

DISCUSSION PAPER SERIES

No. 5003

DYNAMIC MONOPOLY PRICING WITH HERDING

Subir Bose, Gerhard O Orosel,
Marco Ottaviani and Lise Vesterlund

INDUSTRIAL ORGANIZATION



Centre for **E**conomic **P**olicy **R**esearch

www.cepr.org

Available online at:

www.cepr.org/pubs/dps/DP5003.asp

DYNAMIC MONOPOLY PRICING WITH HERDING

Subir Bose, University of Texas at Austin
Gerhard O Orosel, Universität Wien
Marco Ottaviani, London Business School (LBS) and CEPR
Lise Vesterlund, University of Pittsburg

Discussion Paper No. 5003
April 2005

Centre for Economic Policy Research
90–98 Goswell Rd, London EC1V 7RR, UK
Tel: (44 20) 7878 2900, Fax: (44 20) 7878 2999
Email: cepr@cepr.org, Website: www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in **INDUSTRIAL ORGANIZATION**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as a private educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions. Institutional (core) finance for the Centre has been provided through major grants from the Economic and Social Research Council, under which an ESRC Resource Centre operates within CEPR; the Esmée Fairbairn Charitable Trust; and the Bank of England. These organizations do not give prior review to the Centre's publications, nor do they necessarily endorse the views expressed therein.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Subir Bose, Gerhard O Orosel, Marco Ottaviani and Lise Vesterlund

CEPR Discussion Paper No. 5003

April 2005

ABSTRACT

Dynamic Monopoly Pricing with Herding*

This paper studies dynamic pricing by a monopolist selling to buyers who learn from each other's purchases. The price posted in each period serves to extract rent from the current buyer, as well as to control the amount of information transmitted to future buyers. As information increases future rent extraction, the monopolist has an incentive to subsidize learning by charging a price that results in information revelation. Nonetheless in the long run, the monopolist generally induces herding by either selling to all buyers or exiting the market.

JEL Classification: D83, L12 and L15

Keywords: herd behaviour, informational cascade, monopoly, public information and social learning

Subir Bose
Economics Department
University of Texas at Austin
Austin, Texas 78712
USA
Email: bose@eco.utexas.edu

Gerhard O Orosel
Department of Economics
Universität Vienna
Hohensraufengasse 9
A 1010 Wien
AUSTRIA
Tel: (43 1) 2912 8781
Fax: (43 1) 2912 8569
Email: gerhard.orosel@univie.ac.at

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=162377

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=106332

Marco Ottaviani
Department of Economics
London Business School
Sussex Place, Regent's Park
London
NW1 4SA
Tel: (44 20) 7646 0876
Email: mottaviani@london.edu

Lise Vesterlund
Dept of Economics
University of Pittsburgh
4925 WW Posvar Hall
Pittsburgh, PA 15260
USA
Tel: (1 412) 648 1794
Email: vester@pitt.edu

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=137297

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=162376

*This and the companion paper (Bose, Orosel, Ottaviani, and Vesterlund, 2005) generalize the model and the results contained in 'Monopoly Pricing with Social Learning' by Ottaviani (1996) and 'Optimal Pricing and Endogenous Herding' by Bose, Orosel, and Vesterlund (2001). Ottaviani (1996) gave a first formulation of the problem and derived implications for learning and welfare. Independently, Bose, Orosel, and Vesterlund (2001) formulated a similar model but with a different focus on the dependence of the solution on the model's parameters. Bose, Orosel, and Vesterlund's team has then joined forces with Ottaviani to produce this and the companion paper. While this paper analyses the general case with a finite number of signals, the companion paper focuses on the special case with symmetric binary signals.

Submitted 24 March 2005

1 Introduction

This paper studies optimal pricing by a monopolist in a market in which buyers learn from each other's purchases. We consider buyers who have limited first-hand information about the product's value, and who freely can observe the decisions made by other buyers in the past. For example, readers deciding which book to buy often consult the list of best sellers,¹ licensees of a new technology look at the behavior of other potential adopters, and employers assess job applicants on the basis of their employment history. In these situations, buyers decide without taking into account the value of the information revealed to the future buyers who observe these decisions. An *informational externality* is therefore present.

When buyers face *fixed prices*, this externality easily leads to pathological outcomes, such as herd behavior. As shown by Banerjee (1992) and Bikhchandani, Hirshleifer and Welch (1992), henceforth BHW, public information quickly swamps private information, and late buyers end up imitating the decisions of a few initial buyers who act alike. Since these late buyers behave in the same way regardless of their private information, nothing can be inferred from their behavior. These informational cascades result in the loss of a potentially large amount of private information collectively possessed by late buyers.

This paper investigates what happens when *prices* are instead *variable*, being set by a monopolist. The monopolist is a long-run player and can charge different prices over time, and thereby affect the process of information aggregation. We examine what happens when prices are allowed to adjust in response to the buyers' choices; how the monopolist influences the buyers' learning process; and finally, whether monopoly power improves or worsens herd behavior.

In our model, on the demand side of the market there is a sequence of potential buyers, one in each period. All buyers have unit demand and identical (ex post) valuations of the product. The common value of the good is either high or low and is unknown to buyers and seller. Each buyer observes a discrete private signal partially informative about the good's value, as well as the decisions made by all previous buyers. On the supply side, in each period the monopolist posts a price without observing the private signal of the current buyer, who then decides whether to accept or reject the offer. The seller and the subsequent buyers observe the public history of posted prices and purchase decisions, but do not directly observe the private signals on the basis of which these decisions were taken. We assume that the monopolist has no private information about the good's value and is learning alongside with the buyers.

In this setting, prices have two functions: *immediate rent extraction* and *information screening*. On the one hand, the current price allows the monopolist to extract rent from the current buyer. On the other hand, the inference of future buyers about the current buyer's signal de-

¹For a recent empirical analysis see Sorensen (2004).

depends on the price the current buyer is charged. That is, the monopolist also uses the price as a screening device, affecting how much of the current buyer's information is made publicly available, and thus, how much future buyers are willing to pay for the product.

As a building block to our analysis, we isolate the *immediate rent extraction* role of prices by considering static monopoly pricing with a partially informed buyer. In Section 4, we revisit the classic monopoly trade-off between price and quantity, for the case in which the buyer's willingness to pay depends on the realization of a discrete signal about the good's value. Since long-run outcomes depend on what happens when beliefs are *extreme*, we are specifically interested in the situation of extreme prior beliefs. We find that when the prior belief about the good's value is favorable enough, it is optimal for the monopolist to sell even to the buyer with the most unfavorable signal.² When the prior belief is instead sufficiently unfavorable, the monopolist either exits the market by setting a price so high that no buyer type purchases (if the low value of the good is below its cost); or sells to all buyer types (if the low value of the good is above its cost). In the borderline case in which the low value of the good exactly equals its cost, the seller never exits the market and may find it optimal, even at extremely unfavorable beliefs, to demand a price such that only buyers with sufficiently high signals purchase the good.

In Section 5, we analyze the effect of prices on *information screening* and expected future profits. We start by examining how prices affect the information available to future buyers. In each period, any given price induces a bi-partition of the set of signal realizations. Buyers purchase if and only if they observe a private signal above a certain threshold. If the price is so low that all buyer types purchase, or so high that no type purchases, no information is revealed publicly, as in an informational cascade. Intermediate prices instead allow future buyers to infer some information, and result in more information than extreme prices.³ The question of interest is whether the monopolist benefits from charging an information-revealing price. To answer this, we first show that the monopolist's static profit is convex in the prior belief. Since revelation of public information corresponds to a mean preserving spread in the belief distribution, it increases expected profits. Hence, the monopolist's *expected* future profits are maximized at a price which reveals information to future buyers. Of course, the monopolist risks that *ex post* the information turns out to be unfavorable. However, *ex ante* this loss is more than compensated by the gain realized if the information is favorable. Being a long-run player, the monopolist has an incentive to partially internalize the informational externality.

The dynamically optimal prices depend on the interplay of immediate rent extraction and information screening (Section 6). The price that maximizes the expected *present period's profit*

²To understand what happens in this limit, note that exclusion of the buyer with the most unfavorable signal would result in an infinitesimal increase in the price, but in a discrete reduction in quantity.

³But in general, the amount of information revealed by different prices cannot be ranked in the sense of Blackwell.

generally does not also maximize expected *future profits*. The monopolist may therefore sacrifice immediate rent extraction to ensure that information is revealed to future buyers. The resolution of this trade-off naturally depends on the discount factor and the public belief about quality. For any fixed public prior belief, a sufficiently patient seller is exclusively interested in information screening and so demands a price that reveals information. However, given any discount rate, immediate rent extraction dominates when the belief is sufficiently extreme. The reason is that the value of information screening goes to zero when the belief that the value is high is either sufficiently optimistic or pessimistic. This is because the posterior belief is very close to the prior belief when the prior is very high or very low. Since the value of information goes to zero, while its cost is bounded away from zero, the optimal price is the one that maximizes immediate rent extraction.

We then turn to the long-run predictions of the model, for a fixed discount factor.⁴ Provided that the cost does not equal the low value of the good, the monopolist eventually stops the learning process by inducing an informational cascade or exiting, thereby preventing asymptotic learning of the good’s true value. Instead, if the cost equals the low value of the good, herding is still triggered at high public beliefs, but it may be optimal to charge an informative price for all low beliefs. Only in this case and only if the value of the good is actually low, may buyers asymptotically learn the good’s true value. In all other cases, learning stops before the true value is learned.

The paper proceeds as follows. Section 2 discusses the related literature. Section 3 formulates the model, Section 4 analyzes immediate rent extraction, and Section 5 considers the effect of information screening on future rent extraction. Section 6 gives results for the general dynamic model, focusing on the long-term outcomes and consequences for information aggregation. Section 7 concludes by discussing the role played by competition for the efficiency of the process of information aggregation. The Appendix collects the proofs of all the results.

2 Related literature

Despite the pervasiveness of social influence on economic decisions, its implications for pricing in markets have been largely overlooked in economics until recently. Becker (1991) considered a competitive market in which each buyer’s demand for a good depends *directly* on the demand by other buyers. Our model provides a foundation for such dependence based on informational externalities. In our model, the payoff to a buyer depends on the decisions of others only *indirectly*, through the information revealed about the good’s value.

⁴Our companion paper (Bose, Orosel, Ottaviani, and Vesterlund, 2005), provides a full characterization of the seller’s optimal strategy, comprising the short as well as the long run, for the specific case where buyers have binary signals.

In the herding literature, other papers have analyzed the role of prices. Welch (1990) considered pricing by a monopolist (IPO issuer) in a market with a sequence of partially informed buyers (investors). In his model, the monopolist is constrained to choose a fixed price for all buyers and finds it optimal to induce an immediate informational cascade.⁵ In this paper, we instead characterize the optimal dynamic pricing strategy for the monopolist. We find that the monopolist benefits from inducing social learning in the market, and typically delays the occurrence of informational cascades. In the conclusion, we discuss why we obtain such sharply different predictions.

Avery and Zemsky (1995) have also introduced history-dependent prices in the herding model. In their sequential trading model, in each period a privately informed trader places an order on either side of the market. Prices are set competitively by market makers and so incorporate retrospectively the information revealed by past trades. In that setting, herd behavior is impossible when agents have unidimensional information. In our model instead, a monopolist sets prices taking into account the prospective effect of information to be revealed in the future.

The paper closest in spirit to ours is Caminal and Vives (1996).⁶ In their two-period model, there is a continuum of buyers with normally distributed private signals about the quality of two competing products.⁷ Second-period buyers infer quality from the observable quantity sold in the first period, but do not directly observe first-period prices. Our model is instead designed to study the effect of monopoly pricing on informational cascades, and so differs from theirs in a number of ways. First, in our model there is a single buyer in each period with a discrete signal about the binary value of the good, as in BHW.⁸ Second, we consider the case of monopoly rather than duopoly.⁹ Third, we assume that past prices are observed publicly in order to avoid the additional signal jamming effect that they identify.¹⁰

This paper also relates to the literature on learning and experimentation in markets, pioneered by Rothschild (1974) and further developed by Easley and Kiefer (1988) and Aghion, Bolton, Harris, and Jullien (1991), among others. As in those models, in our model the price charged by the monopolist affects what is learnt about the demand curve. But while learning in those models is one sided, in our model the seller as well as the buyers have the opportunity of becoming better

⁵Clearly, the monopolist obtains higher expected profits by conditioning the price charged on the history of purchases by previous buyers. Welch argues that a sufficiently risk averse monopolist would prefer to create a cascade immediately and forgo the gain in expected profits resulting from dynamic pricing.

⁶See also Neeman and Orosel (1999) and Taylor (1999) for models in which a seller of a *unique* object approaches a sequence of bidders.

⁷See also Caminal and Vives (1999) for a more general multi-period version of the model.

⁸As shown by Smith and Sørensen (2000), if instead buyers had unboundedly informative (e.g., normal) signals no cascades would occur even with fixed prices.

⁹See Moscarini and Ottaviani (2001) for an analysis of the static problem of duopoly competition with a privately informed buyer.

¹⁰See Bose, Orosel, and Vesterlund (2001) for an analysis of a version of this model with unobservable past prices.

informed about the quality of the good. The effect of current prices on future demand curves through learning is typically absent in those models, but plays a key role in our model.

Our analysis of the effect of endogenous pricing on the outcomes of bilateral learning is in the spirit of Bergemann and Välimäki (1996 and 2000). While they focused on situations in which ex-post information (such as experience) about product quality is revealed *publicly* over time, in our model buyers make purchase decisions on the basis of pre-existing (ex-ante) *private* information.¹¹ In Bergemann and Välimäki's models the information that is publicly revealed in each period depends on the good that is purchased by the current buyer, because experience is specific to the good. In our model instead, the information that is publicly revealed depends not only on the purchase decision, but also on the price at which the decision is taken. Because of this difference, our analysis of the price setting problem is considerably more involved.

3 Model

A risk-neutral monopolist (or *seller*) offers identical goods to a sequence of risk-neutral potential *buyers* with quasi-linear preferences. In each period $t \in \{1, 2, \dots\}$, a different potential buyer arrives to the market, indexed by the time of arrival. The action space for each buyer is $A = \{0, 1\}$. Action $a_t = 1$ indicates purchase of one unit of the good, and $a_t = 0$ no purchase.

The sequence of events in each period t is as follows. First, the seller and buyer t observe the purchase decisions taken by previous buyers, as well as the prices posted in the past. The price of period $\tau \in \{1, 2, \dots\}$ is denoted by p_τ , and the public history at time t is denoted by $h_{t-1} = (p_1, a_1, \dots, p_{t-1}, a_{t-1})$, with $h_0 = \emptyset$. Second, the seller makes a take-it-or-leave-it price offer p_t for a unit of the good to buyer t . Third, buyer t *privately* observes a signal about the good's value. Fourth, buyer t makes the purchase decision a_t . The set of all possible histories is denoted by \mathcal{H} .

The good's common value is either low or high, $v \in \{L, H\}$, with $0 \leq L < H$. Without loss of generality we choose the monetary unit such that $H - L = 1$. The good's value is unknown to the seller and the buyers. The initial prior belief that the value is high, $v = H$, is commonly known to be equal to $\lambda_1 \in (0, 1)$.

We denote the random signal of buyer t by $\tilde{s}_t \in \mathcal{S}$ and its realization by s_t (or s if time does not matter), where $\mathcal{S} = \{s^0, s^1, \dots, s^K\} \subset \mathbb{R}$ is the signal space, with $K \geq 1$ and finite.¹² The

¹¹See also Bar-Isaac (2003) for a related learning model with a long-run seller in which learning is incomplete when the seller does not have perfect information about own quality. We obtain a similar result in our model.

¹²The number of signal realizations, $K + 1$, may be infinite provided that (i) $\inf_{\sigma \in \mathcal{S}} \{\sigma > s^0\} > s^0$, and (ii) s^0 is not perfectly informative (i.e., $\alpha_0 > 0$ and $\beta_0 > 0$ in the notation below). Smith and Sørensen (2001) have shown that informational cascades cannot occur when signals are of unbounded informativeness. Since in our model an informational cascade implies a constant price, the analysis of Smith and Sørensen (2001) applies and the seller cannot trigger an informational cascade and herding.

conditional probabilities of the signals are denoted by

$$\begin{aligned}\alpha^k &\equiv \Pr(s^k | H) > 0 \\ \beta^k &\equiv \Pr(s^k | L) > 0\end{aligned}$$

for $k = 0, \dots, K$, with $\sum_{i=0}^K \alpha^i = \sum_{i=0}^K \beta^i = 1$. For convenience, we index the signal realizations in such a way that the likelihood ratio $\frac{\alpha^k}{\beta^k}$ is weakly increasing in k , so that the Monotone Likelihood Ratio Property (MLRP) is satisfied.¹³ For simplicity, we further assume that each signal realization s^k has a different likelihood ratio $\frac{\alpha^k}{\beta^k}$, which implies that $\frac{\alpha^k}{\beta^k}$ is strictly increasing in k .¹⁴ A signal s^k such that $\frac{\alpha^k}{\beta^k} < 1$ is bad news for $v = H$, while a signal with $\frac{\alpha^k}{\beta^k} > 1$ is good news (see e.g., Milgrom, 1981).

The payoff of buyer t when the value is v is $[v - p_t] a_t$. That is, a buyer that does not purchase the good receives a payoff of zero, whereas a buyer that purchases the good receives the payoff $v - p_t$. Buyer t then purchases the good if and only if its expected value $E[v | s_t, h_t]$ exceeds the price.¹⁵

The seller has a constant marginal cost c per unit sold, with $0 \leq c < H$.¹⁶ In the analysis, we distinguish three cases, depending on whether the cost c is equal to, below, or above L .¹⁷ The seller's payoff is the discounted sum of profits, $\sum_{t=1}^{\infty} \delta^{t-1} [p_t - c] a_t$, with discount factor $\delta \in [0, 1)$. In each period $t \in \{2, 3, \dots\}$, the seller knows the prices p_τ demanded from previous buyers, and the previous buyers' actions a_τ , $\tau \in \{1, \dots, t-1\}$. A pure strategy for the seller is a function $p : \mathcal{H} \rightarrow (L, L+1)$ that maps every history h_{t-1} into a price p_t , $t \in \{1, 2, \dots\}$. Since it is common knowledge that the seller can always sell at some price above the object's minimum value, L , and is unable to sell at a price at or above the maximum value, H , we restrict attention to prices $p_t \in (L, H) = (L, L+1)$.

The model constitutes a game between the seller and the buyers. The perfect Bayesian equilibrium (PBE) of this game can be derived directly from the seller's optimal strategy. The seller anticipates how the buyers react to her dynamic pricing strategy, and maximizes her expected payoff accordingly.

To each period t we can associate the public history $h_{t-1} = (p_1, a_1, \dots, p_{t-1}, a_{t-1})$ of past prices and actions, which determines the period's public prior belief on the value, $\lambda_t \equiv \Pr(v = H | h_{t-1})$.

¹³In this setting with two states the MLRP is then always satisfied, without loss of generality.

¹⁴This simplification is without loss of generality. Signal realizations with identical likelihood ratios lead to identical posterior beliefs and so can be collapsed into one signal.

¹⁵For technical reasons we assume the tie-breaking rule that a buyer purchases the good when indifferent between purchasing and not.

¹⁶The seller incurs the cost c only when the buyer accepts the offer, but not when the offer is rejected. This can be interpreted as production to order.

¹⁷Although the case $c = L$ may seem non-generic, there are plausible economic circumstances where it occurs. For example, in the case of a license for a patent the seller has zero marginal cost and the patent may be worthless, so that $c = L = 0$.

This belief is the state variable in the seller's dynamic optimization problem. The associated value function is denoted by $V(\lambda_t)$. When focusing on a single period, we drop the time subscript of λ_t and treat λ as the key parameter for comparative statics.

In each period, the seller has to decide whether to stay in the market, and if so what price to demand. When making this decision, the seller has to take into account all the future contingencies and associated decisions. Because information revelation depends on the price demanded, the seller must consider the effect of the price on immediate rent extraction as well as on information screening. Before resolving this trade-off, we investigate these two effects in isolation.

4 Immediate rent extraction

In this section, we set the stage for our dynamic model by characterizing the solution of the static problem of monopoly pricing with a partially-informed buyer. The purpose of this exercise is not only to examine immediate rent extraction, but also to show how monopoly profits depend on the prior belief about the good's value. This dependence is important in the presence of social learning, as the price affects future beliefs about the good's value and, therefore, the demand of future buyers.

Demand and profit functions

When receiving signal realization s^k , the buyer's posterior belief is equal to

$$\lambda^k(\lambda) \equiv \Pr(H | \lambda, s^k) = \frac{\alpha^k \lambda}{\alpha^k \lambda + \beta^k (1 - \lambda)} \quad (1)$$

by Bayes' rule. Using the normalization $H - L = 1$, the willingness to pay for the good is then equal to the conditional expected value

$$E(v | \lambda, s^k) = \lambda^k H + (1 - \lambda^k) L = L + \frac{\alpha^k \lambda}{\alpha^k \lambda + \beta^k (1 - \lambda)}.$$

The probability of selling at price p is equal to the probability that the buyer's willingness to pay is greater or equal to p , i.e., $\Pr(E(v | \lambda, s^k) \geq p)$. This is the monopolist's demand function.

The MLRP implies that $\Pr(H | \lambda, s^k)$ is monotonic in k , so that the buyer's willingness to pay is increasing in the signal. With this discrete signal structure, there are $K + 1$ steps in the demand function, each corresponding to the $K + 1$ posterior expected values. At any price strictly above $E(v | \lambda, s^K)$, the probability of selling is 0. At a price weakly lower than $p^k(\lambda) \equiv E(v | \lambda, s^k)$ but higher than $p^{k-1}(\lambda) \equiv E(v | \lambda, s^{k-1})$ the probability of selling is

$$\psi^k(\lambda) \equiv \Pr(s_t \geq s^k | \lambda) = \lambda \sum_{i=k}^K \alpha^i + (1 - \lambda) \sum_{i=k}^K \beta^i. \quad (2)$$

The *pooling price* $p^0(\lambda) \equiv E(v \mid \lambda, s^0)$ results in a sale with probability one.

Given this demand function, the monopolist who wishes to sell with positive probability has $K + 1$ potentially optimal prices, $p^k(\lambda)$ for $k = 0, \dots, K$. In addition, the monopolist can avoid selling by posting any price $p > p^K(\lambda)$. We denote any such *exit price* by $p^E(\lambda)$, with corresponding sale probability $\psi^E = 0$. Because of their critical role in the analysis, we single out the two extreme prices $p^0(\lambda)$ and $p^E(\lambda)$, and refer to all other prices as intermediate prices.

The expected static profit for the seller when posting price $p^k(\lambda)$ is then

$$\pi^k(\lambda) \equiv E \left[\left(p^k(\lambda) - c \right) a \mid \lambda \right] = \left[p^k(\lambda) - c \right] \psi^k(\lambda) \quad \text{for } k \in \{0, \dots, K\}$$

and $\pi^E(\lambda) \equiv E \left[\left(p^E(\lambda) - c \right) a \mid \lambda \right] = 0$ at price $p^E(\lambda)$. Thus, the static problem of monopoly pricing has been reduced to this simple discrete maximization problem

$$\max_{k \in \{0, \dots, K, E\}} \pi^k(\lambda). \quad (3)$$

Maximized static profit is denoted by $\pi(\lambda)$.

Optimal prices

We now characterize the solution to the static problem (3). In this setting, the buyer's demand function is generated from an arbitrary discrete signal structure with binary state. The probability of selling plays the role of quantity. We begin by revisiting the classic monopoly trade-off between price and quantity, and then focus on the solution for extreme prior beliefs.

To find the optimal policy, it is useful to consider how the expected static profit is affected by incremental changes in the probability of sale. In order to increase the probability of selling from $\psi^k(\lambda)$ to $\psi^{k-1}(\lambda)$, the monopolist needs to reduce the price from $p^k(\lambda)$ to $p^{k-1}(\lambda)$ so that the buyer with signal s^{k-1} prefers to purchase. The change in profit is then

$$\begin{aligned} & \pi^{k-1}(\lambda) - \pi^k(\lambda) \\ &= \left[p^{k-1}(\lambda) - c \right] \left[\psi^{k-1}(\lambda) - \psi^k(\lambda) \right] - \left[p^k(\lambda) - p^{k-1}(\lambda) \right] \psi^k(\lambda) \\ &= \left[L - c + \lambda^{k-1}(\lambda) \right] \left[\lambda \alpha^{k-1} + (1 - \lambda) \beta^{k-1} \right] \end{aligned} \quad (4)$$

$$- \left[\lambda^k(\lambda) - \lambda^{k-1}(\lambda) \right] \left[\lambda \sum_{i=k}^K \alpha^i + (1 - \lambda) \sum_{i=k}^K \beta^i \right]. \quad (5)$$

The first term (4) corresponds to the expected revenue acquired from sales to the marginal buyer type, net of the cost of provision. Notice that the mass of marginal buyers (i.e., the increase in the probability of selling) is always bounded away from zero for any fixed λ . The sign of this *quantity effect* is ambiguous in general. If the low value of the good is higher than the cost, $c < L$, this effect is positive because $p^k(\lambda) - c > 0$ for all λ and k . If instead $c > L$, these

increased sales add to profits when λ is large (since $p^k(1) - c = H - c > 0$ for any k), but detract from profits when λ is small (since $p^k(0) - c = L - c < 0$ for any k).

The second term (5) corresponds to the expected revenue lost on all the inframarginal buyer types and is equal to the difference in the willingness to pay of the marginal buyer type and the preceding type, multiplied by the total mass of inframarginal types. This *price effect* is always negative. Starting with $p^k(\lambda)$, the size of the reduction in price, $p^k(\lambda) - p^{k-1}(\lambda)$, necessary to induce the buyer with signal s^{k-1} to purchase depends on the level of the belief λ and the signal probabilities.

If the prior belief is extreme (either $\lambda \rightarrow 0$ or $\lambda \rightarrow 1$), the difference $\lambda^k(\lambda) - \lambda^{k-1}(\lambda)$ goes to zero, so that the price effect converges to zero. This is because the prior belief is not updated at all if the good's value is already known to be high or low. As a consequence, for *large* enough λ , we always have $\pi^{k-1}(\lambda) > \pi^k(\lambda)$, because the quantity effect is positive and bounded away from zero, while the price effect becomes negligible. In contrast, for *small* enough λ , the price effect again becomes negligible, but the sign of the quantity effect depends on the sign of $L - c$: if $L - c > 0$, the quantity effect is positive so that $\pi^{k-1}(\lambda) > \pi^k(\lambda)$; if instead $L - c < 0$, the quantity effect is negative so that $\pi^{k-1}(\lambda) < \pi^k(\lambda)$; finally, if $L - c = 0$, the quantity effect is positive but also negligible, and so $\pi^{k-1}(\lambda) - \pi^k(\lambda)$ may be positive or negative.

Consider now the solution for extreme prior beliefs, beginning with the case of large λ .

Proposition 1 For any λ large enough, the expected static profit is maximized by the pooling price $p^0(\lambda)$.

When the prior is sufficiently favorable, the willingness to pay of the buyer goes to H for any signal. In this limit, the demand curve becomes flat at a price H and the price effect becomes negligible. Since H is above the marginal cost c , the quantity effect is positive and it becomes optimal to sell to the entire market, $\psi^0(\lambda) = 1$.

At the other extreme, consider the optimal pricing policy in the case of small λ .

Proposition 2 For any λ small enough, the expected static profit is maximized by:

- (i) the pooling price $p^0(\lambda)$, if $c < L$,
- (ii) one of the selling prices $p^k(\lambda)$ with $k = 0, \dots, K$, if $c = L$, and
- (iii) a non-selling price $p^E(\lambda)$, if $c > L$.

When λ tends to 0, demand is perfectly elastic at a price equal to the low value, L . If the low value is above the production cost, it is optimal to sell to all buyer types. If instead the low value is below the cost, the good cannot be sold profitably. Finally, in the borderline case with $c = L$, it is not optimal to exit the market, because all buyer types are willing to pay more than the cost.

5 Information screening

Having examined how prices affect immediate rent extraction, we now examine how they affect information screening. We start by showing how the monopolist can use prices to control what information is revealed and thereby the updating that occurs in the market. We then establish that, ignoring the associated costs, this potential revelation of information is beneficial to the monopolist. The reason is that the monopolist's expected static profit is convex in the prior. The implication for information screening, is that expected future profits are maximized at a price which reveals some information to future buyers.

Public updating

In each period, an intermediate price $p^k(\lambda)$ induces a bi-partition of the set of signal realizations of the buyers, into $b^{<k} \equiv \{s^0, \dots, s^{k-1}\}$ and $b^{\geq k} \equiv \{s^k, \dots, s^K\}$. Buyers who privately observe s^k or higher purchase, while those with lower signal realizations do not purchase. As a result, after a purchase $a = 1$ at price $p^k(\lambda)$ the prior belief is updated favorably to

$$\lambda^{\geq k}(\lambda) \equiv \Pr(H \mid \lambda, p^k(\lambda), a = 1) = \frac{\lambda \sum_{i=k}^K \alpha^i}{\lambda \sum_{i=k}^K \alpha^i + (1-\lambda) \sum_{i=k}^K \beta^i}, \quad (6)$$

and after no purchase $a = 0$ unfavorably to

$$\lambda^{<k}(\lambda) \equiv \Pr(H \mid \lambda, p^k(\lambda), a = 0) = \frac{\lambda \sum_{i=0}^{k-1} \alpha^i}{\lambda \sum_{i=0}^{k-1} \alpha^i + (1-\lambda) \sum_{i=0}^{k-1} \beta^i}. \quad (7)$$

Now consider the public updating following extreme prices. When the pooling price p^0 is posted, all buyer types purchase, so that no updating takes place $\lambda^{\geq 0}(\lambda) = \lambda$. Similarly, at the non-selling price p^E no type purchases, and $\lambda^{<E}(\lambda) = \lambda$. Neither of these two extreme prices reveal public information. In accordance with BHW we say that an informational cascade occurs when either of the extreme prices is charged from some period T onwards. As in BHW, in this model an informational cascade is equivalent to herding.

Definition An *informational cascade* occurs at time T , if all buyers $t \geq T$ make the same purchase decisions regardless of their signal realizations.

Until an informational cascade occurs, prices partition the types of buyers in two non-empty subsets. Once an informational cascade occurs, no signals can be inferred from the actions a_t , for all $t \in \{T, T+1, \dots\}$. There are two types of informational cascades, depending on which extreme price is charged.

First, consider a period T in which the seller finds it optimal to charge the pooling price $p_T = p^0(\lambda_T)$, so as to induce even the buyer with the lowest signal to purchase the good. Since all buyer types are willing to buy at this price, a purchase in period T is uninformative to future buyers. Therefore, the situation facing the seller and buyer $T + 1$ is identical to the one that the seller and buyer T were confronted with in period T , and the price charged in the previous period, p_T , is also optimal in period $T + 1$. Hence, all types of buyers will again purchase the good in period $T + 1$ and the argument can be repeated for all the following periods. In this *purchase cascade*, the monopolist sells to all buyers and effectively covers the market.

Second, a similar reasoning applies when the seller optimally decides in period T to charge a non-selling price, $p_T = p^E(\lambda_T)$, i.e., a price so high that even the buyer with the highest signal realization decides not to purchase. Again, future buyers will not be able to infer any information from the absence of purchase in period T . In this *exit cascade*, the monopolist effectively abandons the market.

Next, consider the updating of the belief following intermediate prices. These prices allow future buyers to infer some information, and therefore necessarily result in more information (in the sense of Blackwell) than extreme prices. In general, the strength of the updating following a purchase or no purchase depends on the level of the price at which the decision was made. Since the probability of a purchase is decreasing in the level of the price, a purchase at a higher price is stronger good news about the product's value than a purchase at a lower price. The resulting updating is then skewed to the right, in the sense that the upward update of the belief following a purchase is stronger the higher the price. On the flip side, the inference following the failure of selling at a low price is more damaging than at a high price. Updating following no purchase is then more skewed to the left for lower prices. Note that the amount of information revealed by two different intermediate prices generally cannot be Blackwell ranked.

We now define the set of public beliefs that can be attained starting from the initial prior λ_1 . Depending on whether buyer t buys or not at price $p_t = p^k(\lambda)$, it is publicly observed whether buyer t 's signal realization lies in the set $b^{\geq k} = \{s^k, \dots, s^K\}$ or in the set $b^{< k} = \{s^0, \dots, s^{k-1}\}$. Denoting by \mathcal{P} the set of all $K + 2$ monotonic bi-partitions $\{b^{< k}, b^{\geq k}\}$ of the support of the signal distribution S , for $k \in \{0, \dots, K, E\}$, we define the set of countable public beliefs that can be attained from $\lambda_1 \in (0, 1)$ as

$$\Lambda(\lambda_1) \equiv \left\{ \lambda \mid \begin{array}{l} \text{there exists an integer } T \text{ and a sequence of signal bi-partitions} \\ (b_1, \dots, b_T) \in \mathcal{P}^T \text{ such that } \lambda = \Pr(H | \lambda_1; b_1, \dots, b_T) \end{array} \right\} \subset (0, 1).$$

Value of information

The next question of interest is whether the monopolist benefits from charging a price that reveals information to future buyers. We start by examining the static model. A key property is

that the monopolist's expected static profit is convex in the prior belief. This implies that the monopolist benefits from making publicly available as much information as possible, provided that information is free.

Proposition 3 Expected static profits $\pi^k(\lambda)$ obtained by targeting the buyer with signal s^k at price $p^k(\lambda)$ are strictly convex in the prior belief λ for $k \in \{0, \dots, K-1\}$ and linear for $k = K$.

Given the important role of this convexity property, we now further investigate the intuition behind it. The first step is to analyze the functional dependence of the posterior beliefs (1) on the prior belief. Clearly, the posterior λ^k for any signal s^k is monotonically increasing in the prior λ . In addition, the posterior λ^k is either concave or convex in λ depending on the magnitude of the corresponding likelihood ratio $\frac{\alpha^k}{\beta^k}$.

Lemma 1 The posterior belief λ^k is convex in the prior belief λ whenever s^k is bad news (i.e., $\frac{\alpha^k}{\beta^k} < 1$), and concave whenever the signal is good news (i.e., $\frac{\alpha^k}{\beta^k} > 1$) for $v = H$.

To understand this property of Bayesian updating, consider the posterior corresponding to the worst possible signal, s^0 . Clearly, we have $\frac{\alpha^0}{\beta^0} < 1$ so that λ^0 is strictly convex in λ .¹⁸ The marginal impact of an increase in the prior on the posterior resulting from bad news is higher, the more optimistic the prior belief. Intuitively, the higher the prior belief, the more it dominates an unfavorable signal.

More generally, the posterior belief conditional on *bad* news is *convex* in the prior. Starting with a low prior, a further reduction in the prior makes little difference to the posterior belief conditional on a bad signal. When the prior is low, this bad signal results in a posterior that is relatively unresponsive to the exact magnitude of the prior, because the bad signal contains little information in addition to that already contained in the prior. The posterior is more responsive to changes in the prior, when the prior is instead higher. Analogously, the posterior conditional on *good* news is *concave* in the prior because good news contradicts the prior when it is low.

Returning now to the intuition for Proposition 3, the second step is to consider how the convexity and concavity of $\lambda^k(\lambda)$ combined with the linearity of $\psi^k(\lambda)$ result in a convex $\pi^k(\lambda)$. Focus first on the extreme case in which the monopolist targets the buyer with the lowest signal by posting the pooling price $p^0(\lambda)$. In this case, $\psi^0(\lambda) = 1$ and so the profit is

$$\pi^0(\lambda) = p^0(\lambda) - c = \lambda^0(\lambda) + [L - c]. \quad (8)$$

The convexity of $\pi^0(\lambda)$ is due to the convexity of $\lambda^0(\lambda)$ shown in Lemma 1. This reflects the fact that an increase in the prior results in an increasingly higher effect on the posterior belief of

¹⁸ $\frac{\alpha^0}{\beta^0} < 1$ follows from $\sum_{i=0}^K \alpha^i = \sum_{i=0}^K \beta^i = 1$ and $\frac{\alpha^k}{\beta^k}$ strictly increasing in k .

receiving an unfavorable signal. Second, consider the case in which the monopolist targets the buyer with the highest signal s^K by posting $p^K(\lambda) = \lambda^K(\lambda) + L$. Even though $\lambda^K(\lambda)$ is concave in λ (by Lemma 1, since $\frac{\alpha^K}{\beta^K} > 1$), the profit for $k = K$ is equal to

$$\pi^K(\lambda) = \lambda\alpha^K + [L - c](\lambda\alpha^K + (1 - \lambda)\beta^K) \quad (9)$$

and so is linear in λ . This is due to the fact that when targeting the buyer with the highest signal, the probability of selling $\psi^K(\lambda)$ (cf. (2)) is equal to the probability that signal s^k is observed, and so cancels out with the denominator in the expression for the posterior belief $\lambda^K(\lambda)$ (cf. (1)). Intuitively, the expected static profit $\pi^k(\lambda)$ becomes “less and less convex” in the prior belief λ when we move from the lowest price, $p^0(\lambda)$, to the highest one, $p^K(\lambda)$. Since even at the highest price, $p^K(\lambda)$, the expected static profit $\pi^K(\lambda)$ is still (weakly) convex, it is intuitive that the expected static profit is also convex for all intermediate prices $p^k(\lambda)$, corresponding to intermediate signals s^k , $k \in \{1, \dots, K - 1\}$.

As an immediate corollary of Proposition 3, we have the following result.

Corollary 1 The maximal expected static profit $\pi(\lambda)$ is convex in λ .

Overall, in this model profits are convex for two reasons. The first source of convexity is the buyer’s pre-existing private information, as identified by Proposition 3. With $K \geq 1$, each $\pi^k(\lambda)$ is strictly convex for $k < K$, regardless of whether c is above or below L . The second source of convexity is the seller’s option of targeting different types of buyers depending on the prior belief. This second effect is present also in the absence of private information, provided that $c > L$.¹⁹

The convexity of profit in beliefs plays a key role in our analysis, when combined with the well-known martingale property of beliefs. Using the definitions (2), (6), and (7), it is immediate to verify that the belief process $\{\lambda_t\}_{t=0}^\infty$ is a martingale, i.e., the expected posterior belief is equal to the prior belief:

$$E(\lambda_{t+1}|\lambda_t) = \psi^k(\lambda_t)\lambda_t^{\geq k} + [1 - \psi^k(\lambda_t)]\lambda_t^{< k} = \lambda_t \sum_{i=k}^K \alpha^i + \lambda_t \sum_{i=0}^{k-1} \alpha^i = \lambda_t. \quad (10)$$

Revelation of public information then induces a mean preserving spread in the belief distribution. By the convexity of Corollary 1, Jensen’s inequality guarantees that the monopolist’s static profits increase in the amount of public information revealed. The monopolist then benefits from revelation of public information and is therefore willing to subsidize information revelation in the

¹⁹To see this, with $K = 0$ the seller obtains $\pi(\lambda) = L + \lambda - c$ for $\lambda \geq c - L$, while it exercises the option of not selling for $\lambda < c - L$, in which case $\pi(\lambda) = 0$. Convexity in models of public learning à la Bergemann and Välimäki (1996) is based on this second effect. In those models, experience is revealed ex post after the buyer purchases the good. Rather than depending on the good purchased, in our model the information revealed depends on the price at which the good is purchased.

static model. This result can also be seen as a corollary of a more general result on the value of public information proven by Ottaviani and Prat (2001).²⁰

Turning to the dynamic problem, we now compare different pricing strategies $\{p_\tau = p(\lambda_\tau)\}_{\tau=t}^\infty$ in terms of the expected future profits $E[\sum_{\tau=t+1}^\infty \delta^{\tau-(t+1)} (p_\tau - c) a_\tau | \lambda_t]$ they generate from period $t + 1$ onwards, conditional on the initial belief λ_t .²¹ This allows us to isolate the informational effect of prices. Assuming that it is not optimal to exit the market, in period t the seller chooses either an uninformative price $p^0(\lambda_t)$ or some informative price $p_t \in \{p^1(\lambda_t), \dots, p^K(\lambda_t)\}$. If the seller optimally charges $p_\tau = p^0(\lambda_\tau)$ in period t , no information is revealed. Then, $p^0(\lambda_\tau)$ continues to be optimal forever after, so that this strategy generates an expected payoff of $\frac{\pi^0(\lambda_t)}{1-\delta}$. The following proposition shows that the seller's expected *future* profits are strictly higher when charging an informative price $p_t \in \{p^1(\lambda_t), \dots, p^K(\lambda_t)\}$ in period t and following the optimal strategy thereafter. Consequently, demanding an informative price generates higher expected future profits.

Proposition 4 In any period t in which exit is not optimal, the seller obtains strictly higher expected future profits from period $t + 1$ onwards by charging any informative price $p_t \in \{p^1(\lambda_t), \dots, p^K(\lambda_t)\}$ in period t and the optimal price in every period $\tau \geq t + 1$ rather than charging $p^0(\lambda_\tau) = p^0(\lambda_t)$ in every period $\tau \geq t$.

To understand this result, note that the seller can always demand an informative price in period t and then charge $p^0(\lambda_{t+1})$ for all $\tau \geq t$. Due to the convexity of $\pi^0(\lambda)$ and the martingale property of beliefs, $E(\lambda_{t+1} | \lambda_t) = \lambda_t$, this strategy results in higher expected future profits compared to the uninformative price $p^0(\lambda_t)$ in period t .²²

Unless it is optimal to exit the market, the seller's expected future payoff is maximized at a price that reveals some information to future buyers. As expected future payoffs are larger when revealing some information rather than none, the seller has an incentive to charge a price at which only buyers with sufficiently high signals purchase the good. However, general results about the relative value of different informative prices cannot be derived because different informative prices reveal different information for which a ranking in the sense of Blackwell is not possible in general.

Although information revelation benefits the seller in the future, often it is not without cost. Whenever the expected immediate profit from charging the pooling price $p^0(\lambda_t)$ exceeds that

²⁰This result holds with more than two states provided that the MLRP is satisfied, but there are examples in which it does not hold if the MLRP is violated. We refer to that paper for a discussion of the connection with Milgrom and Weber's (1982) linkage principle.

²¹Since the public belief λ_t is informationally equivalent to the history h_{t-1} , we can consider the monopolist's (pricing) strategies as functions that map every public belief λ_t into a price p_t rather than as functions that map every history h_{t-1} into a price p_t .

²²The result is reinforced by the seller's option to exit the market by demanding $p_\tau = p^E(\lambda_{t+1})$ for all $\tau \geq t$ if $\frac{p^0(\lambda_{t+1}) - c}{1-\delta} < 0$. Overall, we have $E\left[\max\left(\frac{\pi^0(\lambda_{t+1})}{1-\delta}, 0\right) | \lambda_t\right] \geq E\left(\frac{\pi^0(\lambda_{t+1})}{1-\delta} | \lambda_t\right) > \frac{\pi^0(\lambda_{t+1})}{1-\delta}$.

generated by the price $p^k(\lambda_t)$, demanding the price $p^k(\lambda_t)$ has the cost $\pi^0(\lambda_t) - \pi^k(\lambda_t) > 0$ in period t . Therefore, if the price $p^0(\lambda_t)$ strictly maximizes the expected immediate profit of period t , i.e., if $\pi^0(\lambda_t) > \max_{k \in \{1, \dots, K\}} \pi^k(\lambda_t)$, there is a trade-off in period t between expected immediate profit on the one hand, and expected future profits, on the other.

6 Dynamically optimal pricing

Typically, the expected immediate profit and the expected future profits are maximized at different prices. Thus, the seller may face a trade-off between immediate rent extraction and information revelation. In this section, we first show how this trade-off is solved in two important cases. In one case, we consider a seller so patient that the benefit of revealing information to future buyers swamps the cost of deviating from the myopically optimal price. In the second case, we consider what happens when the belief is very optimistic or pessimistic. With such beliefs, there is little benefit from revealing information to future buyers and the monopolist's dynamically optimal price equals the myopically optimal one. We then examine the overall implications for long-run behavior and learning, and we show that in all but one case herding almost surely occurs before the good's true value is learned. Finally, we conclude the section by describing short-run behavior when the signal has only two, rather than many realizations

Patience

Consider first a seller who is sufficiently patient, so that there are large potential returns from revealing information to future buyers. We now establish that a sufficiently patient seller prefers to charge an informative price today and a pooling price $p^0(\lambda_{t+1})$ thereafter, rather than a pooling price $p^0(\lambda_t)$ today and forever thereafter.²³

Proposition 5 For every $\lambda \in (0, 1)$ there exists a threshold $\bar{\delta} \in (0, 1)$, such that whenever the monopolist's discount factor is higher than that threshold, $\delta > \bar{\delta}$, the pooling price $p^0(\lambda)$ is not optimal for the seller.

A patient enough seller cares mostly about expected future profits, and by Proposition 4, we know that expected future profits cannot be maximized by a pooling price. Regardless of the (fixed) belief $\lambda \in (0, 1)$ the pooling price $p^0(\lambda)$ is never optimal for an infinitely patient seller ($\delta \rightarrow 1$).

Extreme beliefs

Next, consider the resolution of the trade-off when the beliefs are extreme. Given any discount factor, the benefit of revealing information to future buyers goes to zero when the belief

²³Provided the seller does not want to exit.

becomes very optimistic or pessimistic. Starting with belief λ , the *benefit* of information revelation associated to price p^k is given by

$$B^k(\lambda) \equiv \left[\psi^k(\lambda) V(\lambda^{\geq k}(\lambda)) + (1 - \psi^k(\lambda)) V(\lambda^{< k}(\lambda)) \right] - V(\lambda),$$

where $V(\cdot)$ denotes the value function of the seller's dynamic optimization problem. The term within the brackets denotes the seller's expected future payoff when charging p^k and using the subsequent purchase decision to update λ , whereas $V(\lambda)$ denotes the seller's expected payoff absent the information revealed by the price p^k . Thus, the difference is the increase in the seller's expected payoff due to the information revelation associated with p^k .²⁴ As shown in the following lemma, the expected present value of the seller's revenues and hence $V(\cdot)$ is bounded for any given $\delta < 1$. Thus, we conclude that $B^k(\lambda) \rightarrow 0$ when $\lambda \rightarrow 0$ or $\lambda \rightarrow 1$, since $\lim_{\lambda \rightarrow 0} \lambda^{\geq k}(\lambda) = \lim_{\lambda \rightarrow 0} \lambda^{< k}(\lambda) = 0$ and $\lim_{\lambda \rightarrow 1} \lambda^{\geq k}(\lambda) = \lim_{\lambda \rightarrow 1} \lambda^{< k}(\lambda) = 1$.

Lemma 2 In any period $\tau \in \{1, 2, \dots\}$, the expected revenues of the seller in this period are bounded by $\lambda_\tau + L$, i.e., $E(p_\tau a_\tau | \lambda_\tau) \leq \lambda_\tau + L$; and conditional on λ_t , the expected present value of the seller's revenues from period t onwards is bounded by $\frac{1}{1-\delta} [\lambda_t + L]$.

Since for extreme priors the benefit of deviating from a myopically optimal price goes to zero, while the cost is bounded away from zero (see Propositions 1 and 2), the seller will choose the myopically optimal price.

There are two cases to examine. First, consider what happens when λ is high enough. From Proposition 1, we know that the pooling price is myopically optimal. Consequently, the monopolist triggers a purchase cascade whenever the public prior belief that $v = H$ is sufficiently high:

Proposition 6 For every discount factor $\delta \in [0, 1)$, there exists an $\varepsilon_\delta > 0$ such that the seller charges the pooling price $p^0(\lambda_t)$ whenever $\lambda_t > 1 - \varepsilon_\delta$.

Second, consider what happens when λ is small enough. From Proposition 2, we know that the pooling price is myopically optimal if $c < L$, the non-selling price is optimal if $c > L$, and a selling price is optimal if $c = L$.

Start by examining the case in which the low value exceeds the unit cost:

Proposition 7 Assume that $c < L$. For every discount factor $\delta \in [0, 1)$, there exists an $\varepsilon'_\delta > 0$ such that the seller charges the pooling price $p^0(\lambda_t)$ whenever $\lambda_t < \varepsilon'_\delta$.

²⁴Note that $B^k(\lambda) = 0$ at the two uninformative prices (p^0 and p^E) because no information is revealed. Because of Proposition 4, $B^k(\lambda) > 0$ for all other prices, and consequently $B^k(\lambda) \geq 0$ for all prices.

According to Proposition 6 and Proposition 7, when the low value of the good exceeds the cost, the monopolist triggers a purchase cascade both when the public prior belief of $v = H$ is high and when it is low.

Consider next a seller for which the low value of the object is below the unit cost. In this case, the seller will eventually decide to exit when the public belief that $v = H$ is sufficiently low.

Proposition 8 Assume that $c > L$. For every discount factor $\delta \in [0, 1)$, there exists a positive threshold for the belief $\lambda^E \in \Lambda(\lambda_1)$, such that the seller exits the market whenever $\lambda_t < \lambda^E$ and stays in the market whenever $\lambda_t > \lambda^E$.

When the low value of the good lies below the cost, the monopolist triggers a purchase cascade at high λ and an exit cascade for low λ . Interestingly, the dynamically optimal price for the seller may be one which lies below the myopically optimal one. The reason is that while it is never myopically optimal to charge a price below cost it may be dynamically optimal to do so because the monopolist benefits from buyers' learning. Another interesting result in this case is that, regardless of the discount factor δ , it is uniquely optimal for the seller to stay in the market and charge a price above the pooling price for some attainable values of the public prior belief that the good's value is high. The reason is that there exist attainable values of $\lambda \in \Lambda(\lambda_1)$ that satisfy $p^0(\lambda) \leq c < p^k(\lambda)$ since $p^k(\lambda) > p^0(\lambda)$ for all $k > 0$ and all $\lambda \in (0, 1)$. For such beliefs, staying in the market and demanding some information-revealing price is optimal since it generates a positive payoff, in contrast to charging the pooling price $p^0(\lambda)$.

Finally, consider the borderline case $c = L$. In this case, buyers are always willing to pay more than the cost, and hence it is optimal for the monopolist to charge a price at which sales may occur, rather than to exit the market. If in addition the prior belief that $v = H$ is sufficiently low, a purchase cascade is not triggered either.

Proposition 9 Assume that $c = L$. The seller never exits the market. If there exists a $\hat{\lambda} > 0$ such that for each $\lambda \in (0, \hat{\lambda})$ there exists a price $p^k(\lambda)$, $k \geq 1$, that satisfies $[p^k(\lambda) - c] \psi^k(\lambda) \geq p^0(\lambda) - c$, then the seller charges an intermediate price at each $\lambda \in (0, \hat{\lambda})$, and a purchase cascade will not arise.

Long-run outcomes

We now examine the implications for the long-run learning outcomes. In all but one special case, the seller's optimal price is such that an informational cascade eventually arises and the true state is never learnt.

Consider first the exception which occurs in the borderline case, where $c = L$. While the seller's optimal strategy prevents the buyers from learning that the good is of high value, it may

enable them to eventually learn that the value is low.

Proposition 10 Assume that $c = L$. If $\alpha_K > \frac{\alpha^0}{\beta^0}$, there exists a $\hat{\lambda} > 0$ such that at each $\lambda < \hat{\lambda}$ the seller charges an information-revealing price.

If $c = L$ and $\alpha_K > \frac{\alpha^0}{\beta^0}$, at sufficiently low probabilities of $v = H$ the seller does *not* trigger an informational cascade. In this case, if the prior λ_1 is sufficiently low and the true value is low ($v = L$), there is a positive probability that the buyers asymptotically learn that the true value is low. When $c = L$, there are then two possible long-run outcomes: either a purchase cascade occurs, or it is learnt asymptotically that the value is low.

Next, consider the non-borderline cases. For these cases we know that a cascade necessarily occurs whenever prior beliefs λ are extreme enough. As can be shown by example, a cascade could also occur at some interior prior belief, in which case extreme beliefs are never reached. This opens the question of whether it is possible that the belief never settles, so that neither a cascade occurs nor extreme prior beliefs are ever reached. The martingale convergence theorem implies that this is not possible. Since the prior belief λ_t converges almost surely, a cascade and herding will occur almost surely.

Proposition 11 If $c \neq L$, an informational cascade eventually occurs almost surely. If $c = L$ and there exist a $\hat{\lambda} > 0$ and a price $p^k(\lambda)$, $k \geq 1$, such that for each $\lambda \in (0, \hat{\lambda})$ it holds that $[p^k(\lambda) - c] \psi^k(\lambda) \geq p^0(\lambda) - c$, then either an informational cascade occurs almost surely or the seller and the buyers asymptotically learn that the true value is low, $v = L$. Finally, if $c = L$ and the value is high ($v = H$), a cascade eventually occurs almost surely.²⁵

Despite the ability to freely change prices and the positive effect of information revelation on the seller's payoff, if $c \neq L$ it is optimal for the seller to eventually settle on the pooling price or exit the market, thereby stopping the process of information aggregation. Only in the borderline case and only if the value of the good is actually low, may buyers asymptotically learn the good's true value. In all other cases, learning is incomplete.

Short-run outcomes

Further restrictions must be imposed on the signal structure in order to obtain sharper results on the seller's optimal pricing policy in the short run. In a version of this model in which the signal has two realizations and a symmetric distribution (i.e., with $K = 1$ and $\alpha_0 = \beta_1 \equiv \alpha \in (1/2, 1)$),

²⁵Whenever $\alpha_K > \frac{\alpha^0}{\beta^0}$, the condition that there exist a $\hat{\lambda} > 0$ and a price $p^k(\lambda)$, $k \geq 1$, such that $[p^k(\lambda) - c] \psi^k(\lambda) \geq p^0(\lambda) - c$ for each $\lambda \in (0, \hat{\lambda})$, is satisfied for $c = L$ (see the proof of Proposition 10). Thus, if $c = L$ and $\alpha_K > \frac{\alpha^0}{\beta^0}$, either an informational cascade occurs almost surely or the seller and the buyers asymptotically learn that the true value is low.

Bose, Orosel, Ottaviani, and Vesterlund (2005) have fully characterized the perfect Bayesian equilibria. Since the analysis is involved, here we only briefly summarize the main findings.

In the symmetric binary signal model, at any given public belief the seller only needs to consider two possible prices—the pooling price and the separating price—and exit. As only the separating price reveals information to future buyers, Proposition 4 guarantees that expected future profits are maximized by charging the separating price, provided that exit is not optimal. The expected immediate profit from the pooling and the separating prices are therefore crucial for determining which price is charged. Specifically, the separating price is uniquely optimal when it also maximizes the expected immediate profit, and the pooling price can only be optimal (but need not be) when it maximizes that profit.

The properties of the equilibrium depend on α , the precision of the buyers' signals. We distinguish two cases. First, when signals are sufficiently precise (given the other parameters), it is optimal for the seller to charge the separating price if and only if the public belief belongs to a connected interval. When the separating price is charged, the buyer purchases the good if only if the signal realization is high. Observing this action, the seller and future buyers revise their beliefs about the good's value. The seller then continues to charge the separating price as long as the belief stays within this interval, and demands the pooling price or exits the market when the belief exits this interval.

Second, when signals are instead sufficiently noisy, the properties of the equilibrium become sensitive to the interplay between the discount factor and the cost of production. When $c \leq L$, a sufficiently impatient seller triggers herding immediately for all priors. In contrast, a more patient seller charges the separating price for some initial priors (and updated public beliefs). In this case, the separating price is optimal in a connected, but possibly empty, set of beliefs. Instead, when the signals are sufficiently noisy and $c > L$, additional complications can arise. In this case, it is possible that the separating and the pooling prices are each uniquely optimal for two disjoint intervals of public beliefs. As λ increases from $\lambda = 0$, exit is initially optimal for the seller, then the separating price is optimal, then the pooling price is optimal, then the separating price is again optimal, and finally the pooling price becomes again optimal for sufficiently high public beliefs.

7 Conclusion

The prototypical model presented in this paper is highly stylized, in that a simple sequential structure is imposed and perfect observation of past prices and decisions of all previous decision makers is assumed. Nevertheless, we have gained insights on how the monopolist affects the aggregation of private information in markets.

We have shown that the monopolist has an incentive to charge a price higher than the pooling price, in order to allow subsequent buyers to infer the information of the current buyer. As information is revealed, the value of learning is reduced and eventually the seller may stop the learning process and induce an informational cascade.²⁶ When inducing a purchase cascade, the price is reduced to the pooling price. If the price were reduced to the pooling price before the good had become either popular or unpopular, all buyers would purchase the good in any case so that their decisions would not reveal their information.²⁷

We conclude by first comparing our results with those obtained with fixed (rather than path-dependent) prices, and second by examining the welfare effect of having monopolistic rather than competitive pricing.

Fixed vs. variable prices

It is worth comparing the prediction of our dynamic pricing model with those obtained by Welch (1995) for the case in which the monopolist is constrained to post a constant price to all buyers, $p_t = p$ for all $t \in \{1, 2, \dots\}$. Welch found that it is optimal for the seller to induce an immediate purchase cascade by charging a price low enough so that all buyers purchase. With flexible prices instead, we show that the seller is willing to delay the informational cascade and allow buyers to learn.

The sharp contrast in the predictions of these two models is due to the different value the monopolist places on private and public information. In Welch's model, the monopolist cannot condition the price on the information revealed during the social learning process. In this situation, social learning essentially increases the private information possessed by the buyers. Giving *private information* to the buyers results in a spread in the distribution of valuation around the initial valuation of each type. Welch assumed that the seller's opportunity cost is lower than the low value of the good (i.e., $c < L$) and that the signal of each individual buyer is relatively uninformative, so that it is optimal for the monopolist to charge a price that attracts all buyers regardless of their signals. In such a mass market, social learning would cause a reduction in the rent that the monopolist can extract (Lewis and Sappington, 1994 and Johnson and Myatt, 2004). In our model instead, the seller is allowed to condition future prices on the information revealed by the buyers' purchase decisions. As a consequence, social learning provides *public information* which is beneficial to the seller (Ottaviani and Prat, 2001).

²⁶This is always true except when the cost c equals the low value of the good, L , and the good's value is low, in this case there is a positive probability that the belief λ will converge to zero.

²⁷Social learning may explain why young independent professionals, such as doctors and lawyers, are willing to be underemployed and charge high fees relative to their perceived quality, rather than reduce the price for their services. The observed discounts given on books once they are listed as best sellers are consistent with the prediction that the price should be reduced only after the good becomes popular.

Monopoly vs. competition

If there are at least two sellers competing à la Bertrand in each period, the price would be equal to the marginal cost, $p_t = c$ for all $t \in \{1, 2, \dots\}$. This competitive benchmark corresponds to BHW's model with fixed price $p_t = c$. We will use this benchmark to determine how the consumer surplus, producer surplus, and social welfare compare under monopolistic and competitive pricing.

Consider the case where the seller's cost of production is below the low value, i.e., $c < L$. Here, the socially optimal allocation requires that all buyers purchase the good, regardless of the realizations of their private signals. Hence the buyers' private information has no social value. While competitive pricing instantly triggers a purchase cascade, the monopolist may choose to postpone a cascade until more information on the good's value has been gathered. Though long-run behavior is efficient under both monopoly and competition, that need not be the case initially. When the monopolist postpones the purchase cascade to induce social learning and extract more rents in the future, she inefficiently excludes some buyers. This exclusion and learning is socially inefficient, and the monopolist is effectively causing too much learning. As in the textbook treatment of static monopoly, we see that going from competitive to monopolistic pricing increases producer surplus and decreases consumer surplus, and that the increase in producer surplus is lower than the reduction in consumer surplus.

Consider next the case with $c > L$, i.e., when information is socially valuable. Now, the monopolist's incentive to generate learning is partially aligned with the social value of learning. Interestingly, in this case a monopoly can deliver higher social welfare than competition. To illustrate this, focus on a special case in which $c = L + \lambda^K (\lambda_1) + \varepsilon$, with $\varepsilon > 0$ small. That is, in the initial period the cost exceeds (slightly) the willingness to pay of the most optimistic buyer. In this case, an exit cascade occurs immediately in the competitive environment. Information being a public good among competitors, no seller would invest in learning by posting a price below the non-selling price, since there would be no way to subsequently recoup this short-term loss. Consumer, producer, and total surplus would all equal zero. In contrast, a sufficiently patient monopolist would be willing to bear short-term losses to induce learning in the market. In the long run, a purchase or exit cascade will result.²⁸ As a result, the consumer surplus of all buyers is higher under monopoly than under competition. Since the expected producer surplus is also increased, in this example monopoly Pareto dominates (ex ante) competition!

²⁸There are also implications for the occurrence of incorrect cascades, in which all but possibly a finite number of buyers purchase a low-quality good, or do not purchase a high-quality good. In this example with $c > L + \lambda^K (\lambda_1)$, the probability of an incorrect cascade is lower under monopoly than competition.

Appendix

Proofs of Propositions 1–11 follow.

Proof of Proposition 1. Consider the limit of $\pi^{k-1}(\lambda) - \pi^k(\lambda)$ as λ tends to 1. The first term (4) goes to $[L + 1 - c]\alpha^{k-1} = [H - c]\alpha^{k-1} > 0$. The second term (5) goes to 0 for λ large enough. By continuity, we conclude that $\pi^{k-1}(\lambda) > \pi^k(\lambda)$ for all k for large enough λ , hence the pooling price $p^0(\lambda)$ is optimal. *Q.E.D.*

Proof of Proposition 2. In the limit as λ goes to 0, the second term (5) goes to zero. The sign of the first term (4), which converges to $[L - c]\beta^{k-1}$ for $\lambda \rightarrow 0$, depends on the sign of $L - c$. If $L - c > 0$, the first term is positive, so that $\pi^{k-1}(\lambda) > \pi^k(\lambda)$ for all k and the pooling price $p^0(\lambda)$ is optimal. If instead $L - c < 0$, the first term is negative, so that $\pi^{k-1}(\lambda) < \pi^k(\lambda)$ for all k and the non-selling price $p^E(\lambda)$ is optimal when λ is small enough. For $L - c = 0$, it holds for all $k = 0, \dots, K$ that $p^k(\lambda) > c$ for all $\lambda > 0$, and thus part (ii) of the proposition follows. *Q.E.D.*

Proof of Proposition 3. When targeting the buyer with signal realization s^k , the monopolist obtains static profit equal to $[p^k(\lambda) - c]\psi^k(\lambda)$. The second derivative of the profit with respect to the prior belief is equal to

$$\frac{\partial^2 \pi^k(\lambda)}{\partial^2 \lambda} = \frac{2\alpha^k \beta^k}{(\alpha^k \lambda + \beta^k (1 - \lambda))^3} \sum_{i=k}^K (\alpha^i \beta^k - \alpha^k \beta^i),$$

where the second factor

$$\sum_{i=k}^K (\alpha^i \beta^k - \alpha^k \beta^i) = \beta^k \sum_{i=k}^K \beta^i \left(\frac{\alpha^i}{\beta^i} - \frac{\alpha^k}{\beta^k} \right)$$

is always strictly positive by the MLRP whenever $k < K$ and equal to zero for $k = K$, showing the result. *Q.E.D.*

Proof of Lemma 1. From

$$\frac{\partial^2 (\lambda^k)}{\partial^2 \lambda} = -\frac{2\frac{\alpha^k}{\beta^k} \left(\frac{\alpha^k}{\beta^k} - 1 \right)}{(\alpha^k \lambda + \beta^k (1 - \lambda))^3},$$

we see that the posterior beliefs of signals with $\frac{\alpha^k}{\beta^k} < 1$ (respectively $\frac{\alpha^k}{\beta^k} > 1$) are convex (respectively concave) in the prior. *Q.E.D.*

Proof of Corollary 1. This follows immediately from Proposition 3 and the fact that the maximum of convex functions is convex. *Q.E.D.*

Proof of Proposition 4. When following the first strategy of charging the informative price $p_t = p^k(\lambda_t) \in \{p^1(\lambda_t), \dots, p^K(\lambda_t)\}$, in period t there is a sale with probability $\psi^k(\lambda_t)$ and no

sale with probability $1 - \psi^k(\lambda_t)$. The posterior belief in period $t + 1$ is $\lambda^{\geq k}(\lambda_t)$ after a sale and $\lambda^{< k}(\lambda_t)$ after no sale. By charging price $p^k(\lambda_t)$ in period t and the optimal price thereafter, the monopolist obtains $\psi^k(\lambda_t) V[\lambda^{\geq k}(\lambda_t)] + [1 - \psi^k(\lambda_t)] V[\lambda^{< k}(\lambda_t)]$ in period $t + 1$, where $V(\cdot)$ denotes the monopolist's value function.

The second strategy of charging the pooling price $p^0(\lambda_\tau) = p^0(\lambda_t)$ in every period $\tau \geq t$ yields $\pi^0(\lambda_t) = p^0(\lambda_t) - c$ in each period. In this case, expected future profits are $\frac{\pi^0(\lambda_t)}{1-\delta}$.

The result follows from

$$\begin{aligned} & \psi^k(\lambda_t) V[\lambda^{\geq k}(\lambda_t)] + [1 - \psi^k(\lambda_t)] V[\lambda^{< k}(\lambda_t)] \\ \geq & \psi^k(\lambda_t) \max\left\{\frac{\pi^0[\lambda^{\geq k}(\lambda_t)]}{1-\delta}, 0\right\} + [1 - \psi^k(\lambda_t)] \max\left\{\frac{\pi^0[\lambda^{< k}(\lambda_t)]}{1-\delta}, 0\right\} \\ \geq & \psi^k(\lambda_t) \frac{\pi^0[\lambda^{\geq k}(\lambda_t)]}{1-\delta} + [1 - \psi^k(\lambda_t)] \frac{\pi^0[\lambda^{< k}(\lambda_t)]}{1-\delta} \\ > & \frac{\pi^0(\lambda_t)}{1-\delta}, \end{aligned}$$

where the first inequality follows from $V(\lambda_t) \geq \max\left\{\frac{\pi^0(\lambda_t)}{1-\delta}, 0\right\}$, the second from $\max\left\{\frac{\pi^0(\lambda_t)}{1-\delta}, 0\right\} \geq \frac{\pi^0(\lambda_t)}{1-\delta}$, and the third inequality is due to the strictly convexity of $\pi^0(\lambda_t)$ (Proposition 3) and the martingale property of beliefs (10). *Q.E.D.*

Proof of Proposition 5. For every $k \in \{1, \dots, K\}$ define $F_k(\lambda)$ as the difference in the seller's expected payoff from t onwards for $\lambda_t = \lambda$ between the following two strategies: (i) set an informative price, $p_t = p^k(\lambda_t)$ with $k \neq 0, E$, in period t and the pooling price $p^0(\lambda_{t+1})$ from period $t + 1$ onwards; (ii) set the pooling price $p_t = p^0(\lambda_t)$ from period t onwards. Thus,

$$F_k(\lambda) = \pi^k(\lambda) - \pi^0(\lambda) + \tag{11}$$

$$\frac{\delta}{1-\delta} \left\{ \psi^k(\lambda) \pi^0[\lambda^{\geq k}(\lambda)] + (1 - \psi^k(\lambda)) \pi^0[\lambda^{< k}(\lambda)] - \pi^0(\lambda) \right\}. \tag{12}$$

The term (12) is positive because $p^0(\lambda)$ is strictly convex and $\{\lambda_t\}_{t=0}^\infty$ is a martingale. Since this term is positive and the term (11) is independent of δ and finite, $F_k(\lambda) > 0$ whenever δ is sufficiently close to 1. *Q.E.D.*

Proof of Lemma 2. First, since $\sum_{i=k}^K \alpha^i \leq 1$ and $\frac{\beta^k}{\alpha^k} \geq \frac{\beta^i}{\alpha^i}$ for $i \geq k$, we have $\psi^k(\lambda_\tau) = \lambda_\tau \sum_{i=k}^K \alpha^i + (1 - \lambda_\tau) \sum_{i=k}^K \beta^i = \sum_{i=k}^K \alpha^i \left[\lambda_\tau + (1 - \lambda_\tau) \frac{\sum_{j=k}^K \beta^j}{\sum_{j=k}^K \alpha^j} \right] \leq \lambda_\tau + \frac{\beta^k}{\alpha^k} (1 - \lambda_\tau)$, so that $\frac{\psi^k(\lambda_\tau)}{\lambda_\tau + \frac{\beta^k}{\alpha^k} (1 - \lambda_\tau)} \leq 1$. In any period τ , if the seller demands the price $p_\tau = p^k(\lambda_\tau)$, then $E(p_\tau a_\tau | \lambda_\tau) = E[p^k(\lambda_\tau) a_\tau | \lambda_\tau] = p^k(\lambda_\tau) \psi^k(\lambda_\tau) = \lambda_\tau \frac{\psi^k(\lambda_\tau)}{\lambda_\tau + \frac{\beta^k}{\alpha^k} (1 - \lambda_\tau)} + L \psi^k(\lambda_\tau) \leq \lambda_\tau + L \psi^k(\lambda_\tau) \leq \lambda_\tau + L$. Exit in any period τ implies zero revenues in that period: $p_\tau a_\tau = p^k(\lambda_\tau) a_\tau = 0$. Since $\{\lambda_t\}_{t=0}^\infty$ is martingale, $E(\lambda_{t+\tau} | \lambda_t) = \lambda_t$ for all $\tau \in \{1, 2, \dots\}$, and thus

$E(p_{t+\tau} a_{t+\tau} | \lambda_t) \leq E(\lambda_{t+\tau} + L | \lambda_t) = \lambda_t + L$ for all $\tau \in \{1, 2, \dots\}$. Consequently, we have $0 \leq E(\sum_{\tau=0}^{\infty} \delta^\tau p_{t+\tau} a_{t+\tau} | \lambda_t) \leq \frac{1}{1-\delta} [\lambda_t + L]$. *Q.E.D.*

Proof of Proposition 6. If $p_\tau = p^0(\lambda_t)$ for all $\tau \geq t$, the seller's payoff from t onwards is $\frac{p^0(\lambda_t) - c}{1-\delta}$. If $p_t = p^k(\lambda_t)$, $k \geq 1$, it is less than $[p^k(\lambda_t) - c] \psi^k(\lambda_t) + \frac{\delta}{1-\delta} [H - c]$, regardless of the prices the seller charges in the periods after period t , since $p_\tau a_\tau < H$ for all τ . The difference between these two expressions is, after rearranging,

$$p^0(\lambda_t) - p^k(\lambda_t) \psi^k(\lambda_t) - \left[1 - \psi^k(\lambda_t)\right] c + \delta \frac{p^0(\lambda_t) - H}{1 - \delta}.$$

Taking the limit $\lambda_t \rightarrow 1$ and using $\lim_{\lambda_t \rightarrow 1} p^0(\lambda_t) = H$ and $\psi^k(\lambda_t) + \sum_{i=0}^{k-1} \alpha^i$, we obtain

$$\begin{aligned} & \lim_{\lambda_t \rightarrow 1} \left\{ p^0(\lambda_t) - p^k(\lambda_t) \psi^k(\lambda_t) - \left[1 - \psi^k(\lambda_t)\right] c + \delta \frac{p^0(\lambda_t) - H}{1 - \delta} \right\} \\ &= \lim_{\lambda_t \rightarrow 1} \left[1 - \psi^k(\lambda_t)\right] [H - c] = [H - c] \sum_{i=0}^{k-1} \alpha^i > 0, \end{aligned}$$

where the last equality follows from $H > c$ and $k \geq 1$. Thus, only $p^0(\lambda_t)$ maximizes the seller's expected payoff if λ_t is sufficiently close to 1. *Q.E.D.*

Proof of Proposition 7. Since $p_t \geq L > c$ for all t , the seller never exits the market. When charging $p^k(\lambda_\tau)$, $k \in \{0, \dots, K\}$ in period τ , the seller's expected immediate profits satisfy $E[(p^k(\lambda_\tau) - c) a_\tau | \lambda_\tau] \leq \lambda_\tau + [L - c] \psi^k(\lambda_\tau) \leq \lambda_\tau + L - c$ by Lemma 2 and $L - c > 0$. If $p_\tau = p^0(\lambda_t)$ for all $\tau \geq t$, the seller's payoff from t onwards is $\frac{p^0(\lambda_t) - c}{1-\delta}$. If $p_t = p^k(\lambda_t)$ with $k \geq 1$ the seller's payoff from t onwards cannot exceed $[p^k(\lambda_t) - c] \psi^k(\lambda_t) + \frac{\delta}{1-\delta} [\lambda_t + L - c]$, regardless of the prices the seller charges in the periods after period t , because of $E[(p^k(\lambda_\tau) - c) a_\tau | \lambda_t] \leq E[\lambda_\tau + L - c | \lambda_t] = \lambda_t + L - c$ for all $k \in \{0, \dots, K\}$ and all $\tau \geq t$. The difference between these two expressions is

$$p^0(\lambda_t) - p^k(\lambda_t) \psi^k(\lambda_t) - \left[1 - \psi^k(\lambda_t)\right] c + \delta \frac{p^0(\lambda_t) - [\lambda_t + L]}{1 - \delta}.$$

Taking the limit $\lambda_t \rightarrow 0$ and using $\lim_{\lambda_t \rightarrow 0} p^k(\lambda_t) = L$ and $\lim_{\lambda_t \rightarrow 0} \psi^k(\lambda_t) = \sum_{i=0}^{k-1} \beta^i$, we obtain

$$\begin{aligned} & \lim_{\lambda_t \rightarrow 0} \left\{ p^0(\lambda_t) - p^k(\lambda_t) \psi^k(\lambda_t) - \left[1 - \psi^k(\lambda_t)\right] c + \delta \frac{p^0(\lambda_t) - [\lambda_t + L]}{1 - \delta} \right\} \\ &= \lim_{\lambda_t \rightarrow 0} \left\{ [L - c] \left[1 - \psi^k(\lambda_t)\right] \right\} = [L - c] \sum_{i=0}^{k-1} \beta^i > 0, \end{aligned}$$

where the inequality follows from $L > c$ and $k \geq 1$. Thus, only $p^0(\lambda_t)$ maximizes the seller's expected payoff if λ_t is sufficiently close to 0. *Q.E.D.*

Proof of Proposition 8. We begin by showing that the seller's expected payoff increases in λ . This holds because the seller can always adopt the optimal strategy associated with a lower initial λ .

More precisely, suppose that contingent on any history of past sales the seller targets the same buyer types prescribed by the optimal strategy associated with the lower λ . As long as the optimal strategy associated to the lower initial λ does not prescribe immediate exit (implying $V(\lambda) = 0$), this imitating strategy results in strictly higher revenues. This follows from the the following two properties: (i) prices $p^k(\lambda)$ strictly increase in λ for all $k \in \{0, \dots, K\}$ and (ii) probabilities of sale $\psi^k(\lambda)$ are positive for all $k \in \{0, \dots, K\}$ and non-decreasing in λ . Consequently, whenever at some $\bar{\lambda}$ immediate exit is not uniquely optimal, $V(\lambda)$ is strictly increasing in λ for all $\lambda > \bar{\lambda}$.

By Lemma 2, $\max_k E[(p^k(\lambda_\tau) - c) a_\tau | \lambda_\tau] \leq \lambda_\tau + L - c$, i.e., the seller's expected immediate return in any period τ , conditional on λ_τ , is bounded by $\lambda_\tau + L - c$. By the martingale property of the belief process $\{\lambda_\tau\}_{\tau=t}^\infty$, this implies that the seller's expected payoff conditional on λ_t is bounded by $\frac{\lambda_t - [c - L]}{1 - \delta}$, provided that the seller never exits the market. Consequently, whenever $\lambda_t < c - L$ the seller is better off exiting the market than staying in the market forever. Since by assumption $c < H$, immediate exit is not optimal for all $\lambda \in (0, 1)$. Because the seller's expected payoff strictly increases in λ unless immediate exit is uniquely optimal, there must exist a critical $\lambda^E \in \Lambda(\lambda_1)$ such that exit is optimal for $\lambda = \lambda^E$ and uniquely optimal for all $\lambda \in (0, \lambda^E)$, whereas for all $\lambda \in (\lambda^E, 1)$ it is uniquely optimal for the seller to stay in the market. Since exit is optimal for $\lambda = \lambda^E$, it must hold that $p^0(\lambda^E) \leq c$. *Q.E.D.*

Proof of Proposition 9. Since $p_t > L = c$ for all t , the seller never exits the market. If the seller charges $p_\tau = p^0(\lambda_t)$ for all $\tau \geq t$, the seller's payoff from t onwards is $\frac{p^0(\lambda_t) - c}{1 - \delta}$. By charging $p_t = p^k(\lambda_t)$, $k \geq 1$, in period t and $p^0(\lambda_{t+1})$ from period $t+1$ onwards, the seller's expected payoff from t onwards is $[p^k(\lambda_t) - c] \psi^k(\lambda_t) + \delta \left[\psi^k(\lambda_t) \frac{p^0[\lambda^{\geq k}(\lambda_t)] - c}{1 - \delta} + (1 - \psi^k(\lambda_t)) \frac{p^0[\lambda^{< k}(\lambda_t)] - c}{1 - \delta} \right]$. The payoff difference between the second and the first strategy is

$$F_k(\lambda_t) \equiv [p^k(\lambda_t) - c] \psi^k(\lambda_t) - [p^0(\lambda_t) - c] + \tag{13}$$

$$\frac{\delta}{1 - \delta} \left\{ \psi^k(\lambda_t) p^0[\lambda^{\geq k}(\lambda_t)] + (1 - \psi^k(\lambda_t)) p^0[\lambda^{< k}(\lambda_t)] - p^0(\lambda_t) \right\}. \tag{14}$$

By assumption there exists a $k \geq 1$ such that the term (13) is non-negative. The term (14) is positive because $p^0(\lambda_t)$ is strictly convex and λ_t is a martingale. Hence, $F_k(\lambda_t) > 0$ and it cannot be optimal to charge $p_t = p^0(\lambda_t)$. *Q.E.D.*

Proof of Proposition 10. For $\lambda_t = 0$, $E[(p^0(\lambda_t) - c) a_t | \lambda_t] = p^0(0) - c = L - c = 0$ and

$$E[(p^K(\lambda_t) - c) a_t | \lambda_t] = \alpha^K \lambda_t + [L - c] \psi^K(\lambda_t) = \alpha^K \lambda_t = 0.$$

Since $E[(p^0(\lambda_t) - c) a_t | \lambda_t] = p^0(\lambda_t) - c = \frac{\alpha^0 \lambda_t}{\alpha^0 \lambda_t + \beta^0 (1 - \lambda_t)} + L - c = \frac{\alpha^0 \lambda_t}{\alpha^0 \lambda_t + \beta^0 (1 - \lambda_t)}$, the slope at $\lambda_t = 0$ is $\frac{dE[p^0(\lambda_t) a_t | \lambda_t = 0]}{d\lambda_t} = \frac{\alpha^0}{\beta^0}$. For $E[p^K(\lambda) a_t | \lambda] = \alpha^K \lambda$ the slope at $\lambda = 0$ is α^K . Since by assumption $\alpha^K > \frac{\alpha^0}{\beta^0}$, $p^k(\lambda_t) \psi^k(\lambda_t) = \alpha^K \lambda_t \geq p^0(\lambda_t)$ whenever λ_t is small but positive. The result follows from this and Proposition 9. *Q.E.D.*

Proof of Proposition 11. The stochastic process $\{\lambda_t\}_{t=1}^{\infty}$ is a martingale by (10). Since λ_t is bounded, the martingale convergence theorem (see, e.g., Doob 1953, p. 319) implies that $\lim_{t \rightarrow \infty} \lambda_t = \lambda_{\infty}$ exists with probability 1. If $\lambda_{\infty} \in (0, 1)$, it must be that for some finite T it holds that $\lambda_t = \lambda_{\infty}$ for all $t \geq T$, since $\Lambda(\lambda_1) \cap (\varepsilon, 1 - \varepsilon)$ is a finite set for all $\varepsilon > 0$ and $\lambda_{\infty} \in (0, 1)$ implies $\lambda_{\infty} \in (\varepsilon, 1 - \varepsilon)$ for some sufficiently small ε . But this implies that for $t \geq T$ we have a cascade. Thus, the claim follows from Propositions 6-8 for the case $c \neq L$, and from Propositions 6 and 9 for the case $c = L$. Specifically, the last sentence of the proposition follows from Proposition 6 and the fact that if $v = H$ the public belief λ_t cannot converge to 0. *Q.E.D.*

References

- AGHION, P., BOLTON, P., HARRIS, C., AND JULLIEN, B. “Optimal Learning by Experimentation.” *Review of Economic Studies*, Vol. 58 (1991), pp. 621–654.
- AVERY, C. AND ZEMSKY, P. “Multidimensional Uncertainty and Herd Behavior in Financial Markets.” *American Economic Review*, Vol. 88 (1998), pp. 724–748.
- BAR-ISAAC, H. “Reputation and Survival: Learning in a Dynamic Signalling Model.” *Review of Economic Studies*, Vol. 70 (2003), pp. 231–251.
- BANERJEE, A.V. “A Simple Model of Herd Behavior.” *Quarterly Journal of Economics*, Vol. 107 (1992), pp. 797–817.
- BECKER, G.S. “A Note on Restaurant Pricing and Other Examples of Social Influences on Price.” *Journal of Political Economy*, Vol. 99 (1991), pp. 1009–1016.
- BERGEMANN, D. AND VALIMAKI, J. “Learning and Strategic Pricing.” *Econometrica*, Vol. 64 (1996), pp. 1125–1149.
- AND —. “Experimentation in Markets.” *Review of Economic Studies*, Vol. 67 (2000), pp. 213–234.
- BIKHCHANDANI, S., HIRSHLEIFER, D., AND WELCH, I. “A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades.” *Journal of Political Economy*, Vol. 100 (1992), 992–1026.
- BOSE, S., OROSEL, G., AND VESTERLUND, L. “Optimal Pricing and Endogenous Herding.” Department of Economics, University of Vienna, Working Paper No. 0204, September 2001.
- , —, OTTAVIANI, M., AND VESTERLUND, L. “Monopoly Pricing in the Binary Herding Model.” Mimeo, March 2005.
- CAMINAL, R. AND VIVES, X. “Why Market Shares Matter: An Information-Based Theory.” *RAND Journal of Economics*, Vol. 27 (1996), pp. 221–239.

- AND —. “Price Dynamics and Consumer Learning.” *Journal of Economics and Management Strategy*, Vol. 8 (1999), pp. 95–131.
- EASLEY, D. AND KIEFER, N. “Controlling a Stochastic Process with Unknown Parameters.” *Econometrica*, Vol. 56 (1988), pp. 1045–1064.
- JOHNSON, J. AND MYATT, D. “On the Simple Economics of Advertising, Marketing and Product Design.” Cornell and Oxford University, Mimeo, 2004.
- LEWIS, T. AND SAPPINGTON, D. “Supplying Information to Facilitate Price Discrimination.” *International Economic Review*, Vol. 35 (1994), pp. 309–327.
- MILGROM, P. “Good News and Bad News: Representation Theorems and Applications.” *Bell Journal of Economics*, Vol. 12 (1981), pp. 380–391.
- AND WEBER, R. “A Theory of Auctions and Competitive Bidding.” *Econometrica*, Vol. 50 (1982), pp. 1089–1122.
- NEEMAN, Z. AND OROSEL, G.O. “Herding and the Winner’s Curse in Markets with Sequential Bids.” *Journal of Economic Theory*, Vol. 85 (1999), pp. 91–121.
- MOSCARINI, G. AND OTTAVIANI, M. “Price Competition for an Informed Buyer.” *Journal of Economic Theory*, Vol. 101 (2001), pp. 457–493.
- OTTAVIANI, M. “Monopoly Pricing with Social Learning.” in: *Social Learning in Markets*. Ph.D. Thesis, MIT, 1996.
- AND PRAT, A. “The Value of Public Information in Monopoly.” *Econometrica*, Vol. 69 (2001), pp. 1673–1683.
- ROTHSCHILD, M. “A Two-Armed Bandit Theory of Market Pricing.” *Journal of Economic Theory*, Vol. 9 (1974), pp. 185–202.
- SMITH, L. AND SØRENSEN, P. “Pathological Outcomes of Observational Herding.” *Econometrica*, Vol 68 (2000), pp. 371–398.
- SØRENSEN, A.T. “Bestseller Lists and Product Variety: The Case of Book Sales.” Stanford University, Mimeo, 2004.
- TAYLOR, C.R. “Time-on-the-market as a Sign of Quality.” *Review of Economic Studies*, Vol. 66 (1999), pp. 555–578.
- WELCH, I. “Sequential Sales, Learning and Cascades.” *Journal of Finance*, Vol. 47 (1992), pp. 695–732.