + c_S$ stays low, and snobs continue to participate. Thus T^L is large (snobs persist in a quiet niche), establishing $T^H < T^L$ in the reversal region.

Step 7: Existence of growth phase. We show that $\alpha < \bar{\alpha}_{\text{growth}}$ implies $n_{t+1} > n_t$ in Phase

[46] When $\alpha + c_S < 1$, snob adoption declines toward zero without reaching it exactly. This case corresponds to small α; it is compatible with the reversal region only when $\beta(1-\lambda)/\lambda < 1 - c_S$.

I. The equilibrium adoption satisfies $n_t = \Phi(n_t; \hat{\theta}_t)$. By the implicit function theorem:

$$n_{t+1} - n_t = \frac{dn}{d\hat{\theta}} \cdot (\hat{\theta}_{t+1} - \hat{\theta}_t) + O((\Delta\hat{\theta})^2) \tag{38}$$

where $dn/d\hat{\theta} > 0$ (higher beliefs increase adoption, as shown in Step 5). When $\theta = H$, $\mathbb{E}[\hat{\theta}_{t+1} - \hat{\theta}_t | \theta = H] > 0$ (Step 3), so $\mathbb{E}[n_{t+1} - n_t] > 0$.

It remains to verify that $dn/d\hat{\theta}$ is bounded, so that the growth is well-defined. From the IFT expression in Step 5, $dn/d\hat{\theta}$ is finite provided $1 - \Phi'(n_t) > 0$, which holds under Assumption 3. The denominator $1 - \Phi'(n) = 1 + \lambda\alpha g^S - (1-\lambda)\beta g^C$ is bounded below by $1 - \kappa > 0$ where $\kappa = \max_n |\Phi'(n)| < 1$.

The growth condition $\alpha < \bar{\alpha}_{\text{growth}}$ ensures that positive belief drift translates into positive adoption growth throughout Phase I. Under Assumption 3 ($\kappa < 1$), $1 - \Phi'(n) > 0$ uniformly, so the IFT gives bounded $dn/d\hat{\theta}$. The additional restriction $\lambda\alpha g_{\max} < 1$ provides a sufficient condition on α alone:

$$\bar{\alpha}_{\text{growth}} \equiv \frac{\sqrt{2\pi \text{Var}(\mu \mid \hat{\theta}_0)}}{\lambda} \tag{39}$$

□

A.7 Proof of Proposition 5 (Impossibility Under Homogeneous Preferences)

Proof. Consider any adoption model where flow utility is $u_i(\mu, n) = v(\mu) + \gamma_i w(n)$ with $w'(n) > 0$, and lifecycle duration $T(\theta) \equiv \sup\{t : n_t^-(\theta) > 0\}$, where n_t^- denotes adoption by agents with $\gamma_i < 0$ (the "snob" type).

Case (i): $\gamma_i \geq 0$ for all i (positive externality). There are no agents with $\gamma_i < 0$, so $n_t^-(\theta) = 0$ for all t and both states, giving $T^H = T^L = 0$ trivially. For completeness under the alternative definition $T(\theta) = \sup\{t : n_t(\theta) > 0\}$, we show $n_t^H \geq n_t^L$ along expected paths by induction. Agent i adopts iff $v(\mu_i) + \gamma_i w(n) \geq c_i$, i.e., $\mu_i \geq v^{-1}(c_i - \gamma_i w(n)) \equiv \underline{\mu}_i(n)$. Since $\gamma_i \geq 0$, thresholds are decreasing in n (strategic complements). *Base case* ($t = 0$): beliefs are identical ($\hat{\theta}_0 = p$), but MLRP ensures that the distribution of posteriors under H stochastically dominates the distribution under L, so $\Phi(n; \hat{\theta}_0, H) \geq \Phi(n; \hat{\theta}_0, L)$ at every n. Since thresholds are decreasing in n (complements), the fixed point satisfies $n_0^H \geq n_0^L$. *Inductive step*: suppose $n_s^H \geq n_s^L$ for all $s \leq t-1$. By the same logic as Lemma 5, higher realized adoption under H generates more favorable observed signals $\tilde{n}_{t-1}$; combined with MLRP, Bayesian updating yields $\hat{\theta}_t^H \geq \hat{\theta}_t^L$ in expectation. A higher public belief shifts the posterior distribution rightward, so $\Phi(n; \hat{\theta}_t^H, H) \geq \Phi(n; \hat{\theta}_t^L, L)$ at every n. With strategic complements, the fixed point $n_t^H \geq n_t^L$. Since $\{t : n_t^H > 0\} \supseteq \{t : n_t^L > 0\}$ along expected paths, $T^H \geq T^L$.

Case (ii): $\gamma_i \leq 0$ for all i (negative externality). All agents have $\gamma_i \leq 0$, so $n_t^- = n_t$

and $T(\theta) = \sup\{t : n_t(\theta) > 0\}$. Thresholds $\underline{\mu}_i(n) = v^{-1}(c_i + |\gamma_i| w(n))$ are now increasing in n (strategic substitutes). The induction again works through *beliefs*, not through complementarity in n. *Base case*: at $t = 0$, MLRP ensures $\Phi(n; \hat{\theta}_0, H) \geq \Phi(n; \hat{\theta}_0, L)$ at every n – the upward shift in the posterior distribution under H raises the mass above any fixed threshold, regardless of how that threshold moves with n. Because $\Phi(\cdot; \hat{\theta}_0, H)$ lies weakly above $\Phi(\cdot; \hat{\theta}_0, L)$, the highest fixed point satisfies $n_0^H \geq n_0^L$.[47] *Inductive step*: given $n_s^H \geq n_s^L$ for $s \leq t-1$, the observed signal $\tilde{n}_{t-1}$ is stochastically higher under H, yielding $\hat{\theta}_t^H \geq \hat{\theta}_t^L$ in expectation. A higher public belief shifts posteriors rightward, so $\Phi(n; \hat{\theta}_t^H, H) \geq \Phi(n; \hat{\theta}_t^L, L)$ at every n. The fixed point $n_t^H \geq n_t^L$. Strategic substitutes do not break the argument: the channel is $H \to$ more adoption $\to$ more favorable belief updates $\to$ rightward-shifted posteriors $\to$ higher Φ at each n. Since all thresholds move in the same direction (all increase with n), no subset of agents is *attracted* by high adoption; higher quality simply raises the level at which the decreasing best-response crosses the 45-degree line. Hence $T^H \geq T^L$.

Case (iii) follows: with uniform sign of γ_i, either (i) or (ii) applies. The reversal $T^H < T^L$ requires some agents with $\gamma_i > 0$ (whose entry is accelerated by high adoption) and others with $\gamma_j < 0$ (whose exit is triggered by that entry). The key mechanism absent in Cases (i) and (ii) is that under H, the $\gamma_i > 0$ agents' entry raises n, which pushes $\gamma_j < 0$ agents' thresholds above achievable posteriors *faster* than the improved beliefs can compensate, shortening the lifecycle despite higher quality. □

A.8 Belief Dominance Under Higher Quality

Lemma 5 (Stochastic Dominance of Beliefs)**.** *Let $\hat{\theta}_t^H$ and $\hat{\theta}_t^L$ denote the expected public belief at period t under $\theta = H$ and $\theta = L$ respectively, starting from a common prior $\hat{\theta}_0 = p$. Then $\hat{\theta}_t^H \geq \hat{\theta}_t^L$ for all $t \geq 1$, with strict inequality for $t \geq 1$.*

Proof. By induction. At $t = 0$, $\hat{\theta}_0^H = \hat{\theta}_0^L = p$. Suppose $\hat{\theta}_{t-1}^H \geq \hat{\theta}_{t-1}^L$. Under MLRP, $n_H^*(\hat{\theta}) > n_L^*(\hat{\theta})$ for any $\hat{\theta}$ (high quality generates stochastically better signals, hence higher posteriors and more adoption). Higher realized adoption produces a more favorable observed signal $\tilde{n}_{t-1}$, which raises the posterior via Bayes' rule (9). Hence $\mathbb{E}[\hat{\theta}_t \mid \theta = H, \hat{\theta}_{t-1}^H] > \mathbb{E}[\hat{\theta}_t \mid \theta = L, \hat{\theta}_{t-1}^L]$, giving $\hat{\theta}_t^H > \hat{\theta}_t^L$. □

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).

[47] With strategic substitutes, Φ is decreasing in n, so the fixed point is unique and the inequality is strict when the posterior distributions under H and L are distinct.

A.9 Proof of Proposition 4 (Quality-Duration Relationship)

Proof. We prove each part in turn. Part (i) establishes the existence of a threshold via the intermediate value theorem applied to the difference $T^H - T^L$ as a function of α. Part (ii) derives the closed form under the constant-density approximation. Part (iii) verifies the comparative statics hold exactly.

Step 1: Setup. Quality $\theta \in \{L, H\}$ with $v(H) = 1$, $v(L) = 0$. Let T^H and T^L denote cycle durations under high and low quality respectively. We show conditions under which $T^H < T^L$.

Step 2: Adoption under each quality. Under threshold strategies, period-t adoption given public belief $\hat{\theta}_t$ is:

$$n_t(\hat{\theta}_t) = \lambda[1 - G(\alpha n_t + c_S; \hat{\theta}_t)] + (1 - \lambda)[1 - G(c_C - \beta n_t; \hat{\theta}_t)] \tag{40}$$

Higher $\hat{\theta}_t$ shifts the posterior distribution rightward, increasing adoption: $\partial n_t / \partial \hat{\theta}_t > 0$.

Step 3: Belief dynamics under each quality. By Lemma 5, $\hat{\theta}_t^H > \hat{\theta}_t^L$ for all $t \geq 1$:

$$\mathbb{E}[\hat{\theta}_{t+1} - \hat{\theta}_t \mid \theta = H] > 0, \quad \mathbb{E}[\hat{\theta}_{t+1} - \hat{\theta}_t \mid \theta = L] < 0 \tag{41}$$

High quality generates higher adoption on average, which is a favorable signal, causing beliefs to rise. Low quality has the opposite effect.

Step 4: Snob exit condition. Snobs exit when $\underline{\mu}^S(n_t) = \alpha n_t + c_S$ exceeds most posteriors in the population. Define exit time $t^*(\theta)$ as the first period where snob adoption n_t^S begins to decline, i.e., $t^* = \arg\max_t n_t^S$.

Step 5: Two effects of higher quality. Compare paths under $\theta = H$ vs $\theta = L$:

- *Persistence effect*: Under H, beliefs $\hat{\theta}_t$ rise faster, keeping posteriors above the snob threshold longer. This tends to increase t^*.
- *Acceleration effect*: Under H, higher beliefs cause faster conformist entry, raising n_t faster. Since $\underline{\mu}^S(n) = \alpha n + c_S$ is increasing in n, faster adoption raises the threshold faster, triggering earlier snob exit. This tends to decrease t^*.

Step 6: Formalizing the trade-off. Since quality is binary ($\theta \in \{L, H\}$), we compare the two paths directly. Let $\{\hat{\theta}_t^H\}$ and $\{\hat{\theta}_t^L\}$ denote the expected belief paths under $\theta = H$ and $\theta = L$ respectively. Under H, beliefs rise faster: $\hat{\theta}_t^H > \hat{\theta}_t^L$ for all $t \geq 1$ (Step 3).

The key trade-off operates through the equilibrium adoption function $n^*(\hat{\theta})$. Consider the effect of a unit increase in $\hat{\theta}$ on snob adoption. The total effect has two components: (i) the posterior distribution shifts rightward, increasing the mass above $\underline{\mu}^S$; (ii) total adoption n rises, which raises $\underline{\mu}^S = \alpha n + c_S$, reducing snob mass. The net effect on snob adoption depends on whether the belief improvement (i) or the threshold rise (ii) dominates. From the IFT, the equilibrium response of n to beliefs is $dn^*/d\hat{\theta} = (\partial \Phi / \partial \hat{\theta}) / (1 - \Phi'(n^*))$, where

$\partial\Phi/\partial\hat{\theta} > 0$ is the total belief effect on the best-response mapping (Step 5 of the lifecycle proof) and the denominator is positive at stable equilibria.

Step 7: Critical threshold. The persistence effect dominates when snob adoption is increasing in $\hat{\theta}$ (so higher quality, which raises $\hat{\theta}$ faster, keeps snobs active longer). Snob adoption is $n^S(\hat{\theta}) = \lambda[1 - G(\underline{\mu}^S(n^*(\hat{\theta})); \hat{\theta})]$ where $\underline{\mu}^S(n) = \alpha n + c_S$ and $G(\mu; \hat{\theta})$ is the posterior CDF.

Differentiating with respect to $\hat{\theta}$ (suppressing arguments):

$$\frac{dn^S}{d\hat{\theta}} = \lambda \left[\underbrace{-g^S \cdot \alpha \cdot \frac{dn^*}{d\hat{\theta}}}_{\text{threshold effect}(<0)} + \underbrace{\frac{\partial[1 - G(\underline{\mu}^S; \hat{\theta})]}{\partial\hat{\theta}}}_{\text{belief effect}(>0)} \right] \tag{42}$$

The belief effect (the partial derivative of $1 - G$ with respect to $\hat{\theta}$, holding the threshold fixed) is positive: higher $\hat{\theta}$ shifts posteriors rightward, increasing the mass above $\underline{\mu}^S$. Denote the per-unit belief effect as $\lambda \cdot b^S > 0$. Then $dn^S/d\hat{\theta} > 0$ iff $b^S > g^S \cdot \alpha \cdot (dn^*/d\hat{\theta})$.

To obtain a condition on primitives, note that the belief effect on total adoption satisfies $\partial\Phi/\partial\hat{\theta} = \lambda b^S + (1 - \lambda)b^C$, and from the IFT, $dn^*/d\hat{\theta} = (\lambda b^S + (1 - \lambda)b^C)/(1 - \Phi')$. The condition $dn^S/d\hat{\theta} > 0$ then requires:

$$b^S > g^S \cdot \alpha \cdot \frac{\lambda b^S + (1 - \lambda)b^C}{1 - \Phi'} \tag{43}$$

Under the constant-density approximation ($g^S \approx g^C \equiv g^*$ and $b^S \approx b^C \equiv b^*$, where $b^\tau \equiv -\partial G(\underline{\mu}^\tau; \hat{\theta})/\partial\hat{\theta}$), the condition simplifies. With $b^* = g^*$ (the per-unit belief shift equals the density under the approximation) and $\partial\Phi/\partial\hat{\theta} = g^*$, the IFT gives $dn^*/d\hat{\theta} = g^*/(1 - \Phi')$. The condition $b^* > g^* \cdot \alpha \cdot dn^*/d\hat{\theta}$ becomes:

$$1 > \frac{\alpha g^*}{1 - \Phi'} = \frac{\alpha g^*}{1 + \lambda\alpha g^* - (1 - \lambda)\beta g^*} \tag{44}$$

Cross-multiplying ($1 - \Phi' > 0$ at stable equilibria):

$$1 + \lambda\alpha g^* - (1 - \lambda)\beta g^* > \alpha g^* \tag{45}$$

$$1 > (1 - \lambda)g^*(\alpha + \beta) \tag{46}$$

This condition involves the conformist density $g^C = g(\underline{\mu}^C(n^*))$ at the conformist threshold. Reversal occurs when conformists are highly responsive (concentrated near threshold).

Step 8: Deriving $\bar{\alpha}(\lambda)$. Step 7 gives a necessary condition for persistence: $1 > (1 - \lambda)g^C(\alpha + \beta)$, or equivalently $\alpha < 1/((1 - \lambda)g^C) - \beta$. The RHS depends on g^C, which itself depends on α, β through the equilibrium.

To obtain a closed-form approximation, we evaluate at the peak where both thresholds bind for marginal agents. At peak, the derivative of the best-response mapping satisfies $\Phi'(n^*) = -\lambda\alpha g^S + (1-\lambda)\beta g^C \approx 0$, reflecting the approximate balance between marginal snob exit and marginal conformist entry at the adoption maximum. This "flow-balance" condition – which is distinct from the threshold-level condition $\underline{\mu}^S = \underline{\mu}^C$ in Corollary 1 and instead reflects $\Phi'(n^*) \approx 0$ – gives $\lambda\alpha g^S \approx (1-\lambda)\beta g^C$. Under the constant-density approximation ($g^S \approx g^C \equiv g^*$, Assumption 4), this simplifies to:

$$\lambda\alpha \approx (1-\lambda)\beta \tag{47}$$

This balance characterizes the boundary: the aggregate crowding effect on snobs is $\alpha\lambda$, and the aggregate bandwagon effect on conformists is $\beta(1-\lambda)$. When $\alpha\lambda > \beta(1-\lambda)$, crowding dominates and the reversal obtains. Setting them equal:

$$\alpha\lambda = \beta(1-\lambda) \quad \Rightarrow \quad \bar{\alpha}(\lambda) = \beta \cdot \frac{1-\lambda}{\lambda} \tag{48}$$

Step 8a: Exact existence of threshold (part (i)). The density condition from Step 7 ($1 > (1-\lambda)g^C(\alpha+\beta)$) defines an exact threshold. Let $F(\alpha) \equiv (1-\lambda)g^C(n^*(\alpha))(\alpha+\beta)$, where g^C depends on α through the equilibrium $n^*(\alpha)$.

We verify the limiting behavior of F. Note that $n^*(\alpha)$ is continuous in α by the implicit function theorem applied to the fixed-point condition $n^* = \Phi(n^*; \hat{\theta}, \alpha)$, since Φ is C^1 and $|1 - \Phi'(n^*)| > 0$ at stable equilibria. Hence F is continuous. As $\alpha \to 0$: $(\alpha+\beta) \to \beta$, and g^C is bounded (the posterior density at $\underline{\mu}^C(n^*)$ is positive and finite). Thus $F(\alpha) \to (1-\lambda)\beta g^C(n^*(0))$, which is less than 1 when conformist preferences are not too strong (guaranteed by Assumption 3). As $\alpha \to \infty$: $n^*(\alpha) \to 0$ (since $\partial n^*/\partial\alpha < 0$ and adoption is bounded below by 0), so $\underline{\mu}^C(n^*) = c_C - \beta n^* \to c_C$, giving $g^C \to g(c_C; \hat{\theta}) > 0$ (the posterior density at c_C is positive by Assumption 1). Meanwhile $(\alpha+\beta) \to \infty$, so $F(\alpha) \to \infty$. Since F is continuous (by continuity of $n^*(\alpha)$ and g), with $F(0) < 1$ and $F(\alpha) \to \infty$, the intermediate value theorem gives $\bar{\alpha}_{\text{exact}}(\lambda) \in (0, \infty)$ with $F(\bar{\alpha}_{\text{exact}}) = 1$. Below this threshold, persistence dominates ($T^H > T^L$); above it, acceleration dominates ($T^H < T^L$).

Step 8b: Closed form under the constant-density approximation (part (ii)). Under the constant-density approximation ($g^S \approx g^C$), the aggregate balance condition $\alpha\lambda \lessgtr \beta(1-\lambda)$ coincides with the exact condition. This yields the closed form $\bar{\alpha}(\lambda) = \beta(1-\lambda)/\lambda$.

Step 8c: Comparative statics (part (iii)). From the closed form, $\bar{\alpha}(\lambda) = \beta(1-\lambda)/\lambda$:

$$\frac{\partial\bar{\alpha}}{\partial\lambda} = -\frac{\beta}{\lambda^2} < 0, \qquad \frac{\partial\bar{\alpha}}{\partial\beta} = \frac{1-\lambda}{\lambda} > 0 \tag{49}$$

The first says the reversal is easier to trigger (lower $\bar{\alpha}$) when snobs are scarce. The second says stronger conformist attraction raises the bar for reversal, because conformist bandwagon effects offset crowding. These signs hold for the closed-form approximation; for the exact threshold $\bar{\alpha}_{\text{exact}}(\lambda)$, the qualitative direction is the same when the density response is small relative to the direct effect, which holds under Assumption 4. □

A.10 Proof of Proposition 10 (Welfare and Composition)

Proof. We show that social welfare is maximized at an interior market composition.

Step 1: Welfare function. Define social welfare as the population-weighted sum of expected utilities:

$$W(\lambda) = \lambda \cdot EU^S(\lambda) + (1-\lambda) \cdot EU^C(\lambda) \tag{50}$$

where $EU^\tau(\lambda)$ is the expected lifetime utility for type τ given equilibrium dynamics, with outside options normalized to zero.

Step 2: Boundary behavior. As $\lambda \to 0$: the market consists almost entirely of conformists. The conformist-only equilibrium $n_C^{**} = (1-\lambda)[1-G(c_C - \beta n_C^{**})]$ is positive (some conformists adopt based on private signals), but the information externality is absent: without snobs to drive the discovery phase, beliefs update slowly and conformists adopt only when their private signals strongly favor quality. Define $W(0) \equiv EU^C_{\text{no snobs}} > 0$ as the conformist welfare in this no-discovery equilibrium.

As $\lambda \to 1$: nearly all agents are snobs. $W(1) = EU^S(1) > 0$ as before.

Step 3: Interior improvement over boundaries. For intermediate λ, the information externality operates: snobs adopt early, generating informative public signals that allow conformists to assess quality. This improves conformist quality matching relative to the no-snob baseline. Specifically, the expected posterior $\hat{\theta}_{t_1}$ at the time of conformist entry is higher under the snob-conformist lifecycle than under the conformist-only equilibrium, because snob-driven adoption is more informative about quality (snobs adopt based on quality signals, not bandwagon effects). This informational benefit raises $EU^C(\lambda) > EU^C(0)$ for small $\lambda > 0$. Additionally, snobs earn pioneer rents, contributing $\lambda \cdot EU^S > 0$. Thus $W(\lambda) > W(0)$ for small $\lambda > 0$.

Similarly, $W(\lambda^*) > W(1)$ when $\lambda^* = \beta/(\alpha+\beta)$ (the peak-maximizing composition from Proposition 6): the additional surplus from conformist participation $(1-\lambda^*)EU^C$ and the extended discovery phase (improving quality assessment) raise welfare above the snob-only baseline.

Since W is continuous on $[0,1]$ with interior points exceeding both boundary values, the maximum occurs at some $\lambda^{**} \in (0,1)$.

Step 4: First-order condition. At the optimum:

$$\left.\frac{dW}{d\lambda}\right|_{\lambda^{**}} = EU^S - EU^C + \lambda^{**}\frac{dEU^S}{d\lambda} + (1-\lambda^{**})\frac{dEU^C}{d\lambda} = 0 \tag{51}$$

Step 5: Second-order condition. We verify $d^2W/d\lambda^2|_{\lambda^{**}} < 0$. Increasing λ raises total adoption n, which increases crowding costs, so $dEU^S/d\lambda < 0$ near the interior optimum. The compression effect on conformists also gives $dEU^C/d\lambda < 0$ for λ near λ^{**}. Since snobs are directly harmed by their own crowding, $dEU^S/d\lambda < dEU^C/d\lambda$, making the leading terms of $d^2W/d\lambda^2$ negative. Under Assumption 3 ($\kappa < 1$), social preferences are moderate relative to signal precision, ensuring second-order density effects do not overturn this, so λ^{**} is a strict local maximum.[48]

Step 6: Comparative statics of λ^{**}. By the implicit function theorem applied to the FOC $dW/d\lambda|_{\lambda^{**}} = 0$:

$$\frac{d\lambda^{**}}{d\alpha} = -\frac{\partial^2 W/\partial\lambda\partial\alpha}{d^2W/d\lambda^2}\bigg|_{\lambda^{**}} \tag{52}$$

The denominator is negative (Step 6). For the numerator: increasing α raises the crowding externality, which disproportionately harms EU^S and accelerates snob exit. At λ^{**}, $\partial^2 W/\partial\lambda\partial\alpha < 0$ because an additional snob imposes larger crowding costs when α is higher. Thus $d\lambda^{**}/d\alpha < 0$.

For σ: noisier private signals reduce the precision of individual quality assessment, increasing the relative value of the public signal generated by snob-driven adoption. At λ^{**}, $\partial^2 W/\partial\lambda\partial\sigma > 0$ because the marginal snob's informational contribution is more valuable when private signals are noisier. Thus $d\lambda^{**}/d\sigma > 0$. □

[48]The local result suffices: we establish existence of an interior maximum via the FOC and show that corners are dominated.

Online Appendix

Snobs and Conformists: Platform Design and Product Lifecycles

Tommaso Bondi
Cornell University

This Online Appendix contains two extensions to the baseline model – composition-dependent preferences and robustness to convex social preferences – followed by proofs of corollaries and secondary results. Proofs of the main propositions appear in Appendix A.

Extensions

OA.1 Extension: Who Adopts, Not Just How Many

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. A product used by tastemakers carries different connotations than one used by the mainstream, even at the same adoption level. This section extends the model to allow snob utility to depend on adopter composition, showing that the core results survive with quantitative amplification and that new predictions emerge about cycle speed and the possibility of permanently exclusive equilibria.

The baseline is the right starting point: aggregate adoption n is the sole public signal when agents cannot distinguish snob from conformist co-adopters, and all qualitative results survive the extension with only quantitative modification. However, for markets where adopter identity is observable, the composition framework sharpens predictions.

The essential observation is that snob adoption and conformist adoption have asymmetric effects.

To formalize this, we decompose total adoption into its components. Let n_t^S and n_t^C denote snob and conformist adoption at period t, with $n_t = n_t^S + n_t^C$. Define a product's coolness as:[49]

$$C_t \equiv n_t^S - \xi n_t^C \tag{53}$$

where $\xi > 0$ captures how much conformist adoption dilutes cachet. When $\xi = 1$, snob and conformist adoption have symmetric but opposite effects; when $\xi > 1$, conformist adoption is especially damaging to coolness.

[49] We follow the conceptual framework of Warren et al. (2019), who identify "subcultural appeal" and "perceived exclusivity" as components of brand coolness. For formal treatments of composition-dependent status in signaling contexts, see Pesendorfer (1995) and Corneo & Jeanne (1997a).

The coolness formulation unpacks what "disliking popularity" means: a snob who dislikes high n is really saying "I dislike that conformists have arrived."

With coolness in the utility function, snob utility becomes:

$$U^S = v(\theta) - \alpha n_t + \eta C_t = v(\theta) - \alpha n_t + \eta n_t^S - \eta \xi n_t^C \tag{54}$$

where $\eta \geq 0$ captures how much snobs value coolness per se. When $\eta > 0$, snobs now have *two* reasons to dislike conformist entry: the direct crowding effect ($-\alpha n^C$) and the coolness dilution effect ($-\eta \xi n^C$). This amplifies the exit dynamic.

Proposition 13 (Coolness-Amplified Dynamics). *With composition-dependent coolness ($\eta > 0$):*

(i) Each conformist reduces snob utility by $\alpha + \eta\xi$ rather than α alone, increasing the effective crowding cost by $(1 + \eta\xi/\alpha)$.

(ii) Duration satisfies $T_\eta < T_0$ with $T_\eta / T_0 \approx \alpha/(\alpha + \eta\xi)$ for small $\eta\xi$.

(iii) Coolness C_t peaks in Phase I, becomes negative in Phase II ($C_t = 0$ at $n_t^C = n_t^S/\xi$), so products are maximally popular when minimally cool.

The proposition establishes that accounting for adopter composition does not merely add nuance to the baseline – it *amplifies* the core mechanism. If $\eta\xi \approx \alpha$, conformist entry is twice as damaging to snobs as in the baseline, potentially halving lifecycle duration for products where "who uses it" matters intensely (fashion, cultural goods, social platforms).

The composition framework also reveals a possibility absent from the baseline: cool equilibria where products sustain positive coolness indefinitely by limiting conformist entry.

Definition 4 (Cool Equilibrium). A cool equilibrium is a steady state where snobs adopt ($n^S > 0$), conformists do not adopt ($n^C = 0$), and the configuration is stable.

Corollary 7 (Existence of Cool Equilibria). *A cool equilibrium exists if and only if:*

$$\underline{\mu}^C(n_{\max}^S) > \bar{\mu}(\hat{\theta}_\infty) \tag{55}$$

where $n_{\max}^S = \lambda[1 - G(c_S; \hat{\theta}_\infty)]$ is maximal snob adoption and $\bar{\mu}(\hat{\theta})$ is the highest achievable posterior given $\hat{\theta}$. In the composition-dependent model with $\eta > 0$, a snob-only configuration yields effective snob utility $v(\theta) - (\alpha - \eta)n^S$ (since $n^C = 0$ implies $C_t = n_t^S$), so $n_{\max}^S$ depends on $\alpha - \eta$ rather than α: for $\eta > 0$, snobs are more tolerant of other snobs, which tends to increase $n_{\max}^S$. Because $\underline{\mu}^C(n) = c_C - \beta n$ is decreasing, a higher $n_{\max}^S$ makes the existence condition harder to satisfy, all else equal. Cool equilibria are therefore driven primarily by conformist-side parameters: high c_C (high outside option for conformists), low β (weak bandwagon preferences), and large σ (diffuse posteriors that keep $\bar{\mu}(\hat{\theta}_\infty)$ from concentrating near 1). Sufficient: $c_C > \bar{\mu}(\hat{\theta}_\infty) + \beta n_{\max}^S$, or $\beta < (c_C - \bar{\mu})/n_{\max}^S$, or σ large.

Cool equilibria explain markets that sustain exclusivity despite demand pressure: vinyl records (inconvenience screens casual listeners, preserving countercultural cachet), literary fiction (difficulty screens casual readers), and Berghain (door policy screens mainstream attendees).[50] In each case, structural features raise conformist entry costs, sustaining positive coolness through barriers that keep c_C high.

[50]This parallels the "trickle-down" dynamics documented by Bellezza (2023), where luxury products lose status as they diffuse. Cool equilibria formalize a market structure where such trickle-down *never occurs* because conformist entry is perpetually blocked.

OA.2 Robustness to Functional Form

The baseline model assumes linear preferences over adoption: $U^S = \theta - \alpha n$ and $U^C = \theta + \beta n$. This appendix examines robustness to alternative functional forms, focusing on the convex case that generates empirically relevant predictions about cycle shape.

OA.2.0.1 Motivation for Convexity

Several considerations motivate convex crowding costs: crowding effects may exhibit increasing marginal disutility (the shift from "underground" to mainstream feels larger than early growth); Bellezza (2023) documents threshold effects in status signaling; and psychological research on distinctiveness (Berger & Heath, 2007) finds that identity-signaling value drops sharply once outgroup members adopt. For conformists, convex bandwagon benefits capture superlinear returns to coordination.

OA.2.0.2 Specification

Consider the convex utility functions:

$$U^S = \theta - \alpha n^\rho \tag{56}$$

$$U^C = \theta + \beta n^\zeta \tag{57}$$

where $\rho, \zeta > 1$ capture convexity. The baseline linear model is the special case $\rho = \zeta = 1$. The pure quadratic case sets $\rho = \zeta = 2$.

OA.2.0.3 Modified Equilibrium

With convex preferences, threshold strategies still obtain but with modified cutoffs:

$$\underline{\mu}^S(n) = \alpha n^\rho + c_S \tag{58}$$

$$\underline{\mu}^C(n) = c_C - \beta n^\zeta \tag{59}$$

The snob threshold is now convex in n (curving upward faster than linear), while the conformist threshold is concave (curving downward slower than linear when $\zeta > 1$).

At the peak, both marginal types are indifferent:

$$\hat{\theta}^* - \alpha(n^*)^\rho - c_S = 0 \tag{60}$$

$$\hat{\theta}^* + \beta(n^*)^\zeta - c_C = 0 \tag{61}$$

For the symmetric case $\rho = \zeta = 2$ and $c_S = 0$:

$$n^* = \sqrt{\frac{c_C}{\alpha + \beta}}$$

Compare to the linear case: $n^*_{\text{linear}} = \frac{c_C - c_S}{\alpha+\beta}$. For typical parameters, convex crowding costs allow higher peak adoption because snobs tolerate early growth more easily.

OA.2.0.4 Cycle Shape and Asymmetry

The main qualitative difference with convex preferences is in cycle *shape*. The linear model produces relatively symmetric adoption curves. The convex model produces right-skewed curves with gradual ascent and rapid descent.

The asymmetry arises from the marginal crowding cost $\partial U^S / \partial n = -\alpha \rho n^{\rho - 1}$. In early phases when n is small, this marginal cost is negligible – snobs barely notice the first conformists arriving. As n grows, the marginal cost accelerates. Near the peak, even small additional adoption triggers sharp utility drops, causing rapid exit.

This generates several distinctive features. The growth phase extends relative to the linear case because early crowding is tolerable. The peak is more pronounced – a sharper "cliff" rather than a rounded "hill." The decline phase is compressed because exit cascades once crowding costs bite. The asymmetry increases with ρ: higher convexity means more dramatic cliffs.

Empirically, many product lifecycles exhibit this right-skewed pattern. Fashion trends often build gradually over seasons then crash within weeks. Social platforms grow steadily for years then collapse rapidly (MySpace peaked over 2007–2008 but lost 80% of users within 18 months; Vine maintained relevance for three years before sudden shutdown; Clubhouse's entire rise and fall occurred within one year). The convex model captures this asymmetry naturally.

OA.2.0.5 Robustness of Main Results

The core qualitative results survive under convex preferences with modified formulas. The three-phase lifecycle persists because opposing responses to adoption are unchanged. The reversal also survives: the threshold shifts but the non-monotonicity persists whenever acceleration dominates persistence. What changes is primarily quantitative: peak adoption formulas, comparative statics magnitudes, and – most consequentially – cycle shape, with the convex model predicting more asymmetric, cliff-like patterns.

OA.2.0.6 Empirical Discrimination

The linear versus convex specification generates testable predictions. Define skewness Skew = $(t_{\text{peak}} - t_{\text{start}})/(t_{\text{end}} - t_{\text{peak}})$. The linear model predicts Skew ≈ 1 (symmetric); the convex model predicts Skew > 1 (right-skewed). We conjecture that fashion, social media, and cultural products exhibit convex patterns (Skew > 1), while technology and durables exhibit more linear ones.

OA.2.0.7 Worked Example: Quadratic Preferences with Beta Prior

We provide one fully specified non-linear example where the main results hold *exactly*, not under density approximations.

Setup. Let snob utility be $U^S = \theta - \alpha n^2$ and conformist utility be $U^C = \theta + \beta n$ (quadratic crowding, linear bandwagon). Let $c_S = 0$ and $c_C = c > 0$. Suppose posteriors μ are distributed Beta(a_t, b_t) with parameters evolving via Bayesian updating, so the density $g(\mu) = \mu^{a_t-1}(1-\mu)^{b_t-1}/B(a_t, b_t)$ is known in closed form.

Thresholds. Optimal strategies are:

$$\text{Snob adopts} \iff \mu \geq \alpha n^2 \equiv \underline{\mu}^S(n) \tag{62}$$

$$\text{Conformist adopts} \iff \mu \geq c - \beta n \equiv \underline{\mu}^C(n) \tag{63}$$

The snob threshold is convex and increasing; the conformist threshold is linear and decreasing. They intersect at the unique n^* solving $\alpha(n^*)^2 + \beta n^* = c$:

$$n^* = \frac{-\beta + \sqrt{\beta^2 + 4\alpha c}}{2\alpha} \tag{64}$$

Quality-duration reversal (exact). Define the rate of conformist entry and snob exit at the peak. Using the Beta CDF $I_x(a, b)$:

$$\text{Snob exit rate} = \lambda \cdot g(\underline{\mu}^S(n^*)) \cdot 2\alpha n^* \tag{65}$$

$$\text{Conformist entry rate} = (1-\lambda) \cdot g(\underline{\mu}^C(n^*)) \cdot \beta \tag{66}$$

At the peak, both thresholds equal $\mu^* = \alpha(n^*)^2 = c - \beta n^*$, so $g(\underline{\mu}^S(n^*)) = g(\underline{\mu}^C(n^*)) = g(\mu^*)$. The quality-duration reversal occurs when the snob exit rate exceeds the rate at which higher quality can sustain adoption. Since $d\underline{\mu}^S/dn = 2\alpha n$ (compared to α in the linear case), the critical condition is:

$$\lambda \cdot 2\alpha n^* > (1-\lambda) \cdot \beta \quad \iff \quad \alpha > \frac{\beta^2(1-\lambda)(1+\lambda)}{4c\lambda^2} \tag{67}$$

which gives an exact threshold $\bar{\alpha}_Q(\lambda, \beta, c)$ that is finite and well-defined without any density approximation. For the special case $\lambda = 1/3$, $\beta = 1$, $c = 1$, this gives $\bar{\alpha}_Q = (1 \cdot 2/3 \cdot 4/3)/(4 \cdot 1/9) = (8/9)/(4/9) = 2$: any $\alpha > 2$ produces the reversal.

Lifecycle phases (exact). The three-phase structure holds by the same logic as the linear case: at low n, $\underline{\mu}^S < \underline{\mu}^C$, so only snobs adopt; at high n, $\underline{\mu}^S > \underline{\mu}^C$, so only conformists adopt; at intermediate n, both are active. Because the Beta density $g(\mu^*)$ is strictly positive and bounded on $(0,1)$, the best-response mapping $\Phi(n)$ is continuous and satisfies the boundary conditions for Brouwer. The lifecycle characterization (Proposition 3) therefore holds exactly.

This example demonstrates that the main results are not artifacts of the linear-utility, constant-density simplification used in the main text.

Additional Proofs

Proof of Corollary 1 (Peak Adoption)

Proof. At the peak t^*, both snob and conformist thresholds bind for the marginal agents. We derive the peak adoption level.

Step 1: Peak condition. The adoption path satisfies $n_t = \Phi(n_t; \hat{\theta}_t)$ where $\hat{\theta}_t$ evolves via Bayesian updating. The peak occurs at t^* where n_t reaches its maximum over the path. At any fixed $\hat{\theta}$, the equilibrium is determined by $n^*(\hat{\theta}) = \Phi(n^*(\hat{\theta}); \hat{\theta})$. The peak occurs when the snob exit from rising n first offsets the conformist entry from rising beliefs. We do not derive a general peak condition from first principles; instead, we characterize n^* under the density approximation.

Step 2: Exact characterization. The peak adoption n^* satisfies $n^* = \Phi(n^*; \hat{\theta}_{t^*})$. Since Φ is continuous and we are at the peak of n_t, the equilibrium condition pins down n^* as a function of $\hat{\theta}_{t^*}$. In general, n^* must be solved numerically from this fixed point.

Step 3: Closed-form approximation. Under Assumption 4, $g(\underline{\mu}^S(n^*)) \approx g(\underline{\mu}^C(n^*)) \equiv g^*$. We derive the peak adoption level using a direct characterization under the constant-density approximation.

The key observation is that when the posterior density is approximately constant across the threshold region, the adoption masses of both types depend linearly on the distance between each threshold and the boundary of the support. Specifically, snob adoption is approximately $\lambda g^*(1-\underline{\mu}^S(n))$ and conformist adoption is approximately $(1-\lambda)g^*(1-\underline{\mu}^C(n))$, up to an additive constant. The equilibrium adoption level n^* is where total adoption, as a function of n, equals n (the fixed-point condition). Under constant density, the fixed-point condition is linear in n, and the solution equates the two thresholds: $\underline{\mu}^S(n^*) = \underline{\mu}^C(n^*)$. To see why, note that the constant-density best-response mapping $\Phi(n) = \lambda g^*(1 - \alpha n - c_S) + (1-\lambda)g^*(1 - c_C + \beta n) + \text{const}$ is linear in n with slope $\Phi'(n) = g^*[(1-\lambda)\beta - \lambda\alpha]$. The peak of the *lifecycle* path (not the static equilibrium) occurs when the belief-driven growth of total adoption stalls, which under the approximation coincides with the point where the marginal snob exiting balances the marginal conformist entering: $\Phi'(n^*) = 0$, giving $\lambda\alpha = (1-\lambda)\beta$. At this point, the threshold-crossing condition $\underline{\mu}^S = \underline{\mu}^C$ holds as well, yielding:

$$\alpha n^* + c_S = c_C - \beta n^* \quad \Rightarrow \quad n^* = \frac{c_C - c_S}{\alpha + \beta} \tag{68}$$

This is the adoption level at which the two types' thresholds cross. Since this equals $n^\dagger$ from Lemma 1, the peak occurs at the threshold crossing point under the density approximation.

Step 4: When the approximation holds and fails. The constant-density approximation is accurate when: (i) signal noise σ is moderate, so the posterior distribution is spread across the threshold region, and (ii) both thresholds fall within the bulk of the posterior distribution rather than in the tails. When σ is very small (precise signals), the posterior density is concentrated and $g^S \neq g^C$ in general; the exact n^* must be solved from the full fixed-point condition.

Step 5: Comparative statics. Under the constant-density approximation, $n^* = (c_C - c_S)/(\alpha + \beta)$, so both $\partial n^*/\partial\alpha < 0$ and $\partial n^*/\partial\beta < 0$ follow by direct differentiation.

Without the approximation, the exact comparative static for α at the *within-period* equilibrium (holding $\hat{\theta}$ fixed) is:

$$\left.\frac{\partial n^*}{\partial \alpha}\right|_{\hat{\theta}\text{ fixed}} = \frac{-\lambda g(\underline{\mu}^S(n^*)) \cdot n^*}{1 - \Phi'(n^*)} < 0 \tag{69}$$

This requires $1 - \Phi'(n^*) > 0$, i.e., $\Phi'(n^*) < 1$, which holds at any stable equilibrium. The sign result for the dynamic peak follows because higher α depresses the within-period equilibrium at every belief level *and* weakly depresses the equilibrium belief path (through lower adoption generating less favorable signals).

Note: $\partial n^*/\partial\beta|_{\hat{\theta}\text{ fixed}} > 0$ (higher β increases the static equilibrium), but the lifecycle peak $n^* = (c_C - c_S)/(\alpha + \beta)$ decreases in β because the threshold crossing occurs at lower adoption. □

Proof of Corollary 2 (Preference Parameter Effects)

Proof. We derive comparative statics with respect to α and β.

Step 1: Effect on peak adoption. The peak of the lifecycle path n^* is characterized differently from the static equilibrium at a fixed belief. Under the constant-density approximation (Assumption 4), peak adoption occurs at the threshold crossing point $n^* = (c_C - c_S)/(\alpha + \beta)$ (Corollary 1). Direct differentiation gives:

$$\frac{\partial n^*}{\partial \alpha} = -\frac{c_C - c_S}{(\alpha+\beta)^2} < 0, \qquad \frac{\partial n^*}{\partial \beta} = -\frac{c_C - c_S}{(\alpha+\beta)^2} < 0 \tag{70}$$

Both are negative: stronger preferences (of either type) bring the threshold crossing to a lower n, reducing the peak adoption mass.

Without the density approximation, $\partial n^*/\partial\alpha < 0$ holds exactly at any stable equilibrium by the IFT:

$$\left.\frac{\partial n^*}{\partial \alpha}\right|_{\hat{\theta}\text{ fixed}} = \frac{-\lambda g(\underline{\mu}^S(n^*)) \cdot n^*}{1 - \Phi'(n^*)} < 0 \tag{71}$$

For β, the effect on the static equilibrium at fixed $\hat{\theta}$ is positive: $\partial n^*/\partial\beta|_{\hat{\theta}\text{ fixed}} > 0$ (higher β encourages more conformist adoption). The lifecycle peak decreases in β because the threshold crossing point $(c_C - c_S)/(\alpha + \beta)$ falls, even though the equilibrium at any fixed belief rises.

Step 2: Effect of α on duration. Consider the snob threshold $\underline{\mu}^S(n) = \alpha n + c_S$. An increase in α to $\alpha + d\alpha$ raises the threshold by $n \cdot d\alpha$ at every $n > 0$. This means the set of posteriors $\{\mu : \mu \geq \underline{\mu}^S(n)\}$ shrinks for every n. At each period t during the lifecycle, snob adoption $n_t^S = \lambda[1 - G(\alpha n_t + c_S)]$ is lower. The cascade from fewer snobs $\to$ weaker belief updating $\to$ slower conformist entry $\to$ earlier peak means that higher α uniformly compresses the lifecycle. Since peak occurs at lower n^* and is reached from a lower trajectory, T decreases.

More precisely, consider two economies with $\alpha_1 < \alpha_2$. At each period, $n_t(\alpha_2) < n_t(\alpha_1)$ (by induction: $n_0(\alpha_2) < n_0(\alpha_1)$ since the snob threshold is higher, and the subsequent belief path is less favorable). Since T is the first time $n_t^S = 0$ after peak, the lower trajectory reaches this condition sooner: $T(\alpha_2) < T(\alpha_1)$.

Step 3: Effect of β on the reversal threshold. From the closed form $\bar{\alpha}(\lambda) = \beta(1 - \lambda)/\lambda$:

$$\frac{\partial \bar{\alpha}}{\partial \beta} = \frac{1 - \lambda}{\lambda} > 0 \tag{72}$$

Higher β raises the reversal threshold, meaning stronger conformist preferences make the quality-duration reversal harder to trigger. The intuition is that the reversal requires crowding to dominate bandwagon effects ($\alpha\lambda > \beta(1-\lambda)$); increasing β strengthens the bandwagon side, requiring commensurately higher α for reversal.

The effect of β on lifecycle duration T is regime-dependent. In the persistence region ($\alpha < \bar{\alpha}$), higher β extends the lifecycle: faster conformist entry generates more adoption, which improves belief updating, sustaining snob participation. In the reversal region ($\alpha > \bar{\alpha}$), the effect is ambiguous: higher β accelerates conformist entry (which pushes snobs out faster via higher n) but also raises $\bar{\alpha}$ (moving the economy toward persistence). The net sign depends on parameters.

Step 4: Volatility. From the best-response derivative, $|\Phi'(n)| = |(1 - \lambda)\beta g^C - \lambda\alpha g^S|$. Higher α increases $\lambda\alpha g^S$, which increases $|\Phi'|$ when $\lambda\alpha g^S > (1 - \lambda)\beta g^C$ (snob-dominated region). This increases the sensitivity of adoption to shocks. $\square$

Proof of Proposition 6 (Non-Monotonic Composition Effects)

Proof. We show that cycle duration $T(\lambda)$ is non-monotonic in the snob share λ.

Step 1: Boundary behavior as $\lambda \to 0$. When $\lambda \to 0$, initial snob adoption $n_0^S = \lambda[1 - G(c_S)] \to 0$. With negligible early adoption, the observed signal $\tilde{n}_0 \approx \varepsilon_0$ is uninformative. Belief updating stalls: $\hat{\theta}_1 \approx \hat{\theta}_0 = p$. Without rising beliefs, conformist adoption never starts (since $c_C > p$ for typical parameters). The "cycle" consists of vanishingly small snob adoption that generates no information cascade. Thus $T(\lambda) \to 0$ as $\lambda \to 0$.

Step 2: Boundary behavior as $\lambda \to 1$. When $\lambda \to 1$, nearly all agents are snobs. With $(1-\lambda) \approx 0$, conformist adoption $n_t^C = (1-\lambda)[1 - G(\underline{\mu}^C)] \approx 0$ for all t. Without conformist entry, the three-phase lifecycle degenerates: no Phase II occurs and no conformist surge triggers snob exit. Under Definition 2, T is defined as $\min\{t > t^* : n_t^S < \epsilon\}$. In the snob-only limit, snob adoption converges to a stable level and need not fall below ϵ, so $T \to \infty$ under the formal definition.

However, what drives the welfare result is that the *boom-bust interaction* between snobs and conformists – the mechanism generating surplus from preference heterogeneity – vanishes. The economically relevant contribution of the lifecycle to welfare requires both types to participate. As $\lambda \to 1$, the welfare contribution from conformist participation $(1-\lambda)EU^C \to 0$, and the total surplus from snob-conformist interaction vanishes even though snobs persist. The boundary argument for Proposition 6 therefore relies on the welfare function $W(\lambda)$, not on the formal duration T: $W(1) = EU^S(1)$ is finite and interior values exceed it (Step 3).

Step 3: Interior maximum. For intermediate λ, both types contribute: snobs drive growth and generate information, conformists sustain adoption and extend the lifecycle. Duration is maximized when both effects are balanced. Since T is continuous, low at both boundaries, and positive in the interior, the extreme value theorem gives $\lambda^* \in (0, 1)$.

Step 4: Characterizing λ^*. At the optimum, $dT/d\lambda = 0$. Under the constant-density approximation, the reversal condition (Proposition 4) is $\alpha\lambda \lesseqgtr \beta(1-\lambda)$. Duration T depends on λ through the balance of crowding and bandwagon effects. At the critical λ^* where these forces are balanced:

$$\alpha\lambda^* = \beta(1-\lambda^*) \quad \Rightarrow \quad \lambda^* = \frac{\beta}{\alpha+\beta} \tag{73}$$

Below λ^*, increasing λ extends the lifecycle (more snobs improve information generation, and bandwagon effects still dominate crowding). Above λ^*, increasing λ shortens the lifecycle (crowding dominates, triggering the reversal). This is the composition at which the reversal threshold $\bar{\alpha}(\lambda) = \beta(1-\lambda)/\lambda$ exactly equals α. □

Proof of Lemma 3 (Advertising and Lifecycle Duration)

Proof. We show that increasing advertising effectiveness can reduce total lifecycle duration.

Step 1: Advertising model. Let advertising shift each agent's posterior upward by $a \geq 0$ (persuasive advertising that makes agents more optimistic about quality). Under advertising, type τ adopts iff $\mu_i + a \geq \underline{\mu}^\tau(n)$, or equivalently $\mu_i \geq \underline{\mu}^\tau(n) - a$. This is equivalent to lowering the effective threshold by a: adoption at mass n is

$$n_t(a) = \lambda[1 - G(\alpha n_t + c_S - a)] + (1-\lambda)[1 - G(c_C - \beta n_t - a)] \tag{74}$$

Since G is increasing, $n_t(a) > n_t(0)$ for all $a > 0$.

Step 2: Effect on adoption and snob exit. Advertising increases total adoption: by the IFT applied to the fixed-point condition $n = \Phi(n; a)$:

$$\frac{dn^*}{da} = \frac{\partial \Phi / \partial a}{1 - \partial \Phi / \partial n} = \frac{\lambda g(\underline{\mu}^S - a) + (1-\lambda) g(\underline{\mu}^C - a)}{1 - \Phi'(n^*)} > 0 \tag{75}$$

The numerator includes contributions from *both* types: advertising attracts additional snobs (first term) and additional conformists (second term). The denominator is positive at a stable equilibrium ($\Phi'(n^*) < 1$).

The question is whether the adoption increase raises the snob threshold fast enough to trigger earlier exit. The snob threshold at $n_t(a)$ is $\alpha n_t(a) + c_S - a$ (net of advertising). The advertising-induced increase in n raises this by $\alpha \cdot dn^*/da$, while the direct advertising effect lowers it by 1. The threshold rises (making snob exit sooner) when $\alpha \cdot dn^*/da > 1$, i.e.:

$$\alpha \cdot \frac{\lambda g^S + (1-\lambda) g^C}{1 - \Phi'(n^*)} > 1 \tag{76}$$

This holds when α is large (strong crowding aversion) or when the adoption response is large (both types are highly responsive to the posterior shift). To verify this condition holds at $n > n^\dagger$: at the threshold crossing point, conformists are entering and the adoption response dn^*/da is amplified by the conformist bandwagon multiplier. From the IFT expression, $dn^*/da = [\lambda g^S + (1-\lambda) g^C]/(1 - \Phi')$. When $n > n^\dagger$, $\Phi'(n) = (1-\lambda)\beta g^C - \lambda \alpha g^S$ is positive (conformist complementarity dominates), so $1 - \Phi' < 1$ and the multiplier exceeds 1. Combined with $\alpha > \bar{\alpha}(\lambda) = \beta(1-\lambda)/\lambda$ (which ensures the reversal regime), the condition $\alpha \cdot dn^*/da > 1$ holds generically in the reversal regime.[51]

Step 3: Duration comparison. The lemma claims $\partial T / \partial a < 0$ when $n > n^\dagger$. At $n > n^\dagger$, conformists are already entering. Advertising raises n further (Step 2), which raises the snob threshold $\alpha n + c_S$ (recall that the net threshold under advertising is $\alpha n_t(a) + c_S - a$, and under the Step 2 condition $\alpha \cdot dn^*/da > 1$, the net threshold increases with a). Formally, comparing adoption paths with $a > 0$ versus $a = 0$: at each period during Phase II, $n_t(a) > n_t(0)$ and hence the net snob threshold $\alpha n_t(a) + c_S - a > \alpha n_t(0) + c_S$ under the Step 2 condition,

[51] The condition can fail near the boundary $\alpha \approx \bar{\alpha}$ with very precise signals; the lemma applies under non-degenerate posterior densities.

meaning fewer snobs adopt. The snob exit time therefore arrives earlier, and T is shorter.

Step 4: When backfire occurs. The backfire condition requires $\partial T/\partial a < 0$ to dominate any per-period profit increase. This is more likely when α is large (strong crowding), a is applied during Phase II (when conformist entry is already occurring), and $n > \bar{n}^S$ (the snob exit threshold has been crossed). □

Proof of Corollary 5 (Welfare Comparison)

Proof. We decompose the welfare comparison across lifecycle phases.

Step 1: Dynamic expected utility. A randomly drawn snob's ex ante expected lifetime utility from the equilibrium lifecycle is:

$$EU^S = \sum_{t=0}^{T} \delta^t \cdot \mathbb{E}\left[\max\{\mu_i - \alpha n_t, c_S\} - c_S\right] \tag{77}$$

where the expectation is over the posterior distribution $G(\cdot;\hat{\theta}_t)$ and the max captures the adoption decision: the snob adopts iff $\mu_i - \alpha n_t \geq c_S$, earning surplus $\mu_i - \alpha n_t - c_S \geq 0$, and otherwise earns zero (relative to the outside option c_S, which we normalize to zero).

Decomposing: let $\pi_t^S = 1 - G(\underline{\mu}^S(n_t))$ be the probability a random snob adopts at t, and $\bar{\mu}_t^S = \mathbb{E}[\mu_i \mid \mu_i \geq \underline{\mu}^S(n_t)]$ be the average posterior among adopting snobs. Then:

$$EU^S = \sum_{t=0}^{T} \delta^t \pi_t^S [\bar{\mu}_t^S - \alpha n_t - c_S] \tag{78}$$

Each term is non-negative (adopting snobs earn positive surplus by revealed preference). In Phase III, $\pi_t^S \approx 0$ (few snobs adopt), so these terms contribute negligibly.

Step 2: Conformist expected utility. By the same logic:

$$EU^C = \sum_{t=0}^{T} \delta^t \pi_t^C [\bar{\mu}_t^C + \beta n_t - c_C] \tag{79}$$

where $\pi_t^C = 1 - G(\underline{\mu}^C(n_t))$ and $\bar{\mu}_t^C = \mathbb{E}[\mu_i \mid \mu_i \geq \underline{\mu}^C(n_t)]$. In Phase I, $\pi_t^C \approx 0$ (few conformists adopt because $\underline{\mu}^C(n_t)$ is high when n_t is low).

Step 3: Phase-by-phase comparison. During Phase I ($t < t_1$): $\pi_t^S > 0, \pi_t^C \approx 0$. Snobs earn pioneer rents; conformists earn approximately zero. During Phase III ($t > t^*$): $\pi_t^S \approx 0, \pi_t^C > 0$. Conformists earn bandwagon benefits; snobs earn approximately zero. During Phase II ($t_1 \leq t \leq t^*$): both $\pi_t^S > 0$ and $\pi_t^C > 0$.

Step 4: Comparison. Subtracting:

$$\Delta EU = EU^S - EU^C \approx \underbrace{\sum_{t=0}^{t_1-1} \delta^t \pi_t^S [\bar{\mu}_t^S - \alpha n_t - c_S]}_{\text{Pioneer rents (Phase I)}} - \underbrace{\sum_{t=t^*+1}^{T} \delta^t \pi_t^C [\bar{\mu}_t^C + \beta n_t - c_C]}_{\text{Bandwagon benefits (Phase III)}} + R \tag{80}$$

where R captures the surplus difference during Phase II. The sign of ΔEU depends on the discount factor (which downweights Phase III bandwagon benefits) and the relative duration of each phase. □

Proof of Proposition 13 (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:

$$U^S = v(\theta) - \alpha n_t + \eta C_t = v(\theta) - \alpha(n_t^S + n_t^C) + \eta(n_t^S - \xi n_t^C) \tag{81}$$

Rearranging:

$$U^S = v(\theta) + (\eta - \alpha) n_t^S - (\alpha + \eta\xi) n_t^C \tag{82}$$

Step 2: Marginal effect of conformist entry. Differentiating:

$$\frac{\partial U^S}{\partial n^C} = -(\alpha + \eta\xi) < -\alpha \tag{83}$$

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:

$$\Delta C_t = \Delta n_t^S - \xi \Delta n_t^C \tag{84}$$

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 $n_t^C > n_t^S / \xi$. □

Proof of Corollary 3 (Optimal Pricing Path)

Proof. At each period, the monopolist chooses price p_t to maximize revenue subject to participation constraints.

Step 1: Phase I (snobs only). During $t < t_1$, only snobs adopt. A snob with posterior μ adopts iff $\mu - \alpha n_t - p_t \geq c_S$, i.e., $\mu \geq \alpha n_t + p_t + c_S$. Adoption mass is $n_t = \lambda[1 - G(\alpha n_t + p_t + c_S)]$. The monopolist maximizes $p_t \cdot n_t$. Since n_t depends on p_t both directly (through the threshold) and indirectly (through the crowding term αn_t), implicit differentiation of the fixed-point condition gives $\partial n_t/\partial p_t = -\lambda g(z)/(1 + \lambda \alpha g(z))$ where $z = \alpha n_t + p_t + c_S$. The FOC $n_t + p_t(\partial n_t/\partial p_t) = 0$ then yields:

$$p_t = \frac{1 - G(z)}{g(z)} \cdot [1 + \lambda \alpha g(z)] \tag{85}$$

The first factor is the standard inverse-hazard-rate markup; the second factor $[1 + \lambda \alpha g(z)]$ captures the crowding feedback: a higher price reduces adoption, which lowers crowding costs and partially offsets the demand reduction. When $\lambda \alpha g(z)$ is small (moderate crowding relative to signal dispersion), this feedback term is close to 1 and the optimal price simplifies to the approximate expression $p_t^I \approx \hat{\theta}_t - \alpha n_t - c_S$, which extracts the marginal snob's surplus above the participation threshold. As n_t rises during growth, αn_t increases and the optimal price tends to fall.

Step 2: Phase II (both types) with discrimination. With type discrimination, the monopolist sets type-specific prices. The snob participation constraint is $\mu - \alpha n_t \geq p_t^S + c_S$, giving $p_t^S = \hat{\theta}_t - \alpha n_t - c_S$ for the marginal snob. The conformist participation constraint is $\mu + \beta n_t \geq p_t^C + c_C$, giving $p_t^C = \hat{\theta}_t + \beta n_t - c_C$. Since $\beta n_t + \alpha n_t = (\alpha + \beta) n_t > 0$, we have $p_t^C - p_t^S = (\alpha + \beta) n_t - (c_C - c_S)$, which is positive when $n_t > n^\dagger = (c_C - c_S)/(\alpha + \beta)$.

Step 3: Phase II without discrimination. The uniform-price monopolist chooses p_t to maximize:

$$p_t \cdot [\lambda(1 - G(\alpha n_t + p_t + c_S)) + (1 - \lambda)(1 - G(c_C - \beta n_t + p_t))] \tag{86}$$

The optimal uniform price lies between p_t^S and p_t^C, weighted by the relative demand elasticities of each type.

Step 4: Phase III (decline). During decline, snob adoption is negligible and demand is conformist-driven. The conformist participation constraint is $\mu + \beta n_t \geq c_C + p_t$. As n_t falls, the bandwagon benefit βn_t shrinks, and the price must fall to maintain positive demand. □

Proof of Proposition 9 (Optimal Targeting)

Proof. Let $\tau \in [0, 1]$ be the share of advertising directed at snobs. Advertising lowers type-specific thresholds: snob threshold becomes $\underline{\mu}^S(n) - \tau a$ and conformist threshold becomes $\underline{\mu}^C(n) - (1 - \tau)a$, where a is total advertising effectiveness.

The firm maximizes discounted profits over the lifecycle:

$$\Pi(\tau) = \sum_{t=0}^{T(\tau)} \delta^t \pi(n_t(\tau))$$

where period profit $\pi(n_t) = p(n_t) \cdot n_t$ and both n_t and lifecycle duration T depend on τ. Since T is formally a discrete stopping time, we treat it as a smooth function of τ by taking the continuous relaxation $T(\tau) = \inf\{t > t^* : n_t^S(\tau) < \epsilon\}$ and noting that $n_t^S(\tau)$ is C^1 in τ by the implicit function theorem applied to the equilibrium fixed-point condition (Lemma 2). Under this interpretation:

$$\frac{d\Pi}{d\tau} = \underbrace{\sum_{t=0}^{T} \delta^t \pi'(n_t) \frac{\partial n_t}{\partial \tau}}_{\text{per-period profit effect}} + \underbrace{\pi(n_T)\delta^T \frac{\partial T}{\partial \tau}}_{\text{duration effect}} \tag{87}$$

Per-period effect. Shifting advertising from conformists to snobs ($d\tau > 0$) increases snob adoption by $a\lambda g(\underline{\mu}^S)$ and decreases conformist adoption by $a(1-\lambda)g(\underline{\mu}^C)$. The net effect on n_t is $\partial n_t/\partial\tau = a[\lambda g(\underline{\mu}^S) - (1-\lambda)g(\underline{\mu}^C)]$, which can be positive or negative.

Duration effect. More snob advertising sustains the growth phase (snobs adopt longer, beliefs improve more) while less conformist advertising slows the entry that triggers snob exit. Both channels extend duration: $\partial T/\partial\tau > 0$.

FOC and solution. Setting $d\Pi/d\tau = 0$ and solving for τ^*:

$$\tau^* = \min\left\{1, \underbrace{\frac{\alpha}{\alpha+\beta}}_{\substack{\text{static: balances}\\ \text{marginal WTP}}} + \underbrace{\frac{\delta}{1-\delta} \cdot \frac{\lambda}{1-\lambda} \cdot \frac{\partial T/\partial n^S}{|\partial T/\partial n^C|}}_{\substack{\text{dynamic: lifecycle}\\ \text{extension value}}}\right\} \tag{88}$$

The static term $\alpha/(\alpha+\beta)$ arises from the myopic profit-maximization under the constant-density approximation (Assumption 4): when $g^S \approx g^C$ and the posterior densities at the two thresholds are approximately equal, the marginal adoption response to snob-targeted advertising is proportional to α and to conformist-targeted advertising proportional to β, so the optimal allocation equates marginal profit across types. Without the density approximation, the static optimal share depends on λ, $g(\underline{\mu}^S)$, and $g(\underline{\mu}^C)$; the simple form $\alpha/(\alpha+\beta)$ is a parameter-only proxy that captures the qualitative result that targeting tilts toward snobs. The dynamic term captures the present value ($\delta/(1-\delta)$) of lifecycle extension, weighted by the relative population ($\lambda/(1-\lambda)$) and the relative sensitivity of duration to each type's adoption.

When δ is large (patient firm) or $\lambda/(1-\lambda)$ is large (snob-heavy market), the dynamic term dominates and $\tau^* = 1$. □

Proof of Lemma 4 (Equilibrium Response to Visibility)

Proof. Part (i). Under visibility φ, conformists observe a noisy signal $\tilde{n}_t = n_t + \varepsilon_t$ with $\varepsilon_t \sim N(0, \sigma_0^2(1-\varphi)/\varphi)$ as specified in the model. With a Gaussian prior on n_t centered at $\bar{n}$, the conformist's posterior mean is a shrinkage estimator $\mathbb{E}[n_t \mid \tilde{n}_t] = w(\varphi)\tilde{n}_t + (1 - w(\varphi))\bar{n}$, where the weight $w(\varphi)$ on the observation is increasing in φ (higher visibility means lower noise variance, so more weight on the signal). At $\varphi = 1$, $w = 1$ and the signal is exact; as $\varphi \to 0$, $w \to 0$ and conformists rely on the prior. The conformist entry time t_1 is the first period where the perceived adoption, combined with the public belief, crosses the conformist threshold: $\hat{\theta}_{t_1} + \beta\mathbb{E}[n_{t_1} \mid \tilde{n}_{t_1}] \geq c_C$. Increasing φ raises the weight on the true n_{t_1} (which exceeds $\bar{n}$ during Phase I growth), so the threshold is crossed earlier: $\partial t_1/\partial\varphi < 0$.

Part (ii). When $\alpha > \bar{\alpha}(\lambda)$, conformist entry triggers snob exit (Proposition 4). Faster conformist entry (lower t_1) accelerates the Phase II transition, compressing the lifecycle. Since $\partial t_1/\partial\varphi < 0$ and $\partial T/\partial t_1 > 0$ (later conformist entry extends duration), the chain rule gives $\partial T/\partial\varphi = (\partial T/\partial t_1)(\partial t_1/\partial\varphi) < 0$.

Part (iii). Under Assumption 3, the equilibrium mapping $n^*(\hat{\theta}, \varphi)$ is the unique fixed point of $\Phi(n; \hat{\theta}, \varphi)$, which is C^1 in φ by the implicit function theorem (since Φ is C^1 and $|\Phi'| < 1$). Welfare $W(\varphi) = \sum_t \delta^t[\lambda \cdot EU_t^S(\varphi) + (1-\lambda) \cdot EU_t^C(\varphi)]$ inherits C^1 differentiability from the equilibrium path. □

Proof of Proposition 11 (Welfare-Maximizing Visibility)

Proof. We establish each part of the proposition.

Part (i): Interior optimum. We show $W'(0) > 0$ and $W'(1) < 0$ under the stated conditions, which with continuity implies an interior maximum.

At $\varphi = 0$ (complete opacity), conformists cannot observe adoption at all. They must rely entirely on private signals, entering only when their own posterior exceeds c_C. Since conformists have high outside options ($c_C > c_S$), few conformists enter in early periods, even when snob adoption is high. But this means conformists never receive the bandwagon benefits they value. Marginally increasing φ from 0 allows conformists to observe that adoption has begun, lowering their effective threshold and permitting entry. This strictly increases conformist welfare; the induced change in snob welfare is of second order relative to the conformist gain, because snobs' adoption decisions in Phase I depend primarily on their own signals and on n_t (which is little affected by marginal changes in conformists' information when conformist adoption is near zero). Hence $W'(0) > 0$.

At $\varphi = 1$ (full transparency), conformists observe n_t exactly. Consider the marginal effect of reducing φ slightly below 1. This introduces noise in the adoption signal, causing some conformists who would have entered at t_1 (the first period where n_{t_1} crosses their threshold) to delay. Let $\Delta t_1 > 0$ denote the expected delay in conformist entry.

The welfare effect decomposes as follows. Since $\partial t_1/\partial\varphi < 0$ (higher visibility accelerates conformist entry), we write the effect of a marginal *reduction* in φ from full transparency, which delays conformist entry by $\Delta t_1 > 0$:

$$-\frac{dW}{d\varphi}\bigg|_{\varphi=1} = \underbrace{(1-\lambda)\cdot g(\underline{\mu}^C(n_{t_1}))\cdot[\hat{\theta}_{t_1}+\beta n_{t_1}-c_C]\cdot|\Delta t_1|}_{\text{Cost: foregone conformist bandwagon benefits}}$$
$$- \quad \underbrace{\lambda\cdot\sum_{t=0}^{t_1-1}\frac{\partial EU_t^S}{\partial t_1}\delta^t\cdot|\Delta t_1|}_{\text{Benefit: extended snob pioneer rents}}$$
$$-\underbrace{(1-\lambda)\cdot\frac{\partial EU^C}{\partial\hat{\theta}}\cdot\frac{\partial\hat{\theta}}{\partial t_1}\cdot|\Delta t_1|}_{\text{Benefit: improved quality assessment}} \tag{89}$$

The first term is the direct cost of delaying conformist entry: marginal conformists lose Δt_1 periods of bandwagon benefits (positive, since these benefits are foregone). The second term is the benefit to snobs: with conformist entry delayed, snobs face lower crowding costs for Δt_1 additional periods, earning additional pioneer rents. The third term is the benefit to all agents from improved quality assessment: Δt_1 additional periods of snob-dominated adoption generate more informative belief updates, improving $\hat{\theta}_{t_1}$ for the subsequent phases.

When $\alpha > \bar{\alpha}(\lambda)$ (the reversal regime), the second and third terms dominate the first. The intuition is that in this regime, conformist entry triggers a rapid snob exit cascade (Proposition 4). The value of delaying this cascade – preserving the discovery phase – exceeds the forgone bandwagon benefits for marginal conformists. Formally, the condition $\alpha > \bar{\alpha}(\lambda) = \beta(1-\lambda)/\lambda$ ensures that snob exit is sufficiently rapid that the welfare cost of premature conformist entry exceeds its benefit.

Part (ii): Comparative statics. The welfare loss from full transparency is $W(\varphi^*) - W(1)$. Differentiating with respect to α:

$$\frac{\partial[W(\varphi^*)-W(1)]}{\partial\alpha} > 0 \tag{90}$$

because higher α makes snob exit more sensitive to conformist entry, increasing the cost of the cascade that full transparency triggers. Similarly, the loss is decreasing in λ because with more snobs, the conformist mass $(1-\lambda)$ is smaller, so the magnitude of the conformist-acceleration externality under transparency falls mechanically; when snobs are abundant, transparency already preserves substantial snob-driven discovery, reducing the marginal value of extending it through opacity.

Part (iii): First-order condition. At the interior optimum φ^*, the FOC is:

$$\left.\frac{\partial W}{\partial \varphi}\right|_{\text{quality assessment}} + \left.\frac{\partial W}{\partial \varphi}\right|_{\text{conformist acceleration}} = 0 \tag{91}$$

The first term is always positive: more visibility helps all agents distinguish high-quality products from low-quality products, improving matching efficiency. The second term is negative in the reversal regime: more visibility accelerates conformist entry, compressing the lifecycle. At φ^*, these effects balance. □

Proof of Proposition 12 (Platform Patience and Visibility Distortion)

Proof. The platform maximizes $\Pi(\varphi; \delta_P) = \sum_t \delta_P^t r(n_t(\varphi))$ with $r' > 0$.

Part (i): Existence of $\bar{\delta}_P$. Define $M(\varphi, \delta_P) \equiv \partial \Pi / \partial \varphi = \sum_t \delta_P^t r'(n_t(\varphi)) \cdot (\partial n_t / \partial \varphi)$. By Lemma 4(i), higher visibility accelerates conformist entry, raising n_t in early periods but shortening the lifecycle. Decompose M into the per-period adoption effect and the lifecycle truncation effect:

$$M(\varphi, \delta_P) = \underbrace{\sum_{t=0}^{T(\varphi)} \delta_P^t r'(n_t) \frac{\partial n_t}{\partial \varphi}}_{\text{per-period effect}} + \underbrace{r(n_T) \delta_P^T \frac{\partial T}{\partial \varphi}}_{\text{duration effect}(<0)} \tag{92}$$

The per-period effect is positive in early periods (conformist entry raises n_t) and the duration effect is negative (shorter lifecycle, by Lemma 4(ii)). As $\delta_P \to 0$, only early periods matter: $M \to r'(n_0)(\partial n_0 / \partial \varphi) > 0$, so the platform wants to increase φ beyond φ_W^*; hence $\varphi_\Pi^* > \varphi_W^*$ (over-reveals). As $\delta_P \to 1$, the duration effect dominates: the platform values all periods equally, and losing late-lifecycle periods outweighs the per-period adoption gain; $M|_{\varphi_W^*} < 0$, so $\varphi_\Pi^* < \varphi_W^*$ (under-reveals). Since $M(\varphi_W^*, \delta_P)$ is continuous in δ_P, positive at $\delta_P = 0$ and negative at $\delta_P = 1$, the intermediate value theorem gives $\bar{\delta}_P \in (0, 1)$ with $M(\varphi_W^*, \bar{\delta}_P) = 0$. Uniqueness follows from $\partial M / \partial \delta_P < 0$ at the crossing: increasing patience raises the weight on late periods (where the lifecycle compression cost binds), monotonically shifting M downward.

Part (ii): Comparative statics. By the implicit function theorem applied to $M(\varphi_W^*, \bar{\delta}_P; \alpha, \lambda) = 0$:

$$\frac{\partial \bar{\delta}_P}{\partial \alpha} = -\frac{\partial M / \partial \alpha}{\partial M / \partial \delta_P} \tag{93}$$

The denominator $\partial M / \partial \delta_P < 0$ (as established above). For the numerator: increasing α accelerates snob exit, which shortens the lifecycle and reduces the per-period adoption gain (the conformist surge is briefer because snobs flee faster, depressing n_t in later periods). Both channels make $\partial M / \partial \alpha < 0$, so $\partial \bar{\delta}_P / \partial \alpha = (-)/(-) < 0$: higher snob aversion lowers the patience threshold.

For λ: increasing the snob share reduces the conformist mass $(1-\lambda)$, shrinking the per-period adoption gain from over-revealing (fewer conformists means a smaller surge). This makes $\partial M/\partial\lambda < 0$, so $\partial\bar{\delta}_P/\partial\lambda < 0$: more snobs lower the patience threshold. □

Proof of Corollary 6 (Quality and Optimal Visibility)

Proof. Let θ index product quality continuously, with higher θ corresponding to better products. Optimal visibility $\varphi^*(\theta)$ is implicitly defined by the FOC from Proposition 11. Differentiating the FOC with respect to θ:

$$\frac{\partial^2 W}{\partial\varphi^2}\cdot\frac{d\varphi^*}{d\theta}+\frac{\partial^2 W}{\partial\varphi\partial\theta}=0 \tag{94}$$

The second-order condition for a maximum requires $\partial^2 W/\partial\varphi^2 < 0$. The cross-partial $\partial^2 W/\partial\varphi\partial\theta$ captures how quality affects the marginal benefit of transparency. Higher quality has two effects on $\partial W/\partial\varphi$: it increases the acceleration cost (better products attract conformists faster, making conformist entry more damaging to the snob discovery phase) and it may increase the quality-assessment benefit (better matching is more valuable when products are better). In the reversal regime ($\alpha > \bar{\alpha}(\lambda)$), the acceleration channel dominates because the conformist entry cascade is rapid and compounds across periods, while the assessment benefit is bounded by the per-period quality difference $v(H)-v(L)$. Hence $\partial^2 W/\partial\varphi\partial\theta < 0$ in this regime.

Combining: $\frac{d\varphi^*}{d\theta} = -\frac{\partial^2 W/\partial\varphi\partial\theta}{\partial^2 W/\partial\varphi^2} < 0$.

Optimal visibility is decreasing in quality: better products should be less visible. □

Proof of Proposition 7 (Information Precision)

Proof. More precise signals (lower σ) shift the posterior distribution $G(\cdot;\hat{\theta}_t,\sigma)$ toward the true quality θ. When $\theta = H$, this shifts mass rightward, lowering the conformist threshold crossing time.

Conformist entry accelerates. Conformists adopt in significant numbers when the conformist threshold $c_C - \beta n_t$ falls into the bulk of the posterior distribution. More precise signals (lower σ) cause the public belief $\hat{\theta}_t$ to converge faster toward θ (each period's adoption signal is more informative). Under $\theta = H$, beliefs rise faster, shifting posteriors rightward earlier, which raises total adoption n_t and lowers $\underline{\mu}^C$. Conformist entry therefore occurs at a lower t_1: $\partial t_1/\partial\sigma > 0$.

Lifecycle shortens when $\alpha > \bar{\alpha}$. By Proposition 4, when $\alpha > \bar{\alpha}(\lambda)$, faster conformist entry shortens the lifecycle. Since lower σ accelerates conformist entry, $\partial T/\partial\sigma > 0$ in the reversal regime.

Snob welfare. Snob expected utility EU^S includes pioneer rents earned during $[0, t_1)$. When t_1 decreases (conformists enter earlier), the pioneer phase is compressed, reducing EU^S. The quality-assessment benefit ($\bar{v}_t^S$ increases with precision) is bounded, while the compression of pioneer rents compounds across periods when δ is large. □

Proof of Proposition 8 (Artificial Scarcity and Firm Profits)

Proof. Under market clearing, the firm earns discounted profits $\Pi^{\text{clear}}(\delta) = \sum_{t=0}^{T} \delta^t \pi_t$, where π_t is per-period profit and the lifecycle lasts T periods. Under the scarcity strategy with cap $\bar{n} < n^*$, adoption remains at $\bar{n}$ permanently (the snob-dominated phase never ends), yielding per-period profit $\pi^{\text{scarce}}(\bar{n})$ and total discounted profit $\Pi^{\text{scarce}}(\delta) = \pi^{\text{scarce}}(\bar{n})/(1-\delta)$.

The firm prefers scarcity iff $\Pi^{\text{scarce}}(\delta) > \Pi^{\text{clear}}(\delta)$:

$$\frac{\pi^{\text{scarce}}(\bar{n})}{1-\delta} > \sum_{t=0}^{T} \delta^t \pi_t \tag{95}$$

At $\delta = 0$: the LHS equals π^{scarce} while the RHS equals π_0; scarcity is not preferred when $\pi_0 > \pi^{\text{scarce}}$ (the firm would rather take the first-period profit and exit). As $\delta \to 1$: the LHS diverges to $+\infty$ while the RHS converges to $\sum_{t=0}^{T} \pi_t < \infty$; scarcity is strictly preferred. Since $\Pi^{\text{scarce}}(\delta)$ is strictly convex in δ and $\Pi^{\text{clear}}(\delta)$ is a polynomial in δ, there exists a unique crossing point $\bar{\delta}(\bar{n}) \in (0,1)$ such that scarcity is preferred iff $\delta > \bar{\delta}(\bar{n})$. The crossing point is implicitly defined by the indifference condition $\pi^{\text{scarce}}(\bar{n})/(1-\bar{\delta}) = \Pi^{\text{clear}}(\bar{\delta})$, or equivalently $\pi^{\text{scarce}}(\bar{n}) = (1-\bar{\delta})\Pi^{\text{clear}}(\bar{\delta}) \equiv \bar{\pi}^{\text{clear}}(\bar{\delta})$, where $\bar{\pi}^{\text{clear}}(\delta)$ denotes the annualized market-clearing profit. Since $\pi^{\text{scarce}}(\bar{n}) > 0$, we have $\bar{\delta} \in (0,1)$. The threshold is decreasing in π^{scarce} (higher scarcity profits favor scarcity) and increasing in the π_t's (higher market-clearing profits favor clearing). □

Proof of Corollary 7 (Cool Equilibria)

Proof. A cool equilibrium requires $n_t = n^{**} < n^*$ for all t, with snobs adopting and conformists deterred. This requires: (i) snob participation is positive at n^{**}: $\underline{\mu}^S(n^{**}) < 1$, i.e., $\alpha n^{**} + c_S < 1$; (ii) conformist participation is zero at n^{**}: $\underline{\mu}^C(n^{**}) > \bar{\mu}(\hat{\theta}_t)$, i.e., $c_C - \beta n^{**} > \bar{\mu}(\hat{\theta}_t)$; (iii) the snob adoption mass at n^{**} is self-sustaining: $\lambda[1 - G(\underline{\mu}^S(n^{**}))] = n^{**}$.

Condition (iii) defines n^{**} as a fixed point of the snob-only adoption function $\Phi^S(n) = \lambda[1-G(\alpha n + c_S)]$. Since Φ^S is continuous and strictly decreasing in n (because $\underline{\mu}^S(n) = \alpha n + c_S$ is increasing), with $\Phi^S(0) > 0$ and $\Phi^S(n) \to 0$ as $n \to (1-c_S)/\alpha$, it has a unique fixed point n^{**} by the intermediate value theorem. Conditions (i)–(ii) hold jointly when n^{**} is small, which requires λ small enough that the snob-only equilibrium does not cross the conformist entry threshold. Crucially, under $\theta = H$ the public belief $\hat{\theta}_t$ rises over time under snob-only adoption, so the conformist deterrence condition must hold at the *limiting* belief $\hat{\theta}_\infty$ (the highest belief reachable along the snob-only path): the binding condition is $n^{**} < (c_C - \bar{\mu}(\hat{\theta}_\infty))/\beta$. Since n^{**} is increasing in λ (more snobs implies higher snob-only adoption), there exists $\lambda^\dagger \in (0, 1)$ such that this inequality holds iff $\lambda < \lambda^\dagger$. The threshold $\lambda^\dagger$ depends on conformist-side parameters: high c_C, low β, and large σ (which keeps $\bar{\mu}(\hat{\theta}_\infty)$ from concentrating too close to 1) all make cool equilibria easier to sustain. □