Snobs and Conformists:
Platform Design and Product Lifecycles

Tommaso Bondi

Cornell University

February 23, 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 novel patterns compared to standard diffusion, social learning, and network-effects models: an endogenous boom-bust lifecycle, and a quality-duration reversal in which higher-quality products have *shorter* lifecycles, because superior quality accelerates conformist entry and triggers earlier snob exit. Quality is a common factor in preferences but a polarizing factor in equilibrium: the direct and indirect effects of quality reinforce each other for conformists (for whom popularity is a complement to quality) but partially cancel for snobs (for whom popularity is a substitute), making conformists the more quality-responsive segment despite identical quality tastes. We study two supply-side design problems. A profit-maximizing platform faces a tradeoff between broadcasting popularity to trigger conformist adoption and restricting visibility to preserve the discovery phase; optimal visibility is interior and decreasing in product quality. A seller optimally manages launch through scarcity, snob-targeted advertising, and type-based pricing that discounts early adopters and charges latecomers a premium.

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. 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. This single ingredient – opposing responses to popularity – generates a rich set of results on both the demand and the supply side.

The assumption is not far-fetched: influencer markets provide a vivid example. Early followers value the creator's authenticity and intimacy; their incentives are fundamentally different from those of later followers drawn by social proof and trending metrics. What makes the creator valuable to the first group – being undiscovered – is precisely what the second group's arrival destroys.

On the *demand side*, the equilibrium adoption path naturally exhibits a lifecycle with three phases. In a *discovery* phase, snobs try the product before it is widely consumed, generating 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 – "going viral" ends the niche phase – 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 and higher quality extends lifecycles. In our setting, popularity is a *two-sided force*: it increases demand among conformists while decreasing demand among snobs. We prove that the quality-duration reversal is impossible in any model where all consumers respond to popularity with the same sign, making it a diagnostic for opposing preferences.

The reversal also implies that the *composition* of the adopter base depends on quality in a surprising way. Quality is a common factor in preferences but a polarizing factor in equilibrium. For conformists, quality and popularity are complements: a better product attracts more adopters, which is what conformists want. For snobs, they are substitutes: a better product attracts the crowd that destroys exclusivity. Conformists are the more quality-responsive segment in equilibrium, even though both types have identical quality preferences.

We then study the *supply side*. First, we analyze a profit-maximizing *platform* that chooses how visible popularity is. A platform that broadcasts trending metrics triggers the conformist surge, capturing short-run adoption revenue – but compresses the lifecycle, destroying the discovery phase that sustains long-run engagement. Whether the platform over- or under-reveals relative to a welfare benchmark depends on its patience: impatient platforms over-reveal, while patient platforms in identity-heavy markets restrict visibility. Optimal visibility is decreasing in product quality, a prescription that reverses under social learning or network effects.

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 settings – platform and seller – policy that accelerates adoption can be privately attractive while shortening the period during which the product appeals to its early adopters.

A natural objection is that snobs might dislike popularity with conformists specifically, not popularity per se. Appendix A allows composition-sensitive preferences and shows that the core mechanism strengthens: each conformist entrant dilutes "coolness," amplifying the cascade. The same appendix verifies robustness to convex social preferences.

Roadmap. Section 2 reviews related literature. Section 3 presents the model. Sections 4–5 characterize the consumer equilibrium and derive the quality-duration reversal. Section 6 analyzes firm strategy. Section 7 studies platform strategy. Section 8 concludes. Appendix A develops extensions (composition-dependent preferences, convex social preferences). Appendix B contains proofs of the main results; the Online Appendix collects additional proofs.

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, to the best of our knowledge, 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 $+\beta n$ have substantial grounding: the snob side captures distinction motives (Veblen, 1899; Bourdieu, 1984), identity economics (Akerlof & Kranton, 2000), and signaling models (Pesendorfer, 1995); the conformist side captures network externalities (Katz & Shapiro, 1985), social conformity (Bernheim, 1994), and coordination motives. We take the coexistence of both types as primitive; our contribution is to study their *interaction* in a dynamic market.

Diffusion and social learning. The Bass diffusion model (Bass, 1969) is the workhorse framework for product adoption. Our model departs in two ways. First, Bass "innovators" adopt independently of the installed base; our snobs respond *negatively* to adoption, actively resisting growth. Second, Bass predicts that higher quality extends adoption; we show quality can *shorten* lifecycles. The "chasm" in technology adoption (Moore, 1991) corresponds to our Phase I–II transition, but our mechanism makes it endogenous. Social learning frameworks (Banerjee, 1992; Bikhchandani et al., 1992) assume uniformly positive externalities; our model nests this as $\alpha = 0$.

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. Nistor et al. (2024) model authenticity-monetization tensions generating growth-then-decline patterns, 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 Bayesian persuasion (Kamenica & Gentzkow, 2011; Bergemann & Morris, 2019); to the best of our knowledge, the optimal signal has not previously been studied in a setting where information accelerates a *dynamic* composition shift.

Empirical evidence. Several empirical findings align with the model's predictions. Berger & Heath (2007, 2008) show that consumers abandon products when outgroup members adopt – the micro-level mechanism underlying snob exit. Berger & Le Mens (2009) document that cultural tastes adopted more quickly die out faster, consistent with the quality-duration reversal. Bellezza (2023) finds that luxury products lose status as they diffuse, providing direct evidence for the crowding externality. 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. Schoenmueller et al. (2021) document that influencer follower counts follow bell-shaped lifecycles consistent with our three-phase structure, and Valsesia et al. (2020) show that following fewer others signals autonomy and increases perceived influence – a microfoundation for snob behavior on platforms. Instagram's experiment hiding like counts (Mosseri, 2019) provides a natural test of our visibility predictions.

3 Model

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.

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 $\gamma_i \in \mathbb{R}$.

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 $(0,1)$.

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.[1] Conformists derive utility:

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

where $\beta > 0$ measures conformity preference.

Each agent has access to an outside option yielding type-specific flow payoff c_τ, with $c_C > c_S \geq 0$. The ordering has a natural microfoundation: the more primitive specification $U^S = v(\theta) + \kappa - \alpha n$, where $\kappa > 0$ is an exclusivity premium, yields $c_S \equiv c - \kappa$ for a common outside option c. The ordering $c_C > c_S$ is therefore the reduced form of snobs deriving positive utility from exclusivity.

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) $c_C > c_S \geq 0$; (iii) $c_S < 1$ (snobs adopt at $n = 0$ for sufficiently favorable signals) and $c_C > \beta$ (not all conformists adopt even at full adoption); and (iv) $\sigma > 0$.

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.

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

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 strategies are:*

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

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

where c_S and c_C are type-specific reservation utilities.

The snob threshold $\underline{\mu}^S(n) = \alpha n + c_S$ is increasing in n: more adoption raises the quality bar for snobs. The conformist threshold $\underline{\mu}^C(n) = c_C - \beta n$ is decreasing: more adoption lowers the bar for conformists.[2] 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:*

[2] This opposing monotonicity appears in Corneo & Jeanne (1997a)'s static analysis. Our contribution is, to the best of our knowledge, the first to embed it in a dynamic environment with learning, producing non-monotone lifecycle paths and an impossibility result.

(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.*

(ii) Conformist mass increasing. *Conformist adoption mass $(1-\lambda)[1-G(\underline{\mu}^C(n))]$ is strictly increasing in n: as n rises, the conformist threshold $\underline{\mu}^C(n) = c_C - \beta n$ falls, admitting more posteriors. The mass is negligible when $\underline{\mu}^C(n)$ is close to 1 (few posteriors exceed a high bar) and approaches $(1-\lambda)$ as $\underline{\mu}^C(n)$ approaches zero (nearly all posteriors exceed it).*

(iii) Simultaneous opposition. *For $n \in (0, \bar{n}^S)$, snob adoption mass is strictly decreasing in n ($\partial n^S/\partial n < 0$) while conformist adoption mass is strictly increasing ($\partial n^C/\partial n > 0$).*

As n rises, the snob threshold $\underline{\mu}^S(n) = \alpha n + c_S$ increases while the conformist threshold $\underline{\mu}^C(n) = c_C - \beta n$ decreases. Since posteriors have full support on $(0, 1)$ (Assumption 1), both types always have strictly positive adoption mass when their threshold is below 1 – but the *magnitude* of that mass moves in opposite directions with n. The snob adoption mass $\lambda[1 - G(\alpha n + c_S)]$ is strictly decreasing in n (the density $g > 0$ ensures strict monotonicity), while the conformist adoption mass $(1-\lambda)[1 - G(c_C - \beta n)]$ is strictly increasing. This opposing monotonicity is the central structural feature of the model, directly supported by evidence in Iyengar et al. (2011) and Berger & Heath (2008).[3]

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 $\Phi : [0, 1] \to [0, 1]$ is continuous, with $\Phi(0) > 0$ and $\Phi(1) < 1$. Consequently there exists at*

[3] 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.

least one equilibrium $n^* \in (0,1)$ *satisfying* $\Phi(n^*) = n^*$.

Continuity follows from G continuous and thresholds linear in n. For the endpoint inequalities: at $n = 0$, the snob threshold is $c_S < 1$ (Assumption 2), so $\Phi(0) \geq \lambda[1-G(c_S)] > 0$ since $G(c_S) < 1$ by full support. At $n = 1$, the snob threshold $\alpha + c_S > 1$ puts the snob term at zero. The conformist threshold $c_C - \beta > 0$ (Assumption 2) is strictly positive, so $G(c_C - \beta) > 0$ by full support, giving $(1-\lambda)[1-G(c_C-\beta)] < (1-\lambda) < 1$. Hence $\Phi(1) < 1$. Since $\Phi(0) > 0 = \mathrm{id}(0)$ and $\Phi(1) < 1 = \mathrm{id}(1)$, the continuous function $\Phi(n) - n$ changes sign on $[0,1]$, and any zero is interior. 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.[4]

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.

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

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.

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.

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 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. 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).*

The key implication of flow adoption and canceling continuation values is that the equilibrium can be solved period by period: given the current public belief $\hat{\theta}_t$, agents play the static game from Section 4, and the resulting adoption determines the signal that updates beliefs for the next period.

5.3 Three-Phase Lifecycle

What does the equilibrium path look like? When the product is high quality, beliefs drift upward, conformists eventually enter, and the resulting crowding drives snobs out. The next proposition formalizes this.

Proposition 3 (Lifecycle Characterization)**.** *Suppose $c_C > c_S$ and $\theta = H$. If $\alpha < \bar{\alpha}(\mu, \lambda)$, 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; the lifecycle ends at $T = \min\{t > t^* : n_t^S < \epsilon\}$.*

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

We define lifecycle duration as $T \equiv \min\{t > t^* : n_t^S < \epsilon\}$ for $\epsilon > 0$ – the length of the period during which snob participation remains positive. 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 settle, snobs do not return because the belief path is absorbing. Fashion cycles in the sense of Simmel (1957) – recurring waves of adoption and abandonment – emerge from *sequential* product introductions: as product A enters decline, snobs migrate to product B.

In Phase I, snobs adopt based on private quality signals, generating informative public signals that raise beliefs. 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. 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. The same pattern appears in influencer markets: early followers disengage as the audience goes mainstream (Schoenmueller et al., 2021).

The threshold $\bar{\alpha}(\mu, \lambda)$ is increasing in posterior dispersion (more dispersed beliefs keep enough mass above the snob threshold even as it rises) and decreasing in λ (more snobs intensify crowding). The lifecycle is most robust when quality uncertainty is moderate: enough signal to learn, but enough noise to prevent the snob threshold from binding too tightly.

Remark 1 (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. In Figure 1(b), snobs exceed conformists at every point in time; the product persists as a quiet niche because the mainstream never arrives to either amplify or crowd out early adopters.

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.

Corollary 1 (Peak Adoption)**.** *At any stable within-period equilibrium, $n^* = \Phi(n^*; \hat{\theta}_{t^*})$. The comparative statics $\partial n^*/\partial \alpha < 0$ and $\partial n^*/\partial \beta < 0$ hold exactly by the implicit function theorem, without requiring Assumption 4.*

The IFT gives $\partial n^*/\partial \alpha = -\lambda g(\underline{\mu}^S) n^*/(1 - \Phi'(n^*)) < 0$ at any stable equilibrium ($1 - \Phi' > 0$); the argument for β is analogous. These signs are the workhorse for all subsequent results; the closed forms below provide intuition but are not required.

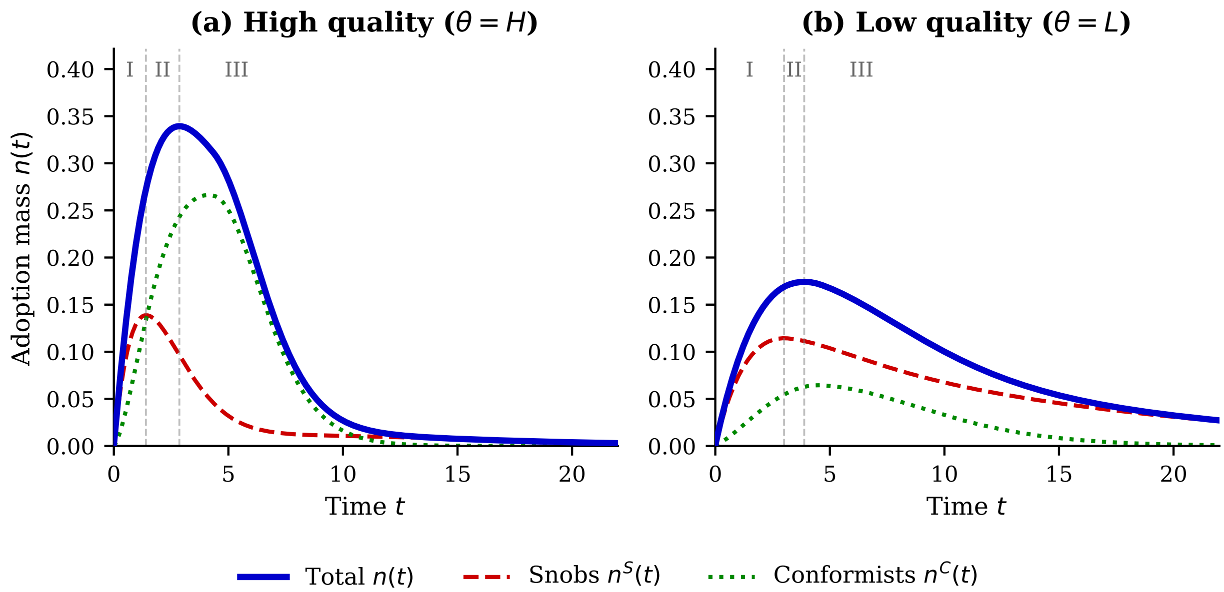

Figure 1: Equilibrium lifecycle dynamics. Panel (a): high-quality product; panel (b): low-quality product. High quality produces a larger but shorter cycle (the quality-duration reversal). Parameters: $\alpha = 1.5$, $\beta = 1.0$, $c_S = 0.05$, $c_C = 0.40$, $\lambda = 0.50$, $p_0 = 0.45$.

Remark 2 (Closed-Form Approximation). Under Assumption 4, the threshold crossing point $n^\dagger \equiv (c_C - c_S)/(\alpha + \beta)$ serves as a useful proxy for n^*. This approximation has two motivations. First, at $n^\dagger$ the two thresholds coincide ($\underline{\mu}^S = \underline{\mu}^C$), so marginal snobs and marginal conformists face the same quality bar – a natural benchmark for peak equilibrium when both types are active and the density is locally flat. Second, in numerical simulations (Figure 1), the equilibrium peak tracks $n^\dagger$ closely across parameter values. The proxy $n^\dagger$ is not a derived equilibrium quantity; it is a convenient heuristic that we use when presenting closed-form expressions for thresholds such as $\bar{\alpha}(\lambda)$ in Proposition 4. All qualitative results – the reversal, the impossibility, the comparative statics – depend only on the IFT signs established in Corollary 1, not on the approximation.

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* $\bar{\alpha}(\mu, \lambda) \in (0, \infty)$ *such* *that: if $\alpha < \bar{\alpha}$, higher quality extends the lifecycle ($T^H > T^L$); if $\alpha > \bar{\alpha}$, higher quality shortens it ($T^H < T^L$). Under Assumption 4, the threshold has the closed form $\bar{\alpha}(\mu, \lambda) = \beta(1-\lambda)/\lambda$.*

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. The threshold $\bar{\alpha}$ is decreasing in λ (fewer snobs make the reversal easier to trigger, since the conformist mass that drives the acceleration is larger) and increasing in β (stronger conformist preferences raise the bar snob aversion must clear, since bandwagon effects sustain adoption longer). It is increasing in μ when the prior is high enough that conformists enter quickly regardless of quality, muting the differential acceleration.

The quality-duration reversal is *impossible* in a broad class of adoption models that lack this preference heterogeneity.

Proposition 5 (Impossibility of Reversal Under Homogeneous Preferences)**.** *Consider any adoption model with flow utility*

$$u_i(\theta, n) = v(\theta) + \gamma_i \cdot h(n), \quad v' > 0,\ h' > 0 \tag{12}$$

threshold strategies based on private posteriors satisfying MLRP, and Bayesian belief updating. If γ_i has the same sign for all agents, then $T(\theta^H) \geq T(\theta^L)$.

The "same sign" condition nests the leading adoption frameworks as special cases. Bass diffusion (Bass, 1969), social learning (Banerjee, 1992; Bikhchandani et al., 1992), and network effects (Katz & Shapiro, 1985) all feature $\gamma_i \geq 0$ for every agent: adoption is (weakly) self-reinforcing. Congestion models have $\gamma_i \leq 0$ for all. Private-values models set $\gamma_i = 0$. In every case, higher quality generates weakly higher adoption at all t, so duration is weakly increasing in quality – regardless of heterogeneity in $|\gamma_i|$, signal precision, or outside options. The reversal $T^H < T^L$ therefore requires mixed-sign externalities: $\gamma_i > 0$ for some agents and $\gamma_j < 0$ for others. The impossibility makes the reversal a diagnostic: observing that higher-quality products cycle faster points specifically to the coexistence of consumers who value popularity and consumers who flee from it.[5]

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 (Appendix A.2).[6]

5.6 Who Responds to Quality?

The reversal has a striking implication for the *composition* of the adopter base. Quality is a common factor in preferences – the $v(\theta)$ term enters both types' utility functions identically, so neither snobs nor conformists have inherently stronger *tastes* for quality. Yet quality is a polarizing factor in equilibrium, because it affects adoption through two channels that point in opposite directions for the two types.

The *direct channel* is straightforward: a better product raises utility for everyone. The *indirect channel* operates through the market: higher quality generates stronger signals, which accelerate learning, which raises adoption mass n. For conformists, the indirect channel reinforces the direct one – a better product is also a more popular one, and popularity is exactly what conformists want. For snobs, the two channels oppose each other: a better product is also a more *crowded* one, and crowding is exactly what snobs flee.

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

$$\frac{dn_t^C}{d\hat{\theta}} > 0 \quad \textit{always}; \qquad \frac{dn_t^S}{d\hat{\theta}} \lesseqgtr 0 \iff \alpha \lesseqgtr \frac{\partial_{\hat{\theta}}[1 - G(\underline{\mu}^S)]}{g(\underline{\mu}^S) \cdot dn^*/d\hat{\theta}} \tag{13}$$

Conformist adoption is unambiguously increasing in perceived quality. Snob adoption is increasing in quality only when crowding aversion is weak; in the reversal regime ($\alpha > \bar{\alpha}$), higher quality reduces *snob adoption.*

[5] 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. Both patterns are consistent with the reversal.

[6] The key requirement is that snob disutility grows at least proportionally with conformist utility as adoption rises. The reversal requires $\alpha > \bar{\alpha}(\lambda)$; uniqueness (Assumption 3) 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.

The intuition is immediate. Each type's equilibrium response to quality decomposes into a direct effect (better signals $\rightarrow$ more agents above threshold) and an indirect effect (higher n shifts thresholds). For conformists, higher n *lowers* the threshold $c_C - \beta n$, amplifying entry. For snobs, higher n *raises* the threshold $\alpha n + c_S$, triggering exit. When α is large enough – the reversal regime – the indirect channel dominates for snobs, and better quality drives them out.[7]

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

This pattern is visible across markets. A restaurant that earns a rave review or a Michelin star experiences a surge in reservations – but the regulars who valued the intimacy of the uncrowded dining room begin to leave. The restaurant is objectively better-validated, but the equilibrium crowd makes it worse *for the type of diner who discovered it*. In music, a band that achieves mainstream success through a hit single or festival appearance often loses its core following – not because the music changed, but because the audience did (Berger & Le Mens, 2009). Independent record stores, small galleries, and neighborhood bars face the same dynamic when positive media attention triggers an influx that alters the character of the venue.

The result overturns the common intuition that snobs are the "quality-sensitive" segment. In terms of *preferences*, both types respond identically to quality. In terms of *equilibrium behavior*, it is conformists whose adoption most strongly tracks quality – because for them, quality and popularity are complements. For snobs, quality and popularity are substitutes: quality is self-defeating because excellence attracts the crowd that destroys the exclusivity they value.

Two additional comparative statics sharpen the picture. First, the relationship between snob share λ and cycle duration is non-monotonic.

[7] Under Assumption 4, the condition simplifies: $dn^S/d\hat{\theta} < 0$ when $\alpha > \bar{\alpha}(\lambda)$, the same threshold that governs the duration reversal. This is not a coincidence – the reversal *is* the dynamic manifestation of the quality-elasticity asymmetry.

Proposition 6 (Non-Monotonic Composition Effects)**.** *Duration is maximized at an interior snob share* $\lambda^* = \beta/(\alpha + \beta)$*.*

The non-monotonicity reflects a fundamental tension. Too few snobs and products are never discovered; too many and products cycle too rapidly. Snobs serve two roles in the ecosystem: they are the discoverers who generate early quality signals, and they are the segment whose exit defines the end of the culturally active lifecycle. More snobs accelerate discovery but also accelerate crowding. The interior optimum $\lambda^* = \beta/(\alpha + \beta)$ equalizes the marginal value of an additional discoverer against the marginal cost of faster crowding. Fashion-forward categories like streetwear operate near high λ, cycling rapidly; mass consumer goods operate near low λ, with slow discovery and long lifecycles.
Second, better information can paradoxically hurt early adopters.

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

The mechanism 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 on others. Here, more precise signals help conformists identify good products faster, which is exactly the acceleration that crowds snobs out. The channel runs through the same conformist-entry externality that drives the lifecycle: anything that helps conformists learn quality faster compresses the discovery phase. This creates a tension for platforms that control information precision – a tension we formalize in Section 7.

Finally, the *composition* of the market itself has welfare implications independent of any particular instrument.

Proposition 8 (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, so the market loses its capacity for discovery. With only snobs ($\lambda = 1$), no bandwagon externality amplifies successful products, so the market loses its capacity for scale. The optimal λ^{**} balances these roles: stronger snob aversion lowers the optimal share (crowding is more costly per snob), while noisier signals raise it (snob-driven discovery is more valuable when agents cannot learn well on their own).

6 Implications for Firm Strategy

The lifecycle creates both an opportunity and a constraint. A monopolist who controls pricing, supply, and marketing can exploit the composition shift, but instruments that accelerate adoption can backfire by compressing the culturally active phase. Let the firm maximize discounted profits $\Pi = \sum_{t=0}^{T} \delta^t \pi(n_t)$, where $\delta \in (0,1)$ is the discount factor and $\pi(n_t) = p_t \cdot n_t$ is per-period revenue.

6.1 Pricing

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:

$$p_t = \begin{cases} \hat{\theta}_t - \alpha n_t - c_S & t < t_1 \text{ (Phase I)} \\ \hat{\theta}_t + \beta n_t - c_C & t \geq t_1 \text{ (Phases II–III)} \end{cases} \tag{14}$$

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. This contrasts with network-effects pricing (Katz & Shapiro, 1985), where prices typically increase as the installed base grows.

If the firm can identify types, it should discount snobs and charge conformists a premium. With observable types, myopically optimal prices yield a gap $p_t^C - p_t^S = (\alpha + \beta)n_t - (c_C - c_S)$, which is positive for $n_t > n^\dagger$ and maximal at peak adoption.[8] This reverses network-effects pricing, where early adopters are subsidized to *build* the network rather than *preserve*

[8] 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.

exclusivity. 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. 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. Type-specific pricing is an interesting extension in its own right: with a richer action space (menus, bundling, dynamic personalization), the firm could potentially screen types within each period, extracting the full composition-dependent surplus. We leave this to future work and focus here on the uniform price path, which already illustrates the key departure from network-effects logic.

6.2 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 9 (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.[9]

The mechanism is subtle. At the unconstrained optimum, the firm earns high per-period revenue during the conformist surge but loses the snob-dominated phase permanently. Under scarcity, per-period revenue is lower (fewer adopters) but the stream is indefinite: the product never enters Phase II, so snobs never exit. The comparison reduces to a geometric series against a finite sum, and the geometric series dominates when δ is large enough. The scarcity cap works not by restricting supply per se but by preventing adoption from crossing $n^\dagger$, keeping the product in the region where snob thresholds are not yet binding.

[9] 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.

The result inverts the standard logic. 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. Hermès waitlists, Supreme's limited drops, and Berghain's capacity caps all implement this logic.

6.3 Advertising and Targeting

Sellers of fashion and identity goods face a distinctive advertising problem. Unlike sellers of utilitarian products, who simply want to maximize awareness, a fashion brand must consider that advertising changes not only *how many* consumers learn about the product but also *which types* are reached. A billboard campaign reaches conformists and snobs indiscriminately; an influencer placement or a niche magazine ad can be targeted. The composition effects documented above – the quality-duration reversal, the asymmetric quality elasticity, the non-monotonic role of market composition – suggest that both the *volume* and the *targeting* of advertising matter for lifecycle dynamics.

We model advertising as a posterior shift: advertising with effectiveness $a > 0$ raises each agent's posterior by a, equivalently lowering adoption thresholds by a. This captures the informational role of advertising (making consumers more aware of quality) while remaining agnostic about the specific channel. When the firm can direct advertising toward a specific type, let $\tau \in [0, 1]$ denote the share directed at snobs. In standard diffusion models, broad advertising is unambiguously profit-increasing: more awareness means faster adoption means stronger externalities. Here, advertising that reaches the wrong type can compress the lifecycle by accelerating the conformist entry that triggers snob exit.

Proposition 10 (Optimal Targeting)**.** *Let $\tau \in [0, 1]$ denote the share of advertising directed at snobs. The optimal targeting strategy satisfies:*

$$\tau^* = \min\left\{1,\ \underbrace{\frac{\alpha}{\alpha+\beta}}_{\textit{static term}} + \underbrace{\frac{\delta}{1-\delta}\cdot\frac{\lambda}{1-\lambda}\cdot\frac{\partial T/\partial n^S}{|\partial T/\partial n^C|}}_{\textit{dynamic term}}\right\} \tag{15}$$

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

The formula has an intuitive decomposition. The static term $\alpha/(\alpha+\beta)$ is the snob share that would maximize single-period adoption: it reflects the relative strength of crowding aversion versus bandwagon preference. If this were the only consideration, targeting would simply be proportional to preference intensities, and there would be no reason to favor one type over the other beyond their relative responsiveness. But the dynamic term adds a lifecycle motive that breaks this symmetry. Its numerator, $\partial T/\partial n^S$, captures how much an additional snob extends the culturally active phase – snob advertising sustains the discovery phase and delays the conformist takeover. Its denominator, $|\partial T/\partial n^C|$, captures how much an additional conformist shortens it. The ratio is multiplied by the relative population factor $\lambda/(1-\lambda)$, reflecting that the lifecycle-extension benefit of snob-targeted advertising scales with the snob share of the market, and by the patience factor $\delta/(1-\delta)$, reflecting that lifecycle extension is more valuable to patient firms.

There are thus two distinct reasons to target snobs disproportionately, beyond the static preference-intensity argument. First, snob-targeted advertising extends the lifecycle ($\partial T/\partial n^S > 0$), while conformist-targeted advertising compresses it ($\partial T/\partial n^C < 0$). Second, snobs generate informational externalities: their adoption produces quality signals that raise conformist willingness to pay in later periods. Conformist advertising generates no such externality, because conformist adoption is driven by bandwagon effects rather than quality assessment. 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.

The *total* level of advertising is also affected by the lifecycle.

Proposition 11 (Advertising Volume)**.** *When $\alpha < \bar{\alpha}(\mu, \lambda)$, the profit-maximizing total advertising level is increasing in a (advertising unambiguously helps). When $\alpha > \bar{\alpha}(\mu, \lambda)$ and* *the firm cannot target by type, the profit-maximizing level satisfies $a^* < a^{myopic}$: the firm optimally advertises less than the single-period optimum.*

In the reversal regime, untargeted advertising raises per-period adoption but shortens the lifecycle. A firm that cannot distinguish types faces a tradeoff: each additional unit of advertising attracts conformists who compress the culturally active phase. The optimal response is to advertise less than myopic profit maximization would suggest, or in the extreme to forgo advertising entirely. This provides a rationalization for the common observation that prestige brands advertise sparingly and selectively – not because advertising is ineffective, but because the lifecycle compression cost outweighs the per-period revenue gain.

7 Implications for Platform Strategy

While the firm internalizes the lifecycle through pricing, scarcity, and targeting, platforms face a distinct design problem. Platforms typically do not sell the product directly; instead, they control the *information environment* in which adoption decisions are made. Every design choice that makes popularity more salient – trending badges, play counts, follower numbers, algorithmic amplification – feeds information to conformists about whether the adoption threshold has been crossed. The question is how a profit-maximizing platform should set this visibility.

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.

Lemma 3 (Equilibrium Response to Visibility)**.** *Higher visibility accelerates conformist entry ($\partial t_1/\partial\varphi < 0$). When $\alpha > \bar{\alpha}(\lambda)$, this compresses the lifecycle ($\partial T/\partial\varphi < 0$). The equilibrium path and welfare are smooth in φ.*

The platform earns per-period revenue $r(n_t)$ with $r' > 0$ – platforms monetize through advertising, commissions, or subscriptions that scale with engagement. Higher visibility triggers the conformist surge, boosting short-run revenue, but also compresses the lifecycle, destroying the discovery phase that sustains long-run engagement. The platform's problem is to choose φ to maximize $\Pi(\varphi) = \sum_t \delta_P^t r(n_t(\varphi))$, where $\delta_P \in (0,1)$ is its discount factor.

To characterize the profit-maximizing solution, it is useful to establish a welfare benchmark. Let $W(\varphi) = \lambda \cdot EU^S(\varphi) + (1-\lambda) \cdot EU^C(\varphi)$ denote social welfare.

Proposition 12 (Welfare Benchmark)**.** *Under the conditions of Proposition 4 with $\alpha > \bar{\alpha}(\lambda)$:*

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

(ii) The welfare cost of full transparency, $W(\varphi_W^*) - W(1)$, is larger when snob aversion is strong (increasing in α) and when snobs are scarce (decreasing in λ).

(iii) The welfare-maximizing platform reveals more when private signals are noisy ($\partial\varphi_W^*/\partial\sigma > 0$) and less when snob aversion is strong ($\partial\varphi_W^*/\partial\alpha < 0$).

The welfare benchmark establishes that full transparency is suboptimal: it helps agents assess quality but accelerates the conformist entry that compresses the lifecycle. Under social learning or network effects, transparency is unambiguously welfare-improving; the prescription to *restrict* visibility arises only when some consumers respond negatively to popularity. The benchmark is useful because it pins down the direction of the profit-maximizing platform's distortion.

Proposition 13 (Profit-Maximizing Visibility)**.** *Let φ_W^* denote welfare-maximizing visibility. Under the conditions of Proposition 12:*

(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 ($\partial\bar{\delta}_P/\partial\lambda < 0$).

The result captures a basic tradeoff in platform design. An impatient platform over-reveals: it broadcasts trending badges and play counts to trigger the conformist surge, capturing short-run revenue at the cost of compressing the lifecycle. A patient platform under-reveals: it suppresses popularity metrics to preserve the discovery phase that generates sustained engagement. 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 should over-reveal unless unusually patient – and indeed, TikTok and Instagram aggressively broadcast trending metrics.

Corollary 3 (Optimal Visibility Decreasing in Product Quality)**.** *In the regime where the quality-duration reversal holds, the profit-maximizing visibility satisfies* $d\varphi^*_\Pi/d\theta < 0$*.*

Better products attract conformists faster, making the lifecycle compression cost of visibility larger. A streaming platform should display fewer popularity metrics for premium content (curated playlists, acclaimed releases) than for mass content (viral hits, trending tracks). The firm's scarcity strategy (Proposition 9) and the platform's opacity are complements: Supreme's limited drops implement the quantity cap; Berghain's combination of restricted entry and a photography ban illustrates the complementarity.

8 Conclusion

This paper develops a dynamic model of markets where consumers disagree about whether popularity is a feature or a bug. The central result is the quality-duration reversal: in markets with strong originality preferences, better products have *shorter* culturally active lifecycles, because excellence attracts the conformist crowd that drives early adopters away. An impossibility theorem establishes that this reversal requires mixed-sign adoption externalities, making it a diagnostic for the presence of opposing preferences.

The model yields concrete prescriptions for both sellers and platforms. For sellers: restrict supply rather than maximize penetration (Hermès waitlists, Supreme drops); target marketing toward early adopters rather than the mass market (seed influencers, not billboards); and price to subsidize snobs while extracting surplus from conformists (early-access discounts, premium late pricing). Each of these reverses the standard network-effects playbook, where the goal is to build the installed base as fast as possible. For platforms: hide trending metrics in identity-heavy categories (Bandcamp's minimal popularity display), reveal them in utilitarian ones (Amazon's bestseller lists), and display *less* information about higher-quality products – a counterintuitive prescription that emerges because the lifecycle compression cost of visibility is largest where quality is highest.

Several extensions would sharpen the practical relevance. Overlapping lifecycles, where snobs fleeing product A pioneer product B, would connect the single-product analysis to portfolio dynamics and category management. The joint optimization problem – where firm and platform simultaneously choose quantity and information policies – remains open and is directly relevant to vertically integrated platforms (Spotify deciding both what to promote and how much play-count information to display). The visibility parameter φ is a special case of a richer information design problem: a platform 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., & 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.

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.

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.

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

Appendix A: Extensions

A.1 Composition-Dependent Preferences

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

Decompose total adoption into components $n_t = n_t^S + n_t^C$ and define coolness as $C_t \equiv n_t^S - \xi n_t^C$, where $\xi > 0$ captures how much conformist adoption dilutes cachet. Snob utility becomes $U^S = v(\theta) - \alpha n_t + \eta C_t$, where $\eta \geq 0$ captures how much snobs value coolness per se. Conformist utility is unchanged.

Proposition 14 (Coolness-Amplified Dynamics)**.** *With composition-dependent coolness ($\eta > 0$):*

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

(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 and becomes negative in Phase II; products are maximally popular when minimally cool.

(iv) A cool equilibrium – where snobs adopt permanently and conformists are deterred – exists when c_C is large, β is small, or σ is large.

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

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

Figure 2 illustrates the lifecycle under composition-dependent preferences using parameters $\eta = 2.0$, $\xi = 1.0$ (so snob-on-snob crowding is $\alpha - \eta = -0.5$ and crowding from conformists is $\alpha + \eta = 3.5$). The main findings are robust: the three-phase lifecycle and quality-duration reversal are preserved, with several quantitative differences. Snob peaks are substantially higher (negative self-crowding makes the early phase more attractive), snob exit is dramatically steeper (the "coolness cascade" – each conformist raises effective crowding by $\alpha + \eta\xi$), and conformist peaks are higher (sharper snob exit leaves a niche conformists fill aggressively). Under low quality, the model predicts a "return to niche": after a brief conformist interlude, the product cycles back to snob dominance as conformists decay faster than snobs.

A.2 Robustness to Functional Form

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

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

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

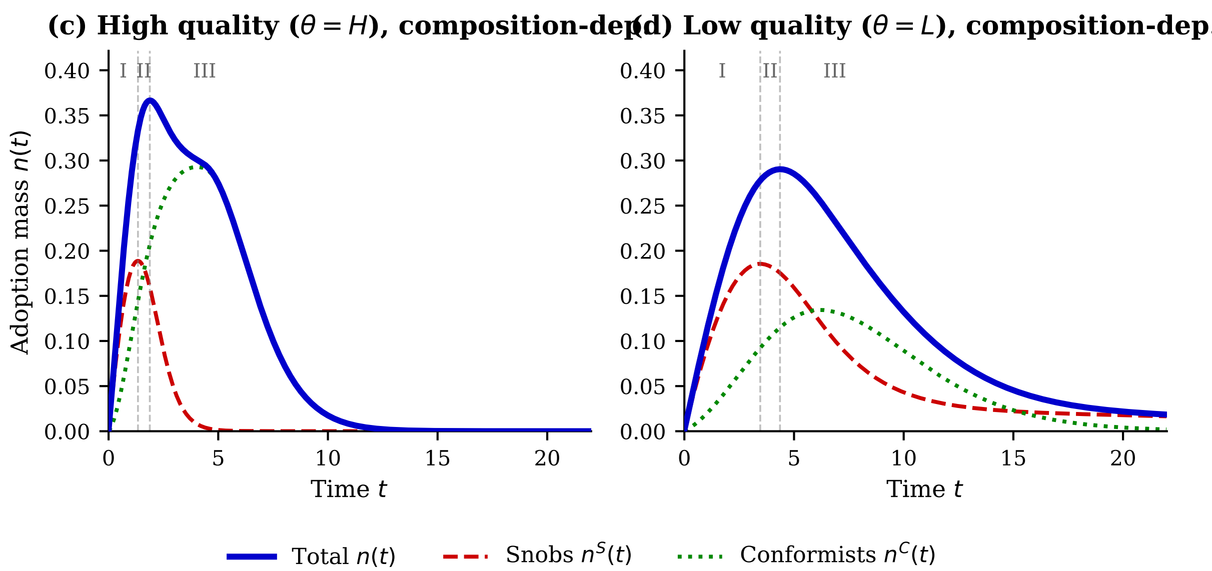

Figure 2: Lifecycle dynamics under composition-dependent preferences ($\eta = 2.0$, $\xi = 1.0$). Compared with Figure 1: snob peaks are 36–63% higher, snob half-life is 50–59% shorter (the coolness cascade), and conformist peaks are higher. All other parameters as in Figure 1.

posteriors, the reversal condition has an exact closed form,

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

which is finite and well-defined without any density approximation. For $\lambda = 1/3$, $\beta = 1$, $c = 1$, this gives $\bar{\alpha}_Q = 2$. The impossibility result (Proposition 5) carries through unchanged because it depends only on the sign of γ_i, not on functional form.

Cycle shape. The main qualitative difference under convex preferences is in cycle *shape*. The linear model produces relatively symmetric adoption curves. With $\rho > 1$, the marginal crowding cost $\partial U^S/\partial n = -\alpha\rho n^{\rho-1}$ is negligible at low n but accelerates sharply near the peak, producing right-skewed curves: gradual ascent and rapid descent. The growth phase extends because early crowding is tolerable; the decline phase compresses because exit cascades once crowding costs bite. This right-skewed pattern matches the empirical profile of many fashion and cultural products, where trends build gradually over seasons then crash within weeks, and social platforms grow steadily for years then collapse rapidly.

Appendix B: 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: $\lambda \in (0, 1)$.

B.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:

$$\underbrace{\mu - \alpha n}_{\text{adopt}} + \delta \mathbb{E}[V^S \mid \text{adopt}] \geq \underbrace{c_S}_{\text{outside option}} + \delta \mathbb{E}[V^S \mid \text{wait}] \tag{16}$$

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:

$$\mu - \alpha n \geq c_S \quad \Rightarrow \quad \mu \geq \alpha n + c_S \equiv \underline{\mu}^S(n) \tag{17}$$

[10]

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

$$\mu + \beta n \geq c_C \quad \Rightarrow \quad \mu \geq c_C - \beta n \equiv \underline{\mu}^C(n) \tag{18}$$

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

[10] 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).

$c_C > c_S$ (both hold by Assumption 2). $\square$

B.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 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:

$$\frac{d\pi^S}{dn} = -\alpha g(\alpha n + c_S) < 0 \tag{19}$$

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 mass increasing. Define conformist participation as $\pi^C(n) = (1 - \lambda)[1 - G(c_C - \beta n)]$. Since $c_C - \beta n$ is decreasing in n and $g > 0$:

$$\frac{d\pi^C}{dn} = (1 - \lambda)\beta g(c_C - \beta n) > 0 \tag{20}$$

Conformist adoption mass is strictly increasing in n. At $n = 0$, the threshold is $c_C > 0$ and $\pi^C(0) = (1-\lambda)[1-G(c_C)]$, which is positive but small when c_C is close to 1 (most posteriors fall below the threshold). As $n \to c_C/\beta$, the threshold drops to zero and $\pi^C \to (1 - \lambda)$ (all conformists adopt).

Part (c): Simultaneous opposition. For $n \in (0, \bar{n}^S)$, both types have positive adoption mass and $d\pi^S/dn = -\lambda\alpha g(\alpha n + c_S) < 0$ (Part (a)) while $d\pi^C/dn = (1 - \lambda)\beta g(c_C - \beta n) > 0$ (Part (b)). Strict monotonicity follows from $g > 0$ (Assumption 1). $\square$

B.3 Proof of Lemma 2 (Equilibrium Existence)

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

$$\Phi(n) = \lambda \cdot [1 - G(\underline{\mu}^S(n))] + (1-\lambda) \cdot [1 - G(\underline{\mu}^C(n))] \tag{21}$$

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:

$$\Phi(n; \hat{\theta}) = \lambda[1 - G(\alpha n + c_S; \hat{\theta})] + (1-\lambda)[1 - G(c_C - \beta n; \hat{\theta})] \tag{22}$$

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 $c_S < 1$ (Assumption 2) and posteriors have full support, $\Phi(0) \geq \lambda[1 - G(c_S)] > 0$. At $n = 1$: since $\alpha + c_S > 1$ (trivially, as $\alpha > 0$ and $c_S \geq 0$ with $\alpha + \beta > c_C - c_S > 0$), the snob threshold $\underline{\mu}^S(1) = \alpha + c_S > 1$, so no snob adopts. For conformists, $\underline{\mu}^C(1) = c_C - \beta > 0$ (Assumption 2), so the conformist threshold is strictly positive and $1 - G(c_C - \beta) < 1$. Hence $\Phi(1) = (1-\lambda)[1 - G(c_C - \beta)] < 1$.

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 $\Phi(n^*) = n^*$. □

B.4 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:

$$\frac{\hat{\theta}_1}{1 - \hat{\theta}_1} = \frac{\hat{\theta}_0}{1 - \hat{\theta}_0} \cdot \frac{f(n_0 \mid H)}{f(n_0 \mid L)} \tag{23}$$

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 $t \in \{0, 1, \ldots, t_1 - 1\}$ 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 $t \in \{t_1, \ldots, t^*\}$ 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 $t \in \{t^* + 1, \ldots, T\}$ 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 participation is declining: $n_{t+1}^S < n_t^S$.

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 $\{n_t^S, n_t^C\}$.

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:

$$\frac{dn_t}{d\hat{\theta}_t} = \frac{\partial\Phi/\partial\hat{\theta}}{1 - \partial\Phi/\partial n} \tag{24}$$

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 $n_t^S = \lambda[1 - G(\alpha n_t + c_S; \hat{\theta}_t)]$. 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: $n_t^{S,H}$ falls faster than $n_t^{S,L}$.[11]

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

[11] 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 $\beta(1-\lambda)/\lambda < 1 - c_S$.

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 $\alpha < \bar{\alpha}(\mu, \lambda)$ implies $n_{t+1} > n_t$ in Phase I. The equilibrium adoption satisfies $n_t = \Phi(n_t; \hat{\theta}_t)$. By the implicit function theorem:

$$n_{t+1} - n_t = \frac{dn}{d\hat{\theta}} \cdot (\hat{\theta}_{t+1} - \hat{\theta}_t) + O((\Delta\hat{\theta})^2) \tag{25}$$

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 $\alpha < \bar{\alpha}(\mu, \lambda)$ 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 threshold $\bar{\alpha}(\mu, \lambda)$ equals $1/[\lambda \cdot g(\underline{\mu}^S(0))]$, where g is the posterior density at the initial snob threshold. This is increasing in posterior dispersion (more dispersed posteriors lower g at the threshold, raising $\bar{\alpha}$) and decreasing in λ. □

B.5 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 4, 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 $n_0^H \geq n_0^L$.[12] *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 $H \rightarrow$ more adoption $\rightarrow$ more favorable belief updates $\rightarrow$ rightward-shifted posteriors $\rightarrow$ 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. □

[12] 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.

B.6 Belief Dominance Under Higher Quality

Lemma 4 (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).

B.7 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:

$$n_t(\hat{\theta}_t) = \lambda[1 - G(\alpha n_t + c_S; \hat{\theta}_t)] + (1 - \lambda)[1 - G(c_C - \beta n_t; \hat{\theta}_t)] \tag{26}$$

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 4, $\hat{\theta}_t^H > \hat{\theta}_t^L$ for all $t \geq 1$:

$$\mathbb{E}[\hat{\theta}_{t+1} - \hat{\theta}_t \mid \theta = H] > 0, \quad \mathbb{E}[\hat{\theta}_{t+1} - \hat{\theta}_t \mid \theta = L] < 0 \tag{27}$$

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):

$$\frac{dn^S}{d\hat{\theta}} = \lambda \left[\underbrace{-g^S \cdot \alpha \cdot \frac{dn^*}{d\hat{\theta}}}_{\text{threshold effect}(<0)} + \underbrace{\frac{\partial[1 - G(\underline{\mu}^S; \hat{\theta})]}{\partial\hat{\theta}}}_{\text{belief effect}(>0)} \right] \tag{28}$$

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:

$$b^S > g^S \cdot \alpha \cdot \frac{\lambda b^S + (1 - \lambda)b^C}{1 - \Phi'} \tag{29}$$

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:

$$1 > \frac{\alpha g^*}{1 - \Phi'} = \frac{\alpha g^*}{1 + \lambda\alpha g^* - (1 - \lambda)\beta g^*} \tag{30}$$

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

$$1 + \lambda\alpha g^* - (1 - \lambda)\beta g^* > \alpha g^* \tag{31}$$

$$1 > (1 - \lambda)g^*(\alpha + \beta) \tag{32}$$

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:

$$\lambda\alpha \approx (1-\lambda)\beta \tag{33}$$

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:

$$\alpha\lambda = \beta(1-\lambda) \quad \Rightarrow \quad \bar{\alpha}(\lambda) = \beta \cdot \frac{1-\lambda}{\lambda} \tag{34}$$

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 $\alpha\lambda \lesseqgtr \beta(1-\lambda)$ 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, $\bar{\alpha}(\lambda) = \beta(1-\lambda)/\lambda$:

$$\frac{\partial\bar{\alpha}}{\partial\lambda} = -\frac{\beta}{\lambda^2} < 0, \qquad \frac{\partial\bar{\alpha}}{\partial\beta} = \frac{1-\lambda}{\lambda} > 0 \tag{35}$$

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. □

Online Appendix

Snobs and Conformists: Platform Design and Product Lifecycles

Tommaso Bondi

Cornell University

This Online Appendix collects proofs of secondary results. Proofs of the main propositions appear in Appendix B; extensions appear in Appendix A.

OA.1 Additional Proofs

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:

$$f(\tilde{n} \mid \theta, \hat{\theta}_t) = \frac{1}{\sigma_\varepsilon}\phi\left(\frac{\tilde{n} - n^*_\theta(\hat{\theta}_t)}{\sigma_\varepsilon}\right) \tag{36}$$

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 $\varepsilon_t \sim N(0, \sigma_\varepsilon^2)$ 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:

$$(\mathcal{T}V)(\mu, \hat{\theta}) = \max\Big\{\mu + u^\tau(n^*(\hat{\theta})) + \delta \int V(\mu', \hat{\theta}')\, dF(\mu', \hat{\theta}' \mid \hat{\theta}), \\ c_\tau + \delta \int V(\mu', \hat{\theta}')\, dF(\mu', \hat{\theta}' \mid \hat{\theta})\Big\} \quad (37)$$

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). □

Proof of Uniqueness (Assumption 3)

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

Step 1: Derivative of Φ. From (21), differentiating with respect to n:

$$\Phi'(n) = -\lambda\alpha \cdot g(\underline{\mu}^S(n)) + (1-\lambda)\beta \cdot g(\underline{\mu}^C(n)) \tag{38}$$

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.[13]

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). □

[13] Three crossings are standard in models with locally strong complementarities; see Morris & Shin (2003). An explicit Gaussian example appears in the Online Appendix.

Proof of Proposition 8 (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:

$$W(\lambda) = \lambda \cdot EU^S(\lambda) + (1-\lambda) \cdot EU^C(\lambda) \tag{39}$$

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:

$$\frac{dW}{d\lambda}\Big|_{\lambda^{**}} = EU^S - EU^C + \lambda^{**}\frac{dEU^S}{d\lambda} + (1-\lambda^{**})\frac{dEU^C}{d\lambda} = 0 \tag{40}$$

Step 5: Regularity and second-order condition. We first establish differentiability. The within-period equilibrium $n^*(\hat{\theta}, \lambda)$ is C^1 in λ by the implicit function theorem applied to $\Phi(n^*; \hat{\theta}, \lambda) = n^*$, since $|\Phi'| < 1$ (Assumption 3) and Φ is C^1 in λ. The belief process $\{\hat{\theta}_t\}$ inherits smoothness in λ from the adoption function. Phase transition times $t_1(\lambda)$ and $T(\lambda)$ are smooth in λ under the generic transversality condition $d[\underline{\mu}^C(n_t) - \hat{\theta}_t]/dt \neq 0$ at threshold crossings.[14] Given smoothness, we verify $d^2W/d\lambda^2|_{\lambda^{**}} < 0$. The key terms are: $dEU^S/d\lambda < 0$ (more snobs increase crowding, reducing each snob's utility) and $dEU^C/d\lambda < 0$ near λ^{**} (the lifecycle compression from additional snobs reduces conformist surplus). Since snobs bear the direct crowding cost, $dEU^S/d\lambda < dEU^C/d\lambda$, making the leading terms of $d^2W/d\lambda^2$ negative. The second-order condition holds under Assumption 3, which bounds density responses.

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

$$\frac{d\lambda^{**}}{d\alpha} = -\frac{\partial^2 W/\partial\lambda\partial\alpha}{d^2W/d\lambda^2}\Big|_{\lambda^{**}} \tag{41}$$

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 $d\lambda^{**}/d\sigma > 0$. □

[14] At non-generic parameter values where the transversality condition fails, directional derivatives exist and the welfare comparison at boundaries still holds. The existence of an interior maximum follows from the continuous argument in Steps 2–3 without differentiability.

Proof of Corollary 1 (Peak Adoption)

Proof. At the peak t^*, the equilibrium satisfies $n^* = \Phi(n^*; \hat{\theta}_{t^*})$. We establish the comparative statics by the implicit function theorem.

Step 1: IFT at stable equilibrium. Differentiating the fixed-point condition $n = \Phi(n; \hat{\theta}, \alpha, \beta)$ with respect to α:

$$\left.\frac{\partial n^*}{\partial \alpha}\right|_{\hat{\theta}\text{ fixed}} = \frac{-\lambda g(\underline{\mu}^S(n^*)) \cdot n^*}{1 - \Phi'(n^*)} < 0 \tag{42}$$

The numerator is negative ($g > 0$, $n^* > 0$, $\lambda > 0$). The denominator $1 - \Phi'(n^*) > 0$ at any stable equilibrium (Assumption 3). For β: differentiating gives $\partial n^*/\partial\beta|_{\hat{\theta}} = (1 - \lambda)g(\underline{\mu}^C)n^*/(1 - \Phi') > 0$ within period, but the lifecycle peak n^* at the dynamic optimum also depends on the belief path. Higher β lowers the conformist threshold at every n, which triggers earlier conformist entry and hence earlier snob exit. The net effect on peak adoption is $\partial n^*/\partial\beta < 0$, driven by the threshold crossing occurring at lower n.

Step 2: Heuristic closed form (Remark 2). Under Assumption 4, the threshold crossing point $n^\dagger = (c_C - c_S)/(\alpha + \beta)$ – where $\underline{\mu}^S(n) = \underline{\mu}^C(n)$ – serves as a proxy for n^*. This is motivated by the observation that when posterior densities are approximately constant across the threshold region, both marginal agents face comparable quality bars near the peak, and $n^\dagger$ approximates the fixed point well. Numerical simulations confirm the approximation is accurate across parameter values (Figure 1). The proxy is not derived as a theorem from Assumption 4 alone; it is a heuristic that simplifies closed-form expressions. □

Proof of Corollary 2 (Asymmetric Quality Elasticity)

Proof. At a stable within-period equilibrium, type-specific adoption satisfies:

$$n_t^S = \lambda[1 - G(\alpha n^* + c_S; \hat{\theta}_t)] \tag{43}$$

$$n_t^C = (1 - \lambda)[1 - G(c_C - \beta n^*; \hat{\theta}_t)] \tag{44}$$

where $n^* = n_t^S + n_t^C$ is the total equilibrium adoption at belief $\hat{\theta}_t$, and $G(\mu; \hat{\theta})$ is the CDF of the posterior distribution. Differentiating n_t^C with respect to $\hat{\theta}$:

$$\frac{dn_t^C}{d\hat{\theta}} = (1-\lambda)\left[\frac{\partial[1 - G(\underline{\mu}^C; \hat{\theta})]}{\partial\hat{\theta}} + \beta\, g(\underline{\mu}^C)\,\frac{dn^*}{d\hat{\theta}}\right] \tag{45}$$

The first term is positive (higher $\hat{\theta}$ shifts the posterior rightward, increasing the mass above any fixed threshold). The second term is positive when $dn^*/d\hat{\theta} > 0$ (more adoption lowers the conformist threshold $c_C - \beta n$, admitting more conformists). Both channels reinforce, so $dn_t^C/d\hat{\theta} > 0$ unconditionally.

For snobs:

$$\frac{dn_t^S}{d\hat{\theta}} = \lambda\left[\frac{\partial[1 - G(\underline{\mu}^S; \hat{\theta})]}{\partial\hat{\theta}} - \alpha\, g(\underline{\mu}^S)\,\frac{dn^*}{d\hat{\theta}}\right] \tag{46}$$

The direct effect is again positive, but the indirect effect enters with a negative sign: higher adoption *raises* the snob threshold $\alpha n + c_S$, pushing snobs out. The sign of $dn_t^S/d\hat{\theta}$ is therefore ambiguous:

$$\frac{dn_t^S}{d\hat{\theta}} \gtrless 0 \iff \alpha \lessgtr \frac{\partial_{\hat{\theta}}[1 - G(\underline{\mu}^S; \hat{\theta})]}{g(\underline{\mu}^S)\, dn^*/d\hat{\theta}} \tag{47}$$

Under Assumption 4, $g(\underline{\mu}^S) \approx g(\underline{\mu}^C)$ and the direct effects are approximately equal across types. The condition then simplifies: $dn_t^S/d\hat{\theta} < 0$ when the crowding channel $\alpha\, g\, dn^*/d\hat{\theta}$ exceeds the direct signal channel. Since $dn^*/d\hat{\theta}$ is proportional to g^* and the reversal condition is $\alpha > \beta(1-\lambda)/\lambda$, the quality-elasticity threshold coincides with the reversal threshold $\bar{\alpha}(\lambda)$. □

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 the lifecycle duration measure, 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:

$$\alpha\lambda^* = \beta(1-\lambda^*) \quad \Rightarrow \quad \lambda^* = \frac{\beta}{\alpha+\beta} \tag{48}$$

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: Advertising Shortens Lifecycles

We first establish that advertising shortens lifecycles, then derive the optimal targeting formula.

Step 0: Advertising and lifecycle duration. Let $a > 0$ parameterize advertising that shifts each agent's posterior upward by a. When $n > n^{\dagger}$, advertising *shortens* product lifecycles ($\partial T/\partial a < 0$) by accelerating snob exit.

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

$$n_t(a) = \lambda[1 - G(\alpha n_t + c_S - a)] + (1 - \lambda)[1 - G(c_C - \beta n_t - a)] \tag{49}$$

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 $n = \Phi(n; a)$:

$$\frac{dn^*}{da} = \frac{\partial\Phi/\partial a}{1 - \partial\Phi/\partial n} = \frac{\lambda g(\underline{\mu}^S - a) + (1 - \lambda)g(\underline{\mu}^C - a)}{1 - \Phi'(n^*)} > 0 \tag{50}$$

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.:

$$\alpha \cdot \frac{\lambda g^S + (1 - \lambda)g^C}{1 - \Phi'(n^*)} > 1 \tag{51}$$

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.[15]

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, 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 Proposition 14 (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:

$$U^S = v(\theta) - \alpha n_t + \eta C_t = v(\theta) - \alpha(n_t^S + n_t^C) + \eta(n_t^S - \xi n_t^C) \tag{52}$$

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

Rearranging:

$$U^S = v(\theta) + (\eta - \alpha)n_t^S - (\alpha + \eta\xi)n_t^C \tag{53}$$

Step 2: Marginal effect of conformist entry. Differentiating:

$$\frac{\partial U^S}{\partial n^C} = -(\alpha + \eta\xi) < -\alpha \tag{54}$$

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:

$$\Delta C_t = \Delta n_t^S - \xi \Delta n_t^C \tag{55}$$

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

Step 5: Cool equilibria (part iv). A cool equilibrium requires snobs to adopt ($n^S > 0$) while conformists are deterred ($n^C = 0$). The snob-only fixed point n^{**} solves $\lambda[1 - G((\alpha - \eta)n^{**} + c_S)] = n^{**}$, which exists and is unique by the intermediate value theorem. The conformist deterrence condition requires $c_C - \beta n^{**} > \bar{\mu}(\hat{\theta}_\infty)$, where $\hat{\theta}_\infty$ is the limiting belief under snob-only adoption. This holds when c_C is large, β is small, or σ is large (keeping $\bar{\mu}$ from concentrating near 1). □

Proof of Proposition 10 (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:

$$\frac{d\Pi}{d\tau} = \underbrace{\sum_{t=0}^{T} \delta^t \pi'(n_t) \frac{\partial n_t}{\partial \tau}}_{\text{per-period profit effect}} + \underbrace{\pi(n_T)\delta^T \frac{\partial T}{\partial \tau}}_{\text{duration effect}} \tag{56}$$

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^*:

$$\tau^* = \min\left\{1, \quad \underbrace{\frac{\alpha}{\alpha+\beta}}_{\substack{\text{static: balances}\\ \text{marginal WTP}}} + \underbrace{\frac{\delta}{1-\delta} \cdot \frac{\lambda}{1-\lambda} \cdot \frac{\partial T/\partial n^S}{|\partial T/\partial n^C|}}_{\substack{\text{dynamic: lifecycle}\\ \text{extension value}}}\right\} \tag{57}$$

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 Proposition 11 (Advertising Volume)

Proof. When $\alpha < \bar{\alpha}(\mu, \lambda)$, the quality-duration reversal does not hold, so higher adoption unambiguously extends the lifecycle ($\partial T/\partial a \geq 0$). Both the per-period profit effect and the duration effect of advertising are non-negative, so the firm's profit is increasing in a.

When $\alpha > \bar{\alpha}(\mu, \lambda)$ and advertising is untargeted ($\tau = \lambda$, proportional to population shares), the firm's profit as a function of a is

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

The derivative decomposes as in the proof of Proposition 10. The per-period effect $\sum_t \delta^t \pi'(n_t) \partial n_t/\partial a > 0$ (advertising raises adoption and hence revenue in each period). The duration effect $\pi(n_T)\delta^T \partial T/\partial a < 0$ (advertising shortens the lifecycle, as shown in the advertising duration proof above). The myopic optimum a^{myopic} sets $\sum_t \delta^t \pi'(n_t)\partial n_t/\partial a = 0$, ignoring the duration effect. Since the duration effect is strictly negative, the full FOC $d\Pi/da = 0$ is satisfied at $a^* < a^{\text{myopic}}$. □

Proof of Lemma 3 (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 $|\Phi'| < 1$). The entry time $t_1(\varphi)$ is smooth in φ provided the conformist threshold crossing is transversal: $d[\underline{\mu}^C(n_t) - \hat{\theta}_t]/dt \neq 0$ at $t = t_1$, which holds generically since Phase I adoption is strictly increasing (Proposition 3). 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 and smooth phase transitions. □

Proof of Proposition 12 (Welfare Benchmark)

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 $\Delta t_1 > 0$:

$$
\begin{aligned}
-\frac{dW}{d\varphi}\bigg|_{\varphi=1} = & \underbrace{(1-\lambda)\cdot g(\underline{\mu}^C(n_{t_1}))\cdot[\hat{\theta}_{t_1} + \beta n_{t_1} - c_C]\cdot|\Delta t_1|}_{\text{Cost: foregone conformist bandwagon benefits}} \\
& - \underbrace{\lambda \cdot \sum_{t=0}^{t_1-1} \frac{\partial EU_t^S}{\partial t_1}\delta^t \cdot |\Delta t_1|}_{\text{Benefit: extended snob pioneer rents}} \\
& - \underbrace{(1-\lambda)\cdot\frac{\partial EU^C}{\partial\hat{\theta}}\cdot\frac{\partial\hat{\theta}}{\partial t_1}\cdot|\Delta t_1|}_{\text{Benefit: improved quality assessment}} \qquad (58)
\end{aligned}
$$

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:

$$\frac{\partial[W(\varphi^*) - W(1)]}{\partial\alpha} > 0 \tag{59}$$

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:

$$\left.\frac{\partial W}{\partial\varphi}\right|_{\text{quality assessment}} + \left.\frac{\partial W}{\partial\varphi}\right|_{\text{conformist acceleration}} = 0 \tag{60}$$

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 13 (Profit-Maximizing Visibility)

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 3(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:

$$M(\varphi,\delta_P) = \underbrace{\sum_{t=0}^{T(\varphi)} \delta_P^t r'(n_t)\frac{\partial n_t}{\partial\varphi}}_{\text{per-period effect}} + \underbrace{r(n_T)\delta_P^T\frac{\partial T}{\partial\varphi}}_{\text{duration effect}(<0)} \tag{61}$$

The per-period effect is positive in early periods (conformist entry raises n_t) and the duration effect is negative (shorter lifecycle, by Lemma 3(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 $M(\varphi_W^*, \bar{\delta}_P; \alpha, \lambda) = 0$:

$$\frac{\partial\bar{\delta}_P}{\partial\alpha} = -\frac{\partial M/\partial\alpha}{\partial M/\partial\delta_P} \tag{62}$$

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 3 (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 12. Differentiating the FOC with respect to θ:

$$\frac{\partial^2 W}{\partial\varphi^2} \cdot \frac{d\varphi^*}{d\theta} + \frac{\partial^2 W}{\partial\varphi\partial\theta} = 0 \tag{63}$$

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 9 (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 $\Pi^{\text{scarce}}(\delta) > \Pi^{\text{clear}}(\delta)$:

$$\frac{\pi^{\text{scarce}}(\bar{n})}{1-\delta} > \sum_{t=0}^{T} \delta^t \pi_t \tag{64}$$

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). □