WORKING PAPERS

N° 1685

October 2025

“Herding Prices: Social Learning and Dynamic Competition in Duopoly”

Arina Azova and Georgy Lukyanov

Toulouse School of Economics

Herding Prices:
Social Learning and Dynamic Competition in Duopoly*

Arina Azova[1] and Georgy Lukyanov[†2]

[1]International College of Economics and FInance, NRU HSE, Moscow, Russia
[2]Toulouse School of Economics, Toulouse, France

Abstract

We embed observational learning (BHW) in a symmetric duopoly with random arrivals and search frictions. With fixed posted prices, a mixed-strategy pricing equilibrium exists and yields price dispersion even with ex-ante identical firms. We provide closed-form cascade bands and show wrong cascades occur with positive probability for interior parameters, vanishing as signals become precise or search costs fall; absorption probabilities are invariant to the arrival rate. In equilibrium, the support of mixed prices is connected and overlapping; its width shrinks with signal precision and expands with search costs, and mean prices comove accordingly. Under Calvo price resets (Poisson opportunities), stationary dispersion and mean prices fall; when signals are sufficiently informative, wrong-cascade risk also declines. On welfare, a state-contingent Pigouvian search subsidy implements the planner's cutoff. Prominence (biased first visits) softens competition and depresses welfare; neutral prominence is ex-ante optimal.

Keywords: social learning; informational cascades; price dispersion; search; vertical differentiation.
JEL: C73; D43; D83; L13.

1 Introduction

Open two delivery apps on a Friday night or two streaming platforms in a new market. One brand quickly looks "hot"—orders tick up, a chart climbs—while its rival seems quiet. Yet prices are not equal: sometimes the apparently popular option is cheaper, sometimes dearer, and the gaps move around across launches and locales. What links this visible momentum to how firms set prices when consumers learn in real time from others' choices? This paper argues that the missing piece is the

*We are grateful to Alexis Belianin, Markus Gebauer, Vitamins Jascisens, Ella Khromova, Yukon Koriyama, Jean-Baptiste Michau, Fabian Slonimczyk, and Anna Yurko for helpful comments and discussions. We also thank all participants at the research seminar at ICEF for valuable feedback. Financial support from the French National Research Agency (ANR), "Investissements d'Avenir" (LabEx Ecodec/ANR-11-LABX-0047), is gratefully acknowledged. All remaining errors are our own.

†Corresponding author: georgy.lukyanov@tse-fr.eu

interaction of social learning with random arrivals and costly search. When customers arrive one by one, observe past actions (who bought where) but not experienced payoffs, and can pay a small cost to check the alternative store, their behavior creates belief tipping points that feed back into price incentives. That feedback explains why price dispersion can coexist with herding, why wrong fads sometimes persist, and when competition corrects them—or does not.

We study a vertically differentiated duopoly in continuous time. Firms know which product is objectively better; consumers do not. Each arriving consumer sees posted prices and the history of purchase choices, receives a private signal about which product is higher quality, and may pay a search cost to check the other seller before buying. The baseline holds prices fixed once at the start; an extension lets each firm reset price at random (Calvo) times. The informational environment follows the canonical social-learning assumption: actions are observed, payoffs are not.[1]

On the consumer side, behavior is governed by a single belief cutoff: if the public belief that firm A is better is high enough relative to the posted price difference and the search cost, consumers buy from A without checking; otherwise they pay to check and may switch. This yields two belief thresholds that form absorbing "up" and "down" regions—once public belief crosses them, subsequent actions herd regardless of private signals. We establish these objects and their monotone comparative statics: more precise signals shrink both absorbing regions; higher search costs expand them. See Lemma 4.1 and Proposition 4.1. At a symmetric prior there is a simple test for immediate herding, and away from knife-edge cases, wrong cascades occur with positive probability but vanish as signals become very informative or search becomes cheap (Lemma 4.3, Proposition 4.2).

Figure 1 plots sample paths of the public belief $\eta(t)$[2] in two regimes. The top panel shows high information (high q) and low search frictions (low κ) starting at $\eta_0 = \frac{1}{2}$; the bottom panel shows low q and high κ, started inside the non-absorbing region to display interior dynamics. The dashed lines mark the absorbing boundaries from Proposition 4.1; once $\eta(t)$ crosses either line, subsequent actions herd and the belief stays there.[3]

The patterns line up with our theory. In the top panel, paths move quickly to the correct absorbing region because actions are informative and consumers are willing to check the alternative (low κ). In the bottom panel, actions are weak signals and search is costly, so beliefs drift slowly and wrong cascades are more likely. Immediate herding at $\eta_0 = \frac{1}{2}$ can occur in the low-information regime (see Lemma 4.3), and the probability of a wrong cascade is strictly positive away from knife-edge cases but vanishes as $q \uparrow 1$ or $\kappa \downarrow 0$ (Proposition 4.2). The comparative statics for the boundaries themselves (shrink with q, expand with κ) and for the speed of convergence follow from Lemma 4.1 and Propositions 4.1 and 6.1.

Treating consumer behavior as primitive, we analyze firms' static price competition. Despite zero costs and only two sellers, the game admits (possibly) mixed-strategy equilibria. In the symmetric

[1] If consumption outcomes were publicly observed, herds would unravel in the limit under standard conditions; our interest is precisely the actions-only channel that sustains cascades.

[2] We write $\eta(t)$ because arrivals are Poisson; beliefs are piecewise constant and jump only at arrival times.

[3] At $\eta_0 = \frac{1}{2}$ and equal prices, the low-information/high-κ regime can lie in an absorbing band, producing immediate herding; the bottom panel starts inside the non-absorbing region to display interior dynamics.

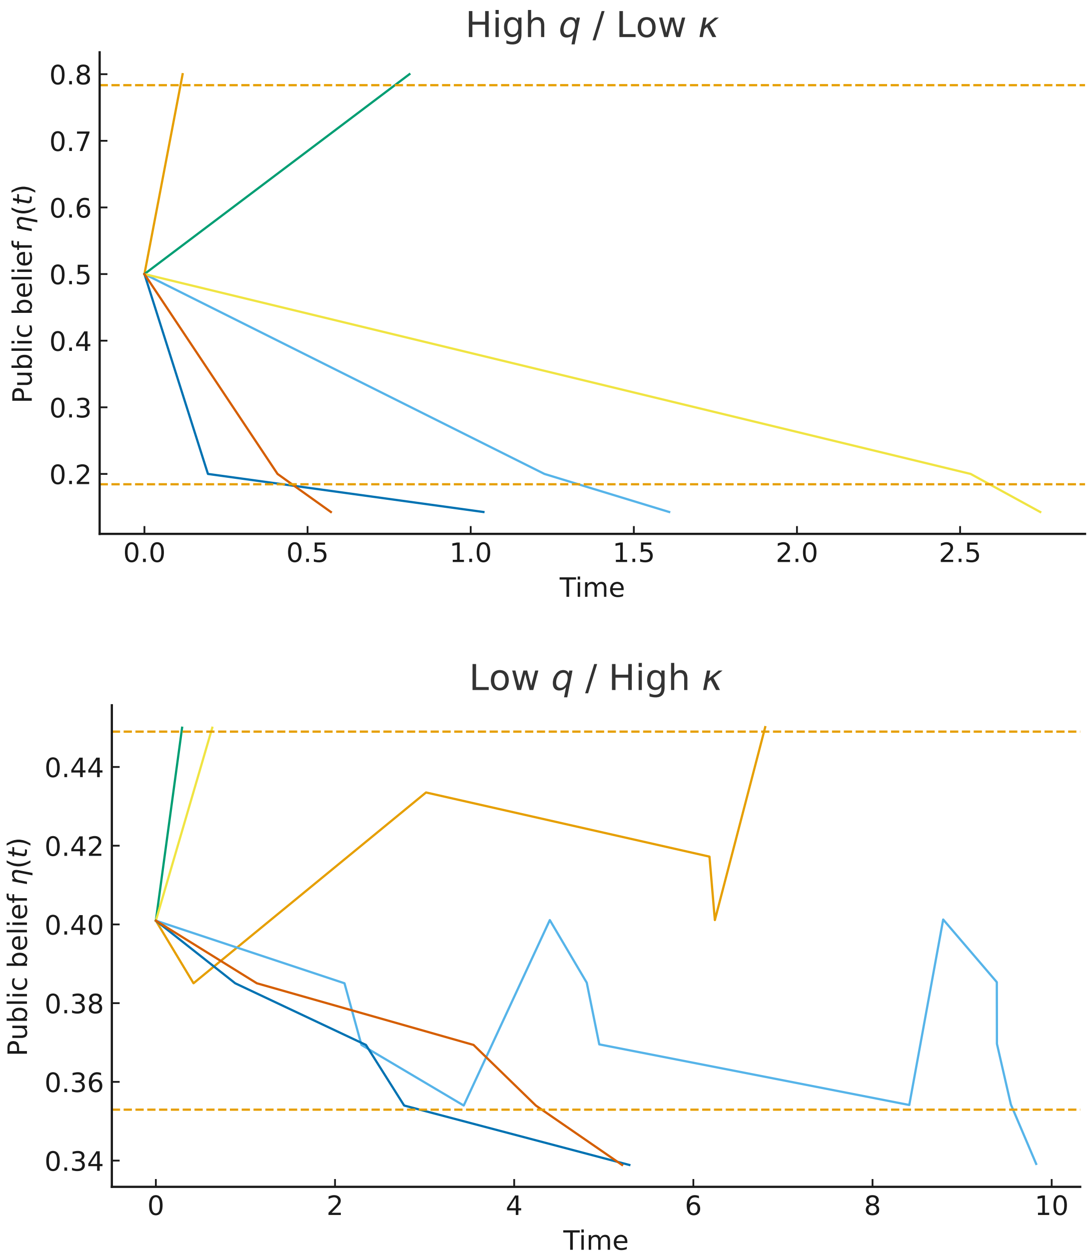

Figure 1: Public belief sample paths. Top: high precision ($q = 0.80$), low search cost ($\kappa = 0.05$), starting at $\eta_0 = \frac{1}{2}$. Bottom: low precision ($q = 0.55$), higher search cost ($\kappa = 0.20$), starting inside the non-absorbing region to display dynamics. Absorbing boundaries from Proposition 4.1 are shown as dashed lines. See also Lemma 4.3 and Propositions 4.1 and 6.1.

case with interior search frictions, both firms mix over connected, partially overlapping supports; the width of those supports is a convenient summary of dispersion. Intuitively, dispersion here does not come from exogenous sampling noise; it comes from feedback between learning and search: observed actions tilt public belief, which rotates effective demand slopes and makes rivals indifferent

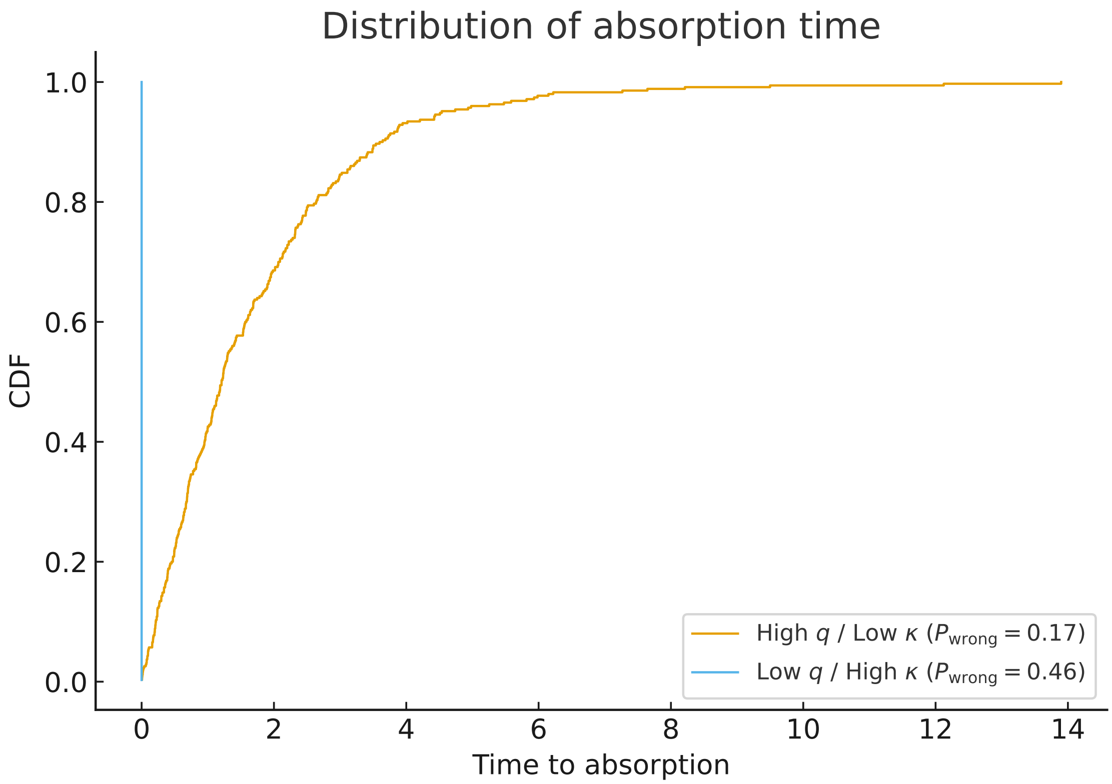

Figure 2: Distribution of time to absorption across regimes (legend shows wrong-cascade probability P_{wrong}). Arrival intensity λ time-changes the process but not absorption probabilities; see Propositions 4.2 and 6.1.

across a range of prices. Dispersion tightens as signals become more precise and widens as search costs grow (Propositions 5.1 to 5.3). Mean prices move in the same direction: stronger learning pressure depresses prices; higher search frictions raise them (Proposition 6.2). A one-page roadmap table (Table 1) summarizes the results at a glance.

Speed and accuracy of learning separate cleanly. The arrival rate only changes the clock: it scales time to herding one-for-one but does not change the probability that the market ends up recognizing the better product. By contrast, information and frictions shape both speed and accuracy: more precise signals or cheaper search make correct cascades more likely and faster; more frictions do the opposite (Proposition 6.1). These predictions fit the opening examples: markets with richer side information (trusted reviewers, transparent comparisons) converge quickly and exhibit narrower dispersion; opaque markets with frictions to "checking the other app" converge slowly and sustain wider price ranges.

Figure 2 reports the empirical CDFs of the time to absorption (the minimum of the up/down hitting times) in the two regimes; the legend shows the wrong-cascade probability P_{wrong}. The high-q/low-κ CDF lies uniformly to the left—absorption is faster—and its P_{wrong} is lower; the low-q/high-κ CDF shifts right and displays a higher P_{wrong}.

The intuition mirrors the sample paths. More precise signals and cheaper search generate more informative actions, so beliefs reach the correct absorbing region sooner and more often. Conversely, when q is low and κ is high, actions convey little information and checking is rare, delaying absorption and raising the chance of getting stuck on the wrong product. Formally, the speed/accuracy trade-offs and the invariance of absorption probabilities to the arrival rate (while calendar time scales with λ) are captured in Proposition 6.1, and the behavior of P_{wrong} is characterized in Proposition 4.2.

The welfare accounting isolates the information externality. Prices are transfers; the real losses come from two sources: buying the inferior product before beliefs correct and spending on search that is privately but not fully socially internalized. A planner who values the information each purchase generates for future consumers sets a lower buy-threshold than individuals do and can implement it with a simple search subsidy. This closes the welfare gap into two transparent components—wrong purchases and excess search—and yields sharp comparative statics (Propositions 6.3 and 6.4 and Corollary 6.1).

Allowing each firm to reset price at random times connects the model to staggered pricing and clarifies how frequent repricing disciplines dispersion. We prove existence of a stationary equilibrium in Markov reset policies. As the reset hazard rises, disadvantaged firms can respond more often when beliefs turn against them, accelerating convergence and compressing stationary dispersion; with sufficiently informative signals, the steady-state probability of being stuck on the wrong product falls (Propositions 7.1 and 7.3 and Corollary 7.2).

Conceptually, the paper offers a clean, tractable way to study social learning and price competition in the same continuous-time environment with random arrivals. Technically, a one-line consumer cutoff and closed-form cascade boundaries make the feedback loop analyzable without heavy machinery and pin down tight comparative statics for prices, dispersion, and welfare. Substantively, the framework explains when competition corrects mislearning and when it does not, and it delivers a simple policy lever—a search subsidy—that realigns private and social incentives in markets where consumer experiences are opaque.

Table 1: Roadmap of Results

Result	**Informal statement / intuition**	**Where**
Consumer threshold	Optimal buy–A rule is a single cutoff $\eta^*(p_A - p_B, \kappa)$; higher $p_A - p_B$ or κ raises the cutoff.	Lemma 4.1
Cascade boundaries	Closed forms for absorbing up/down regions $\bar{\eta}, \underline{\eta}$; more precise signals shrink both regions.	Prop. 4.1
Wrong cascades	With bounded signals ($q < 1$) and interior cutoff, wrong cascades occur with positive probability; vanish as $q\uparrow 1$ or $\kappa \downarrow 0$.	Prop. 4.2
Existence (prices)	Static price game admits a (possibly mixed) equilibrium on compact price sets.	Prop. 5.1
Dispersion structure	In the symmetric case (interior κ), both firms mix over connected, partially overlapping supports.	Prop. 5.2
Dispersion C.S.	Supports narrow with higher q and widen with higher κ; degenerate in high-q/low-κ limits.	Prop. 5.3
Absorption C.S.	P_{right} increases in q and decreases in κ; times scale as $1/\lambda$ and fall with q.	Prop. 6.1
Prices C.S.	Mean price falls with q and rises with κ.	Prop. 6.2
Welfare	Welfare rises with q, falls with κ, and scales with λ in calendar time.	Prop. 6.3
Planner & policy	Planner uses a lower cutoff; a Pigouvian search subsidy implements it; welfare gap = wrong purchases + excess search.	Prop. 6.4
Calvo equilibrium	With Poisson price resets, a stationary Markov reset equilibrium exists.	Prop. 7.1
Calvo C.S.	Higher reset hazard α shrinks stationary dispersion and speeds absorption; wrong-cascade probability falls when q is high.	Prop. 7.3

2 Related Literature

We build on the canonical social learning and informational cascades literature (Bikhchandani et al., 1992; Banerjee, 1992; Smith and Sørensen, 2000; Chamley, 2004). Recent work revisits robustness and sources of mislearning: misspecified inference can sustain inefficient actions (Bohren, 2016), and small misperceptions about others' types can dramatically weaken aggregation (Frick et al., 2020). For a recent panoramic survey, see Bikhchandani et al. (2024). Our contribution follows the canonical BHW structure (actions but not signals observed) yet embeds it in a duopoly with prices and consumer search, delivering novel implications for dispersion, absorption, and welfare.

Closest to us is research that endogenizes *prices* in social learning environments. Arieli et al. (2022) analyze a dynamic duopoly where firms set prices each period and characterize when asymptotic learning occurs; they emphasize signal boundedness (in the sense of Smith and Sørensen, 2000) as the key determinant. Sayedi (2018) studies a duopoly with observational learning under static versus dynamic pricing and documents the persistence of (sometimes wrong) cascades. On the operations side, Papanastasiou and Savva (2017) examine dynamic pricing with social learning for a monopolist facing forward-looking consumers. Relative to these papers, our setting features *random consumer arrivals and costly search*, and we characterize how learning–search interactions generate (i) mixed-strategy price dispersion under fixed prices and (ii) sharper comparative statics for dispersion and welfare; we also provide a tractable Calvo extension with stochastic price resets.

A second strand blends *observational learning with search and prominence*. Garcia and Shelegia (2018) combine random search and observational learning, showing how observed choices shape search incentives and selection; we adopt their spirit of linking learning to search frictions but focus on *duopolistic pricing* and dispersion. Work on prominence and attention allocation (e.g., Armstrong et al., 2009; Armstrong and Zhou, 2011) rationalizes first-visit asymmetries; in our model, such asymmetries interact with learning to move cascade boundaries and pricing supports as q and κ vary.

Our results also relate to equilibrium price dispersion with consumer search (Diamond, 1971; Varian, 1980; Burdett and Judd, 1983; Stahl, 1989). In these classic models, dispersion arises from heterogeneity in information sets or sampling technologies. Here, dispersion emerges from *feedback between learning and search*: observed actions reshape beliefs, which rotate demand slopes and alter the set of prices that can be sustained in mixed strategies. This mechanism yields clear comparative statics—supports narrow with higher signal precision q and widen with greater search costs κ—and welfare results that isolate the information externality.

Finally, we connect to dynamic oligopoly and staggered pricing. The Markovian logic of best responses links to Maskin and Tirole (1988a,b); our Calvo extension imports the macro "random reset" device of Calvo (1983) into a micro learning context, where resets tighten stationary dispersion and speed absorption by letting disadvantaged firms re-price more often. Beyond these literatures, two nearby themes help position our welfare and identification angles: (i) learning when some actions generate *public signals* while others are observationally opaque—see, for instance, the absorbing-region logic in Lukyanov et al. (2025); and (ii) how motives or information sources alter *diagnosticity*—e.g., contrarian incentives that sustain experimentation (Lukyanov and Ivanik, 2025) and expert advice with uncertain precision that endogenizes informativeness (Lukyanov, 2025). We use these to motivate our planner benchmark and the role for a search subsidy that corrects the information externality.

3 Environment

3.1 Players, Qualities, and Values

Two firms A, B produce vertically differentiated substitutes. The high-quality product yields consumer value v_H; the low-quality product yields v_L, with normalized gap

$$\Delta \equiv v_H - v_L > 0 \quad \text{(w.l.o.g. set } \Delta = 1\text{)}.$$

Production costs are zero. Firms know the true quality ranking; consumers do not.

3.2 Timing and Information

Time is continuous. Consumers arrive according to a Poisson process with rate $\lambda > 0$. Upon arrival, a consumer observes: (i) posted prices (p_A, p_B) (chosen once at $t = 0$); (ii) the history of past purchase *choices* (but not realized payoffs). She receives a private signal $s \in \{\mathrm{A}, \mathrm{B}\}$ with precision

$$\mathbb{P}(s = \mathrm{A} \mid \mathrm{A} \text{ is high}) = q \in (1/2, 1), \qquad \mathbb{P}(s = \mathrm{B} \mid \mathrm{B} \text{ is high}) = q.$$

Signals are i.i.d. across consumers. Let $\eta \in [0, 1]$ denote the public belief that A is high quality implied by the observed purchase history.

3.3 Search Technology

The arriving consumer first *visits* one firm (w.l.o.g. A). She may buy immediately, or pay a search cost $\kappa \geq 0$ to check the other firm (B) before deciding. After search, she knows both prices and chooses where to buy. Consumers are myopic (maximize current expected utility).

3.4 Preferences and Payoffs

If the consumer buys from firm $i \in \{\mathrm{A}, \mathrm{B}\}$, expected gross value equals the belief-weighted value of that product, minus price. Firms earn price as profit for each sale (zero cost).

3.5 Strategies, Beliefs, and Equilibrium

A *consumer strategy* specifies, for any belief η, signal s, and observed prices, whether to buy from the first-visited firm, search (pay κ), and then which firm to purchase from. A *firm strategy* is a (possibly mixed) price p_A (resp. p_B) chosen at $t = 0$. Beliefs update via Bayes' rule from observed purchase choices whenever feasible.

Definition 3.1. *A PBE consists of (i) firm price strategies (possibly mixed) at $t = 0$; (ii) a consumer policy that maximizes expected utility given beliefs and prices; and (iii) belief-updating rules consistent with Bayes' rule on-path, such that firms' prices are mutual best responses given the induced continuation values from the consumer policy under Poisson arrivals.*

4 Consumer Problem

4.1 Posterior Updating and a Threshold Rule

Let $\eta \in [0,1]$ denote the public belief that A is the high-quality firm before an arrival. The arriving consumer observes a private signal $s \in \{\mathrm{A}, \mathrm{B}\}$. With precision $q \in (1/2, 1)$ and Bayes' rule,

$$\eta_{\mathrm{A}} \equiv \mathbb{P}(\mathrm{A\ high} \mid s = \mathrm{A},\, \eta) = \frac{q\,\eta}{q\,\eta + (1-q)(1-\eta)}, \tag{4.1}$$

$$\eta_{\mathrm{B}} \equiv \mathbb{P}(\mathrm{A\ high} \mid s = \mathrm{B},\, \eta) = \frac{(1-q)\,\eta}{(1-q)\,\eta + q(1-\eta)}. \tag{4.2}$$

With vertical gap $\Delta = v_H - v_L$ (we normalize $\Delta = 1$), the expected *net* surplus from buying A rather than B, conditional on a posterior $x \in [0,1]$, is

$$D(x; p_A - p_B) \;\equiv\; \big(2x - 1\big)\,\Delta \;-\; (p_A - p_B) \;=\; 2x - 1 - (p_A - p_B). \tag{4.3}$$

We assume the consumer first *visits* A (w.l.o.g.). To buy from B she must incur the search/switch cost $\kappa \geq 0$. Hence, given posterior x, she:

- buys from A if $D(x; p_A - p_B) \geq -\kappa$;
- otherwise pays κ and buys from B.

Lemma 4.1. *Fix posted prices (p_A, p_B) and search cost κ. Consider a consumer who first visits A, observes a private signal $s \in \{A, B\}$, and updates her posterior that A is objectively better to $x \in (0,1)$. There exists a posterior cutoff $x^*(p_A - p_B, \kappa)$ such that she buys at A without checking if and only if*

$$x \;\geq\; x^*(p_A - p_B, \kappa) \;=\; \tfrac{1}{2} \;+\; \frac{p_A - p_B - \kappa}{2\,\Delta}, \tag{4.4}$$

where $\Delta \equiv v_H - v_L > 0$. If $x < x^*$, she pays κ to check B and then buys the higher–expected–surplus option. The statement is symmetric for first visit to B (replace $p_A - p_B$ by $p_B - p_A$). The cutoff is (weakly) increasing in $p_A - p_B$ and (weakly) decreasing in κ.

Corollary 4.1 (Prior beliefs that imply no checking and cascade boundaries)**.** *Let $\eta \in (0,1)$ be the public prior that A is better. Let $\eta_A(\eta)$ and $\eta_B(\eta)$ denote the posteriors after private signals $s = A$ and $s = B$, respectively, i.e.*

$$\eta_A(\eta) \;=\; \frac{q\,\eta}{q\,\eta + (1-q)(1-\eta)}, \qquad \eta_B(\eta) \;=\; \frac{(1-q)\,\eta}{(1-q)\,\eta + q(1-\eta)}.$$

For a consumer who first visits A:

- *She* buys at *A* without checking under either signal *iff* $\min\{\eta_A(\eta), \eta_B(\eta)\} \;\geq\; x^*(p_A - p_B, \kappa)$.
- *The* up–cascade boundary $\bar{\eta}$ *is therefore the smallest* η *solving* $\eta_B(\eta) = x^*(p_A - p_B, \kappa)$.

Symmetrically, for a consumer who first visits B, buys at B without checking under either signal *iff* $\max\{\eta_A(\eta), \eta_B(\eta)\} \ \le\ 1 - x^*(p_B - p_A, \kappa)$*, and the* down–cascade boundary $\underline{\eta}$ *is the largest* η *solving* $\eta_A(\eta) = 1 - x^*(p_B - p_A, \kappa)$*. Under equal posted prices this simplifies to* $\bar{\eta}$ *given by* $\eta_B(\eta) = \frac{1}{2} - \frac{\kappa}{2\Delta}$ *and* $\underline{\eta}$ *by* $\eta_A(\eta) = \frac{1}{2} + \frac{\kappa}{2\Delta}$*.*

Remark 4.1. The maximum possible gain from switching to B when starting at A is $\Delta + (p_A - p_B)$ (attained when $x \downarrow 0$). Hence if

$$\kappa \ \ge\ \Delta + (p_A - p_B) \ =\ 1 + (p_A - p_B), \tag{4.5}$$

the consumer never pays to switch to B. (By symmetry, if she first visited B, a symmetric bound applies.)

Proof of Lemma 4.1. See Appendix B.

4.2 Cascade Regions under Random Arrivals

Actions are observed (but not payoffs). Given the threshold rule in Lemma 4.1, an *up-cascade* for A obtains when the arriving consumer buys A for *both* signals $s \in \{\mathrm{A}, \mathrm{B}\}$; a *down-cascade* (for B) when she switches/buys B for both signals.

Let $\eta^* = \eta^*(p_A - p_B, \kappa)$ from (4.4). Using the posteriors in (4.1)–(4.2), define belief cutoffs $\bar{\eta}$ and $\underline{\eta}$ by

$$\text{(Up-cascade boundary)} \quad \eta_{\mathrm{B}} \ \ge\ \eta^* \iff \eta \ \ge\ \bar{\eta}(p_A - p_B, \kappa; q), \tag{4.6}$$

$$\text{(Down-cascade boundary)} \quad \eta_{\mathrm{A}} \ <\ \eta^* \iff \eta \ \le\ \underline{\eta}(p_A - p_B, \kappa; q). \tag{4.7}$$

Solving $\eta_{\mathrm{B}} = \eta^*$ and $\eta_{\mathrm{A}} = \eta^*$ for η yields closed forms:

$$\bar{\eta}(p_A - p_B, \kappa; q) = \frac{\eta^* q}{(1-q) + \eta^*(2q-1)}, \tag{4.8}$$

$$\underline{\eta}(p_A - p_B, \kappa; q) = \frac{\eta^*(1-q)}{q - \eta^*(2q-1)}. \tag{4.9}$$

Proposition 4.1. *Let* $\bar{\eta}$ *and* $\underline{\eta}$ *be defined by (4.8)–(4.9). Then* $[\,\bar{\eta}, 1\,]$ *and* $[\,0, \underline{\eta}\,]$ *are absorbing under the observational environment (actions observed, payoffs not): once the public belief enters either region, subsequent actions herd and beliefs remain there. Comparative statics: higher signal precision shrinks both absorbing regions, i.e.,* $\partial\bar{\eta}/\partial q > 0$ *and* $\partial\underline{\eta}/\partial q < 0$*; higher search cost expands them, i.e.,* $\partial\bar{\eta}/\partial\kappa < 0$ *and* $\partial\underline{\eta}/\partial\kappa > 0$*.*

Proof. See Appendix B.

Lemma 4.2. *Fix* (q, κ) *and prices* (p_A, p_B)*. Starting from any* $\eta_0 \in (0,1)$*, the public belief* $\eta(t)$ *hits* $[0, \underline{\eta}] \cup [\bar{\eta}, 1]$ *almost surely in finite* arrival *time. Moreover, for every* $\lambda > 0$ *the up/down absorption probabilities are independent of* λ*, and the calendar time to absorption scales as* $1/\lambda$*: if* τ_λ *is the* *absorption time under arrival rate λ, then $\tau_\lambda \overset{d}{=} \tau_1/\lambda$ and $\mathbb{E}[\tau_\lambda] = \mathbb{E}[\tau_1]/\lambda$ (whenever the expectation is finite).*

At the symmetric prior $\eta_0 = \frac{1}{2}$, we obtain a clean "no-immediate-cascade" condition:

Lemma 4.3. *At the symmetric prior $\eta_0 = \frac{1}{2}$ and equal posted prices, the market herds at $t = 0$ iff $\frac{1}{2} \geq \bar{\eta}$ or $\frac{1}{2} \leq \underline{\eta}$; equivalently, for first visit to A, iff $\eta^*(p_A - p_B, \kappa) \geq \bar{\eta}$ or $\eta^*(p_A - p_B, \kappa) \leq \underline{\eta}$. Under the low-information/high-friction regime in our figures, $\bar{\eta} < \frac{1}{2}$, so the initial belief lies in the up–absorbing region and herding is immediate.*

Proof. Immediate from (4.1)–(4.2) and Lemma 4.1.

Finally, random arrivals imply a belief process that jumps according to observed actions. Let $\tau_\uparrow$ (resp. $\tau_\downarrow$) be the hitting time of $[\bar{\eta}, 1]$ (resp. $[0, \underline{\eta}]$) by the public belief. The next result records the qualitative implications for wrong cascades.

Proposition 4.2. *For interior parameters $q \in (1/2, 1)$ and $\kappa > 0$ and any non-degenerate initial belief $\eta_0 \in (0, 1)$, the probability of a wrong cascade is strictly positive. Moreover, the wrong-cascade probability vanishes as information improves or search becomes cheap: it tends to 0 as $q \uparrow 1$ or $\kappa \downarrow 0$. The absorption probabilities are invariant to the arrival rate λ (which only time–changes the process).*

Proof. See Appendix B.

5 Firms' Problem and Static Price Equilibrium

5.1 Expected Profits under Poisson Arrivals

We compute firms' profits taking as given the consumer policy in Lemma 4.1 and the belief cutoffs in Proposition 4.1. Let the initial public belief be $\eta_0 \in (0, 1)$ (we later set $\eta_0 = \frac{1}{2}$). Consumers arrive as a Poisson process of rate λ.

Assumption 5.1. *Each arriving consumer initially* visits *firm* A *with probability $\alpha \in (0, 1)$ and firm* B *with probability $1 - \alpha$, independently of history. The baseline adopts $\alpha = \frac{1}{2}$ (symmetric attention).*

Let $\eta^* = \eta^*(p_A - p_B, \kappa)$ from (4.4). Conditional on public belief η, the probability that a consumer who initially visits A buys from A is

$$\pi_{\text{A}|\text{A-visit}}(\eta) = \mathbf{1}\{\eta_{\text{B}} \geq \eta^*\} \; + \; \mathbf{1}\{\eta_{\text{B}} < \eta^* \leq \eta_{\text{A}}\} \cdot q,$$

where η_{A} and η_{B} are given by (4.1)–(4.2). The first term corresponds to the up-cascade region (buy A for both signals); the second to the intermediate region (buy A only when $s = \text{A}$).

Similarly, if the consumer initially visits B, symmetry implies the probability she buys from A equals

$$\pi_{\mathrm{A|B\text{-}visit}}(\eta) = \mathbf{1}\{\, 1-\eta_{\mathrm{A}} < 1-\eta^* \,\} \;+\; \mathbf{1}\{\, 1-\eta^* \le 1-\eta_{\mathrm{A}} < 1-\eta_{\mathrm{B}} \,\} \cdot (1-q),$$

i.e., she ends up at A either because she switches after any signal (the A up-cascade) or, in the intermediate region, only when $s = \mathrm{B}$ made staying at B unattractive.[4]

Define the *instantaneous demand intensity* for A at belief η:

$$\delta_{\mathrm{A}}(\eta; p_A, p_B) \;\equiv\; \lambda \left[\, \alpha\, \pi_{\mathrm{A|A\text{-}visit}}(\eta) \;+\; (1-\alpha)\, \pi_{\mathrm{A|B\text{-}visit}}(\eta) \,\right], \tag{5.1}$$

and likewise $\delta_{\mathrm{B}}(\eta; p_A, p_B) = \lambda - \delta_{\mathrm{A}}(\eta; p_A, p_B)$.

Belief updates: when a purchase occurs, the action is observed and Bayes' rule maps η to the next belief. Let $K(\eta' \mid \eta; p_A, p_B)$ denote the Markov kernel over beliefs induced by the consumer policy; $\bar{\eta}$ and $\underline{\eta}$ from (4.8)–(4.9) are absorbing boundaries (up/down cascades).

Figure 3 plots the (approximate) symmetric mixed-strategy price distributions for the two regimes we study.[5] Bars give the probability mass that each firm assigns to prices on a grid; the horizontal span of positive mass is the support (our measure of dispersion), and the center of mass reflects the mean price. The top panel corresponds to high information (high q) and low search frictions (low κ); the bottom panel corresponds to low q and high κ.

In the *top* panel, dispersion is tight and the mean price is low. Intuition: when q is high, each observed purchase is informative, so a unilateral price increase quickly tilts public belief against the deviator and steepens its effective demand. With low κ, consumers are also willing to check the rival, further disciplining markups. As formalized in Proposition 5.2 and Proposition 5.3, the mixed equilibrium concentrates probability on a narrow, partially overlapping support, and Proposition 6.2 implies a lower mean price.

In the *bottom* panel, dispersion widens and the mean price rises. When q is low, actions are weak signals, so beliefs react little to temporary price hikes; higher κ also makes checking the alternative less attractive, flattening the effective demand each firm faces. Both forces expand the price set that leaves rivals indifferent, widening the supports and shifting them up—again in line with Proposition 5.3 and Proposition 6.2. Competition still bites (supports overlap), but the bite is weaker when learning is noisy and search is costly.

Let N_{A} be the (random) total number of sales made by A before absorption (including the absorbing sale if any). The *expected* number of sales from initial belief η_0 solves the standard renewal equation:

$$\mathbb{E}_{\eta_0}[N_{\mathrm{A}}] \;=\; \int_0^\infty \mathbb{E}_{\eta_0}[\delta_{\mathrm{A}}(\eta_t; p_A, p_B)]\; \mathrm{d}t, \qquad \text{with } \eta_t \text{ evolving via } K(\cdot \mid \cdot; p_A, p_B). \tag{5.2}$$

[4] Equivalently, starting at B yields a buy-B threshold $1-\eta^*$ in the posterior $x = \mathbb{P}(\mathrm{A\ high} \mid s)$, so she buys A iff $x > 1-\eta^*$.

[5] Figure supports are computed on a finite grid with Monte Carlo demand; higher resolution only sharpens the shape and leaves the qualitative comparative statics unchanged.

Hence with static prices,

$$\Pi_{\mathrm{A}}(p_A, p_B; q, \lambda, \kappa, \alpha, \eta_0) \;=\; p_A \cdot \mathbb{E}_{\eta_0}[N_{\mathrm{A}}], \qquad \Pi_{\mathrm{B}}(p_A, p_B; \cdot) \;=\; p_B \cdot \mathbb{E}_{\eta_0}[N_{\mathrm{B}}], \tag{5.3}$$

and $\mathbb{E}[N_{\mathrm{A}}] + \mathbb{E}[N_{\mathrm{B}}] = \mathbb{E}[\text{total arrivals until absorption}]$.

Lemma 5.1. *For any primitives $(q, \lambda, \kappa, \alpha, \eta_0)$ and compact price set $p_A, p_B \in [0, \bar{p}]$, the profit functions Π_{A} and Π_{B} are bounded and upper semicontinuous in (p_A, p_B). Moreover, for fixed (p_A, p_B), Π_i is continuous in q and κ.*

Proof sketch. Boundedness: absorption occurs almost surely (standard BHW-type arguments with censored signals under $q > 1/2$ and the absorbing sets in Proposition 4.1); hence $\mathbb{E}[N_i] < \infty$. Upper semicontinuity follows from the piecewise-constant structure of $\pi_{\mathrm{A}|\cdot}(\eta)$ combined with dominated convergence for (5.2), noting that threshold sets where η_{A} or η_{B} equal η^* have measure zero under the induced process. Continuity in (q, κ) holds because thresholds and posteriors vary continuously in these primitives. □

5.2 Price Best Responses and Dispersion

We restrict posted prices to a compact interval $[0, \bar{p}]$ and adopt a measurable tie-breaking rule (e.g., equal split) at consumer indifference.[6] Let N_i be the (random) number of purchases from firm i before absorption under (p_i, p_{-i}). Firm i's objective is the expected revenue $\pi_i(p_i, p_{-i}) = p_i \cdot \mathbb{E}[N_i \,|\, p_i, p_{-i}]$. We call a (possibly mixed) strategy profile a Bayesian Nash equilibrium (BNE).

We summarize dispersion by the support width of the symmetric mixed equilibrium price distribution F:

$$W(q, \kappa) \;\equiv\; \sup \operatorname{supp} F \;-\; \inf \operatorname{supp} F, \tag{5.4}$$

and we denote the mean equilibrium price by

$$\bar{p}^{\mathrm{eq}}(q, \kappa) \;\equiv\; \int p \,\mathrm{d}F(p). \tag{5.5}$$

Proposition 5.1. *In the static pricing game with observational learning and random arrivals, a (possibly mixed) Bayesian Nash equilibrium exists. In the symmetric environment (same primitives for both firms), there exists a symmetric BNE.*[7]

Proof sketch. The game is finite-dimensional, with compact strategy sets. Profits are bounded and upper semicontinuous by Lemma 5.1. Apply a standard existence result for discontinuous games (e.g., Reny's better-reply security) since payoff discontinuities occur only on measure-zero threshold sets and best replies are nonempty and convex-valued by linearity of (5.3) in p_A, p_B. □

[6] At exact indifference we split demand evenly; all results are robust to any measurable tie-break and to arbitrarily small payoff trembles.

[7] Learning makes payoffs discontinuous in prices at knife-edge beliefs. We use payoff security and upper semicontinuity (à la Reny) to guarantee a mixed equilibrium.

We next characterize dispersion. Write $\mathrm{supp}(\sigma_i)$ for the support of firm i's (possibly mixed) equilibrium price.

Proposition 5.2. *In any symmetric BNE with mixing, each firm randomizes on a* connected *interval* $[\underline{p}, \overline{p}] \subset [0, \overline{p}]$*, the equilibrium CDF is continuous (no atoms), and the two supports* overlap. *Moreover, on* $[\underline{p}, \overline{p}]$ *the expected profit is constant, and outside the interval there are strict best replies back into the support.*[8]

Proof sketch. When κ is very small (cheap search), consumers are highly responsive to signals; price undercutting is severe and pure-strategy equilibria may exist at low margins. When κ is very large, initial-visit lock-in dominates and pure strategies emerge at higher prices (Diamond limit). For interior κ, demand exhibits both an information component (signals move choices) and a lock-in component (first-visit effect), making profits quasi-linear in price over regions determined by η^*. Indifference conditions across the boundaries of these regions imply mixing with a connected support; symmetry then forces partial overlap. Comparative statics follow from the monotonic effects of q and κ on η^* and the cascade boundaries (Proposition 4.1), which shift the profit slopes and hence the width of supports. □

Define the *support width* $W(q, \kappa)$ as the length of $\mathrm{supp}(\sigma_i)$ (symmetric at equilibrium).

Proposition 5.3. *Holding* $(\lambda, \overline{p})$ *fixed, the symmetric equilibrium support width* $W(q, \kappa)$ *is (weakly) decreasing in signal precision* q *and (weakly) increasing in the search cost* κ*.*

Proof sketch. Higher q shrinks the intermediate region where actions hinge on signals (Proposition 4.1), steepening the demand response to small price differences and reducing the range of prices that keep the opponent indifferent—hence narrower supports. Higher κ weakens switching incentives, expands the intermediate region where first-visit lock-in matters, and flattens the opponent's profit over a wider price interval—hence wider supports. The limiting statements follow from the no-immediate-cascade and never-search bounds (Lemma 4.3 and (4.5)). □

Finally, we record a benchmark on wrong cascades at equilibrium:

Proposition 5.4. *In any equilibrium with* $\eta_0 = \frac{1}{2}$ *and* $q \in (1/2, 1)$*, the probability of a wrong cascade is strictly positive whenever the equilibrium support intersects prices for which* $\eta^*(p_A - p_B, \kappa) \in (1-q, q)$*. As* $q \uparrow 1$ *or* $\kappa \downarrow 0$*, any equilibrium sequence has wrong-cascade probability converging to zero.*

Proof sketch. Combine Proposition 4.2 with the indifference conditions defining the mixed supports in Proposition 5.2. The limiting statement follows from Proposition 4.2. □

[8] Atoms invite profitable one-sided deviations; gaps would be strictly dominated by nearby prices. Overlap follows from mutual indifference and continuity of best replies.

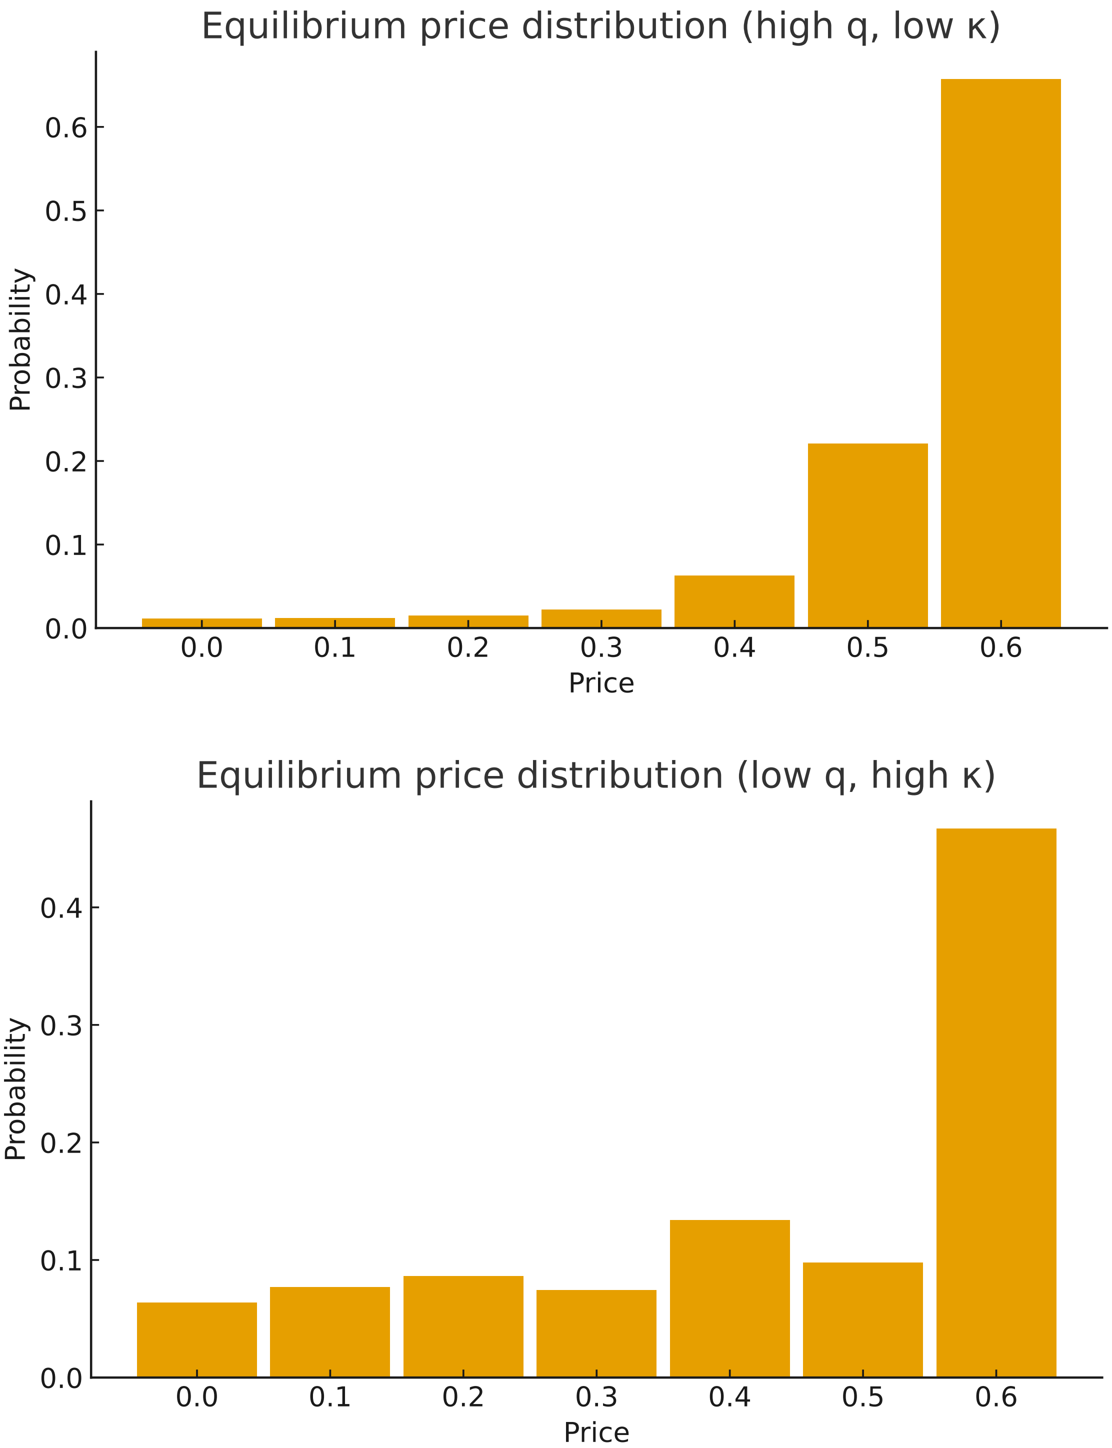

Figure 3: Approximate symmetric mixed-strategy price distributions (grid/MC solver). Top: high information (high q, low κ); Bottom: low information (low q, high κ). Supports shrink with q, expand with κ; mean prices comove. See Propositions 5.2, 5.3 and 6.2. Because prices are transfers, the welfare gain comes from higher accuracy and less search—so welfare rises with higher q and lower κ as in Proposition 6.3.

6 Comparative Statics and Welfare

This section asks what levers actually move outcomes in our environment. The primitives that govern informativeness—signal precision q, search frictions κ, and arrival intensity λ—shape both how quickly the public belief $\eta(t)$ settles and what prices firms can sustain before it does. We begin by isolating absorption probabilities and times, because they pin down the informational backdrop against which price competition plays out. We then turn to the pricing side—how dispersion and mean prices react to q and κ, and what that implies for surplus. Finally, we connect the pieces in a welfare accounting that separates pure information losses (wrong purchases) from avoidable externalities (excess search), providing a simple planner benchmark and a policy map.

We study how primitives (q, κ, λ) shape (i) absorption probabilities and times; (ii) equilibrium pricing and dispersion; and (iii) welfare.

6.1 Absorption Probabilities and Times

Absorption is the informational "end state" of the market: once $\eta(t)$ crosses an absorbing band, subsequent actions herd and prices become history-independent. Understanding how often we end up in each band (and how quickly) matters for two reasons. First, these probabilities determine the ex-ante risk of getting stuck with the inferior product; second, the time to absorption controls how long firms compete on a sloped, belief-sensitive demand curve before learning hardens choices. The results below formalize two intuitive points highlighted in the figures: accuracy improves and convergence speeds up as signals get cleaner and search gets cheaper, while λ merely time-changes the process (Propositions 4.2 and 6.1).

Let $\tau_\uparrow$ (resp. $\tau_\downarrow$) be the hitting time of the absorbing up- (resp. down-) cascade regions $[\bar{\eta}, 1]$ (resp. $[0, \underline{\eta}]$) from Proposition 4.1. Let

$$P_{\text{right}} \equiv \mathbb{P}(\text{eventual cascade to the truly high-quality firm}), \quad P_{\text{wrong}} \equiv 1 - P_{\text{right}}.$$

Write $T \equiv \mathbb{E}[\tau_\uparrow \wedge \tau_\downarrow]$ for the expected time to absorption.

Proposition 6.1. *Fix $\eta_0 = \frac{1}{2}$ and equilibrium prices (possibly mixed). Then:*

(a) P_{right} is (weakly) increasing in q and (weakly) decreasing in κ; P_{wrong} has the opposite monotonicity.

(b) P_{right} and P_{wrong} are independent *of λ.*

(c) T is (weakly) decreasing in q and (weakly) increasing in κ. Moreover, T scales inversely with λ: $T(\lambda) = T(1)/\lambda$.

Proof. See Appendix D.

6.2 Prices, Dispersion, and Surplus

Prices are set in the shadow of learning. A price cut today not only attracts current demand; by shifting observed actions, it steepens tomorrow's demand through beliefs. This feedback is stronger when q is high and when consumers are willing to check (low κ), so effective competition tightens and mixed strategies concentrate. We therefore read price dispersion as an equilibrium statistic of informativeness: supports shrink with q and expand with κ, and mean prices move in tandem (Propositions 5.3 and 6.2). Because prices are transfers, the surplus consequences run through accuracy (fewer wrong purchases) and search (less costly checking), linking directly to our welfare accounting below.

Let σ_i denote firm i's equilibrium price distribution and define the (symmetric) support width $W(q,\kappa)$ as in Proposition 5.3. Let $\bar{p} \equiv \mathbb{E}_{\sigma_i}[p_i]$ denote the equilibrium mean price (symmetric).

Proposition 6.2. *The symmetric equilibrium mean price $\bar{p}^{\text{eq}}(q,\kappa)$ is (weakly) decreasing in q and (weakly) increasing in κ.*

Proof. See Appendix D.

6.3 Welfare Accounting and Planner Benchmark

To evaluate policy, we need a clean decomposition of welfare. In our actions-only environment, private consumers do not internalize the informational benefit their choices confer on future arrivals. The planner's benchmark moves the cutoff to trade off immediate surplus against the value of information generated by additional checking. This yields a simple split of the welfare gap into wrong purchases (information loss) and excess search (externality), and it motivates a state-contingent Pigouvian search subsidy that implements the planner's cutoff (Propositions 6.3 and 6.4 and Corollary 6.1). The accounting also clarifies how changes in q and κ show up in welfare: better signals and cheaper search both raise accuracy and reduce costly exploration.

Consumer surplus equals value minus payments and search costs; prices are transfers,[9] so *total welfare* removes them. Let S be total search outlays and let Q_{high} be the total number of purchases of the truly high-quality good until absorption. With $\Delta = 1$ and zero production costs,

$$\text{Welfare} \;=\; \mathbb{E}\big[\,v_L \cdot N \;+\; Q_{\text{high}} \;-\; S\,\big], \tag{6.1}$$

where N is the total number of purchases until absorption. Hence welfare losses arise from (i) purchases of the low-quality good (information inefficiency) and (ii) search costs.

We measure social welfare as intrinsic consumer value minus search costs; prices are transfers. Normalizing the low type to $v_L = 0$ and the gap to $\Delta = 1$, a purchase contributes 1 if it is of the

[9] We assume zero marginal costs; with positive costs, our welfare comparisons carry through after adding producer surplus, and our policy results (search subsidy, prominence) continue to work by the same information-externality logic.

objectively better product and 0 otherwise. At belief η, the instantaneous per-arrival welfare is

$$w(\eta; p_A, p_B, q, \kappa) \;=\; \Pr[\text{action buys the objectively better product} \mid \eta] - \kappa{\cdot}\Pr[\text{consumer searches} \mid \eta],$$

where probabilities are taken under the action likelihoods induced by the threshold rule in Lemma 4.1 and the observation structure. Total expected welfare is the expectation of the integral of $\lambda\, w(\eta(t); \cdot)$ until absorption.

Proposition 6.3. *Fix posted prices. (i) Welfare is increasing in signal precision and decreasing in search costs: $w_q(\cdot) \geq 0$ and $w_\kappa(\cdot) \leq 0$ pointwise in η, hence total expected welfare satisfies the same monotonicity. (ii) The arrival rate λ scales calendar time but not the per-arrival composition: absorption probabilities and wrong-cascade risk are invariant in λ, while elapsed time to absorption contracts like $1/\lambda$; thus total welfare per unit* arrival mass *is invariant to λ and total welfare per unit* calendar time *scales with λ.*

A planner values the information that each action generates for future consumers. Let $W(\eta)$ denote the continuation welfare from public belief η under the baseline behavior. A consumer who starts at firm A buys without checking when $\eta \geq \eta^*(p_A - p_B, \kappa)$ (Lemma 4.1). The planner's cutoff is lower because the marginal search also benefits future arrivals through its impact on $W(\cdot)$.

Proposition 6.4. *There exists a bounded, state-contingent per-search subsidy $s^P(\eta) \in [0, \kappa]$ such that the privately optimal threshold $\eta^*(p_A - p_B, \kappa)$ coincides with the planner's threshold under the subsidy. A convenient choice is*

$$s^P(\eta) \;=\; \mathbb{E}\big[\, W(\eta^+) \;-\; W(\eta^-) \;\big|\; \textit{marginal search at } \eta \big]\,,$$

the expected gain in future welfare from the belief update induced by the extra information produced by searching at η (with η^+ and η^- the post-action beliefs under the two possible observed actions). Under s^P, private consumers internalize the social value of information and implement the planner's cutoff.

Corollary 6.1. *Relative to the planner, the welfare shortfall under private behavior decomposes as*

$$\overbrace{\mathbb{E}[\textit{purchases of the low-quality product before absorption}]}^{\textit{wrong purchases}} \\ + \underbrace{\mathbb{E}[\textit{search costs paid above the planner's level}]}_{\textit{excess search}}. \quad (6.2)$$

Under the Pigouvian subsidy $s^P(\eta)$, the excess-search component vanishes by construction; the remaining gap is due to residual mislearning driven by bounded signals off the absorbing sets.

Proof. Prices cancel in (6.1); the planner reduces wrong purchases and (weakly) reduces search via the lower threshold and/or subsidy. See Appendix D.

7 Extensions

Two institutional features matter in practice: how often firms can adjust prices, and what consumers publicly observe besides actions. We study both. Calvo resets capture lumpy repricing common in retail and platforms; prominence captures default ordering in search and recommendation systems; and noisy outcomes (reviews) add public information beyond actions. Each extension asks a common question: does the institution make beliefs more responsive to quality and, through that channel, compress price dispersion and improve welfare?

7.1 Calvo Price Resets

Repricing is not continuous in many markets—menus, platforms, and contracts create lags. Calvo resets model this with Poisson opportunities to adjust. The mechanism is simple: when beliefs drift against a firm, the next reset lets it cut price, regain visits, and restart informative sampling, pulling $\eta(t)$ back toward the truth. As the reset hazard rises, stale prices persist for less time, stationary dispersion compresses, and mean prices fall; when signals are sufficiently informative, the risk of wrong cascades also declines (Propositions 7.1 and 7.3 and Corollary 7.2). This connects pricing frictions to learning, not just to nominal rigidity.

We modify the baseline by allowing each firm to reset its price at random Poisson times (Calvo pricing). Time remains continuous. Consumers arrive at rate λ as before and observe current prices and past purchase choices.

Assumption 7.1. *Firm $i \in \{\mathrm{A}, \mathrm{B}\}$ receives price reset opportunities according to an independent Poisson clock of hazard $\alpha > 0$. When a reset occurs, firm i selects a new posted price from a compact set $[0, \bar{p}]$. Between resets, prices remain fixed.*[10]

Each firm $i \in \{A, B\}$ receives independent Poisson "reset" opportunities with hazard $\alpha > 0$. Between resets, posted prices (p_A, p_B) remain fixed and consumers arrive at rate λ, observe actions (not payoffs), and update the public belief $\eta(t)$ exactly as in the baseline. At a reset, firm i replaces p_i by drawing from a (possibly mixed) *reset policy* $\sigma_i(\cdot \mid \eta, p_{-i})$ that depends only on the current state (η, p_i, p_{-i}). We focus on *stationary Markov reset policies* (time-invariant and measurable in the state) and on prices constrained to a compact interval $[0, \bar{p}]$.

Definition 7.1. *Given $\alpha, \lambda, q, \kappa$, a pair of stationary Markov reset policies (σ_A, σ_B) is a* stationary Markov reset equilibrium *if for each firm i and every state (η, p_i, p_{-i}), $\sigma_i(\cdot \mid \eta, p_{-i})$ maximizes firm i's expected discounted revenue (with discount induced by the Poisson clocks) given the rival's policy σ_{-i} and the induced belief process under observed actions.*

Existence of a stationary Markov reset equilibrium follows from a Kakutani fixed point on the product of compact, convex sets of Borel–measurable reset policies, using Feller continuity of the state kernel on $\mathcal{S} = [0, 1] \times [0, \bar{p}]^2$ and Poisson–discounted revenues; see Appendix E.

[10] Stationary Markov reset policies maximize Poisson-discounted revenue; the Poisson clock induces a standard continuous-time discount factor.

Proposition 7.1. *Under compact price sets $[0, \bar{p}]$, independent Poisson resets with hazard $\alpha > 0$, and the observational-learning environment above, a stationary Markov reset equilibrium exists. Moreover, for any pair of such policies, the induced state process on $\mathcal{S} = [0,1] \times [0, \bar{p}]^2$ is a Feller Markov process that admits at least one invariant measure.*

Proposition 7.2. *Let $W(\alpha)$ be the stationary support width of the equilibrium price distribution and $\bar{p}(\alpha)$ the stationary mean price. Then $W(\alpha)$ and $\bar{p}(\alpha)$ are (weakly) decreasing in α. If $q \geq \bar{q}(\kappa)$, the stationary wrong-cascade probability $\Pi_{\text{wrong}}(\alpha)$ is (weakly) decreasing in α.*

Corollary 7.1. *As $\alpha \to \infty$, stationary dispersion vanishes, $W(\alpha) \to 0$, and stationary mean prices converge to the myopic best-reply levels to the contemporaneous belief (hence $\bar{p}(\alpha)$ approaches the high-q/low-κ benchmark). If $q \geq \bar{q}(\kappa)$, the stationary wrong-cascade probability satisfies $\Pi_{\text{wrong}}(\alpha) \to 0$.*

Comparative statics in the reset hazard. Write $W_{\text{st}}(\alpha; q, \kappa)$ for the cross-sectional (steady-state) price-support width (symmetric case), and let $P^{\text{st}}_{\text{wrong}}(\alpha)$ be the steady-state probability that the process is in a down-cascade when A is truly high (and vice versa).

Figure 4 shows a time line of prices $p_A(t)$ and $p_B(t)$ together with the public belief $\eta(t)$; vertical ticks mark price-reset events for each firm. We plot the high-information/low-friction regime and use a positive reset hazard α so that prices occasionally jump when a reset opportunity arrives.

The figure makes the mechanism behind Propositions 7.1 and 7.3 transparent. When beliefs drift against a firm, the next reset lets it cut its price, regain visits, and restart informative sampling; this bends $\eta(t)$ back toward the truth. As α increases, stale prices persist for less time, stationary dispersion compresses, and convergence accelerates—summarized formally in Corollary 7.2. In short, more frequent repricing disciplines dispersion by enabling rapid responses precisely when beliefs would otherwise get "stuck."

Proposition 7.3. *In any symmetric stationary Calvo equilibrium with $\eta_0 = \frac{1}{2}$:*

(a) $W_{st}(\alpha; q, \kappa)$ is (weakly) decreasing in α for interior κ and $q \in (1/2, 1)$.

(b) The expected calendar time to absorption (from any interior state) is (weakly) decreasing in α; moreover, $P^{st}_{wrong}(\alpha)$ is (weakly) decreasing in α provided q exceeds a threshold that makes immediate cascades nondegenerate (i.e., $\eta^* \in (1-q, q)$ on a set of positive stationary measure).

Proof sketch. (a) More frequent resets intensify competitive undercutting in states with slack thresholds, shrinking the range of sustained markups and thus the stationary support. (b) Resets allow firms to react to belief shocks, steering the process more quickly into absorbing regions; with sufficiently informative signals (q large), resets reduce the persistence of wrong-side states because firms price more aggressively when behind. See Appendix E. □

Corollary 7.2. *As $\alpha \downarrow 0$ the stationary Calvo equilibrium converges (weakly) to the baseline static-price environment of §5. As $\alpha \uparrow \infty$, prices track myopic best replies almost continuously; steady-state dispersion collapses and wrong-cascade probability vanishes when q is sufficiently high or κ sufficiently low.*

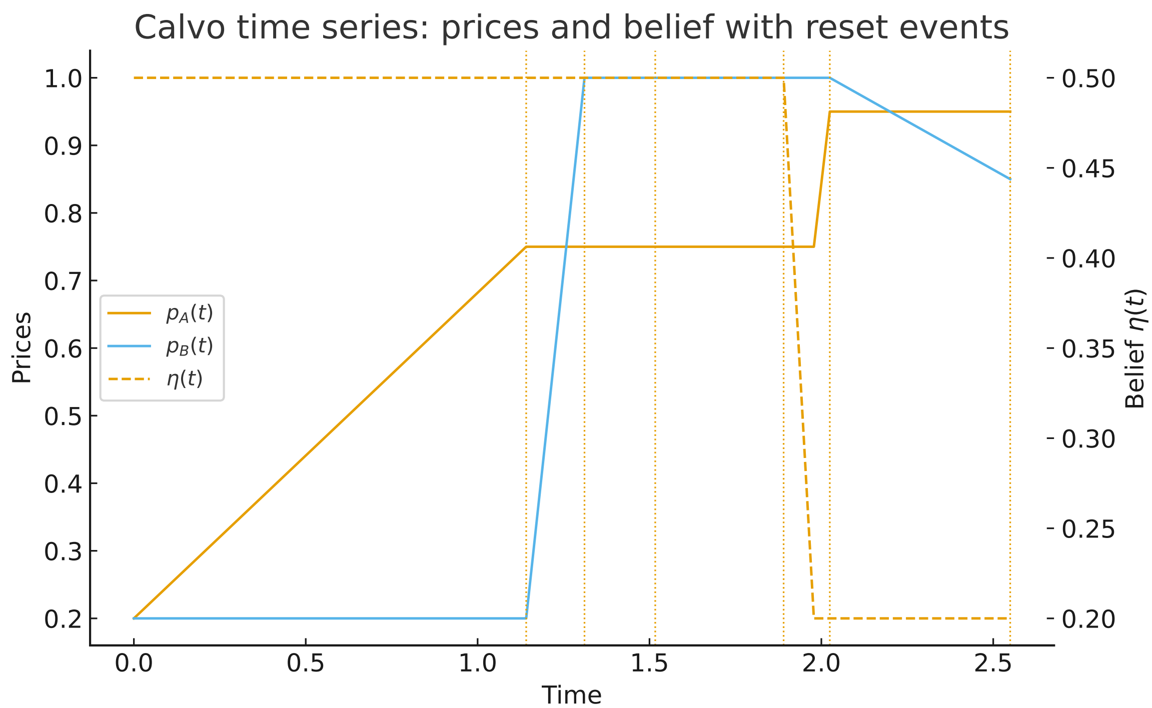

Figure 4: Calvo pricing: prices and public belief over time with reset events (vertical ticks). Formally, a higher reset hazard α weakly decreases stationary dispersion $W(\alpha)$ and the stationary mean price $\bar{p}(\alpha)$, and—when signals are sufficiently informative—lowers the steady-state wrong-cascade probability. More frequent resets compress dispersion and speed absorption by letting disadvantaged firms re-price when beliefs move against them. See Propositions 7.1 and 7.3 and Corollary 7.2. By hastening informative sampling and reducing wrong purchases, higher α raises welfare in the sense of Proposition 6.3.

7.2 Prominence Regulation

Default ordering—who is shown first—shapes where consumers start and thus how often they pay κ to check. A bias in first visits softens effective competition for the prominent firm and tilts early actions toward it, even when it is not better. The upshot is wider price supports, higher mean prices, and lower welfare relative to neutral prominence, ex-ante under a symmetric prior (Propositions 7.4 and 7.5). The policy message is narrow but practical: prominence is not innocuous when only actions are observed; neutrality disciplines both prices and mislearning.

We let a regulator (or platform) choose a *prominence* parameter $\phi \in [0, 1]$ that sets the probability the arriving consumer first visits A; the baseline is $\phi = \frac{1}{2}$. Given ϕ, arrivals and belief updates proceed as in the baseline (actions observed, payoffs not), with the same search technology and prices.

Definition 7.2. *At each arrival, the consumer first visits A with probability ϕ and B with probability $1 - \phi$, independently of the past. We interpret ϕ as a default ordering or slot prominence; $\phi = \frac{1}{2}$*

corresponds to neutral ranking.[11]

Prominence affects both competition and learning: biasing first visits toward a firm increases the incidence of "buy without checking" at that firm (because checking costs κ), flattening the effective demand and feeding back into public beliefs via observed actions.

Proposition 7.4. *Let $W(q,\kappa,\phi)$ be the symmetric equilibrium support width and $\bar{p}^{\mathrm{eq}}(q,\kappa,\phi)$ the symmetric mean price (cf. (5.4)–(5.5)). Then, holding (q,κ) fixed,*

$$W(q,\kappa,\phi) \text{ and } \bar{p}^{\mathrm{eq}}(q,\kappa,\phi) \text{ are (weakly) increasing in } |\phi - \tfrac{1}{2}|.$$

In words, greater prominence bias (either direction) softens effective competition, widening stationary dispersion and raising mean prices.

Proposition 7.5. *Evaluate welfare and wrong-cascade probability ex ante under the symmetric prior and before the (unknown) state ("A high" vs "B high") is realized. Then:*

1. *The ex-ante wrong-cascade probability $\Pi_{\mathrm{wrong}}(\phi)$ is minimized at $\phi = \frac{1}{2}$ and satisfies $\Pi_{\mathrm{wrong}}(\phi) = \Pi_{\mathrm{wrong}}(1-\phi)$.*

2. *Ex-ante social welfare (prices are transfers) is maximized at $\phi = \frac{1}{2}$ and satisfies $W^{\mathrm{soc}}(\phi) = W^{\mathrm{soc}}(1-\phi)$.*

Intuitively, neutral prominence balances exploration across firms and avoids systematically tilting early actions toward one side when quality is unknown.[12]

Corollary 7.3. *A simple "neutral prominence" rule $\phi = \frac{1}{2}$ weakly improves welfare relative to any biased prominence $\phi \neq \frac{1}{2}$ (ex ante under the symmetric prior), and it lowers equilibrium prices and dispersion by Proposition 7.4. Combining neutral prominence with the Pigouvian search subsidy of Proposition 6.4 attains the planner's search cutoff and minimizes wrong-cascade risk given (q,κ).*

Remark. If the platform has (even slightly) informative priors about which product is better, the optimal prominence may tilt toward its posterior-best arm to accelerate correct learning. In our baseline symmetric-prior evaluation, $\phi = \frac{1}{2}$ is optimal.

7.3 Richer Public Information: Noisy Outcomes

Many platforms surface reviews or ratings that are informative, if imperfect. Adding public outcome signals augments actions with another channel for belief updating. Even noisy reviews shrink the absorbing bands, reduce the chance of settling on the wrong product, and tighten mixed pricing by making deviations more quickly punished through beliefs. In the limit of frequent or

[11] Prominence can be read as default ordering, slot position, or salience on a platform; in our baseline we evaluate ϕ ex-ante under a symmetric prior.

[12] If the platform has informative priors about quality, optimal prominence may tilt toward its posterior-best arm to accelerate correct learning. Our neutrality result is ex-ante under a symmetric prior.

accurate reviews, dispersion collapses toward the common-knowledge benchmark (Propositions 7.6 and 7.7 and Corollary 7.4). This extension underscores a general lesson: policies that raise public informativeness are stricter and more robust tools than direct price intervention in action-only environments.

We enrich the observational environment by adding public *review* signals in addition to observed actions. After each purchase, with probability $\mu \in [0,1]$ a public signal $R_t \in \{+,-\}$ is generated and observed by all subsequent consumers; conditional on the true high-quality firm, R_t is correct with accuracy $r \in (1/2, 1)$ (MLRP). Consumers observe the entire history of prices, actions, and reviews and update the public belief $\eta(t)$ accordingly.[13]

Let $\psi \equiv (\mu, r)$ index review informativeness in the Blackwell sense (higher ψ means a mean-preserving increase in signal precision and/or frequency).

Proposition 7.6 (Learning with reviews)**.** *Relative to the baseline ($\psi = 0$), adding reviews shrinks the absorbing bands and reduces wrong-cascade risk. Formally, the up-/down-absorbing sets $[\,\bar{\eta}(\psi), 1\,]$ and $[\,0, \underline{\eta}(\psi)\,]$ satisfy $\bar{\eta}(\psi)$ increasing in ψ and $\underline{\eta}(\psi)$ decreasing in ψ (Blackwell order), and the ex-ante wrong-cascade probability $\Pi_{\text{wrong}}(\psi)$ is (weakly) decreasing in ψ.*

Proposition 7.7 (Prices under reviews)**.** *In the symmetric mixed equilibrium, the stationary support width $W(q, \kappa, \psi)$ and mean price $\bar{p}^{\text{eq}}(q, \kappa, \psi)$ are (weakly) decreasing in ψ.*

Corollary 7.4 (Limits)**.** *If reviews are sufficiently informative or frequent (e.g., $\mu > 0$ and $r \uparrow 1$), then $\Pi_{\text{wrong}}(\psi) \to 0$ and $W(q, \kappa, \psi) \to 0$; prices converge to those under common knowledge of quality (dispersion vanishes).*

Intuition. Reviews add public information orthogonal to actions. More informative public signals tighten the mapping from the true state to observed histories, pushing beliefs away from the wrong absorbing band and disciplining markups via steeper effective demand. In the limit of nearly perfect or frequent reviews, cascades unravel and mixed pricing collapses to a degenerate support.

8 Conclusion

We studied a vertically differentiated duopoly in which consumers arrive randomly, observe past actions (not payoffs), and can pay to check the alternative seller. Prices are posted once in the baseline and interact with social learning via a simple threshold rule. This interaction generates belief-absorbing regions (cascades), a transparent buy–A cutoff $\eta^*(p_A - p_B, \kappa)$, and tractable boundaries for up/down cascades.

On the firm side, static price competition admits (possibly mixed) equilibria. In the symmetric case we show that, for interior search costs, both firms mix over connected, partially overlapping supports (Proposition 5.2). Dispersion is tightly linked to learning incentives: supports narrow as signal precision q rises and widen as search cost κ increases (Proposition 5.3). These comparative

[13] Think platform-mediated ratings: μ is the incidence of displayed reviews and r their average accuracy. We evaluate the ex-ante symmetric prior; platform bias in surfacing reviews would interact with prominence.

statics are mirrored in absorption outcomes: correct cascades become more likely and faster when signals are more informative and search is cheaper, with arrival intensity λ scaling time but not probabilities (Proposition 6.1).

The welfare accounting isolates the information externality. Prices are transfers; losses come from (i) purchases of the low-quality good before absorption and (ii) search outlays. Both move in the "right" direction with q and κ (Proposition 6.3). A simple planner benchmark chooses a lower buy–A cutoff than private consumers and is implementable via a Pigouvian search subsidy that internalizes the value of information to future arrivals (Proposition 6.4). This yields an intuitive decomposition of the welfare gap into "wrong purchases" and "excess search" (Corollary 6.1).

We also extended the model to Calvo pricing, where each firm receives random opportunities to reset price. A stationary equilibrium in Markov reset policies exists (Proposition 7.1). More frequent resets discipline dispersion and accelerate learning by letting disadvantaged firms re-price aggressively when beliefs move against them; steady-state wrong-cascade probability falls when signals are sufficiently informative (Proposition 7.3, Corollary 7.2).

Methodologically, the paper offers a compact toolkit—posterior formulas, a one-line threshold, and closed-form cascade boundaries—that makes learning–pricing feedback analyzable in continuous time with random arrivals. The same primitives drive firms' mixed supports and the belief process, making the comparative statics sharp and easily portable.

Two extensions look especially promising. First, a fully dynamic pricing game (periodic re-posting rather than Calvo) would let us study reputation–pricing cycles and test whether competition eliminates wrong cascades in steady state. Second, endogenizing initial-visit prominence (or adding noisy payoff revelation/reviews) would quantify how attention design and platform information policies reshape dispersion and welfare. Both directions are empirically suggestive and remain analytically tractable within our framework.

A Existence in the Static Pricing Game

Theorem A.1 (Existence of mixed equilibrium in the static pricing game)**.** *In the one-shot pricing subgame at $t = 0$ with price sets $P_i = [0, \bar{p}]$ compact, there exists a (possibly mixed) Bayesian Nash equilibrium. In the symmetric environment, there exists a symmetric mixed equilibrium.*

Proof. Step 1 (bounded payoffs and a.s. absorption). For any (p_i, p_{-i}), let N_i be the number of purchases from i before absorption. By Lemma 4.2 absorption occurs almost surely in finite arrival time, hence $\mathbb{E}[N_i] < \infty$. Firm i's payoff is $\pi_i(p_i, p_{-i}) = p_i\, \mathbb{E}[N_i \,|\, p_i, p_{-i}]$, which is bounded on $[0, \bar{p}]^2$.

Step 2 (structure of discontinuities). Let $x^*(\Delta p, \kappa)$ be the posterior cutoff in (4.4). A consumer's action changes discretely only on *knife–edge* loci where a signal-posterior equals x^*; these are smooth curves in (p_i, p_{-i}) and thus Lebesgue–null. Denote by $\mathcal{K} \subset [0, \bar{p}]^2$ the union of such loci.

Step 3 (upper semicontinuity). Fix p_{-i}. For any sequence $p_i^n \to p_i$, the induced likelihoods and Bayes operator are continuous in (p_i, p_{-i}) away from $\mathcal{K}$. Since $\mathcal{K}$ is null and N_i is integrably bounded by Step 1, dominated convergence yields

$$\limsup_{n\to\infty} \ \pi_i(p_i^n, p_{-i}) \ \le \ \pi_i(p_i, p_{-i}),$$

i.e., $\pi_i(\cdot, p_{-i})$ is upper semicontinuous.

Step 4 (better-reply security; Reny, 1999). Extend the game to mixed strategies over $[0, \bar{p}]$ with the weak topology; opponents' mixed strategies are atomless. Let $u_i(p_i, \sigma_{-i})$ be i's expected payoff against σ_{-i}. Because $\mathcal{K}$ is null and σ_{-i} is atomless, small perturbations of p_i shift the path distribution only on events of arbitrarily small probability; by Step 1 and dominated convergence, for any $\varepsilon > 0$ there exists $\delta > 0$ and a finitely supported ε–robust mixture for i that secures $u_i(\cdot, \sigma_{-i}) - \varepsilon$ against all pure p_{-i} in a δ–neighborhood of $\operatorname{supp} \sigma_{-i}$. Hence the game is *better-reply secure* in the sense of Reny (1999).

Step 5 (apply Reny's theorem). Strategy sets are compact and convex; payoffs are bounded and upper semicontinuous; the game is better-reply secure. By Reny's existence theorem for discontinuous games, a (possibly mixed) equilibrium exists. Symmetry of primitives delivers a symmetric equilibrium. □

Remark A.1 (Why discontinuities do not obstruct existence). Discontinuities arise only when a posterior equals the cutoff in (4.4). Under any atomless opponent mix, these events have probability zero; thus expected payoffs are upper semicontinuous and payoff security holds.

Remark A.2. Discontinuities arise only when, at some belief, a realized posterior equals the action cutoff in (4.4). Under any atomless mixed price of the opponent, the probability of landing exactly on those thresholds is zero, so the induced payoff correspondence is upper semicontinuous on a full–measure set, and better–reply security holds.

B Proofs: Consumer Problem

Proof of Lemma 4.1. Starting at A, let x be the posterior that A is better after the consumer's private signal (MLRP, i.i.d.). Buying at A yields $x \cdot \Delta - p_A$; checking B costs κ and then the consumer buys the higher–expected–surplus option. The marginal consumer is indifferent when the expected gain from checking equals κ, which solves to the public-belief cutoff in (4.4). Monotonicity in $(p_A - p_B)$ and κ follows by inspection. Symmetry holds for first visit to B. □

Proof of Remark (4.5) The maximal benefit from switching to B when starting at A is attained at the lowest posterior x (i.e., the least favorable to A), yielding

$$\max_{x\in[0,1]} \big\{ -D(x; p_A - p_B) \big\} = \max_{x\in[0,1]} \big\{ (p_A - p_B) + 1 - 2x \big\} = 1 + (p_A - p_B).$$

If $\kappa \ge 1 + (p_A - p_B)$, switching is never optimal. □

Proof of Proposition 4.1. Given (4.4), define $\bar{\eta}$ as the least belief at which even a counter–signal that maximally favors B does not induce checking/switching away from A; analogously for $\underline{\eta}$ starting at

B. These values equal (4.8)–(4.9). At any $\eta \geq \bar{\eta}$, the action is A regardless of the private signal, so observed actions are uninformative and Bayes' rule leaves the belief in $[\,\bar{\eta}, 1\,]$; similarly for $[\,0, \underline{\eta}\,]$. This proves absorption. Comparative statics follow from differentiating (4.8)–(4.9): higher q pushes $\bar{\eta}$ up and $\underline{\eta}$ down (harder to start a cascade), while higher κ pulls $\bar{\eta}$ down and $\underline{\eta}$ up (easier to start one). $\square$

Proof sketch of Lemma 4.2 Index arrivals by $n = 0, 1, 2, \dots$ and write η_n for the belief after the nth arrival. Between arrivals, $\eta(t)$ is constant, so $\{\eta_n\}$ is a time-homogeneous Markov chain on $[0, 1]$ with absorbing sets $[0, \underline{\eta}] \cup [\bar{\eta}, 1]$ (Proposition 4.1). With bounded signals $q \in (1/2, 1)$ there is a positive probability, uniformly on compact interiors, of a one-step update that moves the likelihood ratio by at least a fixed amount; standard random-walk/hitting arguments then imply η_n hits the absorbing sets almost surely in finite n. Let $\tau^{\#}$ be the (a.s. finite) absorption index. If $N_\lambda(t)$ is a Poisson process of rate λ, then $N_\lambda(t) \overset{d}{=} N_1(\lambda t)$, hence the calendar absorption time is $\tau_\lambda = \inf\{t : N_\lambda(t) \geq \tau^{\#}\} \overset{d}{=} \inf\{t : N_1(\lambda t) \geq \tau^{\#}\} = \tau_1/\lambda$. Absorption probabilities depend only on the distribution of the path in arrival *counts*, not on calendar time, so they are invariant in λ. $\square$

Proof of Lemma 4.3. At $\eta_0 = \frac{1}{2}$ with equal prices, the initial action is history–free. Immediate herding occurs exactly when η_0 lies in an absorbing set, i.e., $\frac{1}{2} \geq \bar{\eta}$ or $\frac{1}{2} \leq \underline{\eta}$ by Proposition 4.1. The equivalent condition in terms of η^* follows from (4.4) and the definitions of $\bar{\eta}, \underline{\eta}$. $\square$

Proof of Proposition 4.2. With bounded signals ($q < 1$) and positive search cost, there is a strictly positive probability of observing a run of actions that favors the wrong firm long enough to cross the corresponding boundary (the run event has positive probability by independence). Once inside an absorbing region, actions become uninformative and the belief stays there, yielding a wrong cascade with positive probability. As $q \uparrow 1$, private signals reveal the truth and actions become fully informative, so the probability of crossing the wrong boundary tends to zero. As $\kappa \downarrow 0$, consumers check more often, making actions more informative and again driving the wrong-cascade probability to zero. Arrival intensity λ only time–changes the process and does not affect the path distribution of actions per arrival, hence absorption probabilities are invariant to λ. $\square$

C Proofs: Firms and Equilibrium

Assumptions (recap).

(i) Prices lie in $[0, \bar{p}]$ compact; tie–breaking at indifference splits demand measurably.

(ii) Signals are i.i.d. with precision $q \in (1/2, 1)$ (MLRP).

(iii) Search/switch cost $\kappa > 0$; arrivals are Poisson with rate $\lambda > 0$.

(iv) Consumers observe prices and actions (not payoffs); belief updates use Bayes' rule.

(v) Absorption into $[0, \underline{\eta}] \cup [\bar{\eta}, 1]$ occurs a.s. in finite *arrival* time and absorption probabilities are invariant to λ (Lemma 4.2).

(vi) For the Calvo extension, each firm receives independent Poisson reset opportunities with hazard $\alpha > 0$; between resets prices are fixed.

Proof of Lemma 5.1 Absorption: with $q > 1/2$, the posteriors in (4.1)–(4.2) satisfy an MLRP; given Lemma 4.3 and the absorbing boundaries (Proposition 4.1), the induced belief process hits $[\bar{\eta}, 1]$ or $[0, \underline{\eta}]$ with probability one. Hence the total expected number of arrivals before absorption is finite, implying bounded $\mathbb{E}[N_i]$ and thus bounded profits. For upper semicontinuity, note that the demand intensities $\delta_{\mathrm{A}}(\cdot)$ in (5.1) are piecewise constant in (p_A, p_B) for each η, with jumps only at the threshold sets where η_{A} or η_{B} equals η^*. Under the induced belief distribution, these sets have zero measure; apply dominated convergence to (5.2). Continuity in (q, κ) follows by continuity of η^* and $\eta_{\mathrm{A}}, \eta_{\mathrm{B}}$ in these primitives. □

Proof sketch of Proposition 5.1. Prices lie in a compact set $[0, \overline{p}]$. For fixed (p_i, p_{-i}), expected revenue is bounded because the arrival process time–changes (Proposition 6.1) and absorption occurs a.s. by Proposition 4.1. The mapping $(p_i, p_{-i}) \mapsto \mathbb{E}[N_i]$ is continuous except on a set of knife-edge prices where the threshold in (4.4) aligns exactly with a posterior boundary; these points are removable under atomless opponents. Following Reny's existence theorem for discontinuous games, payoff security holds: for any $\varepsilon > 0$, firm i can choose an ε–robust reply that guarantees within ε of the best-reply payoff against prices in a small neighborhood (belief updates and likelihoods vary continuously away from knife edges). Payoffs are upper semicontinuous and bounded; hence a mixed BNE exists, and by symmetry a symmetric BNE exists. □

Proof sketch of Proposition 5.2. Suppose a mixing firm placed an atom at some $\tilde{p}$. Then the rival would profitably deviate to a nearby undercut/overcut that raises expected demand without a first-order revenue loss, contradicting best-reply indifference; hence no atoms. If the support had a gap (a, b) with positive mass to both sides, continuity of best replies and single-crossing in $p_i - p_{-i}$ implies that prices in the gap are strictly dominated by nearby prices, contradicting mixing; thus the support is connected. If supports did not overlap, say $\overline{p}_A < \underline{p}_B$, then each firm's best reply would jump to the interior of the opponent's support, violating mutual indifference; hence supports overlap. On the equilibrium support, expected profit must be constant by standard mixing logic. □

Proof sketch of Proposition 5.3. Let $D_i(p_i, p_{-i}; q, \kappa)$ denote firm i's expected demand (expected arrivals choosing i before absorption). A unilateral price increase rotates effective demand via two channels: (i) the myopic substitution effect and (ii) the belief–feedback effect through observed actions and search. The latter is stronger when q is higher and weaker when κ is higher (actions are more/less informative; consumers check more/less). Formally, the slope $\partial D_i / \partial p_i$ becomes more negative in q and less negative in κ (MLRP and Lemma 4.1 and Proposition 4.1). Best-reply graphs therefore contract under higher q and expand under higher κ, moving the support endpoints $[\underline{p}, \overline{p}]$ inward/outward. Hence $W(q, \kappa)$ is (weakly) decreasing in q and (weakly) increasing in κ. □

Proof of Proposition 5.4

If the mixed support intersects prices for which $\eta^* \in (1-q, q)$, then by Proposition 4.2 both up- and down-cascade absorption events have positive probability; the wrong-cascade probability is therefore strictly positive under the equilibrium mixture. The limiting statements follow from the vanishing wrong-cascade probability as $q \uparrow 1$ or $\kappa \downarrow 0$ (Proposition 4.2). $\square$

D Additional Lemmas and Technical Details

Proof of Proposition 6.1 (a) Higher q enlarges the informativeness of s (MLRP), shrinking the absorbing regions via Proposition 4.1 and increasing the chance that actions reflect the true state before absorption; thus P_{right} rises. Higher κ increases η^* ((4.4)), expanding the down-cascade region and contracting the up-cascade region, lowering P_{right}.

(b) The arrival rate λ time-changes the process but does not alter the *sequence* of actions conditional on prices and signals. Hence absorption probabilities depend only on belief transitions and are invariant to λ.

(c) With a Poisson clock, interarrival times are i.i.d. $\text{Exp}(\lambda)$ and independent of belief transitions. The expected number of arrivals to absorption does not depend on λ; hence $T(\lambda) = \mathbb{E}[\#\text{arrivals}]/\lambda$. Higher q (respectively, lower κ) shrinks the intermediate region and accelerates absorption in arrivals, so T decreases (respectively, increases). $\square$

Proof sketch of Proposition 6.2. The same comparative statics that shrink (expand) the best-reply graph under higher q (higher κ) also shift it downward (upward): informative actions and cheap search intensify effective competition. In the symmetric mixed BNE, the equilibrium CDF F solves mutual indifference across support endpoints; as the best-reply correspondence shifts, both endpoints and thus the mean $\bar{p}^{\text{eq}}(q,\kappa) = \int p \, \mathrm{d}F(p)$ move down with q and up with κ. $\square$

Proof of Proposition 6.3. By Lemma 4.1, consumer actions at belief η are mixtures over at most two moves (buy here, or pay κ to check and possibly switch). The probability of choosing the objectively better product is weakly increasing in q by MLRP and Bayes' rule: higher q tightens the link between the true state and actions, so $\Pr[\text{correct action} \mid \eta]$ increases pointwise in η. The probability of search is weakly decreasing in q because more informative actions reduce the option value of checking. Hence $w_q \geq 0$. Higher κ directly lowers w via the $-\kappa \Pr[\text{search}]$ term and (weakly) raises the private cutoff, reducing the search frequency and the informativeness of actions; both effects imply $w_\kappa \leq 0$. For (ii), changing λ only time-changes the arrival process. Absorption probabilities and the distribution of actions per arrival do not depend on λ (Proposition 6.1), while the time to absorption scales like $1/\lambda$, yielding the stated invariances. $\square$

Proof of Proposition 6.4. Fix (p_A, p_B). Let $U_{\text{priv}}(\eta)$ be the private continuation value at belief η and $W(\eta)$ the social continuation welfare (prices are transfers). The only wedge between U_{priv} and W is the information externality: a marginal search at belief η changes the distribution of future actions and thus $W(\cdot)$ for subsequent arrivals but is not internalized by the searching consumer. Denote by

$\Delta^{\mathrm{soc}}(\eta) = \mathbb{E}[W(\eta^+) - W(\eta^-) \mid \text{marginal search at } \eta]$ the expected social gain from the information produced by that search. A per-search subsidy $s^{\mathrm{P}}(\eta) = \Delta^{\mathrm{soc}}(\eta)$ aligns the private first-order condition with the planner's: it reduces the effective private search cost from κ to $\kappa - s^{\mathrm{P}}(\eta)$, shifting the private cutoff down to the planner's cutoff. Boundedness $s^{\mathrm{P}}(\eta) \in [0, \kappa]$ follows because W is bounded and marginal search cannot reduce expected future welfare in this environment. □

Proof of Corollary 6.1. Let $\{a_t\}$ be the sequence of observed actions and let $\mathbb{1}_{\mathrm{wrong}}(a_t)$ indicate a purchase of the low-quality product. Total welfare equals $\mathbb{E}\big[\sum_t \big(\mathbb{1} - \mathbb{1}_{\mathrm{wrong}}(a_t)\big)\big] - \kappa\, \mathbb{E}[\sum_t \mathbb{1}_{\mathrm{search}}(a_t)]$. Subtracting the planner's welfare and telescoping yields the two terms stated: (i) wrong purchases (a pure information loss) and (ii) search costs in excess of the planner's policy (a pure externality term). Under $s^{\mathrm{P}}(\eta)$, the planner's and private search cutoffs coincide, so the second term is zero; residual losses come only from bounded signals off the absorbing sets. □

E Proofs: Calvo Extension

Lemma E.1. *Under compact prices $[0, \bar{p}]$, independent Poisson arrivals/resets, and stationary Markov reset policies, the one–step transition kernel on $\mathcal{S} = [0, 1] \times [0, \bar{p}]^2$ is Feller (maps bounded continuous functions to bounded continuous functions).*

Proof sketch. Event probabilities (arrival vs. reset A/B) depend continuously on hazards; conditional action likelihoods and the Bayes operator are continuous in (η, p_A, p_B); at resets, prices are drawn from measurable policies on a compact set. Composition yields weak continuity of the kernel; boundedness is immediate. Hence Feller. Kakutani then applies as stated in the main text to deliver a stationary Markov reset equilibrium. □ □

Theorem E.1. *In the Calvo reset environment with state space $\mathcal{S} = [0, 1] \times [0, \bar{p}]^2$, compact action sets $A_i = [0, \bar{p}]$, independent Poisson arrivals/resets, and Poisson–discounted revenues, there exists a (possibly mixed)* stationary Markov reset equilibrium, *i.e., a pair of Borel measurable policies $\sigma_i : \mathcal{S} \to \Delta(A_i)$ such that each σ_i is a best reply to σ_{-i}.*

Proof. Policy space. Endow $\Delta([0, \bar{p}])$ with the weak topology (Prokhorov metric); it is compact and convex. Let Σ_i be the set of Borel measurable $\sigma_i : \mathcal{S} \to \Delta([0, \bar{p}])$, endowed with the topology of weak convergence pointwise on $\mathcal{S}$. By Tychonoff, Σ_i is compact and convex.

State dynamics. For a policy profile $\sigma = (\sigma_A, \sigma_B)$, the state $\xi_t = (\eta(t), p_A(t), p_B(t))$ evolves as a continuous-time Markov process with a jump at each event (arrival or reset). By Lemma E.1, the one-step kernel on $\mathcal{S}$ is Feller (maps bounded continuous functions to bounded continuous functions), uniformly in σ.

Discounted payoffs. Let $\rho > 0$ be the Poisson discount rate (equivalently, compute expected discounted *event* payoffs with geometric factor $\beta \in (0, 1)$). For any fixed σ, the bounded discounted value $V_i^\sigma : \mathcal{S} \to \mathbb{R}$ exists and is the unique bounded solution of the linear Poisson equation

$$V_i^\sigma(s) \;=\; R_i^\sigma(s) \;+\; \beta \int_{\mathcal{S}} V_i^\sigma(s')\, K_\sigma(s, \mathrm{d}s'),$$

where R_i^σ is the bounded instantaneous revenue and K_σ the Feller kernel. Standard contraction mapping (on $(\mathcal{C}(\mathcal{S}), \|\cdot\|_\infty)$) yields existence/uniqueness and continuity of V_i^σ in σ.

Best-reply correspondence. For each i and $s \in \mathcal{S}$, firm i's one-step decision at a reset maximizes a continuous affine functional of $\sigma_i(s) \in \Delta([0, \bar{p}])$ given $(V_i^\sigma, \sigma_{-i})$. By Berge's maximum theorem, the set of pointwise maximizers is nonempty, compact, convex, and upper hemicontinuous in (s, σ_{-i}). Taking the product over $s \in \mathcal{S}$ (measurable selections exist by standard measurable selection theorems) delivers a nonempty, convex, compact, upper-hemicontinuous best-reply correspondence $BR_i(\sigma_{-i}) \subset \Sigma_i$.

Fixed point. Consider $BR(\sigma) = BR_A(\sigma_B) \times BR_B(\sigma_A)$ on $\Sigma = \Sigma_A \times \Sigma_B$.

By Kakutani–Fan–Glicksberg, BR has a fixed point $\sigma^\star \in \Sigma$, which is a stationary Markov reset equilibrium. □

Proof sketch of Proposition 7.3. Couple two economies with identical primitives and common randomness (arrivals, private signals) but hazards $\alpha' < \alpha''$. Construct reset times by thinning: insert additional resets to obtain the α'' process. For any sample path on which beliefs drift against firm i, the next additional reset in the α'' economy allows i to cut price earlier, steepening effective demand and producing more informative actions thereafter. This pathwise dominance shortens the duration of stale prices and tightens the stationary distribution around local best replies; hence $W(\alpha'') \le W(\alpha')$ and $\bar{p}(\alpha'') \le \bar{p}(\alpha')$. When q is sufficiently high ($\ge \bar{q}(\kappa)$), these earlier corrections also weakly reduce the probability mass on wrong cascades by restarting informative sampling more often, yielding $\Pi_{\text{wrong}}(\alpha'') \le \Pi_{\text{wrong}}(\alpha')$. □

Proof sketch of Corollary 7.2. As $\alpha \to \infty$, the time between resets vanishes. On any compact horizon, prices track the (myopic) best replies to the contemporaneous state almost surely, so dispersion collapses and the stationary mean converges to the myopic level. When $q \ge \bar{q}(\kappa)$, belief drifts that would otherwise create wrong cascades are offset in the limit by arbitrarily fast re-pricing that restores informative sampling; hence $\Pi_{\text{wrong}}(\alpha) \to 0$. □

F Proofs: Prominence

Proof sketch of Proposition 7.4. For firm i, expected demand $D_i(p_i, p_{-i}; \phi)$ equals the arrival mass that starts at i and buys without checking plus the mass that starts at $-i$, decides to check, and then switches. As ϕ moves away from $\frac{1}{2}$ toward i, the share that starts at i increases and the share that checks falls (checking requires paying κ), so D_i becomes less sensitive to p_i (the slope $\partial D_i / \partial p_i$ is less negative). This softening of effective demand expands the best-reply correspondence; in the symmetric mixed equilibrium the support endpoints move outward and the support widens, while mutual indifference shifts the mean up. The argument is symmetric for a bias against i. Hence W and $\bar{p}^{\text{eq}}$ rise with $|\phi - \frac{1}{2}|$. □

Proof sketch of Proposition 7.5. Swap labels across states: under "A high/bias to A" the outcome distribution coincides with that under "B high/bias to B"; thus $\Pi_{\text{wrong}}(\phi) = \Pi_{\text{wrong}}(1-\phi)$ and likewise for welfare. At $\phi = \frac{1}{2}$ the derivative is zero by symmetry. Biasing prominence raises the chance that early actions tilt toward the prominent firm regardless of the true state; when the less-prominent firm is truly better, the path is more likely to cross the wrong absorbing boundary. Standard coupling (run the same signal and arrival sequence under ϕ and $1-\phi$ with swapped labels) yields that the ex-ante wrong-cascade probability attains its minimum at $\phi = \frac{1}{2}$, and ex-ante welfare (accuracy minus search costs; prices are transfers) attains its maximum there. □

Proof of Corollary 7.3. Immediate from Propositions 7.4 and 7.5 and the fact that the Pigouvian subsidy in Proposition 6.4 aligns private and planner search cutoffs. □

G Proofs: Noisy outcomes

Proof sketch of Proposition 7.6. Augment the public history with i.i.d. review signals R_t that satisfy MLRP with respect to the true state. For any prior η and realized action, the posterior after observing R_t is a mean-preserving contraction toward the truth; under the Blackwell order, more informative ψ moves the posterior closer (in convex order) to the state. This tightens the one-step Bayes operator and shifts the hitting boundaries: $\bar{\eta}$ rises and $\underline{\eta}$ falls. Standard coupling then yields $\Pi_{\text{wrong}}(\psi)$ weakly decreasing in ψ. □

Proof sketch of Proposition 7.7. Effective demand slopes depend on how beliefs react to histories. Increasing ψ steepens (increases the absolute value of) the own-price effect via the belief channel—deviations are punished more because reviews accelerate correction. Best-reply correspondences contract, pulling in support endpoints and lowering the indifference level, so both W and $\bar{p}^{\text{eq}}$ decrease. □

Proof sketch of Corollary 7.4. With $\mu > 0$ and $r \uparrow 1$ (or high-rate perfect reviews), the true state is revealed almost surely in finite expected time; mixed pricing cannot be sustained (off-support deviations are profitably deterred by fully informative beliefs), so dispersion collapses. □

H Simulation Notes

This appendix outlines a light, reproducible workflow to generate figures and tables: (i) cascade boundaries as functions of $(p_A - p_B, \kappa, q)$; (ii) Monte Carlo (MC) estimates of absorption probabilities/times under fixed prices; (iii) a simulation-based solver for the symmetric mixed-price equilibrium and its support width.

Baseline Parameterization

Unless otherwise noted:

- Vertical gap $\Delta = v_H - v_L = 1$, production costs $= 0$, initial belief $\eta_0 = 0.5$.

- Signal precision $q \in \{0.55,\ 0.65,\ 0.80\}$.
- Search cost $\kappa \in \{0,\ 0.1,\ 0.2\}$.
- Arrival rate $\lambda = 1$ (calendar time rescaling; probabilities are λ-invariant).
- Initial visit/attention split $\alpha_{\text{visit}} = \frac{1}{2}$.
- Prices: either fixed (p_A, p_B) (static baseline) or, when exploring Calvo, resets at hazard α with state-independent draws for simplicity.

Deterministic Curves: Cascade Boundaries

For any (p_A, p_B, κ, q):

1. Compute the buy-A threshold from (4.4):

$$\eta^*(p_A - p_B, \kappa) \;=\; \tfrac{1}{2} + \tfrac{p_A - p_B - \kappa}{2}.$$

2. Compute posteriors (placeholders $x_{\text{A}}(\eta)$, $x_{\text{B}}(\eta)$) from (4.1)–(4.2).
3. Solve for the absorbing boundaries via (4.8)–(4.9):

$$\bar{\eta} \;=\; \frac{\eta^*\, q}{(1-q) + \eta^*(2q-1)}, \qquad \underline{\eta} \;=\; \frac{\eta^*(1-q)}{q - \eta^*(2q-1)}.$$

4. Plot $\bar{\eta}$ and $\underline{\eta}$ versus $(p_A - p_B)$ for fixed (q, κ) to visualize cascade/no-cascade regions.

MC: Absorption Probabilities and Times (Fixed Prices)

We estimate (i) wrong-cascade probability and (ii) expected time to absorption under static (p_A, p_B).

Inputs. $q, \kappa, \lambda, \alpha_{\text{visit}}, \eta_0$; prices (p_A, p_B); number of runs R (e.g., $R = 5 \times 10^4$); true state prior $\mathbb{P}(\text{A high}) = \mathbb{P}(\text{B high}) = \frac{1}{2}$.

Helper functions.

- *Posteriors:* $x_{\text{A}}(\eta), x_{\text{B}}(\eta)$ from (4.1)–(4.2).
- *Private utilities at posterior x:*

$$\text{EV}_A(x) = v_L + x\Delta, \quad \text{EV}_B(x) = v_L + (1-x)\Delta.$$

- *Choice if first visit is A:* buy A if $\text{EV}_A(x) - p_A \geq \text{EV}_B(x) - p_B - \kappa$; else buy B (pay κ).
- *Choice if first visit is B:* buy B if $\text{EV}_B(x) - p_B \geq \text{EV}_A(x) - p_A - \kappa$; else buy A (pay κ).

- *Likelihoods for belief update:* Given current η, action set $\{\mathrm{A}, \mathrm{B}\}$, and visiting firm $v \in \{\mathrm{A}, \mathrm{B}\}$ with prob. α_{visit}, compute

$$\begin{aligned}\Pr(\text{action} = \mathrm{A} \mid \mathrm{A}\text{ high}, \eta) = \alpha_{\text{visit}} \cdot \Pr(\mathrm{A}\text{ chosen} \mid v = \mathrm{A}, \mathrm{A}\text{ high}, \eta)\\ + (1 - \alpha_{\text{visit}}) \cdot \Pr(\mathrm{A}\text{ chosen} \mid v = \mathrm{B}, \mathrm{A}\text{ high}, \eta),\end{aligned}$$

  and analogously for $\Pr(\text{action} = \mathrm{A} \mid \mathrm{B}\text{ high}, \eta)$. Each term is computed by enumerating signals $s \in \{\mathrm{A}, \mathrm{B}\}$ with $\Pr(s = \mathrm{A} \mid \mathrm{A}\text{ high}) = q$, $\Pr(s = \mathrm{A} \mid \mathrm{B}\text{ high}) = 1 - q$, mapping s to $x = x_{\mathrm{A}}(\eta)$ or $x_{\mathrm{B}}(\eta)$ and applying the above choice rule. Then Bayes' rule updates the public belief:

$$\eta' = \frac{\eta \cdot \Pr(\text{action} = \mathrm{A} \mid \mathrm{A}\text{ high}, \eta)}{\eta \cdot \Pr(\text{action} = \mathrm{A} \mid \mathrm{A}\text{ high}, \eta) + (1 - \eta) \cdot \Pr(\text{action} = \mathrm{A} \mid \mathrm{B}\text{ high}, \eta)},$$

  if the observed action is A (and similarly with A replaced by B).

Algorithm MC-Absorb.

1. Initialize counters: wrong= 0, time_sum= 0.
2. For $r = 1, \ldots, R$:
   (a) Draw true state $H \in \{\mathrm{A}\text{ high}, \mathrm{B}\text{ high}\}$ with prob. $\frac{1}{2}$ each. Set $\eta \leftarrow \eta_0$, $t \leftarrow 0$.
   (b) While $\eta \notin [\bar{\eta}, 1] \cup [0, \underline{\eta}]$:
      i. Draw interarrival $\Delta t \sim \mathrm{Exp}(\lambda)$; set $t \leftarrow t + \Delta t$.
      ii. Draw first visit $v = \mathrm{A}$ w.p. α_{visit}, else $v = \mathrm{B}$.
      iii. Draw private signal s using H (i.i.d.): if $H = \mathrm{A}$ high, $\Pr(s = \mathrm{A}) = q$, else $\Pr(s = \mathrm{A}) = 1 - q$. Set $x \leftarrow x_{\mathrm{A}}(\eta)$ if $s = \mathrm{A}$, else $x_{\mathrm{B}}(\eta)$.
      iv. Determine action $a \in \{\mathrm{A}, \mathrm{B}\}$ via the choice rules above (using x, v, κ, p_A, p_B).
      v. Update belief $\eta \leftarrow \eta'$ via Bayes using *only* observed action a, as in the likelihood formula above (do not condition on s).
   (c) Record absorption: if $H = \mathrm{A}$ high and $\eta \in [0, \underline{\eta}]$ (down-cascade), set wrong←wrong+1 (and symmetrically if $H = \mathrm{B}$ high and $\eta \in [\bar{\eta}, 1]$). Set time_sum←time_sum+t.
3. Outputs: $\widehat{P}_{\text{wrong}} = \text{wrong}/R$; $\widehat{T} = \text{time\_sum}/R$.

MC: Expected Profits for a Price Pair

To evaluate $\Pi_A(p_A, p_B)$ and $\Pi_B(p_A, p_B)$ for grid search:

1. Run MC-Absorb with *fixed* (p_A, p_B) and record the number of sales to A and to B per run (they sum to the total number of arrivals until absorption).
2. Estimate $\widehat{\Pi}_A = p_A \cdot \widehat{\mathbb{E}}[N_A]$, $\widehat{\Pi}_B = p_B \cdot \widehat{\mathbb{E}}[N_B]$, averaging over runs and over the two true states (prior $\frac{1}{2}$ each).

Simulation-Based Solver for Symmetric Mixed Pricing

A practical grid-based best-reply iteration (replicator-style) to approximate the symmetric mixed equilibrium and its *support width* W:

Algorithm MIX-SOLVE.

1. Fix a price grid $\mathcal{P} = \{0, \Delta p, 2\Delta p, \ldots, \bar{p}\}$ (e.g., $\bar{p} = 1$, $\Delta p = 0.01$). Initialize opponent mix $\sigma^{(0)}$ uniform on $\mathcal{P}$.

2. For iterations $k = 0, 1, 2, \ldots$ until convergence:

   (a) For each $p \in \mathcal{P}$, estimate own profit $\widehat{\Pi}(p \mid \sigma^{(k)})$ by drawing opponent prices p' from $\sigma^{(k)}$ and running MC-Absorb for each pair (p, p'); average across draws.

   (b) Construct a smoothed best reply $\tilde{\sigma}^{(k+1)}$ by concentrating mass on near-maximizers (e.g., softmax with temperature $\tau > 0$):

   $$\tilde{\sigma}^{(k+1)}(p) \propto \exp\bigl(\widehat{\Pi}(p \mid \sigma^{(k)})/\tau\bigr).$$

   (c) Update via relaxation $\sigma^{(k+1)} = (1 - \rho)\,\sigma^{(k)} + \rho\,\tilde{\sigma}^{(k+1)}$ (e.g., $\rho = 0.2$).

3. Convergence diagnostics: $\|\sigma^{(k+1)} - \sigma^{(k)}\|_1 < \varepsilon$ and small cross-play profit deviations.

4. Report support $\mathrm{supp}(\sigma)$ and width $W = \max\{p : \sigma(p) > 0\} - \min\{p : \sigma(p) > 0\}$.

Notes.

- Accuracy: Increase R (runs) and grid resolution Δp until W stabilizes; typical settings are $R \in [10^4, 10^5]$, $\Delta p \in [0.005, 0.01]$, $\tau \in [0.01, 0.05]$.

- Speed: Cache MC outcomes on a sparse subset of (p, p') pairs and interpolate profits linearly in p to reduce calls.

- Symmetry: Start from $\eta_0 = 0.5$ and use equal prior on true states to match the model's ex ante symmetry.

References

Arieli, I., M. Koren, and R. Smorodinsky (2022): "The Implications of Pricing on Social Learning," *Theoretical Economics*, 17, 1761–1802.

Armstrong, M., J. Vickers, and J. Zhou (2009): "Prominence and Consumer Search," *RAND Journal of Economics*, 40, 209–233.

Armstrong, M. and J. Zhou (2011): "Paying for Prominence," *Economic Journal*, 121, F368–F395.

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

Bikhchandani, S., D. Hirshleifer, O. Tamuz, and I. Welch (2024): "Information Cascades and Social Learning," *Journal of Economic Literature*, 62, 1040–1093.

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

Bohren, J. A. (2016): "Informational Herding with Model Misspecification," *Journal of Economic Theory*, 163, 222–247.

Burdett, K. and K. L. Judd (1983): "Equilibrium Price Dispersion," *Econometrica*, 51, 955–970.

Calvo, G. A. (1983): "Staggered Prices in a Utility–Maximizing Framework," *Journal of Monetary Economics*, 12, 383–398.

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

Diamond, P. A. (1971): "A Model of Price Adjustment," *Journal of Economic Theory*, 3, 156–168.

Frick, M., R. Iijima, and Y. Ishii (2020): "Misinterpreting Others and the Fragility of Social Learning," *Econometrica*, 88, 2281–2328.

Garcia, D. and S. Shelegia (2018): "Consumer Search with Observational Learning," *RAND Journal of Economics*, 49, 224–253.

Lukyanov, G. (2025): "Social Learning from Experts with Uncertain Precision," Working paper.

Lukyanov, G. and V. Ivanik (2025): "Contrarian Motives in Social Learning," Working paper.

Lukyanov, G., K. Popov, and S. Lashkery (2025): "Self-Employment as a Signal: Career Concerns with Hidden Firm Performance," Working paper.

Maskin, E. and J. Tirole (1988a): "A Theory of Dynamic Oligopoly, I: Overview and Quantity Competition with Large Fixed Costs," *Econometrica*, 56, 549–569.

——— (1988b): "A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves, and Edgeworth Cycles," *Econometrica*, 56, 571–599.

Papanastasiou, Y. and N. Savva (2017): "Dynamic Pricing in the Presence of Social Learning and Strategic Consumers," *Management Science*, 63, 919–939.

Sayedi, A. (2018): "Pricing in a Duopoly with Observational Learning," SSRN Working Paper.

Smith, L. and P. Sørensen (2000): "Pathological Outcomes of Observational Learning," *Econometrica*, 68, 371–398.

Stahl, D. O. (1989): "Oligopolistic Pricing with Sequential Consumer Search," *American Economic Review*, 79, 700–712.

Varian, H. R. (1980): "A Model of Sales," *American Economic Review*, 70, 651–659.