

Pricing to Naïve Learners: Belief Manipulation in Dynamic Monopoly*

Efraín Boza[†]

Zachary Breig[‡]

Antonio Rosato[§]

October 27, 2026

Abstract

A long-lived monopolist sells a good of unknown quality to short-lived buyers who observe private signals and all past prices and purchases but neglect that earlier buyers also learned from their predecessors. Once buyers are more optimistic than the prior, the seller can post a price that every buyer accepts: the purchase reveals nothing about quality, yet later buyers read it as favorable private information. Repeating such sales drives buyers' belief toward certainty that the good is valuable. A patient seller who starts from such a belief earns almost the per-period profit from a good known to be valuable, whatever its actual quality, at a discounted cost that grows only logarithmically in her effective horizon. At the prior, manipulation cannot start on its own, and failed attempts can leave buyers permanently pessimistic. With binary signals, a worthless good may eventually be bought by every buyer.

JEL Classification: D82; D83; L12; L15.

Keywords: Dynamic Pricing; Social Learning; Naïve Inference; Information Cascades.

*We acknowledge financial support from the Australian Research Council through Discovery Project DP250100162. Refine.ink and generative-AI tools provided editorial and research assistance; all errors are ours.

[†]University of Queensland (email: e.boza@uq.edu.au).

[‡]University of Queensland (email: z.breig@uq.edu.au).

[§]Deakin University, Università di Napoli Federico II and CSEF (email: antonio.rosato@deakin.edu.au).

1 Introduction

We often judge a product’s quality by how many other people have bought it. Retailers publish bestseller lists and sales ranks, and restaurants point diners to their most popular dishes. Such information affects consumer demand (Cai et al., 2009; Sorensen, 2007), especially for unfamiliar products, such as a debut novel or a newly released film, whose buyers have little experience on which to base their own assessments.¹ For a seller, an early purchase can therefore be valuable beyond the revenue it brings: it may raise what future consumers are willing to pay. Some sellers even go as far as posting fake reviews in order to manufacture favorable social information (Mayzlin et al., 2014; Luca and Zervas, 2016).

Yet, a growing number of purchases need not represent a growing body of independent evidence about quality. Some consumers may have favorable private information, whereas others may simply be responding to earlier buyers’ choices. Indeed, learning from others’ purchases requires accounting for both the prices they paid and the information they already had. A buyer who knows that a product was discounted may still overestimate what its sales reveal if he neglects that earlier customers were themselves following others. Treating each purchase as a separate endorsement is simpler, but counts the same underlying information repeatedly. Experimental evidence suggests that people often suffer from this form of *inferential naïveté*. In Eyster et al. (2018), participants combine their private signals with predecessors’ estimates. When correct inference requires subtracting information repeated across those estimates, participants frequently fail to make the adjustment, and later estimates become excessively extreme. Kübler and Weizsäcker (2005) find that subjects become more willing to follow a cascade as it grows, even though additional decisions within the cascade should reveal no new information. Related evidence concerns the neglect of correlations between information sources and naïve updating in networks (Enke and Zimmermann, 2019; Chandrasekhar et al., 2020).² For a seller, this mistake creates an opportunity: prices influence not only current demand, but also the purchase history from which future buyers learn. Buyers may attribute an observed purchase to favorable private information when it actually reflects optimism that results from even earlier purchases. Thus, we ask whether a seller can exploit this mistake to sustain demand and, if so, how she should set prices.

In this paper, we study optimal pricing by a monopolist facing successive buyers who make this mistake. Her prices determine both who buys today and what later buyers infer from today’s decisions. We show that the seller can manufacture favorable social information through genuine purchases at fully observed prices. Once buyers have become sufficiently optimistic, the seller can guarantee further purchases regardless of buyers’ private information; these purchases reveal nothing about quality, yet later buyers read them as additional favorable private evidence. This *belief pumping* gives the seller a reason to sacrifice current revenue in order to charge higher prices later. We provide conditions under which such pumping is optimal and show that, once initial optimism is established, a patient seller can earn almost as much per period as if her good were known to have the highest possible value—even when it is worthless. Pumping cannot create the initial optimism, however. The seller must first obtain a genuinely informative purchase, and

¹Sorensen (2007) finds larger bestseller-list effects for debut authors, and Moretti (2011) finds stronger social-learning effects for films about which consumers have less precise priors.

²The evidence however is not unanimous. Pooling data from cascade experiments, Weizsäcker (2010) finds that subjects rely too little on others’ choices. In a sequential belief-elicitation experiment, Angrisani et al. (2021) find that overconfidence in one’s own signal fits behavior better than redundancy neglect.

failure to do so can leave buyers permanently pessimistic, even about a valuable good.

These findings have implications for welfare and policy. Purchases of a worthless good transfer money from buyers to the seller and, when production is costly, use resources without creating consumption value. Naïve learning can also produce the opposite inefficiency. After unsuccessful early sales, buyers can remain pessimistic about a good that is valuable, so that some of them permanently forgo it even when it costs nothing to produce. As for policy, the manipulation we study relies only on genuine purchases at prices that every buyer observes. Rules that target fabricated evidence, such as bans on fake reviews, do not reach it. Similarly, requiring sellers to disclose past prices would not help buyers who already see them: the harm arises from how buyers interpret a sales record they observe in full.

We present the model in Section 2, where we introduce inferentially naïve buyers into the dynamic monopoly framework of [Bose et al. \(2006, 2008\)](#). A long-lived seller offers a good that is either valuable or worthless to a sequence of short-lived buyers, one per period; marginal cost is constant and below the good's value in the high state. Each buyer privately observes a signal drawn from a finite set; conditional on the state, signals are independent across buyers. Moreover, signals have full support and satisfy the monotone likelihood ratio property. Each price therefore divides buyers into those whose signals are favorable enough to buy and the rest. Past prices and purchase decisions, but not consumption outcomes, are observable. Following [Eyster and Rabin \(2010\)](#), we assume buyers engage in *best-response trailing naïve inference* (BRTNI): each buyer interprets every predecessor's decision as if it reflected only that predecessor's private signal and the initial prior. A purchase is thus read as evidence that the predecessor's signal was good enough for a buyer holding only the prior to buy at that price, even if he in fact bought because earlier purchases had made him optimistic. Buyers treat a decision as informative only at prices that such a buyer would accept after some signals but not others; in the baseline, decisions at other prices, including purchases at prices above this informative range, leave buyers' belief unchanged. The seller has no private information and updates correctly from actual demand, so her posterior and buyers' public belief can differ even though both are based on the same history. She chooses prices to maximize expected discounted profit, anticipating how buyers purchase and how they read each decision. We show that these two beliefs are the only payoff-relevant state: the seller has a stationary optimal policy that depends only on them, and at each state an optimal price can be found among finitely many candidates.

We begin the analysis in Section 3 by studying how prices affect buyers' beliefs. Pumping works because of the price at which the guaranteed sale occurs: once consumers are more optimistic than the initial prior, the seller can charge a price that the current buyer accepts even with the least favorable signal, but that a buyer relying only on the prior would accept only after a sufficiently favorable one. Later consumers read the purchase as if the buyer had only the prior and his own signal to go on, and so infer favorable information he may never have had. Repeating such sales drives buyers' belief toward certainty that the good is valuable. Had the seller instead kept charging a price that only favorably informed buyers accept, an unfavorable signal would have produced a rejection and lowered buyers' belief; the price cut removes this risk, and the guaranteed purchase is still interpreted as favorable evidence.

The same section also establishes the limits of this strategy and the cheapest way to raise beliefs. At the initial prior, a price that every buyer accepts is also one at which later buyers expect everyone to buy, so a guaranteed purchase cannot generate optimism. Repeated rejections can instead leave buyers so pessimistic that no price can raise their belief, whereas sufficiently

optimistic beliefs can never fall. Per unit of favorable inference, the seller gives up the least revenue by selling at the highest price that buyers still regard as informative, since this sale yields both the most revenue and the strongest inference available from a belief-raising sale.

Section 4 studies the case of a patient seller. Whenever buyers' belief starts above the prior, the seller's per-period profit approaches the profit from selling to every buyer at the full value of a good known to be valuable, irrespective of the good's actual quality; we call this outcome *full exploitation*. With rational buyers, by contrast, expected profit cannot exceed this maximum multiplied by the actual probability that the good is valuable.

Reaching full exploitation takes time and price concessions. As the seller becomes more patient, the profit she forgoes relative to always charging full value grows only logarithmically with her effective horizon: once pumping reaches the highest informative price, each further pump shrinks the gap between full value and the price that every buyer accepts by roughly the same proportion, so a few more pumping periods let her serve many more future buyers at nearly full value. We determine the exact coefficient of this logarithmic cost and show that no pricing policy can improve on it. Holding the prior fixed, the coefficient depends only on the strongest private signal and falls when that signal becomes stronger.

At the common prior, by contrast, profit per period stays bounded away from the maximum however patient the seller is, because her attempts to create optimism may fail. Below the prior, she can learn quality through purchases that leave buyers' belief unchanged, at a negligible cost per period, so that in the limit she earns as much in expectation as if she had known quality from the outset. We characterize her limiting profit at the prior and construct policies that attain it, computing their probability of making buyers optimistic enough for pumping.

Section 5 turns from what a patient seller can achieve to how a seller with a given degree of patience sets prices, assuming symmetric binary signals and zero marginal cost. When buyers are sufficiently optimistic, we give a condition under which the seller pumps until their belief reaches a threshold which rises with her patience, and then charges a constant high price forever. When buyers' beliefs are stuck at a pessimistic level, the seller can still learn about quality by charging a price that only favorably informed buyers accept, and we characterize when she sacrifices current profit to do so.

Under further sufficient conditions, the same pricing rule is optimal at every reachable belief, whatever the seller's posterior: she experiments at the prior with a price that only favorably informed buyers accept, follows a rejection with a cheaper recovery experiment, and follows a purchase with a price cut that starts pumping. We compute the probabilities of the resulting long-run outcomes. A worthless good can end up being bought by every buyer forever, while a valuable good can remain permanently under-purchased, with buyers' belief stuck even as the seller eventually learns its quality. Greater patience can change how long the seller pumps and the eventual price, but not these probabilities.

Section 6 considers two extensions. First, with finite patience and positive marginal cost, we show that the seller may incur losses on early sales that create or amplify optimism, and recover them through later high-price sales. Second, we show that if buyers instead read a purchase at a price above the informative range as evidence of the strongest private signal, the seller can raise prices while continuing to raise beliefs, and her total discounted shortfall from maximal profit remains bounded as she becomes patient. Full exploitation thus survives this alternative updating rule, although its cost changes.

Section 7 concludes by discussing some limitations of our analysis as well as potential direc-

tions for future research. All proofs are in Appendix A. The remainder of this section discusses the related literature.

This paper contributes to the literature on pricing with social learning. [Welch \(1992\)](#) shows that a seller may set prices to trigger purchase cascades among rational buyers, and [Caminal and Vives \(1996, 1999\)](#) study firms and consumers learning from market outcomes. The natural rational benchmark for our model is provided by [Bose et al. \(2006, 2008\)](#), who characterize monopoly pricing when prices control how much information buyers’ decisions transmit. [Arieli et al. \(2022\)](#) show that, under dynamic price competition, rational consumers can learn asymptotically even when private signals are bounded, and [Avery and Zemsky \(1998\)](#) show that, in financial markets, prices set by a competitive, uninformed market maker absorb public information, which limits herding. Other papers analyze dynamic pricing with reviews, ratings, and strategic consumers ([Crapis et al., 2017](#); [Papanastasiou and Savva, 2017](#); [Ifrach et al., 2019](#); [Shin et al., 2023](#); [Carnehl et al., 2024](#); [Chen et al., 2025](#)). In all these papers, prices govern how information about quality is produced, transmitted, or absorbed.³ In ours, the seller strategically uses prices to induce purchases that carry no information but that later buyers read as private evidence.

The paper also belongs to the literature in behavioral industrial organization, which studies how firms respond to consumers’ mistakes; see [Spiegler \(2011\)](#) and [Heidhues and Kőszegi \(2018\)](#) for surveys. Most of this literature concerns how consumers evaluate products or contracts; here, as in [Gagnon-Bartsch and Rosato \(2024\)](#), the mistake is in how they learn from one another. In [Gagnon-Bartsch and Rosato \(2024\)](#), consumers infer quality from earlier purchases and prices but project their own tastes onto others, so their inference about quality rises with the price at which others bought; a monopolist then charges a high price early and lowers it afterwards, which leads to over-adoption. Projection distorts how buyers read purchases that are informative about quality. With BRTNI buyers, our seller can also raise beliefs through purchases that are not informative at all, and once buyers are optimistic, her price rises over a run of such purchases. [Qiu and Whinston \(2017\)](#) find that introductory discounts can prevent behavioral observational learning that would raise later consumers’ willingness to pay; in our model, guaranteed purchases cannot create optimism at the prior, but they amplify it once it exists.

Finally, the paper relates to work on information cascades ([Banerjee, 1992](#); [Bikhchandani et al., 1992](#); [Smith and Sørensen, 2000](#)) and on naïve or misspecified social learning. [Eyster and Rabin \(2010\)](#) introduce BRTNI, building on the cursed equilibrium of [Eyster and Rabin \(2005\)](#). Cursedness has been applied to financial markets by [Eyster et al. \(2019\)](#), and to bilateral trade by [Murooka and Yamashita \(2025\)](#). [Rosato \(2026\)](#) applies BRTNI to English auctions, where bidders decode observed exits much as buyers here decode purchases, and in [Bastianello and Fontanier \(2025\)](#) investors learn from prices while neglecting that prices already embed others’ inferences, which closely parallels BRTNI in a market-clearing setting. [Gagnon-Bartsch and Rabin \(2026\)](#) show that the redundancy neglect underlying BRTNI can lead a society that has learned the truth to abandon it later, because later observers read behavior that reflects inherited beliefs as if it reflected private information alone. In our model, the same neglect pushes buyers’ belief away from the prior in expectation, and belief pumping takes this force to its extreme. [Eyster and Rabin \(2014\)](#) show more generally that learning rules in which people imitate several predecessors without anti-imitating

³Other work gives the seller additional instruments besides the price. In [Chen and Papanastasiou \(2021\)](#), a seller who knows more about quality than consumers can seed observational learning with a fake purchase; in [Mayzlin \(2006\)](#), firms post anonymous promotional messages in online communities. Our seller has no instrument other than a posted price that every buyer observes.

any, as BRTNI buyers do, can lead to confident but wrong beliefs. Related work studies naïve updating in networks (DeMarzo et al., 2003; Golub and Jackson, 2010; Dasaratha and He, 2020; Mueller-Frank and Neri, 2021), coarse inference (Guarino and Jehiel, 2013), and learning under misspecified models (Bohren, 2016; Frick et al., 2020; Bohren and Hauser, 2021).

2 Model

We build on the dynamic monopoly-pricing models of Bose et al. (2006, 2008), but assume that buyers are inferentially naïve. A rational buyer interprets a predecessor’s purchase decision taking into account the price he faced and what he learned from earlier buyers. A BRTNI buyer also conditions on the price, but treats the predecessor as having used only his private signal and the initial prior. As in their models, buyers observe every past price and purchase decision; thus, the seller cannot conceal the terms on which earlier purchases were made, and the manipulation we study does not rely on doing so.⁴

2.1 Environment

A long-lived monopolist sells a good to a sequence of short-lived, unit-demand buyers, one arriving in each period $t = 0, 1, 2, \dots$. The seller and buyers are risk neutral. The good has common value $v \in \{0, 1\}$ and constant marginal cost $c \in [0, 1)$. The seller and buyers have the common prior $\mathbb{P}(v = 1) = \mu_0 \in (0, 1)$. The seller has no private information about v and discounts future profits by $\delta \in [0, 1)$.

Each buyer receives a private signal from the finite set $\mathcal{S} := \{1, \dots, M\}$, where $M \geq 2$. Conditional on v , signals are independent across buyers, with $\mathbb{P}(s_t = s \mid v) = f_v(s)$. We assume *full support*, $f_v(s) > 0$ for every $v \in \{0, 1\}$ and $s \in \mathcal{S}$, and *strict MLRP*: the likelihood ratio $\text{LR}(s) := f_1(s)/f_0(s)$ is strictly increasing in s .⁵ Because $\mathcal{S}$ is finite, full support implies that likelihood ratios are bounded above and bounded away from zero.

At the start of period t , the seller and buyer t observe the public history $h_t := (p_j, a_j)_{j < t}$, with $h_0 := \emptyset$. Let $\mathcal{F}_t$ denote the sigma-field it generates. The seller posts a price $p_t \in [0, 1]$. Buyer t observes the price, privately receives s_t , and chooses $a_t \in \mathcal{A} := \{0, 1\}$, where $a_t = 1$ denotes purchase. His payoff is $(v - p_t)a_t$, and the seller’s current profit is $(p_t - c)a_t$. Buyers purchase when indifferent. The pair (p_t, a_t) is appended to the public history; neither the private signal nor the buyer’s realized consumption payoff is observed by the seller or later buyers.

2.2 Beliefs, inference, and BRTNI decoding

We refer to the belief buyers hold before observing their signal as the *public* belief, while their belief after observing their signal is their *posterior* belief. Let $\mathcal{L}(x) := \log[x/(1 - x)]$ and $\Lambda(y) := e^y/(1 + e^y)$.⁶ For a private signal s , write $\lambda_s := \log[f_1(s)/f_0(s)]$. A buyer with public

⁴Bose et al. (2008, footnote 5) note that if buyers observed purchases but not prices, the seller would want them to believe that a sale had occurred at a high price, and that rational buyers cannot be misled in this way in equilibrium.

⁵With two states, signal realizations can be ordered by their likelihood ratios. Realizations with the same ratio induce the same posterior and can be merged.

⁶We use the boundary extensions $\mathcal{L}(0) := -\infty$, $\mathcal{L}(1) := +\infty$, $\Lambda(-\infty) := 0$, and $\Lambda(+\infty) := 1$.

belief x and signal s has posterior belief and willingness to pay

$$\beta_s(x) := \mathbb{E}[v \mid x, s] = \Lambda(\mathcal{L}(x) + \lambda_s).$$

By strict MLRP, $\beta_1(x) < \dots < \beta_M(x)$ for $x \in (0, 1)$. If exactly the signal types above s buy, a sale reveals $s_t > s$ and a rejection reveals $s_t \leq s$. For $s = 0, \dots, M$, define $F_v(s) := \sum_{j=1}^s f_v(j)$ and $\bar{F}_v(s) := 1 - F_v(s)$, with the empty sum equal to zero. For any belief $x \in [0, 1]$, write

$$F_x(s) := xF_1(s) + (1-x)F_0(s), \quad \bar{F}_x(s) := 1 - F_x(s).$$

For $s = 1, \dots, M-1$, let

$$\ell_s := \log \frac{\bar{F}_1(s)}{\bar{F}_0(s)}, \quad m_s := \log \frac{F_0(s)}{F_1(s)}. \quad (1)$$

The corresponding log-odds increments are ℓ_s after a sale and $-m_s$ after a rejection. The next lemma collects the properties of these objects.

Lemma 1. *Under full support and strict MLRP:*

- (i) $\lambda_1 < 0 < \lambda_M$.
- (ii) $0 < \ell_s \leq \lambda_M$ and $0 < m_s \leq -\lambda_1$ for every $s \in \{1, \dots, M-1\}$.
- (iii) $(\ell_s)_{s=1}^{M-1}$ is strictly increasing with $\ell_{M-1} = \lambda_M$, whereas $(m_s)_{s=1}^{M-1}$ is strictly decreasing with $m_1 = -\lambda_1$.

Lemma 1 orders the informational content of the events that a price can generate. Raising the cutoff makes a sale stronger evidence for $v = 1$ and a rejection weaker evidence against it. The endpoint equalities follow because a sale at cutoff $M-1$ reveals signal M , while a rejection at cutoff 1 reveals signal 1.

From here on, we keep track of two beliefs. Let $r_t := \mathbb{P}(v = 1 \mid \mathcal{F}_t)$ denote the seller's Bayesian posterior. Following [Eyster and Rabin \(2010\)](#), a BRTNI buyer interprets each predecessor's decision as if that predecessor had used only the initial prior μ_0 and his private signal, neglecting his inference from earlier buyers' actions. Let n_t denote the belief buyer t forms from h_t before observing his own signal. Initially, $r_0 = n_0 = \mu_0$. Since the first buyer has no predecessors, his action is decoded correctly and $r_1 = n_1$. From $t = 2$ onward, the seller's posterior and buyers' public belief can differ, despite being based on the same public history. To see how the two beliefs can diverge, suppose buyer 0 receives a favorable signal and buys. Buyer 1, encouraged by that purchase, then buys at a price that he would accept even after an unfavorable signal, but that a buyer holding only the prior would accept only after a favorable one. A rational buyer 2 understands that buyer 1's purchase was guaranteed and learns nothing from it. A BRTNI buyer 2 reads it as a second favorable signal, and so counts buyer 0's information twice: once through buyer 0's own purchase and once through the purchase it induced. This is the redundancy neglect found in several of the experiments discussed in the Introduction, and it is what the seller will exploit.

The purchasing rule attributed to a predecessor with signal s is $\bar{a}(p, s) := \mathbf{1}\{p \leq \beta_s(\mu_0)\}$. This rule induces the *decoding bands*

$$I_s := (\beta_s(\mu_0), \beta_{s+1}(\mu_0)], \quad s = 1, \dots, M-1. \quad (2)$$

Let $I := \bigcup_{s=1}^{M-1} I_s = (\underline{p}, \bar{p}]$, where $\underline{p} := \beta_1(\mu_0)$ and $\bar{p} := \beta_M(\mu_0)$. At a price $p \in I_s$, a BRTNI buyer interprets a sale as the event $\{\text{signal} > s\}$ and a rejection as $\{\text{signal} \leq s\}$. Prices in I are *in-band*; all other prices are *off-band*. Write $n_{t+1} = N(n_t, p_t, a_t)$ for the recursive update of the naïve public belief. At a price $p \in I_s$, a sale adds ℓ_s to $\mathcal{L}(n)$ and a rejection subtracts m_s . Thus, μ_0 determines how the price is decoded, while n is the belief being updated.⁷ Outside I , the imputed purchasing rule predicts a sale with probability one when $p \leq \underline{p}$ and a rejection with probability one when $p > \bar{p}$. The predicted action is uninformative. The opposite action has zero probability under both values of v , so Bayes' rule does not determine its effect on beliefs. We assume that these unexpected actions also leave n unchanged.⁸

Assumption 1. For every $n \in [0, 1]$, $p \notin I$, and $a \in \{0, 1\}$, $N(n, p, a) = n$.

Thus,

$$\mathcal{L}(N(n, p, a)) = \begin{cases} \mathcal{L}(n) + \ell_s, & p \in I_s, a = 1, \\ \mathcal{L}(n) - m_s, & p \in I_s, a = 0, \\ \mathcal{L}(n), & p \notin I. \end{cases} \quad (3)$$

Under Assumption 1, only in-band observations change the naïve public belief. Section 6.2 studies an alternative completion that attributes unexpected off-band actions to extreme signals, showing that this alternative makes manipulation easier.

2.3 Demand, transitions, and the dynamic program

At public belief n , the current buyer purchases if and only if $p \leq \beta_{s_t}(n)$. For $v \in \{0, 1\}$, define the state-conditional sale probability

$$d_v(p, n) := \sum_{j=1}^M f_v(j) \mathbf{1}\{p \leq \beta_j(n)\},$$

and the seller's sale probability

$$D(p; r, n) := r d_1(p, n) + (1 - r) d_0(p, n).$$

When the observed action has positive probability, the seller updates by Bayes' rule:⁹

$$R^1(r, n, p) := \frac{r d_1(p, n)}{D(p; r, n)},$$

$$R^0(r, n, p) := \frac{r [1 - d_1(p, n)]}{1 - D(p; r, n)}.$$

A price *screens* buyers when some signal types purchase and others reject. The next lemma records actual demand, the associated Bayesian updates, and an upper bound on the price at which a sale can occur.

⁷Because the decoded evidence is additive in log odds, updating recursively from $n_0 = \mu_0$ gives the same belief as inference from the full public history. Under the passive off-band rule introduced in Assumption 1, only the counts of sales and rejections in each band matter, not their order.

⁸This completion is obtained from signal-independent trembles: if every signal type takes the unexpected action with the same probability $\eta > 0$, that action has likelihood ratio one.

⁹On a zero-probability branch, set the corresponding posterior equal to r ; this convention does not affect expected payoffs.

Lemma 2. Fix (r, n) and p .

- (i) If $p \leq \beta_1(n)$, a sale occurs surely and r is unchanged.
- (ii) If $\beta_s(n) < p \leq \beta_{s+1}(n)$ for some $s \in \{1, \dots, M-1\}$, a sale occurs with probability $\bar{F}_r(s)$. The seller's posterior becomes $\Lambda(\mathcal{L}(r) + \ell_s)$ after a sale and $\Lambda(\mathcal{L}(r) - m_s)$ after a rejection.
- (iii) If $p > \beta_M(n)$, no sale occurs and r is unchanged.

The naïve transition is given by (3), and a sale requires

$$a_t = 1 \implies p_t \leq \beta_M(n_t). \quad (4)$$

The seller updates from actual demand, which depends on buyers' current public belief n , whereas later buyers decode the same action with bands fixed at the prior. The two beliefs can therefore respond differently to the same event. A purchase that every current buyer would have made leaves r unchanged, but it raises n whenever its price lies in a decoding band. This wedge is what the seller exploits in the rest of the paper.

Let $\mathcal{H}$ denote the set of public histories. A pure pricing policy is a measurable map $\pi : \mathcal{H} \rightarrow [0, 1]$; we also allow measurable randomization over prices after each history. Starting from (r, n) , the seller's expected discounted profit under π and her value are

$$V_\delta^\pi(r, n) := \mathbb{E}_{r,n}^\pi \left[\sum_{t \geq 0} \delta^t (p_t - c) a_t \right], \quad V_\delta(r, n) := \sup_\pi V_\delta^\pi(r, n).$$

We define these continuation values for all $(r, n) \in [0, 1]^2$, keeping the decoding bands fixed at the initial prior μ_0 .

Expected current profit and the conditional distribution of the next state depend on the history only through (r, n) and the chosen price. For a continuation value W , define the Bellman operator

$$\begin{aligned} (TW)(r, n) := \max_{p \in [0, 1]} & \left\{ (p - c)D(p; r, n) \right. \\ & + \delta D(p; r, n)W(R^1(r, n, p), N(n, p, 1)) \\ & \left. + \delta [1 - D(p; r, n)]W(R^0(r, n, p), N(n, p, 0)) \right\}. \end{aligned} \quad (5)$$

Between successive demand or decoding cutoffs, raising the price changes neither demand nor either belief transition, so it weakly increases current profit. The next proposition uses this observation to reduce the price choice to a finite menu and establishes the existence of an optimal stationary policy.

Proposition 1. For each $n \in [0, 1]$, define $\mathcal{T}(n) := \{0, 1\} \cup \{\beta_s(\mu_0) : s \in \mathcal{S}\} \cup \{\beta_s(n) : s \in \mathcal{S}\}$. For every r and bounded continuation value W , the maximum in (5) is attained at a price in $\mathcal{T}(n)$. Moreover, the pair (r, n) is a sufficient payoff-relevant state, and the operator T maps bounded Borel functions into themselves and is a contraction with modulus δ in the sup norm. Its unique bounded fixed point is V_δ , and a Borel-measurable deterministic stationary optimal policy exists. Finally,

$$0 \leq V_\delta(r, n) \leq \frac{1 - c}{1 - \delta}. \quad (6)$$

The reduction preserves the optimization over all history-dependent policies: an optimal price can be chosen using only (r, n) . The seller can compute both beliefs from the public history without observing the current buyer's private signal.

Together with the purchasing rule $a_t = \mathbf{1}\{p_t \leq \beta_{s_t}(n_t)\}$, an optimal stationary policy defines a *naïve-inference equilibrium*. Buyers maximize expected utility under their BRTNI beliefs, and the seller's continuation policy is optimal after every public history given their actual purchasing and updating rules.

3 Belief updating and manipulation

BRTNI beliefs lead to a wedge between actual demand, governed by the current naïve belief n , and its subsequent decoding, governed by the imputed belief μ_0 . The seller can exploit this wedge period by period. This section develops this manipulation technology without making any further assumptions about discount factors or information structures.

3.1 Actual versus perceived information

For an in-band price $p \in I_s$, write the naïve beliefs following a sale and a rejection as

$$n_s^+(n) := \Lambda(\mathcal{L}(n) + \ell_s), \quad n_s^-(n) := \Lambda(\mathcal{L}(n) - m_s). \quad (7)$$

The next proposition relates the expected change in the naïve public belief to the difference between the actual and perceived probabilities of a sale.

Proposition 2. *Suppose $n \in (0, 1)$, $p \in I_s$. Let $k \in \{0, \dots, M\}$ be the actual demand cutoff induced by (p, n) , so that exactly the signal types $j > k$ purchase. Then*

$$\mathbb{E}[n_{t+1} \mid r_t = r, n_t = n, p_t = p] - n = [\bar{F}_r(k) - \bar{F}_n(s)] [n_s^+(n) - n_s^-(n)]. \quad (8)$$

Consequently:

- (i) *if $k = s$, the action is truthfully decoded and the sign of the drift is the sign of $r - n$;*
- (ii) *if $k = 0$, the price is actually pooling and the naïve belief rises deterministically;*
- (iii) *if $k = M$, no sale occurs and the naïve belief falls deterministically.*

Equation (8) compares the actual sale probability $\bar{F}_r(k)$ with $\bar{F}_n(s)$, the probability implied by buyers' beliefs and the imputed cutoff. Since $n_s^+(n) > n_s^-(n)$, the sign of this difference determines the direction of expected belief change. Unlike the seller's posterior, the naïve public belief need not be a martingale.

When $k = s$,

$$\bar{F}_r(s) - \bar{F}_n(s) = (r - n) [\bar{F}_1(s) - \bar{F}_0(s)],$$

so correctly decoded screening raises n on average when $r > n$ and lowers it when $r < n$. If earlier sales have produced $n > r$, such screening erodes optimism. A sure sale avoids this correction because it reveals no new information about v . Conversely, an in-band price rejected by every type lowers n with certainty while leaving r unchanged.

More generally, the difference between actual and perceived sale probabilities in (8) can be written as

$$\bar{F}_r(k) - \bar{F}_n(s) = [\bar{F}_r(k) - \bar{F}_r(s)] + (r - n)[\bar{F}_1(s) - \bar{F}_0(s)].$$

The first term reflects the gap between the cutoff s that buyers impute and the actual cutoff k . Above the prior, every type is more willing to pay than the decoding assumes, so $k \leq s$ and this term is nonnegative; below the prior, $k \geq s$ and it is nonpositive. Naïve decoding therefore pushes buyers' belief away from the prior, even when it coincides with the seller's: if $r = n$, the expected change in n is nonnegative above the prior and nonpositive below it, and strictly so whenever $k \neq s$. The same neglect drives the “unlearning” studied by [Gagnon-Bartsch and Rabin \(2026\)](#), in which later observers read behavior that reflects inherited beliefs as if it reflected private information alone.

3.2 Belief pumping

Belief pumping uses an in-band price accepted by every signal type. The resulting sure sale raises buyers' belief without revealing information about v . Define $p^P(n) := \min\{\beta_1(n), \bar{p}\}$, and

$$n^{\text{cap}} := \Lambda(\mathcal{L}(\mu_0) + \lambda_M - \lambda_1). \quad (9)$$

For $n \in (\mu_0, 1)$, $p^P(n)$ is the highest in-band price accepted by every type.¹⁰ The cap binds exactly when $n \geq n^{\text{cap}}$, or equivalently $\beta_1(n) \geq \bar{p}$. The next proposition shows that, starting above the prior, the seller can drive the naïve belief arbitrarily close to one through repeated sure sales, while her own posterior remains unchanged.

Proposition 3. *Fix any state (r, n) with $n \in (\mu_0, 1)$, and suppose the seller posts $p^P(n)$. Then,*

- (i) $p^P(n) \in I$; let $s(n)$ be the unique index such that $p^P(n) \in I_{s(n)}$;
- (ii) every signal type buys, so a sale occurs surely and the seller's posterior remains r ;
- (iii) the naïve belief becomes $n_{s(n)}^+(n)$, so its log odds rise by $\ell_{s(n)} \geq \ell_1 > 0$;
- (iv) if $n \geq n^{\text{cap}}$, then $p^P(n) = \bar{p} \in I_{M-1}$, $s(n) = M - 1$, and $\mathcal{L}(n_{M-1}^+(n)) = \mathcal{L}(n) + \lambda_M$.

Consequently, repeated pumping drives $n_t \rightarrow 1$ deterministically while r_t remains fixed, and the number of pumps needed to reach n^{cap} from n is at most

$$K_0(n) := \max \left\{ 0, \left\lceil \frac{\mathcal{L}(n^{\text{cap}}) - \mathcal{L}(n)}{\ell_1} \right\rceil \right\}. \quad (10)$$

For illustration, let $M = 2$, $\mu_0 = 1/2$, and $f_1(2) = f_0(1) = 7/10$. The decoding band is $(3/10, 7/10]$. At $n = 7/10$, even the low-signal buyer is willing to pay $1/2$. Posting $p = 1/2$ therefore sells surely and leaves r unchanged. Later buyers, however, believe the purchaser had only the prior and his own signal to go on, and at the prior a buyer pays $1/2$ only after signal 2. They therefore read the sale as evidence of signal 2 and raise their public belief to $49/58$, at which the maximum sure-sale price rises to $7/10$.

¹⁰A sufficiently small reduction in $p^P(n)$ keeps the price in the same band and makes all types strictly prefer buying.

Once $n \geq n^{\text{cap}}$, pumping at $\bar{p}$ combines the highest in-band price with the largest log-odds increase. Charging more requires leaving the decoding region, in which case the purchase decision leaves n unchanged. Above the cap, every pump therefore sells at $\bar{p}$, below the price $\beta_1(n)$ that every buyer would accept. Section 3.3 measures this cost.

3.3 The limits and cost of belief manipulation

At sufficiently low beliefs, no in-band price produces a sale; at sufficiently high beliefs, every in-band price does. Define the lower threshold

$$\underline{n} := \Lambda(\mathcal{L}(\mu_0) + \lambda_1 - \lambda_M) \in (0, \mu_0). \quad (11)$$

The next lemma shows that beliefs cannot rise in the *frozen region* $n \leq \underline{n}$ or fall in the *ratcheted region* $n \geq n^{\text{cap}}$. These restrictions on belief updating give corresponding bounds on the seller's profit.

Lemma 3. *For every $r \in [0, 1]$, the following hold.*

- (i) *If $n_t \leq \underline{n}$, every in-band price is rejected and $n_{t'} \leq n_t$ for every $t' \geq t$ under every policy. Every future sale occurs at a price no greater than $\underline{p}$, and*

$$V_\delta(r, n_t) \leq \frac{(\underline{p} - c)^+}{1 - \delta}. \quad (12)$$

- (ii) *If $n_t \geq n^{\text{cap}}$, every in-band price sells surely and $n_{t'} \geq n_t$ for every $t' \geq t$ under every policy. Moreover,*

$$V_\delta(r, n_t) \geq \frac{(\beta_1(n_t) - c)^+}{1 - \delta} \geq \frac{(\bar{p} - c)^+}{1 - \delta}. \quad (13)$$

Both regions are absorbing.

For $c > 0$, exit is optimal whenever $n \leq \underline{n}$ and $\beta_M(n) \leq c$. Beliefs cannot rise, so every future sale has a nonpositive margin. Hence $V_\delta(r, n) = 0$ for every r , attained by posting $p = 1$ forever. In the ratcheted region, the lower bound follows by posting the fixed price $\beta_1(n)$ forever when $\beta_1(n) \geq c$, and by not selling otherwise. These policies are feasible, but need not be optimal.

The thresholds are symmetric around the prior in log odds: $\mathcal{L}(n^{\text{cap}}) - \mathcal{L}(\mu_0) = \lambda_M - \lambda_1 = \mathcal{L}(\mu_0) - \mathcal{L}(\underline{n})$. Between them lies the *contestable region* $\underline{n} < n < n^{\text{cap}}$, where some in-band prices screen buyers. At these prices, a sale raises n and a rejection lowers it. In the ratcheted region, any screening price is off band and leaves n unchanged.

At or below the prior, every sure-sale price satisfies $p \leq \beta_1(n) \leq \underline{p}$ and leaves n unchanged. Raising n therefore requires an in-band sale at a price some signal types reject. Once $\mu_0 < n < 1$, Proposition 3 instead allows beliefs to rise through sure sales alone. Starting from the prior, the seller must obtain an informative purchase before pumping can begin.

We now measure the cost of raising beliefs relative to the benchmark flow profit $1 - c$. An in-band sale at $p \in I_s$ yields profit $p - c$ and raises naïve log odds by ℓ_s , so its profit shortfall is $1 - p$. Let $\chi := (1 - \bar{p})/\lambda_M > 0$. The next lemma shows that χ is the minimum shortfall per unit increase in log odds across in-band sales.

Lemma 4. For every $s \in \{1, \dots, M-1\}$,

$$\frac{1 - \beta_{s+1}(\mu_0)}{\ell_s} \geq \chi, \quad (14)$$

with equality if and only if $s = M-1$. Consequently, every in-band price $p \in I_s$ satisfies

$$1 - p \geq \chi \ell_s. \quad (15)$$

Equality holds for a sale at $p = \bar{p}$: the shortfall is $1 - \bar{p}$ and the log-odds increment is λ_M . For $n \in [n^{\text{cap}}, 1)$, this sale is certain, so the seller can attain the bound through pumping. The seller's marginal cost does not enter χ , because it cancels from the profit difference $(1 - c) - (p - c)$. Section 4 uses this bound to characterize the seller's optimal profit as $\delta \uparrow 1$.

4 The patient limit

In this section, we study the limit as $\delta \uparrow 1$, using normalized profits $(1 - \delta)V_\delta(r, n)$. This normalization expresses profit on a per-period basis, so that earning $1 - c$ in every period has normalized value $1 - c$. Patience allows the seller to recover the cost of pumping through later sales, but does not eliminate the risk of unsuccessful initial experimentation. We first characterize the seller's profit and the cost of manipulation when buyers' beliefs exceed the prior. We then ask how much she can earn when beliefs start at or below the prior, and how likely experimentation is to make buyers optimistic enough for pumping. Finally, we examine buyers' payments and welfare when $v = 0$.

4.1 Full exploitation and the rational benchmark

We first bound profits when buyers learn rationally. Let $V_\delta^R(\mu)$ denote the seller's value with Bayesian buyers and common public posterior μ .

Lemma 5. For every $\mu \in [0, 1]$ and $\delta \in [0, 1)$,

$$V_\delta^R(\mu) \leq \frac{(1 - c)\mu}{1 - \delta}. \quad (16)$$

With Bayesian buyers, individual rationality bounds expected current profits by $(1 - c)\mu_t$, and the martingale property of μ_t gives (16). Under BRTNI, buyers' valuations instead depend on n_t . Raising n_t without changing r_t can therefore yield profits above this rational-benchmark bound.

Consider a pump-and-harvest policy. Starting from $n \in (\mu_0, 1)$, the seller pumps at $p^P(n_t)$ for $K_0(n)$ periods, then at $\bar{p}$ for another $N \geq 1$ periods. Let n^* be the belief after the last pump. The seller then charges $\beta_1(n^*)$ forever. Every buyer purchases throughout. The next proposition bounds this policy's shortfall from $(1 - c)/(1 - \delta)$. Since optimal pricing yields at least as much profit as this feasible policy, the same upper bound applies to the optimal profit shortfall.

Proposition 4. For every $n \in (\mu_0, 1)$, $r \in [0, 1]$, $\delta \in (0, 1)$, and integer $N \geq 1$,

$$\frac{1 - c}{1 - \delta} - V_\delta(r, n) \leq K_0(n) + N(1 - \bar{p}) + \frac{e^{-[\mathcal{L}(n) + \lambda_1]} e^{-N\lambda_M}}{1 - \delta}. \quad (17)$$

The first two terms bound forgone profits during pumping; the last bounds the discounted loss from subsequent sales below a price of one. Increasing N raises the pumping cost but reduces the subsequent price gap. The next proposition shows that, as $\delta \uparrow 1$, these losses become negligible in normalized terms, so the seller's profit approaches its upper bound $1 - c$.

Proposition 5. *Normalized profit converges to $1 - c$, uniformly over the seller's posterior and all naïve beliefs above the prior:*

$$\lim_{\delta \uparrow 1} \sup_{r \in [0,1], n \in (\mu_0,1)} |(1 - \delta)V_\delta(r, n) - (1 - c)| = 0. \quad (18)$$

Consequently, for every $r \in [0, 1]$ and $n \in (\mu_0, 1]$,

$$\liminf_{\delta \uparrow 1} (1 - \delta)[V_\delta(r, n) - V_\delta^R(r)] \geq (1 - c)(1 - r) > 0. \quad (19)$$

We call this outcome *full exploitation*. It is exploitation because the seller uses buyers' distorted beliefs to obtain payments exceeding their objectively expected consumption value: for each fixed $r < 1$, buyers' normalized expected surplus is negative in the patient limit. It is full because normalized profit approaches $1 - c$, the largest feasible per-period profit.¹¹

With rational buyers, normalized profit is at most $(1 - c)r$ and therefore remains tied to the probability that the good is valuable. Under BRTNI, whenever $n > \mu_0$, a patient seller's profit approaches $1 - c$ regardless of r . Even a seller certain that $v = 0$ attains the same limit as one certain that $v = 1$, whereas the rational benchmark yields zero profit at $r = 0$. The seller gains by raising buyers' willingness to pay without learning anything new about the good or increasing its value.

4.2 The price of manufactured evidence

Proposition 5 shows that normalized profit approaches $1 - c$ whenever $n > \mu_0$. The next proposition quantifies the remaining difference between discounted profits and $(1 - c)/(1 - \delta)$. For each fixed $n \in (\mu_0, 1)$, this difference grows at rate $\chi \log[1/(1 - \delta)]$, where χ is the minimum cost per unit of log odds identified in Lemma 4.

Proposition 6. *Let $\chi := (1 - \bar{p})/\lambda_M > 0$ as before.*

(i) *For every $n \in (\mu_0, 1)$, $r \in [0, 1]$, and $\delta \in (0, 1)$,*

$$\frac{1 - c}{1 - \delta} - V_\delta(r, n) \leq \chi \log \frac{1}{1 - \delta} + C(n), \quad (20)$$

where

$$C(n) := K_0(n) + (1 - \bar{p}) + e^{-[\mathcal{L}(n) + \lambda_1]}. \quad (21)$$

(ii) *For every $n \in (0, 1)$,*

$$\liminf_{\delta \uparrow 1} \inf_{r \in [0,1]} \frac{\frac{1-c}{1-\delta} - V_\delta(r, n)}{\log \frac{1}{1-\delta}} \geq \chi. \quad (22)$$

(iii) *For every fixed $n \in (\mu_0, 1)$,*

$$\lim_{\delta \uparrow 1} \sup_{r \in [0,1]} \left| \frac{\frac{1-c}{1-\delta} - V_\delta(r, n)}{\log \frac{1}{1-\delta}} - \chi \right| = 0. \quad (23)$$

¹¹The convergence is uniform even as $n \downarrow \mu_0$, but excludes $n = \mu_0$.

Part (i) uses the pump-and-harvest policy, which trades a lower earlier price during pumping for a higher price on subsequent sales, with a number of top-band pumps that grows as the seller becomes more patient. Once $n \geq n^{\text{cap}}$, each pump costs $1 - \bar{p}$ relative to the benchmark flow and multiplies buyers' odds by e^{λ_M} . The price that every buyer then accepts falls short of one by an amount roughly inversely proportional to those odds. This gap is paid in every later period, so its discounted cost is about the gap divided by $1 - \delta$, and it pays to keep pumping until the gap is of order $1 - \delta$. Because each pump shrinks the gap by approximately the factor $e^{-\lambda_M}$, this takes about

$$N = \left\lceil \frac{1}{\lambda_M} \log \frac{1}{1 - \delta} \right\rceil$$

pumps, at a total cost of about $\chi \log[1/(1 - \delta)]$. Part (ii) shows that no policy improves on this coefficient: by Lemma 4, every unit of log odds that buyers gain costs the seller at least χ , and a policy that accumulates fewer units loses more by selling below a price of one in later periods. The discounted shortfall therefore grows without bound, but only logarithmically in the effective horizon $1/(1 - \delta)$, and the normalized shortfall vanishes at rate $(1 - \delta) \log[1/(1 - \delta)]$.

For a fixed prior, χ depends on the signal structure only through λ_M , the log likelihood ratio of the most favorable signal. A larger λ_M lowers χ for two reasons: it raises the top-band price $\bar{p}$, and it raises the evidence that buyers credit to a top-band sale. Along a sequence of finite signal structures with μ_0 fixed and $\lambda_M \rightarrow \infty$,

$$\chi(\lambda_M) = \frac{1 - \Lambda(\mathcal{L}(\mu_0) + \lambda_M)}{\lambda_M} \sim \frac{e^{-\mathcal{L}(\mu_0)} e^{-\lambda_M}}{\lambda_M} \rightarrow 0.$$

The likelihood ratio that makes a genuinely selective purchase strong evidence of quality also makes a pooling purchase persuasive when buyers mistake it for that evidence. A more informative top signal therefore lowers the cost of each unit of manufactured evidence.¹²

4.3 Manipulation cannot start itself

The preceding maximal-profit result requires $n > \mu_0$. At the prior, no in-band price is accepted by every type, so raising beliefs requires a purchase that is not guaranteed. To bound the resulting risk, define

$$K_{\text{fr}} := \left\lceil \frac{\lambda_M - \lambda_1}{m_{M-1}} \right\rceil, \quad \Delta_c := 1 - c - (\underline{p} - c)^+ = 1 - \max\{c, \underline{p}\}. \quad (24)$$

Starting from the prior, K_{fr} consecutive in-band rejections suffice to reach the frozen region, regardless of the prices chosen. The constant Δ_c is the difference between $1 - c$ and the upper bound on flow profits in that region. The next proposition shows that this risk prevents normalized profit from approaching $1 - c$, however patient the seller is.

¹²Bounded likelihood ratios play two roles in our model. As in [Smith and Sørensen \(2000\)](#), where bounded private beliefs allow rational herds to settle on the wrong action, they allow buyers' belief to become permanently stuck: in the frozen region, even a buyer with the most favorable signal rejects every in-band price, so no purchase can be read as good news, and buyers' belief cannot recover even if the good is valuable (Lemma 3). They also keep manufactured evidence costly: χ is positive because λ_M is finite, and it vanishes as $\lambda_M \rightarrow \infty$.

Proposition 7. For every $r \in [0, 1]$ and $\delta \in (0, 1)$,

$$(1 - \delta)V_\delta(r, \mu_0) \leq (1 - c) - \Delta_c \left[r f_1(1)^{K_{\text{fr}}} + (1 - r) f_0(1)^{K_{\text{fr}}} \right]. \quad (25)$$

The gap from $1 - c$ is uniformly positive in r and δ .

Conditional on v , the probability that every attempted in-band price is rejected before the seller either stops experimenting or reaches the frozen region is at least $f_v(1)^{K_{\text{fr}}}$. On these paths, any sale occurs at a price no greater than $\underline{p}$. Together with Proposition 5, this implies a discontinuity in limiting normalized profit at $n = \mu_0$. An arbitrarily small amount of initial optimism permits sure-sale pumping; at the prior, obtaining the first favorable purchase entails a risk that patience cannot eliminate.

4.4 Experimentation at and below the prior

Strictly below the prior, a price at or below the lower edge of the decoding band still screens buyers without being decoded, so the seller can learn about v while leaving n unchanged; Proposition 8 shows that, for a patient seller, this learning is essentially free. The natural benchmark is therefore a seller who knows v . Suppose that $r = v$, while buyers' purchasing and inference rules remain unchanged. At $n \in (0, \mu_0]$, her best stationary off-band expected profit per period is

$$h_v(n) := \max \left\{ 0, \max_{p \in [0, \underline{p}]} (p - c) d_v(p, n) \right\}. \quad (26)$$

The zero term allows the seller to stop selling. Alternatively, she can choose an in-band price. By Proposition 1, it suffices to consider $A(n) := \mathcal{T}(n) \cap I$.

For $W : (0, \mu_0] \rightarrow [0, 1 - c]$, let $\widehat{W}(x) = W(x)$ for $0 < x \leq \mu_0$ and $\widehat{W}(x) = 1 - c$ for $\mu_0 < x < 1$. Define

$$(Q_v W)(n, p) := d_v(p, n) \widehat{W}(N(n, p, 1)) + [1 - d_v(p, n)] \widehat{W}(N(n, p, 0)). \quad (27)$$

The continuation value above the prior is $1 - c$ by Proposition 5. The next lemma characterizes the known-state limiting value as a choice between stopping and further in-band experimentation.

Lemma 6. For each $v \in \{0, 1\}$, there is a unique bounded function $U_v : (0, \mu_0] \rightarrow [0, 1 - c]$ satisfying

$$U_v(n) = \begin{cases} h_v(n), & n \leq \underline{n}, \\ \max \left\{ h_v(n), \max_{p \in A(n)} (Q_v U_v)(n, p) \right\}, & n \in (\underline{n}, \mu_0]. \end{cases} \quad (28)$$

Moreover, for every $n \in (0, \mu_0]$,

$$\lim_{\delta \uparrow 1} (1 - \delta) V_\delta(v, n) = U_v(n). \quad (29)$$

Stopping yields the stationary flow $h_v(n)$; continued experimentation can raise n above the prior, where the limiting value is $1 - c$, or lead to the frozen region, from which recovery is

impossible. Since the seller already knows v , these experiments matter only through their effects on buyers' beliefs.

We now return to the original model, in which the seller does not observe v . For $0 < n < \mu_0$, write $U(r, n) := rU_1(n) + (1 - r)U_0(n)$ for the posterior-weighted average of the known-state limiting values. The next proposition distinguishes beliefs strictly below the prior from the prior itself. For each band $s \in \{1, \dots, M - 1\}$, let $n_s^- := \Lambda(\mathcal{L}(\mu_0) - m_s)$ and $r_s^- := rF_1(s)/F_r(s)$ denote the beliefs after a rejection at $n = \mu_0$.

Proposition 8. *For every $r \in [0, 1]$:*

(i) *For every $n \in (0, \mu_0)$,*

$$\lim_{\delta \uparrow 1} (1 - \delta)V_\delta(r, n) = U(r, n). \quad (30)$$

(ii) *At the prior,*

$$\lim_{\delta \uparrow 1} (1 - \delta)V_\delta(r, \mu_0) = \max \left\{ (\underline{p} - c)^+, \max_{s=1, \dots, M-1} [(1 - c) - F_r(s)((1 - c) - U(r_s^-, n_s^-))] \right\}. \quad (31)$$

For $0 < n < \mu_0$, the price $\min\{\beta_M(n), \underline{p}\}$ screens buyers while leaving n unchanged. Repeating this price allows the seller to learn v with arbitrarily high accuracy. As $\delta \uparrow 1$, the learning phase can occupy a vanishing fraction of the discounted horizon, yielding the mixture of known-state values in (30).

At $n = \mu_0$, every off-band price is uninformative. An in-band sale raises n above the prior, whereas a rejection leads to the below-prior recovery problem. Among experiments, (31) therefore minimizes the probability of rejection multiplied by the loss in continuation value after a rejection. Lower bands make a rejection less likely but produce a larger downward revision in n when it occurs. Current profit disappears after multiplication by $1 - \delta$, so the limiting choice depends only on these continuation payoffs. When she experiments, her limiting payoff is thus the expected value of a lottery: in state v , she obtains $1 - c$ if experimentation eventually lifts buyers' belief above the prior, and the stationary flow h_v at the belief where experimentation stops otherwise. By Bayes' rule, each band's limiting payoff is affine in r , making the prior value continuous, convex, and piecewise affine.

We next ask how often policies that attain the limiting profit succeed in making buyers optimistic enough for belief pumping. For each known state v , fix an optimal stationary experimentation rule σ_v for (28). At each belief it chooses an in-band price or stops experimenting and earns $h_v(n)$ thereafter; it stops in the frozen region.¹³ Let $\phi_v^\sigma(n)$ be this rule's probability, in state v , of reaching $n > \mu_0$ before stopping. The next proposition combines the initial experiment with any subsequent recovery attempts.

Proposition 9. *Fix $r \in [0, 1]$ and a band s whose experimental payoff attains the overall maximum in (31). There is a sequence of policies whose normalized profits converge to the limiting value at (r, μ_0) and whose crossing probabilities in each state $v \in \{0, 1\}$ converge to*

$$\pi_v(s; \sigma_v) = \bar{F}_v(s) + F_v(s)\phi_v^\sigma(n_s^-). \quad (32)$$

¹³Different profit-optimal rules can have different crossing probabilities. We fix one rule for each state, including how it resolves ties.

The first term is the probability of an initial purchase; the second accounts for recovery after an initial rejection. In state $v = 0$, this is the probability that the constructed policies reach the region where full exploitation of a worthless good becomes feasible. If the stopping alternative in (31) is optimal, repeating the best off-band action instead attains the limiting value with zero crossing probability. The proof gives the recursion for ϕ_v^σ .

4.5 Buyer payments and welfare

Full exploitation can occur even when buyers receive no value from the good. The next proposition describes their payments, the frequency of trade, and total surplus conditional on $v = 0$. Here each normalized quantity is the corresponding expected discounted sum multiplied by $1 - \delta$.

Proposition 10. *Fix $r \in (0, 1)$ and $n \in (\mu_0, 1)$, and let (π_δ) be any sequence of policies with $(1 - \delta)V_\delta^{\pi_\delta}(r, n) \rightarrow 1 - c$; such sequences exist by Proposition 5. Conditional on $v = 0$:*

- (i) *normalized seller profit converges to $1 - c$;*
- (ii) *normalized trade frequency converges to one;*
- (iii) *normalized buyer payments converge to one, and buyers' normalized payoff converges to -1 ;*
- (iv) *normalized total surplus converges to $-c$.*

Moreover, if T_δ is a random date drawn independently of the state and the entire market history, with $\mathbb{P}(T_\delta = t) = (1 - \delta)\delta^t$, then for every $\varepsilon \in (0, 1)$,

$$\mathbb{P}(a_{T_\delta} = 1, p_{T_\delta} \geq 1 - \varepsilon \mid v = 0) \rightarrow 1, \quad (33)$$

and on this event $n_{T_\delta} \geq \Lambda(\mathcal{L}(1 - \varepsilon) - \lambda_M)$.

When $v = 0$ and $c = 0$, buyers' losses are transfers to the seller: normalized buyer payoff converges to -1 and seller profit to one, while total surplus remains zero. With $c > 0$, each sale also uses resources worth c without creating any consumption value. Normalized total surplus therefore converges to $-c$, rather than the zero attainable through no trade.

The final claim describes a date sampled according to the seller's discount weights. As $\delta \uparrow 1$, a purchase occurs at a price arbitrarily close to one with probability approaching one, even conditional on $v = 0$. The sale-price bound then implies that the buyer's public belief is also arbitrarily close to one.

5 Finite patience with symmetric binary signals

We now study optimal pricing at a fixed discount factor. With symmetric binary signals, a favorable and an unfavorable update cancel exactly. This allows us to characterize when the seller stops raising buyers' willingness to pay and when she sacrifices current profit to learn about quality. We first solve the high- and low-belief regions, then give sufficient conditions for an optimal policy at every reachable belief.

Throughout this section, we assume $c = 0$, $\delta \in (0, 1)$, and $f_1(G) = f_0(B) = \alpha \in (1/2, 1)$, with $f_1(B) = f_0(G) = q := 1 - \alpha$. Thus, G and B correspond to signals 2 and 1 in the general model. Let $\ell := \log(\alpha/q)$, $\theta := e^\ell$, and

$$g(x) := \frac{\alpha x}{\alpha x + q(1-x)}, \quad b(x) := \frac{qx}{qx + \alpha(1-x)}$$

denote the posteriors after favorable and unfavorable evidence. Then, $\mathcal{L}(g(x)) = \mathcal{L}(x) + \ell$, $\mathcal{L}(b(x)) = \mathcal{L}(x) - \ell$, and $b(g(x)) = g(b(x)) = x$. Screening produces a sale with probability $\rho(r) := q + (2\alpha - 1)r$ according to the seller's posterior r . Starting from the common prior, both beliefs take values in

$$x_k := \Lambda(\mathcal{L}(\mu_0) + k\ell), \quad k \in \mathbb{Z}, \quad (34)$$

with $g(x_k) = x_{k+1}$ and $b(x_k) = x_{k-1}$. The decoding band is $I = (x_{-1}, x_1]$, the high-belief cutoff is $n^{\text{cap}} = x_2$, and the frozen boundary is $\underline{n} = x_{-2}$.

5.1 When to stop pumping

At $n \geq x_2$, every in-band price sells surely, so buyers' beliefs cannot fall. The seller can pump at x_1 , sell to everyone at the highest sure-sale price $b(n)$, or screen at $g(n)$. Screening reveals information to the seller but, being off band, leaves buyers' public belief unchanged. We give a sufficient condition under which the seller prefers to sell to everyone, choosing only when to stop pumping.

Let $V^{PH}(n)$ be the value when only pumping and harvesting are allowed. At $n = x_2$, the price $b(x_2) = x_1$ still pumps; harvesting without changing beliefs is available when $n > x_2$. Define

$$\bar{\delta}_H(n) := \frac{b(n) - x_1}{n - x_1}, \quad n \in [x_2, 1). \quad (35)$$

Since $b(g(n)) = n$, pumping once raises the subsequent harvest price from $b(n)$ to n . For $n > x_2$, $\bar{\delta}_H(n)$ is the discount factor at which this investment just compensates for the current price concession.

Proposition 11. *Suppose*

$$\delta x_2 + (1 - \delta)x_1 \geq \alpha. \quad (36)$$

For every $r \in [0, 1]$ and $n \in [x_2, 1)$, off-band screening is strictly suboptimal and $V_\delta(r, n) = V^{PH}(n)$. There is a unique $n^*(\delta) \in (x_2, 1)$ satisfying $\bar{\delta}_H(n^*(\delta)) = \delta$. Pumping at x_1 is optimal below this threshold and harvesting at $b(n)$ is optimal above it; both actions are optimal at the threshold.

Condition (36) compares the return from a sure sale followed by harvesting with the largest possible screening revenue. Under this condition, even learning the true quality cannot make screening preferable. Pumping and harvesting both sell surely, so their profits depend on buyers' willingness to pay, not on the seller's assessment of quality. Pumping sacrifices $b(n) - x_1$ today for higher prices later. As beliefs rise, $\bar{\delta}_H(n)$ increases, so the seller eventually prefers to collect the higher price. Condition (36) holds for every δ when $\mu_0 \geq 1/2$, because then $x_1 \geq \alpha$ and $x_2 > x_1$.

Choosing harvesting whenever the seller is indifferent, the number of pumps is

$$K^*(n, \delta) := \max \left\{ 0, \left\lceil \frac{\mathcal{L}(n^*(\delta)) - \mathcal{L}(n)}{\ell} \right\rceil \right\}.$$

The resulting value is

$$V^{PH}(n) = \frac{x_1[1 - \delta^{K^*(n, \delta)}] + \delta^{K^*(n, \delta)}b(g^{K^*(n, \delta)}(n))}{1 - \delta}. \quad (37)$$

A more patient seller sets a higher stopping threshold and makes weakly more pumps from any given belief. As $\delta \uparrow 1$, $n^*(\delta) \uparrow 1$, consistent with the patient-limit analysis.

5.2 Learning when buyers' beliefs are frozen

At $n = x_j$ with $j \leq -2$, no price can raise buyers' public belief. The seller can nevertheless learn about quality by charging a price accepted only after a favorable signal. The following lemma shows that deliberately lowering buyers' beliefs is never worthwhile, leaving a choice between pooling and screening.

Write $W_j(r) := V_\delta(r, x_j)$ and let $p_j^P := x_{j-1}$ and $p_j^S := x_{j+1}$. Pooling at p_j^P sells surely and leaves both beliefs unchanged. Screening at p_j^S reveals the buyer's signal to the seller but leaves n unchanged since $p_j^S \leq x_{-1}$ is off band.

Lemma 7. *For every $r \in [0, 1]$ and $j \leq -3$, $V_\delta(r, x_j) \leq V_\delta(r, x_{j+1})$. At every $j \leq -2$, an in-band price, which is rejected surely and lowers the naïve belief, is strictly suboptimal.*

Lowering n reduces the prices buyers accept without revealing anything about quality. Under an optimal policy, n therefore stays fixed. Only the seller's posterior changes, and only when she screens. To compare the two remaining prices, let $d_1 := \alpha$, $d_0 := q$, and

$$\Delta_v(j) := d_v p_j^S - p_j^P, \quad v \in \{0, 1\}. \quad (38)$$

This is the difference between screening and pooling revenue in state v . Since $\Delta_1(j) > \Delta_0(j)$, screening is more profitable when the good is valuable. The next proposition characterizes the seller's choice for every α and δ in the ranges specified above.

Proposition 12. *Fix $j \leq -2$.*

- (i) *If $\Delta_1(j) \leq 0$, pooling forever is optimal at every seller posterior.*
- (ii) *If $\Delta_0(j) \geq 0$, screening forever is optimal at every seller posterior.*
- (iii) *If $\Delta_0(j) < 0 < \Delta_1(j)$, there is a cutoff $r_j^* \in (0, 1)$. Pooling is uniquely optimal when $r < r_j^*$ and screening is uniquely optimal when $r > r_j^*$; both are optimal at equality. An optimal policy screens while $r > r_j^*$ and pools forever once $r \leq r_j^*$.*

In the first two cases, knowing v would not change the optimal price: the same action maximizes revenue in both states. In the third, a seller who knew quality would pool when $v = 0$ and

screen when $v = 1$. Screening then has an additional benefit: its sales and rejections help the seller decide which price to charge subsequently. The expression for the cutoff is given in the proof.

Let

$$\hat{r}_j := \frac{-\Delta_0(j)}{\Delta_1(j) - \Delta_0(j)},$$

denote the myopic cutoff, the unique posterior at which the expected current flows from the two actions coincide.

Corollary 1. *In case (iii) of Proposition 12, $r_j^* < \hat{r}_j$. Moreover, $r_j^* \rightarrow \hat{r}_j$ as $\delta \downarrow 0$ and $r_j^* \rightarrow 0$ as $\delta \uparrow 1$.*

The gap between the two cutoffs is where the seller pays for information. For $r \in (r_j^*, \hat{r}_j)$, Proposition 12 prescribes screening even though its expected current flow falls short of the safe flow: the per-period premium is $p_j^P - \rho(r)p_j^S$, strictly positive on this interval and vanishing exactly at $\hat{r}_j$, above which screening earns more today as well. Because the screening price is off band, it leaves the naïve belief fixed, so this sacrifice buys nothing but the option value of learning the state. The two limits in the corollary show how the value of information changes with patience: as $\delta \downarrow 0$, the seller becomes unwilling to sacrifice current profit for information and the experimentation region vanishes; as $\delta \uparrow 1$, experimentation becomes optimal at every interior posterior.

5.3 Optimal pricing and purchase cascades

We now give sufficient conditions under which the seller experiments at the prior, makes one recovery attempt after a rejection, and switches to pumping after a purchase at the prior. The conditions allow the same pricing rule to be optimal whether the seller knows that the good is valuable, knows that it is worthless, or remains uncertain about quality.

Let $\Gamma_v(n) := V_\delta(v, n)$ denote the seller's value when she knows v , keeping buyers' purchasing and inference rules unchanged. Revealing quality can only benefit the seller, so

$$V_\delta(r, n) \leq r\Gamma_1(n) + (1 - r)\Gamma_0(n). \quad (39)$$

If the same feasible pricing rule attains both known-state values without conditioning on v , its expected payoff attains this bound for every r . It is therefore optimal even without knowledge of quality. This observation lets us verify the proposed prices separately in the two states.

Assume (36) holds. Under the proposed policy, the continuation value after a purchase at the prior is $J^1 := x_0 + \delta V^{PH}(x_2)$: the seller pumps at x_0 and then follows Proposition 11. For a known-state sale probability $d \in \{q, \alpha\}$, let J_d^F , J_d^0 , and J_d^{-1} be the unique solution to

$$J_d^F = \frac{dx_{-1}}{1 - \delta}, \quad (40)$$

$$J_d^0 = dx_1 + \delta[dJ^1 + (1 - d)J_d^{-1}], \quad (41)$$

$$J_d^{-1} = dx_0 + \delta[dJ_d^0 + (1 - d)J_d^F]. \quad (42)$$

Here J_d^F is the value of screening forever at $n = x_{-2}$. The other values correspond to screening at x_1 when $n = x_0$ and at x_0 when $n = x_{-1}$. At $n = x_{-1}$, a purchase returns buyers' belief to the prior, while a rejection leads to x_{-2} .

For screening to remain optimal throughout the frozen region even when $v = 0$, it is enough that

$$x_{-3} \leq qx_{-1}. \quad (43)$$

This compares screening and pooling revenue at $n = x_{-2}$ and implies the same comparison at every lower belief. Equivalently,

$$\alpha > \frac{\sqrt{5} - 1}{2} \quad \text{and} \quad \mu_0 \leq \mu^{\text{fr}}(\alpha) := \frac{\alpha^2 + \alpha - 1}{\alpha(2\alpha - 1)}. \quad (44)$$

The bound μ^{fr} increases with α . At $\mu_0 = 1/2$, the restriction reduces to $\alpha \geq 2/3$. The remaining restrictions compare the proposed prices with pooling and with alternative screening prices at the three intermediate beliefs.

Condition 1.

(a) *Inequality (43) holds.*

(b) *For each $d \in \{q, \alpha\}$,*

$$J^1 \geq \max \left\{ \frac{x_{-1}}{1 - \delta}, \frac{dx_2}{1 - \delta}, dx_1 + \delta [dV^{PH}(x_2) + (1 - d)J_d^0] \right\}, \quad (45)$$

$$J_d^0 \geq \frac{x_{-1}}{1 - \delta}, \quad (46)$$

$$J_d^{-1} \geq \max \left\{ \frac{x_{-2}}{1 - \delta}, J_d^F \right\}. \quad (47)$$

Inequality (45) makes a sure-sale pump preferable to pooling or further screening after an initial purchase. Inequalities (46) and (47) make the initial experiment and the recovery attempt worthwhile. The next proposition combines these choices with the policies in the two absorbing regions.

Proposition 13. *Suppose (36) and Condition 1 hold. The following stationary policy is optimal at every (r, x_j) with $r \in [0, 1]$ and $j \in \mathbb{Z}$:*

(i) *for $j \geq 2$, follow the pump-and-harvest policy of Proposition 11;*

(ii) *at $j = 1$, pump at x_0 ;*

(iii) *at $j = 0$, screen in band at x_1 ;*

(iv) *at $j = -1$, screen in band at x_0 ;*

(v) *for $j \leq -2$, screen off band at x_{j+1} forever.*

If all inequalities are strict and $n^*(\delta) \neq x_j$ for every $j \geq 2$, the prescribed price is uniquely optimal at each such state, and the same action pattern remains optimal in a neighborhood of the parameter vector.

Under this policy, starting from $n = \mu_0$, the probability in state v of eventually entering the high-belief region is

$$\varpi_v = \frac{d_v}{1 - d_v(1 - d_v)}. \quad (48)$$

The frozen region is reached with complementary probability.

An initial purchase prompts a price cut from x_1 to x_0 . The lower price guarantees a sale, which buyers mistake for further favorable evidence. The seller then pumps until the higher future price no longer compensates for the current concession. After an initial rejection, she instead offers one lower-priced recovery experiment. A purchase restores the initial beliefs when starting from the common prior; another rejection ends attempts to raise buyers' beliefs, but not trade. The seller continues selling only to buyers with favorable signals.

The probability in (48) includes repeated recovery attempts: a rejection followed by a purchase restarts the problem at the prior. Conditional on quality, these probabilities do not depend on the seller's posterior or discount factor within the stated parameter region. Her ex ante probability of reaching the high-belief region still depends on her assessment of quality, through $r\varpi_1 + (1-r)\varpi_0$.

The following example illustrates the price path under optimal manipulation and compares the probability of a purchase cascade for valuable and worthless goods. It also shows how greater patience can raise the eventual price without changing the probability of a cascade.

Example: purchase cascades on a worthless good. Return to $(\mu_0, \alpha) = (1/2, 7/10)$, as in the illustration following Proposition 3. The inequalities in (36) and Condition 1 hold strictly for every $\delta \in [4/5, 1)$, as verified in Appendix A. Thus, Proposition 13 applies throughout this interval. At each such discount factor, the same intermediate- and low-belief actions remain optimal for nearby parameter values.

At $\delta = 9/10$, an initial sale at $7/10$ is followed by a sure sale at $1/2$. Starting from the resulting belief x_2 , the seller pumps three more times at $7/10$, reaches x_5 , and thereafter charges $x_4 = 2401/2482 \approx 0.9674$ to every buyer. Her value at the common prior is $V_\delta(1/2, 1/2) \approx 5.4814$. At the same prior and discount factor, Lemma 5 bounds profit with rational buyers by 5. The profit gain from naïve inference therefore survives the risk of unsuccessful initial experimentation.

After two consecutive rejections in the experimentation phase, buyers' public beliefs instead stay at $x_{-2} = 9/58$, and the seller charges $x_{-1} = 3/10$ forever. These purchases and rejections reveal independent private signals, so the seller's posterior converges almost surely to the true state. Buyers' public belief nevertheless remains fixed. Even learning that the good is valuable does not allow the seller to restore their willingness to pay.

For every $\delta \in [4/5, 1)$,

$$\varpi_0 = \frac{30}{79}, \quad \varpi_1 = \frac{70}{79}.$$

With the common prior $1/2$, the unconditional probabilities are $15/79$ for universal purchase when $v = 0$, $35/79$ for universal purchase when $v = 1$, and $29/79$ for frozen beliefs. A valuable good is more likely to reach permanent universal purchase, but a worthless good does so with probability $30/79 \approx 0.380$. Misleading early purchases make this outcome possible; the seller's subsequent sure-sale pricing amplifies those purchases without producing further information. Patience changes how long she pumps and the final price, rather than the probabilities of the two long-run regimes.

Failure to establish optimism need not shut down the market. It can nevertheless leave a valuable good permanently under-purchased. Conditional on $v = 1$, buyers' beliefs freeze with probability $9/79$. Thereafter, a fraction $3/10$ of buyers forgo a good worth one that costs nothing to produce. Averaging over quality and the two long-run regimes, the expected per-period welfare

loss in the long run, relative to efficient trade, is

$$\frac{1}{2} \cdot \frac{9}{79} \cdot \frac{3}{10} = \frac{27}{1580}.$$

Zero marginal cost therefore does not eliminate all allocation losses from naïve learning.¹⁴

On successful paths, the initial discount is followed by higher prices. The seller can charge more because pumping raises buyers' willingness to pay through purchases that reveal no new information about quality.

6 Extensions

We now return to the general finite-signal model with marginal cost $c \in [0, 1)$. We first examine when the seller incurs losses on early sales; then, we consider how buyers' interpretation of unexpected off-band actions affects the cost of manipulation.

6.1 Below-cost penetration pricing

With positive marginal cost, the seller may incur losses while obtaining the first favorable purchase or while amplifying the optimism it creates. To identify below-cost sales along the pump-and-harvest construction, for $c \in (\underline{p}, \bar{p})$ define

$$n_c^P := \Lambda(\mathcal{L}(c) - \lambda_1), \quad (49)$$

as the belief at which the highest sure-sale price equals cost: $\beta_1(n_c^P) = c$. This cost restriction implies $\mu_0 < n_c^P < n^{\text{cap}}$. When instead $c > \bar{p}$, every price that can generate an initial sale is below cost. The next proposition shows that a sufficiently patient seller nevertheless prefers to experiment.

Proposition 14. *Distinguish two cases:*

(i) *Suppose $c \in (\underline{p}, \bar{p})$. For every $n \in (\mu_0, 1)$,*

$$p^P(n) < c \iff n < n_c^P. \quad (50)$$

Starting at $n_0 \in (\mu_0, n_c^P)$, repeated pricing at $p^P(n_t)$ produces a finite, nonempty initial sequence of below-cost sales. The pump-and-harvest construction attaining the sharp coefficient χ in Proposition 6 has this initial phase whenever it starts in that interval.

(ii) *Suppose $c \in (\bar{p}, 1)$. For all sufficiently high $\delta < 1$, $V_\delta(\mu_0, \mu_0) > 0$, and every optimal initial price at the common prior belongs to I . The seller therefore earns strictly negative expected current profit from her initial experiment.*

¹⁴With rational buyers, [Bose et al. \(2006, 2008\)](#) show that a valuable good eventually generates a purchase cascade almost surely when marginal cost equals the low value.

In part (i), the loss-making sales are guaranteed. They raise future buyers' willingness to pay without revealing anything about quality, leaving the seller's posterior unchanged. The losses end once $n_t \geq n_c^P$, although pumping may continue; the eventual harvest price exceeds $\bar{p} > c$.

In part (ii), the initial experiment does reveal information. A purchase raises buyers' belief above the prior and makes sure-sale pumping possible. The seller accepts a negative expected current margin because this successful outcome opens a sufficiently profitable continuation.

Low introductory prices also arise without inferential naïvety: in [Shapiro \(1983\)](#), a monopolist sets a low introductory price for an experience good whose quality consumers initially underestimate, because they learn it by using the good. The loss in part (ii) has a similar motive, since the initial sale is informative. The losses in part (i) are different: those sales reveal nothing about quality, but later buyers mistakenly interpret them as informative.

6.2 Optimistic off-band beliefs

Under passive decoding, sales above $\bar{p}$ cannot raise buyers' beliefs. Following the approach in [Eyster and Rabin \(2010, footnote 12\)](#), we now attribute unexpected off-band actions to extreme signals. The following assumption replaces Assumption 1; in-band decoding is unchanged.

Assumption 2. For $p \notin I$, a sale at $p > \bar{p}$ is decoded as signal M , so $\mathcal{L}(N(n, p, 1)) = \mathcal{L}(n) + \lambda_M$. A rejection at $p \leq \underline{p}$ is decoded as signal 1, so $\mathcal{L}(N(n, p, 0)) = \mathcal{L}(n) + \lambda_1$. All other off-band observations leave n unchanged.

Let V_δ^E denote the seller's value under Assumption 2, again allowing any $c \in [0, 1)$. The next proposition shows that, starting above the prior, the optimal discounted profit shortfall remains bounded as $\delta \uparrow 1$.

Proposition 15. For every $n \in (\mu_0, 1)$, $r \in [0, 1]$, and $\delta \in (0, 1)$,

$$1 - \max\{c, \beta_M(n)\} \leq \frac{1-c}{1-\delta} - V_\delta^E(r, n) \leq K_0(n) + 1 + \frac{e^{-[\mathcal{L}(\mu_0) + \lambda_M]}}{1 - e^{-\lambda_M}}. \quad (51)$$

In particular, the shortfall remains bounded as $\delta \uparrow 1$.

The upper bound comes from a policy that first pumps above n^{cap} and then charges $\beta_1(n_t) > \bar{p}$ in every period. Each purchase is certain, but buyers interpret it as evidence of the strongest signal. Their willingness to pay therefore rises, allowing the seller to increase her price again. These prices approach one fast enough that the total profit forgone relative to charging one in every period remains finite, even without discounting.

Proposition 6 shows that under passive decoding the optimal discounted shortfall grows at rate $\chi \log[1/(1-\delta)]$; under Assumption 2 it remains bounded. Full exploitation above the prior therefore holds under both completions, and only its logarithmic cost depends on the passive rule. Because the optimistic completion lets the seller raise both her price and buyers' belief with the same sale, the passive rule is the more conservative choice for our purposes: above the prior, it makes manipulation more expensive.

7 Concluding remarks

We have studied a monopolist selling to buyers who learn from earlier purchases but neglect that those purchases already reflect what earlier buyers had learned. Such buyers let the seller turn optimism into more optimism: once they are more optimistic than the prior, a price that every current buyer accepts yields a purchase that tells the seller nothing but that later buyers read as favorable private evidence. A patient seller who starts from such a state earns, on average, close to the profit she would make selling a good known to have the highest possible value, irrespective of the good’s actual quality, at a cost that grows only with the logarithm of the effective horizon. In the extreme case of a good that is worthless for consumers but costly to produce, the manipulation we identified ends up lowering total surplus, not just transferring it. At the prior, optimism must instead be seeded by an informative purchase, and a run of unfavorable signals can leave buyers’ belief permanently low. With symmetric binary signals, we characterize the optimal policy at fixed discount factors under sufficient conditions on primitives.

The model yields predictions that separate BRTNI from rational social learning. At the level of individual inference, a BRTNI buyer’s response to an earlier purchase depends on the price at which it was made, but not on how optimistic the earlier buyer already was; a rational buyer discounts purchases by buyers who would have bought regardless of their signals. A laboratory design in the spirit of [Eyster et al. \(2018\)](#) could hold fixed the price and choice of an early subject while varying what that subject had observed. At the market level, prices can rise during a run of purchases that every buyer would have made whatever his signal, as on the successful paths of Section 5, whereas with rational buyers such purchases leave willingness to pay unchanged.

From a policy perspective, the manipulation we study uses only genuine purchases at prices that buyers observe, so rules against fake reviews, such as the one adopted by the U.S. Federal Trade Commission in 2024, do not reach it.¹⁵ The harm comes from how buyers interpret the sales record. Hiding others’ purchases would prevent this mislearning, but it would also withhold the information that genuine purchases convey.

We end the paper by discussing some of its limitations, as well as avenues for future research. In our analysis, we take BRTNI as given and assume that all buyers engage in it. With a mix of naïve and rational buyers, the seller would need to account for how rational buyers read her prices. Moreover, buyers in our model observe the entire history of prices and decisions, whereas many platforms report summaries such as sales ranks or purchase counts, often without the prices at which purchases were made. How naïve buyers read such summaries, and how a seller prices when only summaries are visible, is a natural next question. Our model features a single seller and a good with a pure common value; competition among sellers and heterogeneous tastes, as in [Gagnon-Bartsch and Rosato \(2024\)](#), could change both the incentive to manipulate and its cost. Finally, at fixed discount factors, we characterized optimal pricing only for symmetric binary signals and under sufficient conditions, whereas for general signal structures, our analysis focused on the patient limit; furthermore, our arguments relied on bounded likelihood ratios. We hope that future research will be able to address these gaps.

¹⁵Federal Trade Commission, “Trade Regulation Rule on the Use of Consumer Reviews and Testimonials,” 16 CFR Part 465, 89 Fed. Reg. 68034 (August 22, 2024).

A Proofs

Proofs for Sections 2 and 3

Proof of Lemma 1. By full support and MLRP, $h(s) := f_1(s)/f_0(s) = e^{\lambda s}$ is well defined and strictly increasing. Since $\mathbb{E}_0[h(S)] = 1$, we have $h(1) < 1 < h(M)$, and hence $\lambda_1 < 0 < \lambda_M$.

For $s \in \{1, \dots, M-1\}$,

$$\frac{\bar{F}_1(s)}{\bar{F}_0(s)} = \mathbb{E}_0[h(S) \mid S > s], \quad \frac{F_1(s)}{F_0(s)} = \mathbb{E}_0[h(S) \mid S \leq s].$$

Both conditional means are strictly increasing in s . Moreover, because h is strictly increasing and has unconditional mean one, the lower-tail mean is below one and the upper-tail mean is above one. Hence $\ell_s, m_s > 0$. The upper-tail mean is bounded above by $h(M)$, with equality only at $s = M-1$, while the lower-tail mean is bounded below by $h(1)$, with equality only at $s = 1$. Taking logs gives $\ell_s \leq \lambda_M$ with $\ell_{M-1} = \lambda_M$, and $m_s \leq -\lambda_1$ with $m_1 = -\lambda_1$. Their stated monotonicity follows from that of the two conditional means, completing the proof. ■

Proof of Lemma 2. The purchasing rule selects an upper set of signals. If all types buy or all reject, the realized action has probability one in both states and leaves r unchanged. Otherwise, for some $s \in \{1, \dots, M-1\}$, a sale reveals $\{s_t > s\}$ and a rejection reveals $\{s_t \leq s\}$. Bayes' rule gives the posterior formulas in part (ii), including at $r = 0, 1$. The naïve update is (3). Finally, a sale requires the highest-signal type to buy, giving (4). ■

Proof of Proposition 1. Fix (r, n) and order the distinct elements of $\mathcal{T}(n)$ as $0 = y_0 < y_1 < \dots < y_K = 1$. For $p \in (y_{k-1}, y_k]$, the inequalities $p \leq \beta_j(n)$ and $p \leq \beta_j(\mu_0)$ have the same truth values as at $p = y_k$. Actual demand and both belief transitions are therefore constant on this interval. The objective in (5) is affine in p , with slope $D(p; r, n) \geq 0$, so replacing p by y_k cannot lower it. Since $0 \in \mathcal{T}(n)$, the maximum is attained in $\mathcal{T}(n)$.

Expected current profit and the conditional distribution of the next state depend on the history only through (r, n) and the chosen price. The candidate-price functions $0, 1, \beta_s(\mu_0)$, and $\beta_s(n)$, together with the associated demand and transition maps, are Borel measurable. Hence T maps bounded Borel functions into themselves. For any two such functions,

$$\|TW - TW'\|_\infty \leq \delta \|W - W'\|_\infty.$$

Thus T has a unique bounded Borel fixed point W . A fixed ordering of the candidate-price functions gives a Borel-measurable maximizing price $\pi^*(r, n)$.

For any history-dependent policy π , iterating the Bellman inequality gives

$$W(r, n) \geq \mathbb{E}_{r,n}^\pi \left[\sum_{t=0}^{H-1} \delta^t (p_t - c) a_t + \delta^H W(r_H, n_H) \right].$$

Equality holds under π^* . Since W and flow profits are bounded, letting $H \rightarrow \infty$ yields $W \geq V_\delta^\pi$ for every π and $W = V_\delta^{\pi^*}$. Therefore $W = V_\delta$, and π^* is optimal among all admissible policies.

Finally, posting $p = c$ forever earns zero profit, whereas no policy earns more than $1 - c$ per period. These observations give (6). ■

Proof of Proposition 2. Under the true state distribution induced by r , a sale occurs with probability $\bar{F}_r(k)$, where $F_r(k) := 1 - \bar{F}_r(k)$. BRTNI buyers nevertheless update to $n_s^+(n)$ after a sale and $n_s^-(n)$ after a no-sale. Therefore,

$$\mathbb{E}[n_{t+1}] = \bar{F}_r(k)n_s^+(n) + F_r(k)n_s^-(n).$$

A Bayesian with prior n who correctly interpreted the cutoff s would satisfy the martingale identity

$$n = \bar{F}_n(s)n_s^+(n) + F_n(s)n_s^-(n).$$

Subtracting gives (8). By Lemma 1, $\ell_s, m_s > 0$, so $n_s^+(n) > n > n_s^-(n)$ and the second factor is strictly positive. If $k = s$, strict MLRP implies $\bar{F}_1(s) > \bar{F}_0(s)$, so the sign is the sign of $r - n$. The cases $k = 0$ and $k = M$ are immediate. ■

Proof of Proposition 3. Since $n > \mu_0$ and β_1 is strictly increasing, $\beta_1(n) > \beta_1(\mu_0) = \underline{p}$. Thus, $p^P(n) = \min\{\beta_1(n), \bar{p}\}$ lies in $(\underline{p}, \bar{p}] = I$, proving (i). Since $p^P(n) \leq \beta_1(n) \leq \beta_s(n)$ for every s , the price is weakly below every type's willingness to pay: the sale is certain and leaves r unchanged, proving (ii). It is decoded according to the band containing $p^P(n)$, so log odds rise by $\ell_{s(n)} \geq \ell_1 > 0$, proving (iii).

The inequality $n \geq n^{\text{cap}}$ is equivalent to $\mathcal{L}(n) + \lambda_1 \geq \mathcal{L}(\mu_0) + \lambda_M$, or $\beta_1(n) \geq \bar{p}$. Hence $p^P(n) = \bar{p} \in I_{M-1}$ and the increment is $\ell_{M-1} = \lambda_M$, proving (iv). After K pumps from n , log odds are at least $\mathcal{L}(n) + K\ell_1$. Since $\ell_1 > 0$, repeated pumping therefore sends the naïve belief to one while leaving r unchanged. Moreover,

$$\mathcal{L}(n) + K_0(n)\ell_1 \geq \mathcal{L}(n^{\text{cap}}),$$

so at most $K_0(n)$ pumps are needed to reach n^{cap} , as claimed. ■

Proof of Lemma 3. By construction, $\beta_M(\underline{n}) = \Lambda(\mathcal{L}(\mu_0) + \lambda_1) = \underline{p}$. Hence, if $n \leq \underline{n}$, then $\beta_M(n) \leq \underline{p}$. Every in-band price is strictly above $\underline{p}$, so no type buys and the naïve belief weakly decreases, while off-band observations leave it unchanged. Thus n can never rise. Any future sale must occur at a price no greater than $\beta_M(n_t) \leq \underline{p}$, so flow profit is at most $(\underline{p} - c)^+$. Summing discounted profits gives (12).

If instead $n \geq n^{\text{cap}}$, then $\beta_1(n) \geq \bar{p}$. Every in-band price therefore sells surely and satisfies $N(n, p, 1) \geq n$, while off-band observations leave n unchanged. Hence n can never fall and the region $\{n \geq n^{\text{cap}}\}$ is absorbing. Starting from n , posting the fixed price $\beta_1(n)$ forever induces a sale at every date and yields $(\beta_1(n) - c)/(1 - \delta)$. Combining this with $V_\delta(r, n) \geq 0$ from Proposition 1 gives (13). ■

Proof of Lemma 4. Lemma 1 gives $\ell_s \leq \lambda_M$, while monotonicity of willingness to pay gives $\beta_{s+1}(\mu_0) \leq \beta_M(\mu_0) = \bar{p}$. Therefore

$$\frac{1 - \beta_{s+1}(\mu_0)}{\ell_s} \geq \frac{1 - \bar{p}}{\lambda_M}.$$

Both inequalities bind only for $s = M - 1$. If $p \in I_s$, then $p \leq \beta_{s+1}(\mu_0)$, giving (15). ■

Proofs for Section 4

Proof of Lemma 5. At a rational public history, a purchasing buyer pays no more than his posterior expected value. Iterated expectations therefore give

$$\mathbb{E}[(p_t - c)a_t \mid \mathcal{F}_t] \leq \mathbb{E}[(v - c)a_t \mid \mathcal{F}_t] \leq (1 - c)\mu_t,$$

where the second inequality uses $(v - c)a_t \leq (1 - c)v$ in every realization. Since $\mathbb{E}[\mu_t] = \mu$ under every pricing policy, summing discounted expected profits and taking the supremum over policies proves (16). $\blacksquare$

For any policy π and initial state (r, n) ,

$$\frac{1 - c}{1 - \delta} - V_\delta^\pi(r, n) = \mathbb{E}_{r,n}^\pi \left[\sum_{t \geq 0} \delta^t [(1 - c) - (p_t - c)a_t] \right]. \quad (52)$$

The period- t flow shortfall equals $1 - p_t$ after a sale and $1 - c$ after a rejection, and lies in $[0, 1]$.

Proof of Proposition 4. Consider the policy that pumps at $p^P(n_t)$ for $K_0(n)$ periods, reaching $n_t \geq n^{\text{cap}}$ by Proposition 3, pumps N further periods at $\bar{p}$, and thereafter posts $\beta_1(n^*)$ forever, where n^* is the naïve belief after the last pump. Every one of these prices sells with probability one: the pumps by Proposition 3, and the harvest price because it equals the lowest willingness to pay. Since the first phase ends at or above n^{cap} and each top-band pump adds λ_M ,

$$\mathcal{L}(n^*) \geq \max\{\mathcal{L}(n), \mathcal{L}(n^{\text{cap}})\} + N\lambda_M,$$

so with $N \geq 1$, $\beta_1(n^*) = \Lambda(\mathcal{L}(n^*) + \lambda_1) \geq \Lambda(\mathcal{L}(\mu_0) + (N + 1)\lambda_M) > \bar{p}$: the harvest price is off band and the naïve belief stays fixed under Assumption 1.

The flow shortfall is $1 - p_t \leq 1$ in each of the first $K_0(n)$ periods (the bound 1 cannot be replaced by $1 - c$, because early pump prices may lie below cost), $1 - \bar{p}$ in each of the N top-band pumps, and $1 - \beta_1(n^*)$ in every harvest period. Bounding discount factors by one and using $1 - \Lambda(y) \leq e^{-y}$,

$$1 - \beta_1(n^*) \leq e^{-[\mathcal{L}(n^*) + \lambda_1]} \leq e^{-[\mathcal{L}(n) + \lambda_1]} e^{-N\lambda_M},$$

and summing the three phases gives (17). $\blacksquare$

Proof of Proposition 5. Let

$$N(\delta) := \left\lceil \frac{1}{\lambda_M} \log \frac{1}{1 - \delta} \right\rceil.$$

Then $e^{-N(\delta)\lambda_M} \leq 1 - \delta$. Uniformly over $n > \mu_0$,

$$K_0(n) \leq \left\lceil \frac{\lambda_M - \lambda_1}{\ell_1} \right\rceil, \quad e^{-[\mathcal{L}(n) + \lambda_1]} \leq e^{-[\mathcal{L}(\mu_0) + \lambda_1]}.$$

Multiplying (17) by $1 - \delta$ and using these bounds, the right-hand side is at most

$$(1 - \delta) \left\lceil \frac{\lambda_M - \lambda_1}{\ell_1} \right\rceil + (1 - \delta)N(\delta)(1 - \bar{p}) + e^{-[\mathcal{L}(\mu_0) + \lambda_1]}(1 - \delta),$$

which vanishes because $(1 - \delta)N(\delta) \rightarrow 0$. Together with $(1 - \delta)V_\delta \leq 1 - c$ from (6), this proves (18). Combining (18) with (16) gives (19) for $n \in (\mu_0, 1)$. At $n = 1$, posting price one forever yields $(1 - c)/(1 - \delta)$, so the same comparison holds. $\blacksquare$

The lower bound in Proposition 6 uses two lemmas: a pathwise accounting inequality that prices every period of every path—growth periods through Lemma 4, stagnation periods through the sale-price cap (4)—and a deterministic estimate showing that no sequence of belief increments escapes the resulting tradeoff. For a fixed path, write

$$L_t := \mathcal{L}(n_t), \quad G_t := (L_{t+1} - L_t)^+, \quad A_t := \sum_{\tau < t} G_\tau.$$

Lemma A.1. Fix an initial $n \in (0, 1)$ and any pricing policy. Let

$$\kappa_c := \min \left\{ 1 - c, \frac{1}{2} e^{-[\mathcal{L}(n) + \lambda_M]^+} \right\} > 0.$$

Then $G_t \in \{0\} \cup [\ell_1, \lambda_M]$, and the realized flow shortfall satisfies

$$(1 - c) - (p_t - c)a_t \geq \chi G_t \mathbf{1}\{G_t > 0\} + \kappa_c e^{-A_t} \mathbf{1}\{G_t = 0\}. \quad (53)$$

Proof. Under passive off-band beliefs, $G_t > 0$ only after an in-band sale. Then $G_t = \ell_s$ for some s , $a_t = 1$, and Lemma 4 gives

$$(1 - c) - (p_t - c) = 1 - p_t \geq \chi G_t.$$

If $G_t = 0$ and a sale occurs, (4) gives

$$(1 - c) - (p_t - c) = 1 - p_t \geq 1 - \beta_M(n_t).$$

If no sale occurs, the shortfall is $1 - c$. The logistic bound

$$1 - \Lambda(y) \geq \frac{1}{2} e^{-y^+}$$

and $L_t \leq L_0 + A_t$ imply

$$1 - \beta_M(n_t) \geq \frac{1}{2} e^{-[L_t + \lambda_M]^+} \geq \frac{1}{2} e^{-[L_0 + \lambda_M]^+} e^{-A_t}.$$

A no-sale shortfall also exceeds $(1 - c)e^{-A_t}$. Taking the smaller constant yields (53). ■

The next lemma bounds the resulting sum uniformly over all possible sequences of belief increases.

Lemma A.2. Fix $\chi, \kappa > 0$ and $0 < \ell_1 \leq \lambda_M < \infty$. For every $\eta \in (0, 1/2)$ there exists $\bar{\delta}(\eta) < 1$ such that, for every $\delta \geq \bar{\delta}(\eta)$ and every sequence $G_t \in \{0\} \cup [\ell_1, \lambda_M]$,

$$\sum_{t \geq 0} \delta^t [\chi G_t \mathbf{1}\{G_t > 0\} + \kappa e^{-A_t} \mathbf{1}\{G_t = 0\}] \geq (1 - \eta)^2 \chi \log \frac{1}{1 - \delta}. \quad (54)$$

Proof. Let $\bar{A} := \log[1/(1 - \delta)]$, $\bar{A}_\eta := (1 - \eta)\bar{A}$, $t_\eta := \lfloor \eta/(1 - \delta) \rfloor$, and $\tau := \inf\{t : A_t \geq \bar{A}_\eta\}$. For $t \leq t_\eta$, Bernoulli's inequality gives $\delta^t \geq 1 - \eta$.

If $\tau \leq t_\eta$, then

$$\sum_{t < \tau} \delta^t \chi G_t \geq (1 - \eta) \chi A_\tau \geq (1 - \eta)^2 \chi \bar{A}.$$

If $\tau > t_\eta$, then $A_t < \bar{A}_\eta$ for every $t \leq t_\eta$. At most $(\bar{A}_\eta + \lambda_M)/\ell_1$ of those dates have $G_t > 0$. Hence the contribution from dates with $G_t = 0$ is at least

$$\kappa(1 - \eta) \left[t_\eta + 1 - \frac{\bar{A}_\eta + \lambda_M}{\ell_1} \right] (1 - \delta)^{1 - \eta}.$$

Divided by $\bar{A}$, this expression diverges as $\delta \uparrow 1$. It therefore exceeds the right-hand side of (54) for all sufficiently large δ . The threshold $\bar{\delta}(\eta)$ produced by this argument depends only on $(\eta, \chi, \kappa, \ell_1, \lambda_M)$, not on the sequence, which is what permits taking expectations over random paths in the proof of Proposition 6. $\blacksquare$

Proof of Proposition 6. For part (i), apply Proposition 4 with

$$N = \left\lceil \frac{1}{\lambda_M} \log \frac{1}{1 - \delta} \right\rceil.$$

Then $N(1 - \bar{p}) \leq \chi \log[1/(1 - \delta)] + 1 - \bar{p}$ and $e^{-N\lambda_M} \leq 1 - \delta$. Substitution into (17) gives (20), with $C(n)$ as in (21), for every $r \in [0, 1]$.

For part (ii), fix $n \in (0, 1)$ and $\eta \in (0, 1/2)$. Lemma A.1 bounds the realized discounted shortfall below by the left-hand side of (54), with $\kappa = \kappa_c$. This constant depends on n, c , and the signal structure, but not on r , the policy, or the realized path. Lemma A.2 therefore gives, for all sufficiently large δ , uniformly over all these choices,

$$\sum_{t \geq 0} \delta^t [(1 - c) - (p_t - c)a_t] \geq (1 - \eta)^2 \chi \log \frac{1}{1 - \delta}.$$

Taking expectations and then the infimum over policies in (52) yields

$$\inf_{r \in [0, 1]} \frac{\frac{1-c}{1-\delta} - V_\delta(r, n)}{\log \frac{1}{1-\delta}} \geq (1 - \eta)^2 \chi.$$

Taking the lower limit as $\delta \uparrow 1$ and then letting $\eta \downarrow 0$ proves (22). For fixed $n \in (\mu_0, 1)$, the upper bound in part (i) and the uniform lower bound in part (ii) give (23). $\blacksquare$

Proof of Proposition 7. Fix a deterministic stationary optimal policy, which exists by Proposition 1. Let $T_1 < T_2 < \dots$ be its successive in-band dates, let $J \in \{0, 1, \dots\} \cup \{\infty\}$ be their number, and let

$$E := \{a_{T_k} = 0 \text{ for every } k \leq \min\{J, K_{\text{fr}}\}\}.$$

On E , no in-band sale ever occurs: the first $\min\{J, K_{\text{fr}}\}$ in-band dates are no-sales by definition; and if $J \geq K_{\text{fr}}$, those no-sales lower log odds by at least $K_{\text{fr}} m_{M-1} \geq \lambda_M - \lambda_1$, placing the belief in the frozen region, where Lemma 3 makes every further in-band date a sure no-sale. Since off-band dates leave the naïve belief fixed, $n_t \leq \mu_0$ throughout on E .

At every in-band date with $n \leq \mu_0$, the lowest-signal buyer rejects, because the price is strictly above $\underline{p} = \beta_1(\mu_0) \geq \beta_1(n)$; conditional on v and the history up to that date, a no-sale therefore

occurs with probability at least $f_v(1)$. Induction over the first K_{fr} in-band dates, conditioning at each $\mathcal{F}_{T_k}$, gives

$$\mathbb{P}(E \mid v) \geq f_v(1)^{K_{\text{fr}}}.$$

On E , every in-band date yields zero profit. An off-band sale requires $p_t \leq \beta_M(n_t) \leq \beta_M(\mu_0) = \bar{p}$ by (4), so its price cannot lie above the band and must be at most $\underline{p}$. Normalized profit on E is therefore at most $(\underline{p} - c)^+$; outside E it is at most $1 - c$. Hence,

$$(1 - \delta)V_\delta(r, \mu_0) \leq (1 - c) - \Delta_c \mathbb{P}(E),$$

which gives (25). The uniform gap follows from full support. $\blacksquare$

Proof of Lemma 6. Fix v and write $W_{\delta,v}(n) := (1 - \delta)V_\delta(v, n)$. Because h_v includes the option of zero profit, prices below cost can be omitted when computing it. At every remaining price, expected current profit is nondecreasing in n ; hence h_v is nondecreasing. If $n \leq \underline{n}$, Lemma 3 implies that beliefs cannot rise. Every future in-band action then yields zero profit, and every off-band action yields expected current profit at most $h_v(n)$. Repeating an off-band action attaining this flow therefore gives $W_{\delta,v}(n) = h_v(n)$.

For an in-band price, define

$$(Q_{v,\delta}W)(n, p) := (1 - \delta)(p - c)d_v(p, n) + \delta \mathbb{E}_v[W(N(n, p, a))],$$

where $a = 1$ with probability $d_v(p, n)$. An off-band action leaves the known-state problem unchanged, so its Bellman self-loop can be replaced by stopping at $h_v(n)$. Thus, on $(\underline{n}, \mu_0]$,

$$W_{\delta,v}(n) = \max \left\{ h_v(n), \max_{p \in A(n)} (Q_{v,\delta}W_{\delta,v})(n, p) \right\}.$$

At any in-band date before crossing the prior, the lowest signal induces rejection. A block of K_{fr} such draws therefore reaches the frozen region. Under any experimentation rule, let K count experiments before stopping, crossing the prior, or reaching that region. For every $k \geq 0$,

$$\mathbb{P}_v(K > k) \leq [1 - f_v(1)^{K_{\text{fr}}}]^{\lfloor k/K_{\text{fr}} \rfloor}. \quad (55)$$

Starting from $U_v^{(0)} = h_v$, iterate the right-hand side of (28). The k th iterate allows at most k experiments before stopping. The iterates increase and are bounded by $1 - c$. The tail bound (55) gives uniform convergence to a solution U_v ; unrolling any two bounded solutions for k experiments gives uniqueness by the same bound. Each iterate is Borel measurable, so U_v is as well.

For a fixed k and initial belief, the experiment tree is finite. Its current profits vanish after multiplication by $1 - \delta$, its discount factors converge to one, and its continuation values above the prior converge uniformly to $1 - c$ by Proposition 5. Backward induction therefore gives convergence of the k -experiment normalized value to $U_v^{(k)}$. Allowing additional experiments changes either value by at most $(1 - c)$ times the right-hand side of (55). Taking $\delta \uparrow 1$ and then $k \rightarrow \infty$ proves (29). $\blacksquare$

Proof of Proposition 8. Write $W_\delta(r, n) := (1 - \delta)V_\delta(r, n)$ and $W_{\delta,v}(n) := W_\delta(v, n)$. Revealing v to the seller can only increase value, so

$$W_\delta(r, n) \leq rW_{\delta,1}(n) + (1 - r)W_{\delta,0}(n).$$

For the reverse limiting inequality in part (i), fix $n \in (0, \mu_0)$ and repeatedly post $q(n) := \min\{\beta_M(n), \underline{p}\}$. Since $\beta_1(n) < q(n) \leq \beta_M(n)$, this off-band price screens buyers while leaving n fixed. Conditional on v , purchases are independent Bernoulli observations with distinct state-dependent probabilities. After $T_\delta := \lceil (1 - \delta)^{-1/2} \rceil$ observations, let $\hat{v}_\delta$ be a likelihood-ratio-test estimate, with $\varepsilon_{\delta,v} := \mathbb{P}_v(\hat{v}_\delta \neq v) \rightarrow 0$. Then $(1 - \delta)T_\delta \rightarrow 0$ and $\delta^{T_\delta} \rightarrow 1$.

After estimating state w , follow an optimal policy for the known-state- w problem. On a correct estimate its normalized continuation payoff is $W_{\delta,v}(n)$; on an incorrect estimate it is at least $-c$. Since $0 \leq W_{\delta,v}(n) \leq 1 - c$, the conditional expected continuation payoff is at least $W_{\delta,v}(n) - \varepsilon_{\delta,v}$. Hence

$$W_\delta(r, n) \geq -c(1 - \delta)T_\delta + \delta^{T_\delta} \left[rW_{\delta,1}(n) + (1 - r)W_{\delta,0}(n) - \max_v \varepsilon_{\delta,v} \right].$$

Combining this inequality with the upper bound and Lemma 6 proves part (i).

At $n = \mu_0$, every off-band price leaves both beliefs unchanged, and the best stationary flow is $(\underline{p} - c)^+$. This self-loop can again be replaced by stopping. By Proposition 1, each in-band choice can be represented by $p = \beta_{s+1}(\mu_0)$. Its normalized current profit is $(1 - \delta)(\beta_{s+1}(\mu_0) - c)\bar{F}_r(s)$, which vanishes in the limit. A sale leads to $n_s^+ := \Lambda(\mathcal{L}(\mu_0) + \ell_s) > \mu_0$ and $r_s^+ := r\bar{F}_1(s)/\bar{F}_r(s)$; a rejection leads to (r_s^-, n_s^-) with $n_s^- < \mu_0$. Proposition 5 and part (i) therefore give the limiting experimental payoff

$$(1 - c)\bar{F}_r(s) + F_r(s)U(r_s^-, n_s^-) = (1 - c) - F_r(s)[(1 - c) - U(r_s^-, n_s^-)].$$

The maximum is over finitely many bands, so the limit passes through it, proving part (ii). Finally, Bayes' rule gives

$$F_r(s)U(r_s^-, n_s^-) = rF_1(s)U_1(n_s^-) + (1 - r)F_0(s)U_0(n_s^-),$$

which also establishes that each band's limiting payoff is affine in r . ■

Proof of Proposition 9. For $n \in (\underline{n}, \mu_0]$, let

$$\mathcal{A}_v^*(n) := \{\text{stop} : h_v(n) = U_v(n)\} \cup \{p \in A(n) : (Q_v U_v)(n, p) = U_v(n)\}, \quad (56)$$

and set $\mathcal{A}_v^*(n) = \{\text{stop}\}$ in the frozen region. The iterates defining U_v in the proof of Lemma 6 are Borel measurable and converge uniformly. Thus U_v is Borel measurable. Since $A(n)$ consists of a finite family of Borel candidate-price functions, a Borel measurable rule $\sigma_v(n) \in \mathcal{A}_v^*(n)$ exists.

Let K be the number of experiments under this rule before crossing above the prior or stopping. The bound (55) applies, so K has a uniformly bounded mean. Iterating the equality for the chosen action in (28) shows that its expected terminal payoff is $U_v(n)$, where the terminal payoff is $1 - c$ upon crossing and h_v upon stopping. Hence the rule is optimal for the limiting known-state problem.

To compute its crossing probability, extend any $\phi : (0, \mu_0] \rightarrow [0, 1]$ by

$$\tilde{\phi}(x) := \begin{cases} 1, & x > \mu_0, \\ \phi(x), & x \leq \mu_0. \end{cases} \quad (57)$$

Conditioning on the first action gives

$$\phi_v^\sigma(n) = \begin{cases} 0, & \sigma_v(n) = \text{stop}, \\ d_v(p, n) \tilde{\phi}_v^\sigma(N(n, p, 1)) \\ \quad + [1 - d_v(p, n)] \tilde{\phi}_v^\sigma(N(n, p, 0)), & \sigma_v(n) = p. \end{cases} \quad (58)$$

If ϕ and ϕ' are two bounded solutions, iterating for k experiments gives

$$|\phi(n) - \phi'(n)| \leq \|\phi - \phi'\|_\infty \mathbb{P}_v(K > k).$$

The right-hand side vanishes uniformly by (55), proving uniqueness.

Now start from (r, μ_0) with price $\beta_{s+1}(\mu_0)$. An initial sale crosses above the prior with probability $\bar{F}_v(s)$. After a rejection, use the off-band learning phase from the proof of Proposition 8 at n_s^- , then follow $\sigma_{\hat{v}_\delta}$, where $\hat{v}_\delta$ is the estimated state. When the rule stops, repeat an off-band price attaining $h_{\hat{v}_\delta}$, or exit if its zero branch binds. After any crossing, use the pump-and-harvest policy from the proof of Proposition 5.

The probability of estimating v incorrectly vanishes. The learning phase has length T_δ with $(1 - \delta)T_\delta \rightarrow 0$, and the remaining number of experiments has bounded mean by (55). Their normalized profit contributions therefore vanish, and the expected discount factor applied to terminal payoffs converges to one. Uniform convergence above the prior then gives the state- v limiting payoff

$$(1 - c)\bar{F}_v(s) + F_v(s)U_v(n_s^-),$$

and the conditional crossing probability converges to (32). Taking the r -weighted average of these payoffs and applying Bayes' rule gives the band- s term in (31). Since s maximizes that expression, the constructed policies attain the limiting optimum. $\blacksquare$

Proof of Proposition 10. Let $\Pi_{v,\delta}$ denote normalized seller profit conditional on state v under π_δ . Each lies weakly below $1 - c$, and $r\Pi_{1,\delta} + (1 - r)\Pi_{0,\delta} \rightarrow 1 - c$; because $r \in (0, 1)$, both conditional components converge to $1 - c$, proving (i). Let

$$Q_{0,\delta}^a := (1 - \delta)\mathbb{E}\left[\sum_t \delta^t a_t \mid v = 0\right], \quad Q_{0,\delta}^p := (1 - \delta)\mathbb{E}\left[\sum_t \delta^t p_t a_t \mid v = 0\right]$$

denote normalized trade frequency and buyer payments. Since $(p_t - c)a_t \leq (1 - c)a_t$, we have $\Pi_{0,\delta} \leq (1 - c)Q_{0,\delta}^a$, so $Q_{0,\delta}^a \rightarrow 1$ while $Q_{0,\delta}^a \leq 1$, proving (ii). The identity $Q_{0,\delta}^p = \Pi_{0,\delta} + cQ_{0,\delta}^a$ gives $Q_{0,\delta}^p \rightarrow 1$; conditional on $v = 0$ buyers receive no value, so their normalized payoff is $-Q_{0,\delta}^p \rightarrow -1$, proving (iii); and normalized total surplus is $-cQ_{0,\delta}^a \rightarrow -c$, proving (iv).

For the final claim, $\mathbb{E}[p_{T_\delta} a_{T_\delta} \mid v = 0] = Q_{0,\delta}^p \rightarrow 1$. Since $p_t a_t \leq 1$,

$$1 - \mathbb{E}[p_{T_\delta} a_{T_\delta} \mid v = 0] \geq \varepsilon \mathbb{P}(p_{T_\delta} a_{T_\delta} \leq 1 - \varepsilon \mid v = 0),$$

which proves (33). On a sale, $p_t \leq \beta_M(n_t)$ by (4), so $p_{T_\delta} \geq 1 - \varepsilon$ forces $\beta_M(n_{T_\delta}) \geq 1 - \varepsilon$, that is, $n_{T_\delta} \geq \Lambda(\mathcal{L}(1 - \varepsilon) - \lambda_M)$. $\blacksquare$

Proofs for Section 5

The proof of Proposition 11 uses the following solution of the restricted pump-and-harvest problem.

Lemma A.3. The function $\bar{\delta}_H$ is strictly increasing from zero to one on $[x_2, 1)$. Hence $\bar{\delta}_H(n^*(\delta)) = \delta$ has a unique solution in $(x_2, 1)$. An optimal stationary policy pumps below $n^*(\delta)$ and harvests above it, with both actions optimal at equality. Choosing harvesting at equality gives (37).

Proof. Direct differentiation gives

$$\bar{\delta}'_H(n) = \frac{(1 - \theta^{-1})[x_1(1 - n)^2 + \theta^{-1}(1 - x_1)n^2]}{(n - x_1)^2[1 - (1 - \theta^{-1})n]^2} > 0.$$

Moreover, $\bar{\delta}_H(x_2) = 0$ and $\bar{\delta}_H(n) \rightarrow 1$ as $n \uparrow 1$.

The restricted discounted problem has a deterministic stationary optimum by the contraction argument of Proposition 1. At $n > x_2$, harvesting under such a policy continues forever, since it leaves the state fixed. Pumping forever earns $x_1/(1 - \delta)$, strictly less than pumping once and then harvesting. Thus an optimal stationary policy makes finitely many pumps before harvesting.

The value of K pumps followed by harvesting is

$$\Pi_K(n) = \frac{x_1(1 - \delta^K) + \delta^K b(g^K(n))}{1 - \delta},$$

where $K \geq 1$ is required at $n = x_2$. Since $b(g(x)) = x$,

$$\Pi_{K+1}(n) - \Pi_K(n) = \frac{\delta^K [g^K(n) - x_1]}{1 - \delta} [\delta - \bar{\delta}_H(g^K(n))].$$

The bracket decreases strictly with K . The smallest maximizing pump count is therefore $K^*(n, \delta)$, which is at least one at x_2 . This proves the policy and value formula. If a reachable belief equals $n^*(\delta)$, harvesting there or pumping once more gives the same value. ■

Proof of Proposition 11. By Lemma 3, $\mathcal{H} := [0, 1] \times [x_2, 1)$ is absorbing. For $n > x_2$, Proposition 1 reduces the price choice to pumping at x_1 , harvesting at $b(n)$, or screening at $g(n)$. At x_2 , $b(x_2) = x_1$ is a pump, and off-band pooling at x_{-1} must also be considered.

The restricted value is feasible. Pumping once and then harvesting implies

$$(1 - \delta)V^{PH}(n) \geq (1 - \delta)x_1 + \delta n \geq (1 - \delta)x_1 + \delta x_2 \geq \alpha.$$

Screening leaves n fixed and earns $g(n)\rho(r) \leq \alpha g(n) < \alpha$. Thus, evaluated at the continuation value V^{PH} , its Bellman payoff satisfies

$$g(n)\rho(r) + \delta V^{PH}(n) < V^{PH}(n).$$

The restricted value also satisfies $(1 - \delta)V^{PH}(n) \geq b(n)$, so off-band pooling at x_2 and prices rejected surely cannot improve it. Consequently, V^{PH} , viewed as a function independent of r , satisfies the full Bellman equation on $\mathcal{H}$. Contraction gives $V_\delta = V^{PH}$ there and the strict inequality excludes screening. Lemma A.3 supplies the threshold policy. Its monotonicity in δ follows from the strict increase of $\bar{\delta}_H$. ■

Proof of Lemma 7. At $j \leq -2$, Proposition 1 reduces the relevant prices to pooling at p_j^P , screening at p_j^S , and an in-band price that is rejected surely and moves j to $j - 1$. An off-band rejection leaves both beliefs fixed and is dominated by pooling.

For $j \leq -3$, compare any policy using these actions at (r, x_j) with the policy at (r, x_{j+1}) that chooses the same action type after each sequence of outcomes. Couple the private signals. Purchase decisions and the seller's posterior are identical on the two paths, while the second naïve-belief index remains one step higher. Pooling and screening prices are therefore higher on that path, and in-band rejections yield zero on both. Taking suprema proves $W_j(r) \leq W_{j+1}(r)$.

An in-band price at $j \leq -2$ consequently yields $\delta W_{j-1}(r) \leq \delta W_j(r) < W_j(r)$, where strictness uses $W_j(r) \geq p_j^P/(1 - \delta) > 0$. ■

Proof of Proposition 12. Fix $j \leq -2$ and abbreviate $\Delta_v := \Delta_v(j)$. By Lemma 7, the Bellman equation is

$$W_j(r) = \max \{ p_j^P + \delta W_j(r), p_j^S \rho(r) + \delta [\rho(r) W_j(g(r)) + (1 - \rho(r)) W_j(b(r))] \}.$$

Revealing quality gives the upper bound

$$W_j(r) \leq \frac{r \max\{p_j^P, \alpha p_j^S\} + (1 - r) \max\{p_j^P, q p_j^S\}}{1 - \delta}.$$

Pooling forever attains it in case (i), and screening forever attains it in case (ii). Their values are respectively $p_j^P/(1 - \delta)$ and $p_j^S \rho(r)/(1 - \delta)$.

For case (iii), let τ_m be the number of screening periods needed for the posterior's log odds to fall by $m\ell$ from its starting point, with $\delta^\infty := 0$. Conditional on v , first-step analysis and the strong Markov property give

$$\mathbb{E}_v[\delta^{\tau_m}] = z_v^m, \quad z_v := \frac{1 - \sqrt{1 - 4\delta^2 \alpha q}}{2\delta d_v}.$$

Indeed, the one-step expectation is the unique root in $(0, 1)$ of $z = \delta[(1 - d_v) + d_v z^2]$, and successive downward passages multiply their discount factors. The roots satisfy $0 < z_1 < z_0 < 1$ and $z_0 = \theta z_1$.

Define

$$r_j^* := \frac{-\Delta_0(1 - z_0)}{\Delta_1(1 - z_1) - \Delta_0(1 - z_0)}. \quad (59)$$

This is the unique zero of the strictly increasing function $B(r) := r\Delta_1(1 - z_1) + (1 - r)\Delta_0(1 - z_0)$. Also, $B(\hat{r}_j) = (1 - \hat{r}_j)(-\Delta_0)(z_0 - z_1) > 0$, so $r_j^* < \hat{r}_j$.

Consider the policy that screens for $r > r_j^*$ and otherwise pools forever. Write $\underline{W} := p_j^P/(1 - \delta)$. For $r \leq r_j^*$ its value is $\underline{W}$; for $r_j^* < r < 1$, let $m(r) := \lceil [\mathcal{L}(r) - \mathcal{L}(r_j^*)]/\ell \rceil$. The expected number of discounted screening periods in state v is $(1 - z_v^{m(r)})/(1 - \delta)$, giving the policy value

$$W_j^*(r) = \frac{p_j^P + r\Delta_1(1 - z_1^{m(r)}) + (1 - r)\Delta_0(1 - z_0^{m(r)})}{1 - \delta}. \quad (60)$$

At $r = 1$ the value is $\alpha p_j^S/(1 - \delta)$. We verify this bounded candidate in the Bellman equation.

If $r \leq r_j^*$ and $g(r) \leq r_j^*$, screening's excess value over pooling is $r\Delta_1 + (1-r)\Delta_0 < 0$. If $r \leq r_j^* < g(r)$, a favorable update puts the posterior one step above its stopping point. Let $S(r)$ be the screening Bellman payoff evaluated at W_j^* . Bayes' rule and the root equations give

$$(1-\delta)[S(r) - \underline{W}] = \frac{\delta q}{z_1} B(r) \leq 0,$$

with equality exactly at $r = r_j^*$.

For $r_j^* < r < 1$, screening reproduces the candidate by policy evaluation. To compare it with pooling, write $\Omega(r) := r/(1-r)$ and $m := m(r)$. Then $\Omega(r) > \Omega(r_j^*)\theta^{m-1}$ and

$$\frac{1-z_0^m}{1-z_1^m} = \frac{1-z_0}{1-z_1} \frac{\sum_{k=0}^{m-1} \theta^k z_1^k}{\sum_{k=0}^{m-1} z_1^k} \leq \theta^{m-1} \frac{1-z_0}{1-z_1}.$$

Since $\Omega(r_j^*) = -\Delta_0(1-z_0)/[\Delta_1(1-z_1)]$, these inequalities imply

$$(1-\delta)W_j^*(r) - p_j^P = r\Delta_1(1-z_1^m) + (1-r)\Delta_0(1-z_0^m) > 0.$$

Thus screening strictly dominates pooling above the cutoff. The endpoint posteriors follow directly from the known-state comparison. Contraction now gives $W_j = W_j^*$ and the stated optimal actions. On the prior-generated lattice the corresponding index is $i_j^* = 1 + \lfloor [\mathcal{L}(r_j^*) - \mathcal{L}(\mu_0)]/\ell \rfloor \in \mathbb{Z}$. ■

Proof of Corollary 1. The preceding proof establishes $r_j^* < \hat{r}_j$. Rationalizing the root gives

$$z_v = \frac{2\delta\alpha q}{d_v(1 + \sqrt{1 - 4\delta^2\alpha q})}.$$

As $\delta \downarrow 0$, both roots tend to zero, so (59) converges to $\hat{r}_j$. As $\delta \uparrow 1$, $z_1 \rightarrow q/\alpha < 1$ and $z_0 \rightarrow 1$. The numerator in (59) then tends to zero and its denominator to $\Delta_1(1 - q/\alpha) > 0$, proving the second limit. ■

Proof of Proposition 13. Let $\omega_j := \mu_0\theta^j/(1-\mu_0)$. Then

$$\frac{x_{j-1}}{x_{j+1}} = \theta^{-2} \frac{1 + \omega_{j+1}}{1 + \omega_{j-1}}$$

is increasing in j . Hence (43) implies $x_{j-1} \leq qx_{j+1}$ for every $j \leq -2$. Proposition 12(ii) then makes screening forever optimal in both known states throughout the frozen region. Rearranging (43) also gives $\omega_0 \leq (\alpha^2 - q)/q^2$, which is equivalent to (44); at $\mu_0 = 1/2$ it reduces to $\alpha \geq 2/3$. The derivative $(\mu^{\text{fr}})'(\alpha) = (1-\alpha)(3\alpha-1)/[\alpha^2(2\alpha-1)^2] > 0$ gives its stated monotonicity.

Proposition 11 supplies the common policy for $j \geq 2$. Fix $d \in \{q, \alpha\}$. At the remaining beliefs, the proposed policy has value J^1 at $j = 1$ and satisfies (41)–(42). Solving these equations gives

$$J_d^0 = \frac{dx_1 + \delta dJ^1 + \delta d(1-d)x_0 + \delta^2(1-d)^2 J_d^F}{1 - \delta^2 d(1-d)}.$$

The denominator is positive, so these equations uniquely determine the candidate values.

At $j = 1$, pooling at x_{-1} and off-band screening at x_2 leave the known-state problem unchanged. Their Bellman inequalities are therefore equivalent to $J^1 \geq x_{-1}/(1-\delta)$ and $J^1 \geq$

$dx_2/(1 - \delta)$. In-band screening at x_1 gives the third comparison in (45). At $j = 0$, the alternative is pooling at x_{-1} , yielding (46). At $j = -1$, pooling at x_{-2} and off-band screening at x_{-1} give (47). An in-band sure rejection yields $\delta J_d^F < J_d^F \leq J_d^{-1}$ and cannot improve the value. Thus the proposed bounded values satisfy each known-state Bellman equation. The policy, which depends only on n , attains (39) and is optimal for every r .

Strict inequalities imply unique prices at $j \leq 1$. If $n^*(\delta)$ is not a lattice point, the high-belief prices are also unique. Continuity of the finite comparisons and local constancy of the high-belief stopping index give the neighborhood claim.

Finally, from the prior a purchase leads to pumping with probability d_v . A rejection followed by a recovery purchase returns n to the prior with probability $(1 - d_v)d_v$; two rejections lead to the frozen region. Consequently,

$$\varpi_v = d_v + (1 - d_v)d_v\varpi_v,$$

which gives (48). The probability of repeating the recovery cycle indefinitely is zero because $d_v(1 - d_v) < 1$. ■

Verification of the example in Section 5.3. For $(\mu_0, \alpha) = (1/2, 7/10)$, inequality (43) is strict because $\alpha > 2/3$. Also $x_1 = \alpha$ and $x_2 > x_1$, so (36) is strict for every $\delta > 0$.

Write $U := (1 - \delta)V^{PH}(x_2)$ and $\hat{J} := (1 - \delta)J$ for each candidate value. Pumping twice from x_2 and then harvesting at $x_3 = 343/370$ gives, for $\delta \geq 4/5$,

$$U \geq x_1(1 - \delta^2) + \delta^2 x_3 \geq \frac{7819}{9250} =: U_0. \quad (61)$$

Each normalized gap in (45)–(47) is increasing in U . For the in-band screening comparison at $j = 1$, its derivative is

$$\frac{\delta(1 - d)[1 - \delta^2 d(2 - d)]}{1 - \delta^2 d(1 - d)} > 0;$$

the other derivatives are immediately positive from (41)–(42) and $\hat{J}^1 = (1 - \delta)x_0 + \delta U$. It therefore suffices to set $U = U_0$.

Put $t := 5\delta - 4 \in [0, 1)$. Substitution into the ten comparisons other than in-band screening at $j = 1$, followed by clearing positive denominators, gives polynomials in t with strictly positive coefficients. The two remaining gaps have common positive denominator $925000(2164 - 168t - 21t^2)$ and numerators

$$\begin{aligned} 709685496 - 79621628t - 29246532t^2 - 1724961t^3, & \quad d = 3/10, \\ 98408984 - 35990012t - 20942628t^2 - 1809969t^3, & \quad d = 7/10. \end{aligned}$$

Both decrease on $[0, 1]$ and at $t = 1$ equal 599092375 and 39666375, respectively. All comparisons are therefore strict throughout $[4/5, 1)$. Continuity gives the stated neighborhood property for the intermediate- and low-belief actions.

At $\delta = 9/10$, the high-belief indifference equation is $360n^2 - 302n - 49 = 0$. Its positive root lies strictly between $x_4 = 2401/2482$ and $x_5 = 16807/17050$. Thus the seller pumps at $j = 2, 3, 4$ and then harvests at x_4 . Substitution into (37) and (40)–(42) gives

$$J^1 = \frac{53078699}{6205000}, \quad V_\delta(1/2, 1/2) = \frac{282267428}{51495295} \approx 5.4814.$$

Finally, (48) with $d_0 = 3/10$ and $d_1 = 7/10$, weighted by the common prior, gives the reported outcome probabilities.

On frozen paths, subsequent purchases reveal independent signals with state-conditional probabilities α and q . The law of large numbers therefore implies that the seller's posterior converges almost surely to v , while buyers' public belief remains x_{-2} .

For every fixed $\delta \in [4/5, 1)$, the expected long-run per-period surplus is

$$\lim_{t \rightarrow \infty} \mathbb{E}[(v - c)a_t] = \frac{1}{2}[\varpi_1 + (1 - \varpi_1)\alpha] = \frac{763}{1580}.$$

Efficient trade yields $1/2$, so the difference is $27/1580$. ■

Proofs for Section 6

Proof of Proposition 14. For part (i), since $c < \bar{p}$, $p^P(n) < c$ if and only if $\beta_1(n) < c$, which is equivalent to $n < n_c^P$. Each pump raises log odds by at least $\ell_1 > 0$ by Proposition 3, so repeated pricing at $p^P(n_t)$ reaches n_c^P after finitely many periods.

For the claim about the sharp coefficient, use the construction in Proposition 4 with $N(\delta) = \lceil \lambda_M^{-1} \log[1/(1 - \delta)] \rceil$. The policy calculation in its proof bounds the shortfall by $\chi \log[1/(1 - \delta)] + C(n_0)$. The policy-independent lower bound in Proposition 6(ii) then shows that this construction attains the coefficient χ .

For part (ii), let $d := \mu_0 f_1(M) + (1 - \mu_0) f_0(M) > 0$. Starting from the common prior, post $\bar{p}$, exit after a rejection, and continue optimally after a sale. A sale reveals signal M and takes both beliefs to $\beta_M(\mu_0) = \bar{p} > \mu_0$. Hence

$$V_\delta(\mu_0, \mu_0) \geq d[\bar{p} - c + \delta V_\delta(\bar{p}, \bar{p})].$$

By Proposition 5, the right-hand side is positive for all sufficiently high $\delta < 1$.

At the common prior, every off-band price leaves both beliefs unchanged and yields nonpositive current profit when $c > \bar{p}$. Its Bellman payoff is therefore at most $\delta V_\delta(\mu_0, \mu_0) < V_\delta(\mu_0, \mu_0)$. Every optimal initial price must consequently lie in I . Such a price is strictly below cost and induces a sale with positive probability, giving strictly negative expected current profit. ■

Proof of Proposition 15. The shortfall identity (52) does not depend on the off-band completion, so it applies under Assumption 2. In period zero, any realized sale satisfies $p_0 \leq \beta_M(n)$ by (4), so flow profit is at most $\max\{0, \beta_M(n) - c\}$ and the period-zero shortfall is at least $1 - \max\{c, \beta_M(n)\}$; all shortfalls are nonnegative, which gives the lower bound.

For the upper bound, set $T := K_0(n) + 1$ and pump in band for T periods. In-band updating is unchanged, so Proposition 3 gives $n_T > n^{\text{cap}}$. Thereafter post $\beta_1(n_t) > \bar{p}$ in every period. Each sale is certain and, under Assumption 2, adds λ_M to buyers' log odds. Thus, for every $k \geq 0$, $\mathcal{L}(n_{T+k}) = \mathcal{L}(n_T) + k\lambda_M$. Using $\mathcal{L}(n^{\text{cap}}) + \lambda_1 = \mathcal{L}(\mu_0) + \lambda_M$,

$$1 - \beta_1(n_{T+k}) \leq e^{-[\mathcal{L}(n_T) + \lambda_1]} e^{-k\lambda_M} \leq e^{-[\mathcal{L}(\mu_0) + \lambda_M]} e^{-k\lambda_M}.$$

Each of the first T periods has shortfall at most one. Bounding all discount factors by one and summing the geometric bound for the remaining periods gives the right-hand side of (51). ■

References

- ANGRISANI, M., A. GUARINO, P. JEHIEL, AND T. KITAGAWA (2021): “Information Redundancy Neglect versus Overconfidence: A Social Learning Experiment,” *American Economic Journal: Microeconomics*, 13, 163–197.
- ARIELI, I., M. KOREN, AND R. SMORODINSKY (2022): “The Implications of Pricing on Social Learning,” *Theoretical Economics*, 17, 1761–1802.
- EVERY, C. AND P. ZEMSKY (1998): “Multidimensional Uncertainty and Herd Behavior in Financial Markets,” *American Economic Review*, 88, 724–748.
- BANERJEE, A. (1992): “A Simple Model of Herd Behavior,” *Quarterly Journal of Economics*, 107, 797–817.
- BASTIANELLO, F. AND P. FONTANIER (2025): “Expectations and Learning from Prices,” *Review of Economic Studies*, 92, 1341–1374.
- 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.
- BOHREN, J. A. AND D. N. HAUSER (2021): “Learning with Heterogeneous Misspecified Models: Characterization and Robustness,” *Econometrica*, 89, 3025–3077.
- BOSE, S., G. OROSEL, M. OTTAVIANI, AND L. VESTERLUND (2006): “Dynamic Monopoly Pricing and Herding,” *RAND Journal of Economics*, 37, 910–928.
- (2008): “Monopoly Pricing in the Binary Herding Model,” *Economic Theory*, 37, 203–241.
- CAI, H., Y. CHEN, AND H. FANG (2009): “Observational Learning: Evidence from a Randomized Natural Field Experiment,” *American Economic Review*, 99, 864–882.
- CAMINAL, R. AND X. VIVES (1996): “Why Market Shares Matter: An Information-Based Theory,” *RAND Journal of Economics*, 27, 221–239.
- (1999): “Price Dynamics and Consumer Learning,” *Journal of Economics & Management Strategy*, 8, 95–131.
- CARNEHL, C., A. STENZEL, AND P. SCHMIDT (2024): “Pricing for the Stars: Dynamic Pricing in the Presence of Rating Systems,” *Management Science*, 70, 1755–1772.
- CHANDRASEKHAR, A. G., H. LARREGUY, AND J. P. XANDRI (2020): “Testing Models of Social Learning on Networks: Evidence from Two Experiments,” *Econometrica*, 88, 1–32.

- CHEN, C.-H., K.-P. CHEN, AND J. ISHIDA (2025): “Social Learning and Strategic Pricing with Rating Systems,” *American Economic Journal: Microeconomics*, 17, 147–180.
- CHEN, L. AND Y. PAPANASTASIOU (2021): “Seeding the Herd: Pricing and Welfare Effects of Social Learning Manipulation,” *Management Science*, 67, 6734–6750.
- CRAPIS, D., B. IFRACH, C. MAGLARAS, AND M. SCARSINI (2017): “Monopoly Pricing in the Presence of Social Learning,” *Management Science*, 63, 3586–3608.
- DASARATHA, K. AND K. HE (2020): “Network Structure and Naive Sequential Learning,” *Theoretical Economics*, 15, 415–444.
- DEMARZO, P. M., D. VAYANOS, AND J. ZWIEBEL (2003): “Persuasion Bias, Social Influence, and Unidimensional Opinions,” *Quarterly Journal of Economics*, 118, 909–968.
- ENKE, B. AND F. ZIMMERMANN (2019): “Correlation Neglect in Belief Formation,” *Review of Economic Studies*, 86, 313–332.
- EYSTER, E. AND M. RABIN (2005): “Cursed Equilibrium,” *Econometrica*, 73, 1623–1672.
- (2010): “Naïve Herding in Rich-Information Settings,” *American Economic Journal: Microeconomics*, 2, 221–243.
- (2014): “Extensive Imitation is Irrational and Harmful,” *Quarterly Journal of Economics*, 129, 1861–1898.
- EYSTER, E., M. RABIN, AND D. VAYANOS (2019): “Financial Markets Where Traders Neglect the Informational Content of Prices,” *Journal of Finance*, 74, 371–399.
- EYSTER, E., M. RABIN, AND G. WEIZSÄCKER (2018): “An Experiment on Social Mislearning,” Discussion Paper 73, CRC TRR 190 Rationality and Competition.
- FRICK, M., R. IJIMA, AND Y. ISHII (2020): “Misinterpreting Others and the Fragility of Social Learning,” *Econometrica*, 88, 2281–2328.
- GAGNON-BARTSCH, T. AND M. RABIN (2026): “Naïve Social Learning, Mislearning, and Unlearning,” Working paper.
- GAGNON-BARTSCH, T. AND A. ROSATO (2024): “Quality Is in the Eye of the Beholder: Taste Projection in Markets with Observational Learning,” *American Economic Review*, 114, 3746–3787.
- GOLUB, B. AND M. O. JACKSON (2010): “Naive Learning in Social Networks and the Wisdom of Crowds,” *American Economic Journal: Microeconomics*, 2, 112–149.
- GUARINO, A. AND P. JEHIEL (2013): “Social Learning with Coarse Inference,” *American Economic Journal: Microeconomics*, 5, 147–74.
- HEIDHUES, P. AND B. KŐSZEGI (2018): “Behavioral Industrial Organization,” in *Handbook of Behavioral Economics: Applications and Foundations I*, ed. by B. D. Bernheim, S. DellaVigna, and D. Laibson, North-Holland, vol. 1, 517–612.

- IFRACH, B., C. MAGLARAS, M. SCARSINI, AND A. ZSELEVA (2019): “Bayesian Social Learning from Consumer Reviews,” *Operations Research*, 67, 1209–1221.
- KÜBLER, D. AND G. WEIZSÄCKER (2005): “Are Longer Cascades More Stable?” *Journal of the European Economic Association*, 3, 330–339.
- LUCA, M. AND G. ZERVAS (2016): “Fake It Till You Make It: Reputation, Competition, and Yelp Review Fraud,” *Management Science*, 62, 3412–3427.
- MAYZLIN, D. (2006): “Promotional Chat on the Internet,” *Marketing Science*, 25, 155–163.
- MAYZLIN, D., Y. DOVER, AND J. CHEVALIER (2014): “Promotional Reviews: An Empirical Investigation of Online Review Manipulation,” *American Economic Review*, 104, 2421–2455.
- MORETTI, E. (2011): “Social Learning and Peer Effects in Consumption: Evidence from Movie Sales,” *Review of Economic Studies*, 78, 356–393.
- MUELLER-FRANK, M. AND C. NERI (2021): “A General Analysis of Boundedly Rational Learning in Social Networks,” *Theoretical Economics*, 16, 317–357.
- MUROOKA, T. AND T. YAMASHITA (2025): “Optimal Trade Mechanisms with Adverse Selection and Inferential Naïvety,” *American Economic Journal: Microeconomics*, 17, 33–67.
- PAPANASTASIOU, Y. AND N. SAVVA (2017): “Dynamic Pricing in the Presence of Social Learning and Strategic Consumers,” *Management Science*, 63, 919–939.
- QIU, L. AND A. B. WHINSTON (2017): “Pricing Strategies under Behavioral Observational Learning in Social Networks,” *Production and Operations Management*, 26, 1249–1267.
- ROSATO, A. (2026): “A Winner’s Blessing? Inferential Naïvety in English Auctions with Interdependent Values,” Working paper.
- SHAPIRO, C. (1983): “Optimal Pricing of Experience Goods,” *Bell Journal of Economics*, 14, 497–507.
- SHIN, D., S. VACCARI, AND A. ZEEVI (2023): “Dynamic Pricing with Online Reviews,” *Management Science*, 69, 824–845.
- SMITH, L. AND P. SØRENSEN (2000): “Pathological Outcomes of Observational Learning,” *Econometrica*, 68, 371–398.
- SØRENSEN, A. T. (2007): “Bestseller Lists and Product Variety,” *Journal of Industrial Economics*, 55, 715–738.
- SPIEGLER, R. (2011): *Bounded Rationality and Industrial Organization*, Oxford: Oxford University Press.
- WEIZSÄCKER, G. (2010): “Do We Follow Others When We Should? A Simple Test of Rational Expectations,” *American Economic Review*, 100, 2340–2360.
- WELCH, I. (1992): “Sequential Sales, Learning, and Cascades,” *Journal of Finance*, 47, 695–732.