

WORKING

PAPERS

Subir Bose[†], Gerhard O. Orosel[‡] and Lise Vesterlund[§]

Optimal Pricing and Endogenous Herding

September 2001

Working Paper No: 0204



DEPARTMENT OF ECONOMICS

UNIVERSITY OF VIENNA

All our working papers are available at: <http://mailbox.univie.ac.at/papers.econ>

Optimal Pricing and Endogenous Herding*

Subir Bose[†], Gerhard O. Orosel[‡], and Lise Vesterlund[§]

September 2001

Abstract

We consider a monopolist who sells identical objects of common but unknown value in a herding-prone environment. Buyers make their purchasing decisions sequentially, and rely on a private signal as well as previous buyers' actions to infer the common value of the object. The model applies to a variety of cases, such as the introduction of a new product or the sale of licenses to use a patent. We characterize the monopolist's optimal pricing strategy and its implications for the temporal pattern of prices and for herding. The analysis is performed under alternative assumptions about observability of prices. We find that when previous prices are observable, herding may but need not arise. In contrast, herding arises immediately when previous prices are unobservable and the seller's equilibrium strategy is a pure Markov strategy. While the possibility of social learning is present in the first case, it is absent in the second. Finally, we examine the seller's incentive to manipulate the buyers' evaluation of the object when buyers are naive. Using secret discounts the seller successfully interferes with social learning, and herding occurs in finite time.

JEL CLASSIFICATION NUMBERS: D8, D42

KEYWORDS: Herding, Informational Cascades, Optimal Pricing

* Acknowledgements to be added.

[†] Iowa State University

[‡] University of Vienna

[§] University of Pittsburgh

1. Introduction

Markets for objects of common but unknown value are prone to herding when observable purchasing decisions are made sequentially and buyers receive private signals about the object's value. To infer the common value of the object, buyers use the information from previous buyers' actions as well as their own private signal. As a result, buyers' inferences may give rise to informational cascades and herding, where the available public information swamps the buyers' private information and induces them to behave identically.¹ We analyze a monopolist's optimal pricing strategy in such a market under alternative assumptions about observability of prices.

The model applies to a variety of cases. One case is the sale of licenses to use a patent. The patent owner is a monopolist who is selling a product that has uncertain value in the early stages of its use, and hence buyers may try to infer the value of the license from the sequence of previous purchases. Another case is the introduction of a significantly new product, such as computer software, new building materials, new medical equipment, or a new type of vehicle. Frequently the value of a significantly new product is uncertain to the buyers for quite some time and the innovative firm enjoys a temporary monopoly.

For simplicity, we assume that the common value of the object sold by the monopolist is either "high" or "low." In each period the seller approaches a randomly chosen buyer, who observes a private signal about the value of the object, and decides whether to buy one unit of the object at the price demanded by the monopolist. We rely on the perfect Bayesian equilibrium (PBE) as our equilibrium notion.² For each buyer the optimal strategy is to buy the object, if the price is less than the buyer's expected value of the object. The buyers' actions are publicly observed. After each buyer the monopolist decides which price to charge the next buyer. Due to the particular nature of the buyers' strategies, the PBE can be derived directly from the seller's optimal policy.

To demonstrate the monopolist's optimization problem, consider the case where the price demanded by the seller is publicly observed and buyers are rational. At any point in time the seller must decide which price to charge the next potential buyer. Given the structure of the model, the seller will find it optimal to charge either a low price or a high price relative to the current public evaluation of the object. If the monopolist demands the high price, a prospective buyer will buy only if he has observed a signal that indicates high value, whereas if she charges a low price, the buyer will buy independent of his signal. While a sale at the

¹An *informational cascade* occurs when the (observable) actions reveal no information about private signals. *Herding* occurs when agents behave identically, independently of their private signals. In our model these two phenomena can only occur together.

²We use PBE to denote the singular, i.e., perfect Bayesian equilibrium, as well as the plural, i.e., perfect Bayesian equilibria.

low price reveals no additional information about the object's value, a sale at the high price enables future buyers to infer the present buyer's positive signal, and hence the monopolist can demand even higher prices in the future. On the other hand, failure to sell at the high price reveals that the respective buyer has received a signal that indicates low value, and the monopolist is forced to charge lower prices in the future. Thus, demanding a high price is like an investment with an uncertain outcome. We say that herding occurs when the seller continues to charge the same price and buyers purchase the product independent of their signals. In such a situation future buyers cannot infer anything from the earlier buyers' behavior, i.e., there is an informational cascade and social learning stops. The aim of our analysis is to characterize the monopolist's optimal pricing strategy, and its implication for the temporal pattern of prices and herding.

We show that the monopolist's optimal pricing strategy is very sensitive to whether the price demanded by the seller is publicly observed. If prices are publicly observed, herding will occur with positive probability. However, a patient seller will trigger herding only if the true state is known with high probability. Moreover, for certain parameter constellations the seller triggers herding only when public information indicates that the expected value of the object is high. This implies that while herding occurs when the expected value of the object reaches a sufficiently high level, this critical level may never be reached and in this case herding will not occur. Interestingly, the seller continues to demand the high price as the public evaluation of the object decreases, and the seller will not trigger herding in order to prevent the buyers from learning that the true value of the object is low.

The monopolist's optimal pricing strategy is quite different when buyers don't observe the price at which previous purchasing decisions were made. We show that PBE in pure strategies in which the seller's strategy is Markov have the unique outcome that the low price always is charged and hence herding arises immediately, provided the seller is patient or the quality of the buyer's signal is poor. Although previous prices are not observed these prices are perfectly inferred in a pure strategy equilibrium. Since buyers' beliefs are consistent with the seller's strategy, a purchase at the low price does not reveal any information. Thus, all subsequent periods will be identical to the first. Therefore, once it is uniquely optimal for the seller to charge the low price, the seller's optimal price continues to be the low price, and buyers make no inferences from earlier sales. As a result, the seller would like to commit to observable prices. The analysis reveals that a monopolist with the option of secretly setting prices will be unable to manipulate the aggregation of information. This has large implications for the possibility of social learning. While buyers may aggregate information when prices are observable, this possibility is entirely absent when previous prices are not observed.

Bikhchandani, Hirshleifer, and Welch (1998) suggest that in order to strategically mislead

buyers, “a seller may be tempted to cut price secretly for early buyers, so that later buyers will attribute the popularity of the product to high quality rather than low price (p. 165).” Our analysis shows that rational buyers cannot be deceived by secret discounts. Thus, such an argument implicitly assumes that buyers are naïve in the sense that they do not expect secret price cuts.

To analyze the seller’s incentive to manipulate learning we consider the case where the monopolist posts a publicly observable price at each point in time, but may grant a buyer an unobservable discount. In contrast to the previous two cases we relax the rationality assumption and assume that buyers are naïve in the sense that they believe that previous buyers paid the posted price. As suggested by Bikhchandani, Hirshleifer, and Welch (1998), the seller may manipulate information aggregation when she can offer secret discounts and buyers are naïve. Specifically, the seller will post the high price and charge the low price when the expected value of the object is sufficiently low, causing naïve buyers to increase their evaluation of the object. When the expected value of the object is sufficiently large the seller both posts and charges the low price. As a result, herding will occur in finite time, provided the seller is patient or the quality of the buyer’s signal is poor. Surprisingly we find that generally it will not be optimal for the seller to always offer secret discounts until she triggers herding.

Related to the present paper is the literature on optimal experimentation, where buyers learn from the experiments of earlier buyers. The objective of this literature is to determine how the equilibrium level of experimentation compares to the efficient level. See for example Bergeman and Välimäki (2000) who determine the equilibrium in a model where an entering firm offers a good of uncertain value and engages in price competition with existing firms. Buyers receive a publicly observable and noisy signal on the product’s quality, and sellers account for this accumulation of information when choosing the optimal price. The primary difference between the optimal experimentation literature and our analysis is that we assume that buyers only observe previous purchasing decisions and not the actual signal. We can therefore analyze a situation in which a monopolist optimally may choose to charge a price such that herding occurs and learning stops. In the framework of the optimal experimentation literature learning only stops if competitors choose to undercut the price of the incumbent. Our analysis is therefore more closely related to that of the herding literature than to the literature on optimal experimentation.

Since the pioneering articles of Banerjee (1992) and Bikhchandani et al. (1992) a large number of herding models has appeared.³ In contrast to our model the standard assumption of this literature is that all buyers face the same price. Welch (1992), Avery and Zemsky

³References can be found, for example, in Devenow and Welch (1996), Gale (1996), Vives (1996), and Bikhchandani et al. (1998). For a general analysis see Smith and Sørensen (2000).

(1998), Neeman and Orosel (1999), and Ottaviani (1999) are four exceptions. The article by Welch (1992) provides the first rigorous discussion of strategic pricing in a herding-prone environment. Whereas path-dependent pricing is at the center of our analysis, in Welch (1992) it is considered only briefly as a supplement to the main analysis, and in particular with respect to a risk-averse issuer (p. 708-9). The reason is that Welch aims to explain underpricing of initial public offerings (IPO's). Since the S.E.C. has banned variable-price sales, the issuer can offer shares to the public only at a fixed common price. Consequently path-dependent pricing is not an interesting issue in the context of IPO's.

Avery and Zemsky (1998) relax the fixed-price assumption in a market where a posted *competitive* price reflects the aggregated information. They show that a cascade cannot occur in such an environment unless an additional dimension of uncertainty is introduced.

While in our model the monopolist can sell arbitrary many units of an object, Neeman and Orosel (1999) consider the case where the monopolist is selling one single unit. Potential buyers bid sequentially for this unit. Because of the winner's curse they bid below the object's expected value conditional on public information, and the bids increase in value until herding occurs. At this point the monopolist sells the object for the bid that has been reached. In Neeman and Orosel the monopolist must decide whether to accept the current offer. In our model this is reversed: the buyer must decide whether to accept the monopolist's offer.

To the best of our knowledge, Ottaviani (1999) is the only paper that is closely related to our work.⁴ Similar to our study, he considers monopoly pricing with social learning. While the two papers are complementary in their objectives, they differ in their methods. Furthermore, Ottaviani provides a thorough analysis only of the case where prices are observable, and focuses his interest on welfare analysis and applications. In contrast, we analyze not only the case of observable prices (with a more general treatment of cost), but we also examine rigorously the case of unobservable prices and the case of secret discounts with naive buyers.

The paper proceeds as follows. The model is presented in Section 2. Section 3 deals with the case of observable prices, and Section 4 with the case of unobservable prices. In Section 5 we relax the assumption that buyers are rational and examine the case of secret discounts and naïve buyers. Finally, we summarize the results and discuss extensions in Section 6. Most of the proofs are collected in the Appendix.

2. The Model

We consider a risk-neutral monopolist who sequentially sells identical, indivisible objects to a countably infinite set of risk-neutral buyers. Each buyer buys at most one unit of the

⁴We are indebted to Christophe Chamley for making us aware of Ottaviani's paper. At that time we had already finished our own independent research on the subject.

object. The object's value is common to all buyers, and depends on the state of nature $\omega \in \Omega = \{G, B\}$. In the "good" state $\omega = G$ the value, $\hat{v}(G)$, is "high," whereas in the "bad" state $\omega = B$ the value, $\hat{v}(B)$, is "low," i.e., $\hat{v}(B) < \hat{v}(G)$. Without loss of generality we can choose the monetary unit such that $\hat{v}(G) - \hat{v}(B) = 1$. The state of nature is unknown to the seller and the buyers, however it is common knowledge among the seller and the buyers that the a priori probability of the good state G is $\lambda_1 \in (0, 1)$. We assume that the seller has a constant marginal cost $c \geq 0$. The seller has to incur this cost whenever she sells a unit of the object, but if her offer to sell the object is rejected, the cost c does not accrue. This can be interpreted, for example, as production to order. Time is measured in discrete periods $t = 1, 2, \dots$. In each period the seller approaches a randomly chosen buyer that she has not previously approached. The buyer observes a private signal about the state of nature, and decides whether to buy one unit of the object at the price demanded by the monopolist. We refer to the potential buyer in period t as buyer t and denote his random signal by $S_t \in \mathcal{S}$ with realization s_t , where $\mathcal{S} = \{g, b\}$ denotes the signal space. A signal realization g indicates the good state G and is called the "good signal," and a realization b indicates the bad state B and is called the "bad signal." Conditional on the true state ω , buyers' signals are independent and identically distributed for all buyers, and they are imperfectly informative. The signals are correct with probability $\alpha \in (\frac{1}{2}, 1)$, and incorrect with probability $1 - \alpha$. That is, $\Pr[s_t = g \mid \omega = G] = \Pr[s_t = b \mid \omega = B] = \alpha$ and $\Pr[s_t = g \mid \omega = B] = \Pr[s_t = b \mid \omega = G] = 1 - \alpha$.

In each period $t \in \{1, 2, \dots\}$ the seller demands a price \hat{p}_t from buyer t . For the analysis it is useful to transform the variables $\hat{v}(\omega)$, $\omega \in \{G, B\}$, and \hat{p}_t , $t \in \{1, 2, \dots\}$, by subtracting $\hat{v}(B)$. Therefore, we define $v(\omega) \equiv \hat{v}(\omega) - \hat{v}(B)$, $\omega \in \{G, B\}$, and $p_t \equiv \hat{p}_t - \hat{v}(B)$, $t \in \{1, 2, \dots\}$. For simplicity we will call these transformed variables the value of the object and its price, respectively. Whenever it matters, it will be clear from the context whether we refer to the transformed or to the untransformed variables. Due to the normalization $\hat{v}(G) - \hat{v}(B) = 1$, the transformation gives $v(B) = 0$ and $v(G) = 1$. Moreover, since it is common knowledge that the seller can always sell at some price above the object's minimum value and will never be able to sell at a price at or above the object's maximum value, only prices $p_t \in (0, 1)$ need be considered.

For each buyer the action space is $A = \{0, 1\}$. The action of buyer t is denoted by a_t , where $a_t = 1$ means that he buys one unit of the object, and $a_t = 0$ that he does not buy the object. Each buyer t observes the actions of all the previous buyers a_τ , $\tau \in \{1, \dots, t - 1\}$. The payoff for any buyer t from purchasing the object at a price \hat{p}_t is $\hat{v}(\omega) - \hat{p}_t = v(\omega) - p_t$, and it is zero if the buyer refrains from buying the object. Each buyer t updates the probability that the true state is the good state according to the actions of previous buyers, the prices he observes or believes that previous buyers have been charged, and his own private signal.

The buyer buys the object if and only if the expected payoff from doing so is non-negative; that is, if and only if the price does not exceed the expected value of the object conditional on his updated beliefs.⁵

The seller discounts future revenues according to a discount factor $\delta \in [0, 1)$. Her payoff is the discounted sum of her profits, $\sum_{t=1}^{\infty} \delta^{t-1} (\hat{p}_t - c) a_t = \sum_{t=1}^{\infty} \delta^{t-1} [p_t + \hat{v}(B) - c] a_t$. A particular case is the one where c equals $\hat{v}(B)$. In this case, $\sum_{t=1}^{\infty} \delta^{t-1} (\hat{p}_t - c) a_t = \sum_{t=1}^{\infty} \delta^{t-1} p_t a_t$ and thus the seller's payoff is determined by the transformed prices and the buyers' actions, independent of the cost c and the objects minimum value $\hat{v}(B)$, provided they are equal. In each period $t \in \{2, 3, \dots\}$, the seller knows the prices p_τ she has demanded from previous buyers, and the previous buyers' actions a_τ , $\tau \in \{1, \dots, t-1\}$. The *full history* of demanded prices and buyers' actions at time $t \in \{1, 2, \dots\}$ is denoted by $H_t = (p_1, a_1, \dots, p_t, a_t)$. The history H_0 is given by the empty set. The set of all possible full histories is denoted by \mathcal{H} . A pure strategy for the seller is a function $P : \mathcal{H} \rightarrow (0, 1)$ that maps every full history H_{t-1} into a price p_t , $t \in \{1, 2, \dots\}$.

The model constitutes a game between the seller and the buyers. For this game *herding* is defined as follows.

Definition. *Herding occurs at time T , if for all $t \geq T$ the seller charges a constant price $p_t = p_T$ and all buyers $t \geq T$ purchase the object regardless of their signal realizations.*

Herding implies an informational cascade: no signals can be inferred from the actions a_t , $t \in \{T, T+1, \dots\}$. The alternative conceivable herding situation, where the price is constant and all buyers *refuse* to buy regardless of their signal realizations, is not in the seller's interest and thus cannot occur in equilibrium.⁶

Although the model constitutes a game, this game is relatively simple and the PBE can be derived almost directly from the seller's optimal strategy. The seller can easily anticipate how the buyers react to her dynamic pricing strategy, and she maximizes her expected payoff accordingly. In particular, the seller observes the full history and knows what buyers can observe and what their beliefs about her own strategy are. Corresponding to the two signal realizations $s_t = b$ and $s_t = g$, the seller considers two types of buyer t , say type b and type g , respectively. If she finds it optimal to charge a high price p_t such that only type g buys the good, then buyer t 's signal s_t is revealed to be high if he buys, and low if he declines to buy the object. If the seller finds it optimal to charge a low price p_t such that type b buys the good, then type g will buy as well. In the case where prices are observable and buyers are rational, charging the low price renders a purchase in period t uninformative. Therefore, the situation facing the seller and buyer $t+1$ is identical to the one that the seller and buyer

⁵For technical reasons we assume the tie-breaking rule that a buyer purchases the good when he is indifferent.

⁶When $c > \hat{v}(B)$ the seller has the option of exiting the market.

t were confronted with in period t , and the previous low price p_t is also optimal in period $t + 1$. Hence both types of buyers will purchase the good in period $t + 1$ and the argument can be repeated for all the following periods. For observable prices this shows that herding arises, if in any period t the seller finds it optimal to charge a price p_t such that buyer t buys the good regardless of his type. Consequently, along the equilibrium path the posted prices separate the two types of buyers until herding occurs, and agents can correctly infer the private signals until herding arises.

We use the letter λ to denote the seller's updated probability of the good state. Specifically, let $\lambda_t \equiv \Pr(\omega = G \mid \lambda_1; H_{t-1})$ denote the seller's probability of the good state conditional on the full history H_{t-1} . Whenever the seller charges the high price she can perfectly infer all the buyers' signals. Therefore, any λ_t must be an element of a countable set of λ 's that is defined as follows. For any prior $\lambda_1 \in (0, 1)$ we define the set of λ 's that can be "attained" from λ_1 by

$$\Lambda(\lambda_1) \equiv \left\{ \lambda \mid \begin{array}{l} \text{there exists an integer } T \text{ and a sequence of signal realizations} \\ (s_1, \dots, s_T) \in \{g, b\}^T \text{ such that } \lambda = \Pr(\omega = G \mid \lambda_1; s_1, \dots, s_T) \end{array} \right\} \\ \subset (0, 1)$$

Furthermore, for any given $\lambda^l \in \Lambda(\lambda_1)$, we define for all $k \in \{1, 2, \dots\}$, $\lambda^{l+k} \equiv \Pr(\omega = G \mid \lambda^l, k \text{ signals } s = g)$ and $\lambda^{l-k} \equiv \Pr(\omega = G \mid \lambda^l, k \text{ signals } s = b)$, where $s \in \{g, b\}$ denotes the signal realization. For simplicity we let $\lambda^+ \equiv \Pr(\omega = G \mid \lambda, s = g)$ and $\lambda^- \equiv \Pr(\omega = G \mid \lambda, s = b)$. Thus, a buyer who learns λ and receives a good signal (type g) assigns the updated probability λ^+ to the good state. Similarly, a buyer who learns λ and receives a bad signal (type b) assigns the updated probability λ^- to the true state being good.

As noted above there are two types of buyer t , corresponding to the two signal realizations $s_t = g$ and $s_t = b$. The seller's associated conditional probabilities of the buyer's type are determined by λ_t and given by $\Pr(s_t = g \mid H_{t-1}) = \lambda_t \alpha + (1 - \lambda_t)(1 - \alpha) \equiv \varphi(\lambda_t)$ and $\Pr(s_t = b \mid H_{t-1}) = \lambda_t(1 - \alpha) + (1 - \lambda_t)\alpha = 1 - \varphi(\lambda_t)$, respectively. Denote the probability of the good state that buyer t infers from the history that he observes or perceives by μ_t , where $\mu_t \in \Lambda(\lambda_1)$. This probability does not include the information that buyer t derives from his signal. It is buyer t 's probability of the good state "before" he observes his private signal. When prices are observable, the seller's and the buyers' inference from the history coincide, that is, $\mu_t = \lambda_t$ for all t . If prices are unobservable, this holds along any pure strategy equilibrium path (because each buyer can deduce all previous prices from the seller's pure strategy), but not if the seller deviates. When the seller may mislead buyers by offering secret discounts, buyers' beliefs evolve as if posted prices were actual prices and thus the seller can always infer the (possibly mistaken) inference of each buyer from the perceived

history.

We can simplify the analysis significantly, if we first assume that the marginal cost c equals the minimum value $\hat{v}(B)$ of the object. At the end of the paper we discuss separately the cases $c < \hat{v}(B)$ and $c > \hat{v}(B)$. If $c = \hat{v}(B)$, the game can be analyzed in terms of the transformed variables $v(\omega)$, $\omega \in \{G, B\}$, and p_t , $t \in \{1, 2, \dots\}$, without any reference to the parameters c and $\hat{v}(B)$, which can both be ignored. The assumption $c = \hat{v}(B)$ is not only analytically convenient, for some cases it is also quite plausible. For example, in the case of a license for a patent the seller has no marginal cost and the patent may be without value in the bad state, i.e., $c = \hat{v}(B) = 0$.

3. Observable Prices

In this section we characterize the PBE in the case where the demanded price is publicly observed. That is, each buyer $t \in \{1, 2, \dots\}$ is informed of the *full* history H_{t-1} , and the seller and buyer t draw the same inference λ_t , which is common knowledge.

Buyer t 's strategy is simple: buy if and only if the seller demands a price $p_t \leq E[v(\omega) | \lambda_t, s_t] = \Pr(\omega = G | \lambda_t, s_t)$. Consequently, at every λ_t the seller need only consider two possible prices: she either charges a high price such that only type g buys the object, or she charges a low price such that both buyer types buy the object. Specifically, the seller must choose between:

$$p^H(\lambda_t) \equiv \lambda_t^+ = \frac{\lambda_t \alpha}{\lambda_t \alpha + (1 - \lambda_t)(1 - \alpha)}$$

and

$$p^L(\lambda_t) \equiv \lambda_t^- = \frac{\lambda_t(1 - \alpha)}{\lambda_t(1 - \alpha) + (1 - \lambda_t)\alpha}$$

which we will call the ‘‘high price’’ and the ‘‘low price,’’ respectively. It is important to keep in mind that these prices are both functions of $\lambda_t \in [0, 1]$ and not constants.

Given λ_t , no other prices can be optimal for the seller. The reason is that no sale will occur at a price $p > p^H(\lambda_t)$, and the seller unnecessarily loses rent by charging a price p such that $p^H(\lambda_t) > p > p^L(\lambda_t)$, or $p < p^L(\lambda_t)$. Given buyers' strategies it is clear that whenever either $p^H(\lambda_t)$ or $p^L(\lambda_t)$ is uniquely optimal, λ_t determines the seller's optimal price p_t in period t . If $p^H(\lambda_t)$ and $p^L(\lambda_t)$ are both optimal for some $\lambda_t \in \Lambda(\lambda_1)$, the seller may condition the price p_t on aspects of the history H_{t-1} that are not reflected in λ_t . However, in all three cases the seller's maximum expected payoff from period t onwards is uniquely determined by λ_t . For any $t \in \{1, 2, \dots\}$ and $\lambda_t \in \Lambda(\lambda_1)$, we denote this payoff by $V(\lambda_t)$. That is, $V : \Lambda(\lambda_1) \rightarrow \mathbb{R}$ is the seller's value function in the case where all buyers $t \in \{1, 2, \dots\}$ learn the seller's information λ_t .

If buyer t is charged the low price, then he will purchase the object regardless of his signal and consequently $\lambda_{t+1} = \lambda_t$. For the seller this implies that when $p^L(\lambda_t)$ is optimal at λ_t , then $V(\lambda_t) = p^L(\lambda_t) + \delta V(\lambda_{t+1}) = p^L(\lambda_t) + \delta V(\lambda_t)$. That is,

$$V(\lambda_t) = \frac{p^L(\lambda_t)}{1 - \delta}.$$

As argued above, this shows that whenever the low price $p^L(\lambda_t)$ is uniquely optimal for some λ_t , herding is triggered or continued in period t .⁷ On the other hand if the high price $p^H(\lambda_t)$ is uniquely optimal, then the seller's expected payoff exceeds the one she would get if she triggered herding and thus

$$V(\lambda_t) > \frac{p^L(\lambda_t)}{1 - \delta}.$$

The stochastic process of the updated probabilities $\{\lambda_t\}_{t=1}^{\infty}$ is a martingale. If the seller charges the low price $p^L(\lambda_t)$ at some t , $E(\lambda_{t+1} | \lambda_t) = \lambda_t$ because no information is revealed. If the seller demands the high price $p^H(\lambda_t)$, buyer t 's signal realization s_t will be revealed by his action and therefore

$$E(\lambda_{t+1} | \lambda_t) = \varphi(\lambda_t) \lambda_t^+ + [1 - \varphi(\lambda_t)] \lambda_t^- = \lambda_t.$$

It can be shown that in contrast to the stochastic process of the updated probabilities $\{\lambda_t\}_{t=1}^{\infty}$, the stochastic process of the seller's optimal price is a martingale only after the seller has triggered herding.

For the analysis of the seller's optimal decision it is useful to distinguish between her expected *immediate return* and her expected *future return*. First we examine the seller's expected immediate return, that is, her expected return in period t . The seller's immediate return from charging buyer t the price $p_t \in \{p^L(\lambda_t), p^H(\lambda_t)\}$ is

$$E[p_t a_t | \lambda_t] = \begin{cases} p^L(\lambda_t) = \lambda_t^- = \frac{\lambda_t(1-\alpha)}{\lambda_t(1-\alpha) + (1-\lambda_t)\alpha} & \text{for } p_t = p^L(\lambda_t) \\ p^H(\lambda_t) \Pr(s_t = g | \lambda_t) = \alpha \lambda_t & \text{for } p_t = p^H(\lambda_t) \end{cases}$$

since $p^H(\lambda_t) = \lambda_t^+$ and $\Pr(s_t = g | \lambda_t) = \lambda_t \alpha + (1 - \lambda_t)(1 - \alpha)$. The difference between the expected immediate return from the two prices is

$$\alpha \lambda_t - p^L(\lambda_t) = \frac{\lambda_t [\alpha^2(1 - \lambda_t) - (1 - \alpha)(1 - \alpha \lambda_t)]}{1 - \varphi(\lambda_t)}$$

where $\varphi(\lambda_t) \equiv \Pr(s_t = g | \lambda_t) < \alpha < 1$.

⁷In the event that buyer t does not purchase the object at the low price $p^L(\lambda_t)$ the belief about his type (i.e., about his signal s_t) by the seller and later buyers is irrelevant for the equilibrium outcome.

For $\lambda \in [0, 1]$ the associated low price $p^L(\lambda) = \frac{(1-\alpha)}{(1-\alpha) + \frac{(1-\lambda)}{\lambda}\alpha}$ is a strictly convex function with $p^L(0) = 0$ and $p^L(1) = 1$, whereas $\alpha\lambda$, the expected immediate return from the high price, is a linear function of $\lambda \in [0, 1]$ with $\alpha\lambda = 0$ for $\lambda = 0$ and $\alpha\lambda = \alpha < 1$ for $\lambda = 1$. Thus, either $\alpha\lambda < p^L(\lambda)$ for all $\lambda > 0$ or the two curves have a unique intersection for $\lambda > 0$. Figure 1 illustrates the two cases, with the first case in panel (a) and the second in panel (b).




Figure 1: Immediate expected return from the high and low price, respectively.

If $\alpha^2 \leq (1 - \alpha)$, the term $\alpha^2(1 - \lambda_t) - (1 - \alpha)(1 - \alpha\lambda_t) \leq -(1 - \alpha)^2\lambda_t < 0$ and thus $\alpha\lambda_t - p^L(\lambda_t) < 0$ for all $\lambda_t \in (0, 1)$. When $\alpha^2 > (1 - \alpha)$,

$$\alpha^2(1 - \lambda) - (1 - \alpha)(1 - \alpha\lambda) = \begin{cases} \alpha^2 - (1 - \alpha) > 0 & \text{for } \lambda = 0 \\ -(1 - \alpha)^2 < 0 & \text{for } \lambda = 1 \end{cases}$$

Note that the term $\alpha^2(1 - \lambda) - (1 - \alpha)(1 - \alpha\lambda)$ is continuous in $\lambda \in [0, 1]$, and because $\alpha > \frac{1}{2}$ it decreases in λ . Hence for $\alpha^2 > (1 - \alpha)$ there is a unique $\bar{\lambda}_\alpha \in (0, 1)$ such that

$$\alpha^2(1 - \bar{\lambda}_\alpha) - (1 - \alpha)(1 - \alpha\bar{\lambda}_\alpha) = 0. \quad (3.1)$$

Thus, if $\alpha^2 > (1 - \alpha)$,

$$\alpha\lambda_t - p^L(\lambda_t) \begin{cases} > 0 & \text{for } \lambda_t < \bar{\lambda}_\alpha \\ = 0 & \text{for } \lambda_t = \bar{\lambda}_\alpha \\ < 0 & \text{for } \lambda_t > \bar{\lambda}_\alpha \end{cases}$$

for some $\bar{\lambda}_\alpha \in (0, 1)$. This proves the following lemma.

Lemma 1. *If $\alpha^2 \leq (1 - \alpha)$, the immediate return from the low price $p^L(\lambda_t)$ is larger than the expected immediate return from the high price $p^H(\lambda_t)$ for all $\lambda_t \in \Lambda(\lambda_1)$. If $\alpha^2 > (1 - \alpha)$, there exists a $\bar{\lambda}_\alpha \in (0, 1)$ such that the expected immediate return from the high price $p^H(\lambda_t)$ is identical to the immediate return from the low price $p^L(\lambda_t)$ for $\lambda_t = \bar{\lambda}_\alpha$, whereas it is larger for $\lambda_t < \bar{\lambda}_\alpha$ and smaller for $\lambda_t > \bar{\lambda}_\alpha$.*

Next we examine the seller's expected future return, which is the sum of the expected discounted returns from the next period onwards. The expected *future* return to the seller is the sum of the expected discounted returns from the next period onwards. Thus, if the seller charges the low price $p^L(\lambda_t)$ in t and triggers herding, then in period t her discounted future return is $\delta \frac{p^L(\lambda_t)}{1-\delta}$. The fact that $p^L(\lambda)$ is strictly convex for $\lambda \in [0, 1]$ implies that the expected future return from charging the high price always exceeds that of the low price.

Lemma 2. *For all $\lambda_t \in \Lambda(\lambda_1)$ the expected future return to the seller from charging the high price $p^H(\lambda_t)$ in period t strictly exceeds that from charging the low price $p^L(\lambda_t)$ in period t .*

Proof: The function $p^L(\lambda) = \frac{(1-\alpha)}{(1-\alpha)+\frac{(1-\lambda)}{\lambda}\alpha}$, where $\lambda \in [0, 1]$, is strictly convex; and $V(\lambda_t) \geq \frac{1}{1-\delta} p^L(\lambda_t)$ for all $\lambda_t \in \Lambda(\lambda_1)$. If the seller charges the high price $p^H(\lambda_t)$ in t , the expected future return is

$$\begin{aligned} \delta \varphi(\lambda_t) V(\lambda_t^+) + \delta [1 - \varphi(\lambda_t)] V(\lambda_t^-) &\geq \delta \varphi(\lambda_t) \frac{p^L(\lambda_t^+)}{1-\delta} + \delta [1 - \varphi(\lambda_t)] \frac{p^L(\lambda_t^-)}{1-\delta} \\ &> \delta \frac{p^L(\lambda_t)}{1-\delta} \end{aligned}$$

because $\lambda_t = \varphi(\lambda_t) \lambda_t^+ + [1 - \varphi(\lambda_t)] \lambda_t^-$ and $p^L(\lambda)$ is strictly convex. ■

Combining the expected immediate and future return, Lemma 1 and 2 imply that when $\alpha^2 > (1 - \alpha)$ and $\lambda_t \leq \bar{\lambda}_\alpha$ the seller always charges the high price $p^H(\lambda_t)$, since that price maximizes both the immediate and the future return. Thus, for $\alpha^2 > (1 - \alpha)$ herding will never arise at low λ_t 's. Surprisingly, the seller will not trigger herding in order to prevent buyers from asymptotically learning that the bad state is the true state.⁸ For $\alpha^2 \leq (1 - \alpha)$ demanding the high price is a risky investment as the seller sacrifices some immediate return for a higher expected future return. A sufficiently patient seller will undertake such an investment, but an impatient seller will trigger herding.

Although it is possible that buyers learn that the true state is bad with an arbitrarily large probability, the converse is not true. The following lemma shows that the seller triggers herding when buyers believe that the true state is good with a sufficiently large probability. That is, given λ_1 the probability λ_t is bounded away from 1 along the equilibrium path.

⁸This conclusion is due to the result that $\alpha^2 > (1 - \alpha)$ implies $\varphi(\lambda_t) p^H(\lambda_t) > p^L(\lambda_t)$ for all $\lambda_t \in (0, \bar{\lambda}_\alpha)$ and is not driven by the fact that $\varphi(\lambda) p^H(\lambda) = p^L(\lambda) = 0$ for $\lambda = 0$.

Lemma 3. *Whenever the seller's updated probability of the good state, λ_t , is sufficiently high, the seller charges the low price. That is, for every discount factor $\delta \in (0, 1)$ there exists an $\epsilon_\delta > 0$ such that $p_t = p^L(\lambda_t)$ whenever $\lambda_t \in (1 - \epsilon_\delta, 1)$.*

Proof: The seller's payoff from charging $p^L(\lambda_t)$ in period t equals $\frac{p^L(\lambda_t)}{1-\delta}$. Since the seller's price always is less than 1, her expected payoff from charging the high price $p^H(\lambda_t)$ is less than $\alpha\lambda_t + \frac{\delta}{1-\delta}$. The difference, $p^L(\lambda_t) - \alpha\lambda_t + \delta \frac{p^L(\lambda_t)-1}{1-\delta}$, converges to $1 - \alpha > 0$ for $\lambda_t \rightarrow 1$. Thus, $p^L(\lambda_t)$ generates a higher expected payoff than $p^H(\lambda_t)$ whenever λ_t is sufficiently large. ■

Lemma 3 is intuitive. If λ_t is already high, the potential increase in the seller's expected future return from an increase of λ_t even to its upper limit of 1 is small. On the other hand, if λ_t is high, the expected immediate return from the high price, $\alpha\lambda_t$, is significantly smaller than the immediate return from the low price, $p^L(\lambda_t)$, because $p^L(\lambda_t)$ is almost identical to λ_t for high λ_t . Therefore, when the probability of the good state is sufficiently high the seller prefers to trigger herding rather than to aim at a further increase of this probability. Not surprisingly the point at which herding is triggered is sensitive to the seller's degree of patience. In fact, given any fixed probability $\lambda_t \in (0, 1)$ a sufficiently patient seller will charge the high price $p^H(\lambda_t)$ at that λ_t (see Lemma 4, Appendix A.1).

The quality of the buyers' signal, α , is critical for the seller's optimal strategy. For example, if α is close to 1, that is, if the signal is almost perfect, even an extremely impatient seller will demand the high price unless λ_t is close to 1. This follows directly from Lemma 1, Lemma 2, and the fact that $\bar{\lambda}_\alpha \rightarrow 1$ for $\alpha \rightarrow 1$. The significance of signal quality α for the seller's optimal strategy (and thus for the PBE) has the consequence that we need to distinguish three cases that differ with respect to the quality of the signal. These are the cases (i) $\alpha^2 > 1 - \alpha$ or equivalently, $\alpha > \frac{1}{2}(\sqrt{5} - 1) \sim 0.618$; (ii) $\alpha^2 = 1 - \alpha$ or equivalently, $\alpha = \frac{1}{2}(\sqrt{5} - 1)$; and (iii) $\alpha^2 < 1 - \alpha$ or equivalently, $\alpha < \frac{1}{2}(\sqrt{5} - 1)$. We say that the signal is "strong" in case (i) and "weak" in case (iii). Case (ii) we call the borderline case. In the borderline case the likelihood ratio $\frac{1-\alpha}{\alpha}$ equals the probability α that the signal is correct.⁹

3.1. The Seller's Optimal Strategy When the Signal is Strong

In the case of strong signals the seller's optimal price has the following characteristics.

Proposition 1. *Assume that prices are observable. If $\alpha^2 > 1 - \alpha$, there exists a critical*

⁹Note that the likelihood ratio $\frac{1-\alpha}{\alpha}$ determines the function $p^L(\lambda_t)$, which gives the immediate return associated with charging the low price as a function of λ_t . The probability α determines the expected immediate return $\alpha\lambda_t$ associated with charging the high price as a function of λ_t . For $\lambda_t = 0$ the respective derivatives are $\frac{dp^L(\lambda_t)}{d\lambda_t} = \frac{1-\alpha}{\alpha}$ and $\frac{d(\alpha\lambda_t)}{d\lambda_t} = \alpha$. Incidentally, $\alpha^2 = 1 - \alpha$ is the equation for the "golden section."

probability $\mu^* \geq \bar{\lambda}_\alpha$, $\mu^* \in \Lambda(\lambda_1)$, such that it is uniquely optimal for the seller to demand

$$p_t = \begin{cases} p^H(\lambda_t) & \text{whenever } \lambda_t < \mu^* \\ p^L(\lambda_t) & \text{whenever } \lambda_t > \mu^* \end{cases}$$

For $\lambda_t = \mu^*$, $p^L(\lambda_t)$ is optimal, but $p^H(\lambda_t)$ may be optimal as well. Moreover, $\mu^* > \bar{\lambda}_\alpha$ for $\delta > 0$, $\mu^* = \min_{\lambda \in \Lambda(\lambda_1) \cap [\bar{\lambda}_\alpha, 1]} \lambda$ for $\delta = 0$, and $\mu^* \rightarrow 1$ for $\delta \rightarrow 1$.¹⁰

Proof: Appendix A.1. ■

Proposition 1 shows that the set of attainable λ 's can be partitioned such that for low λ 's the high price is optimal and for high λ 's the low price is optimal. The intuition is that at high λ 's there is little to gain and much to lose from demanding the high price (Lemma 3). At low λ 's the converse holds: there is little to lose and much to gain from demanding the high price. The reason is that when there is a large probability that there will be no sale at the high price, the (high) price and thus the loss from not selling is small.

If the prior λ_1 is sufficiently large, the seller triggers herding immediately. Otherwise she demands the high price $p^H(\lambda_1)$ in period 1. Buyer 1 buys if and only if he has received the good signal. His action is publicly observed, and the seller and future buyers update their beliefs accordingly. If buyer 1 bought the object (which reveals $s_1 = g$) and $\lambda_1^+ = \mu^*$, then the seller will charge $p_2 = p^L(\lambda_1^+) = \lambda_1 < \lambda_1^+ = p^H(\lambda_1) = p_1$ and trigger herding in period 2.¹¹ If buyer 1 bought the object but $\lambda_1^+ < \mu^*$, then the seller charges $p^H(\lambda_1^+)$. In the remaining case where buyer 1 did not purchase the object, the seller charges $p^H(\lambda_1^-)$ in period 2. In this way the process continues. Unless the seller triggers herding in $t = 2$, she charges $p^H(\lambda_t)$ in each period $t = \{3, 4, \dots, T\}$ until, if ever, λ_t hits μ^* at some $t = T$. At this point she lowers the price from $p^H(\lambda_{T-1}) = \lambda_{T-1}^+$ to $p^L(\lambda_T) = \lambda_T^- = \lambda_{T-1} < \lambda_{T-1}^+$ and triggers herding. However, λ_t may never hit μ^* and consequently herding may never arise. If herding does not occur, λ_t will converge to zero due to the martingale convergence theorem. That is, it will asymptotically be revealed that the bad state is the true state.

3.2. The Seller's Optimal Strategy in the Borderline Case

In the borderline case the seller's patience, as measured by her discount factor δ , determines the pattern of her optimal pricing strategy. An impatient seller triggers herding immediately, whereas a patient seller follows a strategy that is analogous to the optimal strategy in the case of strong signals. With a patient seller herding may but need not arise. Specifically, the following proposition holds.

¹⁰In general only $p^L(\mu^*)$ will be optimal at $\lambda_t = \mu^*$ because $\Lambda(\lambda_1)$ is a discrete set.

¹¹This assumes that in the improbable case that $p^L(\mu^*)$ and $p^H(\mu^*)$ are both optimal, the seller chooses $p^L(\mu^*)$. The intuitive explanations below are also based on this simplifying assumption.

Proposition 2. *Assume that prices are observable. If $\alpha^2 = 1 - \alpha$, there exists a discount factor $\delta^* \in [0, 1)$ such that for all $\delta \in [0, \delta^*]$ the uniquely optimal prices are given by $p_t = p^L(\lambda_1)$ for all $t \in \{1, 2, \dots\}$. For each $\delta \in (\delta^*, 1)$ there exists a critical probability $\mu^* \in \Lambda(\lambda_1)$ such that it is uniquely optimal for the seller to demand*

$$p_t = \begin{cases} p^H(\lambda_t) & \text{whenever } \lambda_t < \mu^* \\ p^L(\lambda_t) & \text{whenever } \lambda_t > \mu^* \end{cases}$$

For $\lambda_t = \mu^$, $p^L(\lambda_t)$ is optimal, but $p^H(\lambda_t)$ may be optimal as well. Finally, $\mu^* \rightarrow 0$ for $\delta \rightarrow \delta^*$, and $\mu^* \rightarrow 1$ for $\delta \rightarrow 1$.*

Proof: Appendix A.1. ■

In contrast to the case of strong signals, in the borderline case the high price $p^H(\lambda_t)$ generates a lower expected immediate return than the low price $p^L(\lambda_t)$ even for small λ_t 's (Lemma 1). For an impatient seller the higher expected future return that is associated with the high price $p^H(\lambda_1)$ is not sufficient to compensate for the lower immediate return. Therefore, for an impatient seller the low price $p^L(\lambda_1)$ is always uniquely optimal and herding arises immediately. For a patient seller the borderline case is similar to the case where the signal is strong. The difference is that $\bar{\lambda}_\alpha$, as defined by (3.1), is zero in the borderline case, but positive in the case of strong signals.

3.3. The Seller's Optimal Strategy When the Signal is Weak

For the case of weak signals we show that the seller's optimal strategy has the following characteristics.

Proposition 3. *Assume that prices are observable. If $\alpha^2 < 1 - \alpha$, there exist discount factors $\delta^{**} \in (0, 1)$ and $\delta^{***} \in [\delta^{**}, 1)$ such that*

- *for all $\delta \in [0, \delta^{**})$, the uniquely optimal prices are given by $p_t = p^L(\lambda_1)$ for all $t \in \{1, 2, \dots\}$;*
- *for $\delta \in [\delta^{**}, \delta^{***}]$, $p_t = p^L(\lambda_t)$ is optimal for all $\lambda_t \in \Lambda(\lambda_1)$, but $p_t = p^H(\lambda_t)$ is optimal as well for at least one $\lambda_t \in \Lambda(\lambda_1)$;*
- *for each $\delta \in (\delta^{***}, 1)$ there exist critical probabilities $\mu^* \in \Lambda(\lambda_1)$ and $\mu^{**} \in \Lambda(\lambda_1) \cup \{0\}$, $0 \leq \mu^{**} < \mu^* < 1$, such that it is uniquely optimal for the seller to demand*

$$p_t = \begin{cases} p^H(\lambda_t) & \text{whenever } \lambda_t \in (\mu^{**}, \mu^*) \\ p^L(\lambda_t) & \text{whenever } \lambda_t \in (0, \mu^{**}) \cup (\mu^*, 1); \end{cases}$$

*for $\lambda_t \in \{\mu^{**}, \mu^*\}$, $p^L(\lambda_t)$ is optimal, but $p^H(\lambda_t)$ may be optimal as well; for $\delta \rightarrow 1$, $\mu^{**} \rightarrow 0$ and $\mu^* \rightarrow 1$.*

Proof: Appendix A.1. ■

An impatient seller will always choose the low price $p^L(\lambda_1)$ and trigger herding immediately, because $p^L(\lambda_1)$ generates a higher immediate return (Lemma 1). If the seller is patient, herding will not occur immediately for priors λ_1 that lie within some range (μ^{**}, μ^*) . However, if $\mu^{**} > 0$, herding will arise eventually. Finally, the case where the seller is patient and $\mu^{**} = 0$ is analogous to the case of strong signals. If the prior λ_1 is not too high, there is a positive probability that herding will not arise. In this event λ_t converges to zero and it is asymptotically revealed that the bad state is the true state.

3.4. Summary

When prices are public information herding may but need not arise. Depending on the parameter constellation, the seller either initiates herding immediately or starts with a high price, relative to the current public evaluation of the object. In the latter case she continues to do so as long as the updated public evaluation of the object is within a certain interval. Along this path the price follows a stochastic process where herding constitutes an absorbing barrier. The absorbing barrier is optimally chosen by the seller and thus the seller's problem can also be seen as one of optimal stopping. As soon as the price exceeds a critical level and the buyer actually buys at this price, the seller reduces the price somewhat and triggers herding. However, the price may never hit this critical level. Instead, the price may converge to zero and thereby reveal that the common value of the object is low. Surprisingly, except when the quality of the signals is poor, the seller will not trigger herding in order to prevent the buyers from learning that the true value of the object is low.

A decrease in signal quality α increases the likelihood of herding. For any given probability λ that the object's value is high, a decrease in α increases the low price $p^L(\lambda)$ whereas it decreases $\alpha\lambda$, the expected immediate return from the high price. Moreover, the effect of any given sequence of revealed signal realizations on the price that the seller can achieve, decreases with α . Therefore, it is intuitively plausible that a decrease in α decreases the expected immediate and future return associated with the high price relative to the return from herding. Consequently, the seller's incentive to trigger herding increases when α decreases.

A decrease in the seller's degree of patience δ also increases the likelihood of herding, but the reason is different. Whenever the high price $p^H(\lambda)$ generates a lower expected immediate return than the low price $p^L(\lambda)$, to demand the high price is an investment where some immediate return is sacrificed for a higher expected future return. Thus, only a sufficiently patient seller will charge the high price $p^H(\lambda)$, whereas a less patient seller will demand the low price $p^L(\lambda)$ and trigger herding. Consequently, herding is more likely when the seller is less patient.

4. Unobservable Prices

In this section we examine the case where it is common knowledge that buyers observe the history of purchasing decisions but do not observe the price at which these decisions were made. The price demanded by the seller may be unobservable to later buyers simply because it is not revealed publicly.¹² Or, it may be that prices effectively are unobservable because buyers rationally recognize that the seller has an incentive to manipulate information aggregation by offering secret discounts.

The seller observes the full history, and in period t she believes that the likelihood of the good state equals $\lambda_t \equiv \Pr(\omega = G \mid \lambda_1; H_{t-1})$. For simplicity our notation omits that λ_t depends on the buyers' strategies and beliefs. Buyer t only observes the public history $h_{t-1} = (a_1, \dots, a_{t-1})$, where h_0 is the empty set. Given the seller's optimal strategy $P : \mathcal{H} \rightarrow (0, 1)$, buyer t updates $\mu_t \equiv \Pr(\omega = G \mid P; \lambda_1; h_{t-1}, p_t)$ from the public history h_{t-1} , where for simplicity our notation omits that μ_t depends also on the strategies and beliefs of the buyers $\tau \in \{1, \dots, t\}$. The seller's beliefs are omitted because given the seller's strategy they are not important for the buyers' inference. Common knowledge of the prior λ_1 implies that $\mu_1 = \lambda_1$. In addition, along the equilibrium path of any PBE in pure strategies, buyers can infer the unobserved prices from the seller's equilibrium strategy and from the observable actions of previous buyers. Consequently, in equilibrium the seller and the buyers make the same inferences from history, and along the equilibrium path $\mu_t = \lambda_t$ for all $t \in \{1, 2, \dots\}$. If in any period t the seller deviates from her equilibrium strategy P^* , the buyers $\tau \geq t + 1$ cannot detect this as long as the observable action a_t is consistent with the seller's equilibrium strategy and buyers will still infer $\mu_\tau \equiv \Pr(\omega = G \mid P^*; \lambda_1; h_{\tau-1}, p_\tau)$ from the public history. Hence λ_t need not equal μ_t off the equilibrium path.¹³

We will show that immediate herding is an equilibrium outcome, and moreover this outcome is unique when we restrict the seller's equilibrium strategy (but not her deviation strategies) to be a pure Markov strategy. We define a Markov strategy as follows.

Definition. *A pure strategy $P : \mathcal{H} \rightarrow (0, 1)$ of the seller is Markov, if for any history $H_{t-1} \in \mathcal{H}$ it prescribes a price p_t that depends only on $\lambda_t \equiv \Pr(\omega = G \mid \lambda_1; H_{t-1})$ and*

¹²For example, neither the owner of a patent nor the licensee may reveal the price at which the license to use the patent was sold.

¹³In the case of unobservable prices our notion of a PBE consists of the requirements for a weak PBE as defined by Mas-Colell, Whinston, and Green (1995, p. 285) plus the requirement that players' beliefs are consistent with common knowledge of the structure of the game and of rationality of all players. This implies that a purchase by some previous buyer cannot be interpreted as indicating that the respective buyer has observed the bad signal. The reason is that any rational buyer who is willing to accept the seller's offer after having observed the bad signal would accept that offer had he instead observed the good signal.

$$\mu_t \equiv \Pr(\omega = G \mid P^*; \lambda_1; h_{t-1}, p_t).^{14}$$

We consider only equilibria in pure strategies because equilibria in mixed strategies are intractable. If the seller uses a mixed strategy at some t , future buyers know only the probability distribution of the price that was demanded at t , whereas the seller knows the actual price. Consequently, the seller and the future buyers update their probabilities of the good state differently after they have observed the action a_t of buyer t . Since the price demanded by the seller depends on her updated probability of the good state, which the buyers cannot infer perfectly, buyers cannot determine previous prices and the analysis soon gets intractable. In contrast, buyers can perfectly infer all previously quoted prices in a pure strategy equilibrium.

We first show that immediate herding is an equilibrium outcome. In proving existence of such an outcome we specify the buyers' strategies and beliefs, and show that for any $\lambda_t = \mu_t$ it is a best response for the seller to charge $p_t = p^L(\mu_t)$. Consistent beliefs of a buyer are that the low price was charged when the object was sold, and that the high price was demanded when it was not sold. That is, buyers $\tau \in \{2, 3, \dots\}$ believe that the price \hat{p}_t that was demanded in period $t \in \{1, \dots, \tau - 1\}$ is

$$\hat{p}_t = \begin{cases} p^L(\mu_t) & \text{if } a_t = 1 \\ p^H(\mu_t) & \text{if } a_t = 0 \end{cases} \quad (4.1)$$

Thus, whenever buyers unexpectedly observe that no sale has taken place, they infer that the seller has deviated from her equilibrium strategy by demanding the high price. Moreover, buyers believe that after the deviation the seller continues to quote the low price associated with the buyers' updated beliefs. Since buyers always believe that $\mu_t = \lambda_t$, in which case $p^L(\mu_t)$ is optimal, these beliefs are consistent with the seller's strategy. The beliefs (4.1) of buyer τ are about *previous* prices and are independent of the present price p_τ that the seller demands from buyer τ . If the seller deviates and unexpectedly demands the high price $p^H(\mu_\tau)$ instead of the low price $p^L(\mu_\tau)$ from any buyer τ , the deviation does not influence the respective buyer's beliefs about the prices that have been previously charged. Similarly, the observation that no sale has taken place does not induce buyers to revise their belief that in the past the seller has charged the low price whenever there was a sale. This is consistent with the buyer's information. Each buyer's strategy is to purchase the object, if and only if the price demanded by the monopolist does not exceed the object's expected value conditional on the respective buyer's information and beliefs.

We show that given the buyers' strategies and beliefs it is a best response for the seller

¹⁴Along the equilibrium path $\mu_t = \lambda_t = \Pr(\omega = G \mid \lambda_1; H_{t-1}) = \Pr(\omega = G \mid \lambda_{t-1}, p_{t-1}, a_{t-1})$ for all $t \in \{2, 3, \dots\}$ is a Markov process.

to charge $p_t = p^L(\mu_t)$ whenever $\mu_t = \lambda_t$. Therefore, it is optimal for the seller to charge the constant price $p_t = p^L(\lambda_1)$ for all λ_1 . If the seller optimally charges $p_1 = p^L(\lambda_1)$ in period 1, the buyer's purchase does not reveal any information and the strategic situation is the same in period 2, hence $p_2 = p^L(\lambda_1)$ is optimal in period 2. Applying the same argument repeatedly, implies that $p_t = p^L(\lambda_1)$ is optimal for all $t \in \{1, 2, \dots\}$, if $p^L(\lambda_1)$ is optimal at $t = 1$. Hence, to show that immediate herding is an equilibrium outcome, it is sufficient to show that it is not optimal for the seller to deviate and charge $p_1 = p^H(\lambda_1)$. Note that given buyers' beliefs as specified by (4.1), a deviation in any period t to the high price implies $\mu_{t+1} \leq \mu_t$, whereas λ_{t+1} either increases or decreases relative to λ_t . Furthermore, it must be that $\lambda_t \geq \mu_t$, $t \in \{1, 2, \dots\}$.

If the seller deviates to $p_1 = p^H(\lambda_1)$, this has three consequences: (i) instead of $p^L(\lambda_1)$, the seller's expected immediate return in period 1 is $\varphi(\lambda_1)p^H(\lambda_1)$, where as before $\varphi(\lambda_t) \equiv \lambda_t\alpha + (1 - \lambda_t)(1 - \alpha)$ denotes the seller's probability that buyer t observes the good signal; (ii) instead of $\lambda_2 = \lambda_1$, the seller's updated probability of the good state in period 2 is $\lambda_2 = \lambda_1^+ > \lambda_1$ if there was a sale, and $\lambda_2 = \lambda_1^- < \lambda_1$ if there was no sale in period 1; (iii) instead of $\mu_2 = \lambda_1$, the future buyers' inference is $\mu_2 = \lambda_1$ only if there was a sale, whereas it is $\mu_2 = \lambda_1^-$ if there was no sale in period 1. Clearly, the third consequence is disadvantageous for the seller. The first consequence is good for the seller if $\varphi(\lambda_1)p^H(\lambda_1) > p^L(\lambda_1)$ and bad if $\varphi(\lambda_1)p^H(\lambda_1) < p^L(\lambda_1)$. Since in any period t the return from the low price is $p^L(\mu_t)$, which is independent of λ_t , the second consequence is relevant for the seller only if she considers demanding the high price at some point in the future. However, because of the second consequence we cannot rule out that the seller deviates even when this reduces the immediate return in period 1, i.e., when $\varphi(\lambda_1)p^H(\lambda_1) < p^L(\lambda_1)$.

Since buyers' beliefs imply $\mu_t \leq \lambda_1$ and thus $p^H(\mu_t) \leq p^H(\lambda_1)$, it follows that $\varphi(\lambda_t)p^H(\mu_t) \leq \varphi(\lambda_t)p^H(\lambda_1) < \alpha p^H(\lambda_1)$, that is, the seller's expected immediate return from the high price is always less than $\alpha p^H(\lambda_1)$. Moreover, because $p^L(\mu_t) \leq p^L(\lambda_1)$ for all $t \in \{1, 2, \dots\}$, the seller's expected immediate return from the low price can never exceed $p^L(\lambda_1)$. Consequently, if $p^L(\lambda_1) \geq \alpha p^H(\lambda_1)$, any deviation from $p_t = p^L(\lambda_1)$ to $p_t = p^H(\lambda_1)$ will reduce the seller's expected payoff. Simple calculation shows that the condition $p^L(\lambda_1) \geq \alpha p^H(\lambda_1)$ is equivalent to $\lambda_1 \geq \frac{\alpha^3 - (1-\alpha)^2}{\alpha^3 - (1-\alpha)^3} \equiv \bar{\lambda}_\alpha$. Thus, whenever $\lambda_1 \geq \bar{\lambda}_\alpha$ immediate herding is an equilibrium outcome, irrespective of the seller's degree of patience. This (sufficient) condition is violated for small priors λ_1 . However, in the case of weak or borderline signals (i.e., $\alpha^2 \leq 1 - \alpha$) we can extend the result to low priors (i.e., $\lambda_1 < \bar{\lambda}_\alpha$). The following proposition collects these results.

Proposition 4. *Assume that previous prices are unobservable for buyers. If $\alpha^2 \leq 1 - \alpha$ or $\alpha^2 > 1 - \alpha$ and $\lambda_1 \in [\bar{\lambda}_\alpha, 1)$, there exists a PBE such that $p_t = p^L(\lambda_1)$ for all $t \in \{1, 2, \dots\}$. Thus, there is a PBE where the seller always charges the price $p^L(\lambda_1)$ and herding arises*

immediately at $t = 1$.

Proof: Appendix A.2. ■

The intuition for existence of a PBE that irrespective of the seller's degree of patience has immediate herding as outcome rests on two arguments. First, a deviation to the high price at best leaves future buyers' evaluation of the object unchanged and reduces it with positive probability. Second, under the assumptions of the proposition the expected immediate return from the low price exceeds that from the high price. The only benefit the deviation has for the seller is that she learns the respective buyer's signal, but under the assumptions of the proposition that turns out to be without value.

When the expected immediate return from the high price exceeds that of the low price the situation is different. In this case a sufficiently impatient seller will demand the high price $p_1 = p^H(\lambda_1)$ in period 1. However, a patient seller will give more weight to the fact that future buyers reduce their evaluation of the object whenever no sale occurred. Given buyers' beliefs (4.1) the seller's expected return in the far future is certainly maximized by always charging the low price. Thus, immediate herding is an equilibrium outcome provided the seller is sufficiently patient. This is confirmed by Proposition 5, which covers the cases not considered in Proposition 4.

Proposition 5. *Assume that previous prices are unobservable for buyers. If $\alpha^2 > 1 - \alpha$ and $\lambda_1 \in (0, \bar{\lambda}_\alpha)$, there exists a $\bar{\delta} \in (0, 1)$ such that for all $\delta > \bar{\delta}$ there exists a PBE where $p_t = p^L(\lambda_1)$ for all $t \in \{1, 2, \dots\}$. Thus, there is a PBE where the seller always charges the price $p^L(\lambda_1)$ and herding arises immediately at $t = 1$, provided the seller is sufficiently patient.*

Proof: Appendix A.2. ■

Next we address the question of uniqueness. We will prove that if the seller's equilibrium strategy is restricted to be a pure Markov strategy, then immediate herding is the unique equilibrium outcome of all pure strategy PBE.¹⁵ Proving uniqueness is complicated by the fact that a buyer may reinterpret the public history when confronted with a deviation by the seller. In particular, a buyer who is charged the low rather than the high price may believe that the seller has deviated in the past, and thus revise his updating from the public history. For example, observing a deviation to the low price may convince the buyer that, with the exception of the cases where no object was sold, the seller has never previously demanded the high price. Consequently, the respective buyer's evaluation of the object may decrease drastically, with the effect that the seller can sell the object only for an extremely low price.

¹⁵ Alternative pure Markov strategy PBE differ only in the seller's strategy at nodes that are not reached in equilibrium.

Such off-the-equilibrium-path beliefs act like a punishment of the seller and may support other equilibria.

We show first that there does not exist an equilibrium where the seller always demands the high price independent of the past history. This result is not only useful to prove uniqueness, but is also interesting because the strategy to demand the high price is the one that maximizes social learning, whereas immediate herding implies that there is no social learning at all. Thus, if there existed a PBE where the seller always demands the high price, maximal and minimal social learning could both be equilibrium outcomes. Lemma 10 shows that this is not the case.

Lemma 10. *If either (i) $\alpha\lambda_1 \leq \lambda_1^-$, i.e., signals are weak or borderline, or signals are strong and $\lambda_1 \in [\bar{\lambda}_\alpha, 1)$, or (ii) $\delta > \hat{\delta}$ for some sufficiently large $\hat{\delta} \in (0, 1)$, then there does not exist a PBE where for all signal realizations the price is $p_t = p^H(\lambda_t)$ for all $t \in \{1, 2, \dots\}$, that is, where the seller always demands the high price.*

Proof: Appendix A.3. ■

The reason why we cannot sustain an equilibrium where the high price always is charged is that the seller has an incentive to deviate to the low price. This deviation cannot be detected by future buyers and thus will increase the seller’s expected future return. Furthermore, in the weak and in the borderline case the seller’s expected immediate return also increases when she deviates to the low price in period 1, hence it is not possible to sustain an equilibrium where the high price is always charged. In the case of strong signals the seller must sacrifice some immediate return to increase the buyer’s evaluation of the object and the argument is more complicated. However, at some nodes that are reached with positive probability the loss in the expected immediate return is sufficiently low to be outweighed by expected future gains

To analyze uniqueness we consider only PBE where the seller’s equilibrium strategy is a pure Markov strategy, that is, in each period t the two conditional probabilities of the good state, λ_t and μ_t , determine the seller’s optimal price p_t . Notice that the seller’s deviation strategies are not required to be Markovian. Whereas the two probabilities λ_t and μ_t provide information that is “intrinsically” relevant for the seller, other aspects of the history H_{t-1} can be relevant only if the seller and the buyers have somehow “coordinated on making them relevant.” It is conceivable that in addition to λ_t and μ_t some other aspects of the history H_{t-1} are relevant for the seller because buyer t expects the seller to condition the price p_t on these aspects and “punishes” her if she deviates to a different price. Perhaps such a “coordination on intrinsically irrelevant aspects of the history” can implement a PBE where the seller does not trigger herding immediately. If such a PBE exists, then the respective “coordination” acts as a commitment device for the seller not to “cheat” by deviating and

charging a lower price than prescribed by her equilibrium strategy in order to deceive future buyers. In our analysis we choose to focus solely on strategies that rely only on information that is intrinsically relevant.

We consider first the case where either $\alpha^2 \leq 1 - \alpha$ or $\alpha^2 > 1 - \alpha$ and $\lambda_1 \geq \bar{\bar{\lambda}}_\alpha$. Both cases have in common that along the equilibrium path the expected immediate return from the low price exceeds that from the high price. The intuition for the uniqueness result is therefore straightforward. Suppose there is an equilibrium where $p_1 = p^H(\lambda_1)$. If the seller deviates and charges $p_1 = p^L(\lambda_1)$ instead, buyer 1 purchases the object regardless of his signal and future buyers falsely update $\mu_2 = \lambda_1^+$. Such a deviation is beneficial for the seller because her future as well as her immediate expected return increase.

Proposition 6. *Assume that previous prices are unobservable for buyers. If either $\alpha^2 \leq 1 - \alpha$ or $\alpha^2 > 1 - \alpha$ and $\lambda_1 \in [\bar{\bar{\lambda}}_\alpha, 1)$, then any pure strategy PBE where the seller's equilibrium strategy is Markov has immediate herding, i.e., $p_t = p^L(\lambda_1)$ for all $t \in \{1, 2, \dots\}$, as equilibrium outcome. That is, under these conditions immediate herding is the unique equilibrium outcome.*

Proof: Appendix A.4. ■

Finally, we consider the remaining case where $\alpha^2 > 1 - \alpha$ and $\lambda_1 < \bar{\bar{\lambda}}_\alpha$. In this case the expected immediate return from the high price exceeds that from the low price in period 1.¹⁶ We show that there does not exist a PBE where the seller's strategy is Markovian and $p_1 = p^H(\lambda_1)$, provided the seller is sufficiently patient. To do this we conjecture a PBE that has $p_1 = p^H(\lambda_1)$ and show that this leads to a contradiction if the seller is sufficiently patient. Since deviating from $p_1 = p^H(\lambda_1)$ to $p_1 = p^L(\lambda_1)$ is now costly for the seller and, moreover, may have only temporary beneficial effects, the intuition that underlies Proposition 6 does not generally apply even for a very patient seller. In addition, if the seller deviates at some $t > 1$ from $p_t = p^H(\lambda_t)$ to some other price, buyer t will necessarily notice this deviation and may conclude that the seller has also deviated previously. In that case buyer t may not be willing to buy the object for the low price $p_t = p^L(\lambda_t)$ because he revises his beliefs about the true state. The lowest possible probability of the good state that a rational buyer t may infer from the public history h_{t-1} is the one that is based on the assumption that when there was no sale the respective previous buyer had observed a bad signal, whereas whenever there was a sale the previous buyer has bought the object only because the seller had deviated to a sufficiently low price. This makes the proof of uniqueness as well as the intuition for it more complicated. However, it can be shown that for any conjectured PBE that starts with

¹⁶Actually this is not true for $\lambda_1 \in [\bar{\lambda}_\alpha, \bar{\bar{\lambda}}_\alpha)$ and therefore the uniqueness proof of Proposition 6 could be applied for $\lambda_1 \geq \bar{\lambda}_\alpha$ as well (and not only for $\lambda_1 \geq \bar{\bar{\lambda}}_\alpha$). However, Proposition 4 on existence relies on $\lambda_1 \geq \bar{\bar{\lambda}}_\alpha$ and because of this we also have to assume $\lambda_1 \geq \bar{\bar{\lambda}}_\alpha$ in Proposition 6.

$p_1 = p^H(\lambda_1)$ there is some node that is reached with positive probability where the seller gets a permanent increase in the expected future revenues from deceiving future buyers by deviating to a sufficiently low price. At such a node the immediate loss from charging a sufficiently low price to induce a sale is outweighed by the associated permanent increase in the expected future revenues, provided the seller is sufficiently patient. Therefore, such a seller will deviate at this node and the conjectured equilibrium unravels. Consequently, if the seller is sufficiently patient there can be no PBE where the seller's strategy is Markovian and $p_1 = p^H(\lambda_1)$.

Proposition 7. *Assume that previous prices are unobservable for buyers. If $\alpha^2 > 1 - \alpha$ and $\lambda_1 \in (0, \bar{\lambda}_\alpha)$, there exists a $\bar{\delta} \in (0, 1)$ such that for each $\delta > \bar{\delta}$ any pure strategy PBE where the seller's equilibrium strategy is Markov has immediate herding, i.e., $p_t = p^L(\lambda_1)$ for all $t \in \{1, 2, \dots\}$, as equilibrium outcome. That is, under these conditions immediate herding is the unique equilibrium outcome.*

Proof: Appendix A.4. ■

Essentially the intuition for uniqueness is that the seller cannot commit not to “cheat” by charging a lower price than the one prescribed by her equilibrium strategy. With unobservable prices the seller has the option to cheat, except when there is immediate herding; and in order to mislead future buyers a sufficiently patient seller will, in fact, cheat at some node that is reached with positive probability. The only case in which the seller will be unable to cheat is when there is immediate herding, at least when the seller's strategy is Markov. Thus only immediate herding can be sustained as an equilibrium outcome.

5. Secret Discounts and Naïve Buyers

As mentioned in the Introduction, it has been suggested that by offering secret discounts the seller can strategically mislead buyers to increase their evaluation of the object. The previous section clearly demonstrates that rational buyers cannot be misled in this way. In this section, we therefore relax the assumption that buyers are rational and examine the seller's incentive to manipulate the aggregation of information when faced with a population of naïve buyers.

In contrast to our earlier analysis we assume that the monopolist posts a publicly observable price p_t , but may secretly offer a discount d_t , where $d_t \in [0, 1)$. The offer to the buyer still has the form of a “take it or leave it offer,” where the price demanded by the monopolist equals $p_t - d_t$. In each period t the seller knows the *full* history given by $H_{t-1} = (p_1, d_1, a_1, \dots, p_{t-1}, d_{t-1}, a_{t-1})$. Given this history she chooses an action $(p_t, d_t) \in (0, 1) \times [0, 1)$.

Buyers observe previously posted prices and the associated actions, but they do not

observe the secret discounts. We assume that buyers are naïve, in the sense that they are unaware of even the possibility of secret discounts as long as the seller’s posted prices are consistent with the optimal strategy of an “honest” seller. Therefore, the seller’s posted prices must be consistent with the optimal price in the observable prices case. If, incorrectly, the seller’s posted price is high when it should be low in the observable prices case, the buyers become aware of the discount possibility and forever thereafter they will believe that the price actually charged by the seller is low. Being offered a discount does not cause the same type of belief revision. Rather, when a naïve buyer is offered a discount, he is convinced by the seller that he is a special customer and that no one before him has ever received a discount.

Naïve buyers update their beliefs according to the *perceived* history. The perceived history is defined as the history of posted prices and actions together with the belief that posted prices are actually charged as long as the high price is not posted when the low price is optimal in the observable prices case.¹⁷ For each $t \in \{1, 2, \dots\}$ the *perceived* history is given by $\chi_t = (p_1, 0, a_1, \dots, p_t, 0, a_t)$, which is the full history with all discounts being replaced by zero. We define χ_0 to be the empty set. Buyer t ’s associated updated probability is denoted by $\mu_t \equiv \Pr(\omega = G \mid \lambda_1; \chi_{t-1})$, and the seller’s updated probability by $\lambda_t \equiv \Pr(\omega = G \mid \lambda_1; H_{t-1})$. The seller perfectly infers μ_t and posts either the associated high price $p^H(\mu_t) = \mu_t^+$ or the associated low price $p^L(\mu_t) = \mu_t^-$. Given buyers’ beliefs the seller will post the price that is optimal in the observable prices case for the realization $\lambda_t = \mu_t$.¹⁸ Whenever μ_t assumes a value such that $p^L(\mu_t)$ would be the seller’s uniquely optimal price in the observable prices case, the seller posts (and charges) $p^L(\mu_t)$ and triggers herding.

Since it is optimal for the seller to actually charge either the high price $p^H(\mu_t)$ or the low price $p^L(\mu_t)$, only discounts that reduce the high price $p^H(\mu_t)$ to the low price $p^L(\mu_t)$ need be considered. We show first that when the expected immediate return from the low price exceeds that from the high price, then the seller never actually charges the high price. In particular, she gives a discount $d_t = p^H(\mu_t) - p^L(\mu_t)$ whenever she posts the high price $p^H(\mu_t)$, and thus buyer t actually pays only the low price and buys the object for sure. Charging the low price causes the seller’s return to increase for two reasons. First, the immediate return is the low price, which in the case considered exceeds the expected immediate return from demanding the posted high price; and second, by charging $p^L(\mu_t)$ the seller benefits from the fact that future buyers erroneously infer the good signal realization $s_t = g$ from buyer t ’s purchase. The seller’s only potential cost from charging the low price

¹⁷If the seller posts the low price when the high price should have been posted, buyers believe the posted price is charged. The buyers maintain this belief in the event that no sale occurred at the posted low price.

¹⁸Here λ_t refers to the updated probability of the good state that determines the equilibrium price *in the observable prices case*, not to the λ_t of this section. For simplicity we assume that a seller who is indifferent between posting the high and the low price always posts the low price.

is that she doesn't learn buyer t 's signal, but this foregone knowledge has no value. The reason is that the seller, by secretly charging the low price, can increase buyers' beliefs that the object is of high value, and effectively secure that μ_t reaches (in finitely many steps) a level where the uniquely optimal price is the low price. At this stage all buyers purchase the object and the true state of the world is irrelevant to the seller.

The condition that in period 1 the immediate expected return from the low price exceeds that from the high price is that either $\alpha^2 \leq 1 - \alpha$ or $\alpha^2 > 1 - \alpha$ and $\lambda_1 \geq \bar{\lambda}_\alpha$. For these parameter constellations the following proposition shows that whenever the high price is optimal in the observable prices case, the seller posts the high price and secretly grants a discount in order to deceive future buyers.

Proposition 8. *Assume that the seller may grant secret discounts and buyers are naïve. Let μ^* denote the critical probability of Proposition 1, 2, and 3, respectively, and $\mu^{*+} \equiv \Pr(\omega = G | \mu^*, s = g)$. If (i) $\alpha^2 \leq 1 - \alpha$ or (ii) $\alpha^2 > 1 - \alpha$ and $\lambda_1 \in [\bar{\lambda}_\alpha, 1)$, the seller immediately posts and actually charges the low price $p^L(\lambda_1)$ and triggers herding whenever this is uniquely optimal in the observable prices case; otherwise she posts the high price $p^H(\mu_t)$ and grants a secret discount $d_t = p^H(\mu_t) - p^L(\mu_t)$ for the first T periods $t \in \{1, \dots, T\}$, where T is a finite, deterministic integer. In the latter situation, $\mu_{T+1} = \mu^*$ if $p^L(\mu^*)$ is uniquely optimal at μ^* when prices are observable, and $\mu_{T+1} = \mu^{*+}$ if $p^H(\mu^*)$ is also optimal at μ^* when prices are observable. In period $T + 1$ the seller posts and charges $p^L(\mu^*)$ or $p^L(\mu^{*+})$, respectively, and triggers herding.*

Proof: Appendix A.5. ■

Finally, we consider the case when the signals are strong and the prior is sufficiently low (i.e., $\alpha^2 > 1 - \alpha$ and $\lambda_1 < \bar{\lambda}_\alpha$). We know that whenever μ_t gets sufficiently large the seller will post and charge the low price. Furthermore, secret discounts allow the seller to increase the buyers' public evaluation of the object to the level where herding occurs. However, for the parameter constellations considered now, the expected immediate return from the high price exceeds that from the low price, and hence there is a cost associated with offering a discount and deceiving future buyers. As a result, the price charged by the seller depends on her discount factor. In particular, a seller that doesn't value the future will never offer a discount, whereas a sufficiently patient seller will have an optimal strategy that prescribes her to offer discounts at least at some nodes. Notice, however, that in general it will not be optimal for a patient seller to *always* offer discounts until she triggers herding. The reason is that although demanding the high price without a discount may result in a decrease of future buyers' evaluations of the object (because there is no sale), the probability that there is a sale and thus that future buyers' evaluations increase is still positive and the seller need not incur the cost of the discount. Moreover, the seller always has the option of granting a secret discount at a later stage if necessary. Interestingly, there are histories in which even

a patient seller never cheats.

However, a sufficiently patient seller will always use secret discounts to prevent buyers' updated probabilities of the good state from becoming too low. Consequently, buyers' beliefs μ_t will almost surely reach in finite time a level where the seller triggers herding. This implies the following result for the parameter constellations that are not covered by Proposition 8. Provided the seller is sufficiently patient, herding will occur in finite time with probability 1. The intuition behind this result can be seen when considering the case where herding may realize with a positive probability that is less than 1. This implies that there must be a positive probability of signal realizations such that the seller becomes increasingly pessimistic about reaching the true value of this object and thus about her future revenues. However, by offering secret discounts the seller can prevent buyer's beliefs from falling below some threshold. Since from any such threshold there are only finitely many steps to the herding region, herding will realize with probability 1.

Proposition 9. *Assume that the seller may grant secret discounts and buyers are naïve. Let μ^* denote the critical probability of Proposition 1 and $\mu^{*+} \equiv \Pr(\omega = G \mid \mu^*, s = g)$. If $\alpha^2 > 1 - \alpha$ and $\lambda_1 \in (0, \bar{\lambda}_\alpha)$, there exists a discount factor $\bar{\delta} \in (0, 1)$ such that for all $\delta \in (\bar{\delta}, 1)$ the seller posts (and charges) the price $p^L(\mu^*)$ in finite time with probability 1 if $p^L(\mu^*)$ is uniquely optimal at μ^* when prices are observable, and posts (and charges) $p^L(\mu^{*+})$ in finite time with probability 1 if $p^H(\mu^*)$ is also optimal at μ^* when prices are observable. Thus, herding arises in finite time with probability 1, provided the seller is sufficiently patient.*

Proof: Appendix A.5. ■

If the seller can offer secret discounts and buyers are naïve, then the seller can manipulate information aggregation. This has the effect that with probability 1 herding will occur in finite time. In contrast to the case with observable prices, the bad state will never be asymptotically revealed if the seller is sufficiently patient. Rather than allowing social learning to reveal the bad state, the seller will post high prices relative to the public information and secretly grant discounts in order to deceive future buyers. In this way the seller makes sure that eventually the buyers' updated probability that the object is of high value exceeds a critical level. Then, as in the observable prices case, she reduces the posted price somewhat, triggers herding, and stops giving discounts.¹⁹ Some buyers may get discounts from the

¹⁹The situation is different if, in contrast to our assumption, buyers *always* believe that the seller charged the posted price, irrespective of which price is optimal in the observable prices case. It can be shown that in this case the seller makes sure that eventually the buyers' updated probability of the object being of high value exceeds a critical level, and from then on proceeds by always posting the high price and charging the low price. Consequently, herding never occurs and buyers' probability μ_t of the good state will converge to 1 for $t \rightarrow \infty$.

posted prices because in this way the seller can trick future buyers to make false positive inferences about earlier buyers' signals. Consequently, whenever the seller initially posts the high price $p^H(\lambda_1)$ the path of posted prices rises, at least eventually. As soon as it exceeds a critical level, the price drops somewhat and herding occurs.²⁰

6. Summary and Extensions

In addition to the complete characterization of the seller's optimal pricing policy for the case of observable prices, our analysis has provided three general insights:

1. If prices are observable, herding will occur with positive probability. However, a sufficiently patient seller will trigger herding only if the true state is known with high probability.
2. If previous prices are unobservable to buyers, herding will be more common than when prices are observable.
3. If the seller can grant secret discounts and buyers are naïve, a sufficiently patient seller will always trigger herding at a relatively high price (and for certain parameter constellations every seller is sufficiently patient).

While there may be social learning when both actions and prices are observed, this does not imply that sellers can manipulate aggregation of information by secretly changing their prices or hiding them altogether. The only way the seller can manipulate information aggregation is to trigger herding and thus to end learning. Rather than enabling deception the seller's opportunity to cheat will merely inhibit learning. Rational buyers cannot be fooled. For misleading manipulation to be successful it is necessary that buyers are not rational.

Next we examine the extent to which our results are influenced by the assumption that the seller's constant marginal cost c equals the minimum value $\hat{v}(B)$ of the object. We deal separately with the case $c < \hat{v}(B)$ and the case $c > \hat{v}(B)$. In our context the main difference between these two cases is that if $c > \hat{v}(B)$ it is optimal for the seller to exit the market when rational buyers have a sufficiently low estimate of the good state. Another difference is that

²⁰Our analysis of naïve buyers has focused exclusively on the case where the seller is restricted to making "take it or leave it offers." This is not the seller's preferred strategy when she is faced with buyers who never anticipate discounts. In such an environment the seller is strictly better off if she instead first quotes the high price, and then decreases the price to the low price if and only if the respective buyer rejects the offer. Since by assumption the buyer does not expect discounts to be given, he will not strategically decline to buy at the high price. Clearly, if the seller follows such a strategy, the path of posted prices always increases, until at some point herding occurs.

if $c \leq \hat{v}(B)$, then from an efficiency point of view, every potential buyer should purchase the object. Thus, if $c \leq \hat{v}(B)$ herding is efficient, and from an efficiency point of view the seller should never charge the high price as long as there are no other reasons to learn the true state.²¹ This is so in spite of the fact that herding prevents information aggregation, and the reason is simply that when $c \leq \hat{v}(B)$ information aggregation is useless. In contrast, if the cost c exceeds $\hat{v}(B)$ but is less than $\hat{v}(G)$, information aggregation is socially valuable.²²

Consider first the case where $c < \hat{v}(B)$ and prices are observable to buyers. For any $\lambda \in (0, 1)$ the immediate return net of cost from the low price is now $p^L(\lambda) + \bar{p}$, where $\bar{p} \equiv \hat{v}(B) - c > 0$. This is again a strictly convex function of λ . The expected immediate return net of cost from the high price is now $\varphi(\lambda) [p^H(\lambda) + \bar{p}] = [\alpha + (2\alpha - 1)\bar{p}]\lambda + (1 - \alpha)\bar{p}$, which again is a linear function of λ . For the difference $\Delta(\lambda) \equiv [p^L(\lambda) + \bar{p}] - \varphi(\lambda) [p^H(\lambda) + \bar{p}] = p^L(\lambda) + \bar{p} - \{[\alpha + (2\alpha - 1)\bar{p}]\lambda + (1 - \alpha)\bar{p}\}$ it holds that $\Delta(\lambda) > p^L(\lambda) - \varphi(\lambda)p^H(\lambda) = p^L(\lambda) - \alpha\lambda$ for $\bar{p} > 0$, and thus the low price is more attractive than in the case $\bar{p} = 0$, i.e., the case we analyzed in Section 3 where $c = \hat{v}(B)$. For λ sufficiently close to 1 the difference $\Delta(\lambda)$ is positive since $\Delta(1) = (1 - \alpha)(1 + \bar{p}) > 0$, which corresponds to the situation in Section 3. However, the difference $\Delta(\lambda)$ is positive for sufficiently small λ 's as well because $\Delta(0) = \alpha\bar{p} > 0$, whereas the respective difference was zero at $\lambda = 0$ in Section 3 where $\bar{p} = 0$. Consequently, for any $\delta \in [0, 1)$ the low price is optimal at sufficiently low as well as at sufficiently high values of λ . That is, in contrast to our previous result the seller will always trigger herding whenever λ is sufficiently low and prices are observable.

Since $p^L(\lambda) + \bar{p}$ is strictly convex and $\varphi(\lambda) [p^H(\lambda) + \bar{p}]$ is linear in λ , these two curves intersect at most twice. Given $\bar{p} > 0$, they will not intersect if α is sufficiently small. Thus, for those α 's the immediate return net of cost from the low price is always larger than the expected immediate return net of cost from the high price, and that corresponds to the case of weak signals in Section 3. On the other hand, if α is sufficiently large the two curves will intersect twice. This follows from the combination of two arguments. First, $\Delta(1) = (1 - \alpha)(1 + \bar{p})$ implies that at $\lambda = 1$ the point on $\varphi(\lambda) [p^H(\lambda) + \bar{p}]$ approaches the point on $p^L(\lambda) + \bar{p}$ from below as $\alpha \rightarrow 1$; second, for α close to 1 the slope of $p^L(\lambda) + \bar{p}$ is steeper than the (constant) slope of $\varphi(\lambda) [p^H(\lambda) + \bar{p}]$ because $dp^L(\lambda)/d\lambda = \alpha^2/(1 - \alpha)^2$ for $\lambda = 1$ and thus $dp^L(1)/d\lambda \rightarrow \infty$ for $\alpha \rightarrow 1$. Those α 's for which the two curves intersect twice correspond to the case of strong signals in Section 3. Given $\bar{p} > 0$, there also exists an α such that the two curves share a tangential point at some $\lambda \in (0, 1)$. However, this borderline case differs from the one in Section 3 because there the corresponding tangential point is at $\lambda = 0$, which is not attainable.

²¹This will not be the case if the efficient usage of the object depends on the state.

²²If $c \geq \hat{v}(G)$, the good should not, and will not, be produced.

Applying arguments of the analysis of Section 3 gives the following results for $\bar{p} > 0$. If, given $\bar{p} > 0$, the signals are strong or borderline in the sense that there is at least one $\lambda \in (0, 1)$ such that the expected immediate return net of cost from the low and the high price is the same, then the low price is optimal for the seller for sufficiently high and low λ 's, and the high price is optimal for all intermediate λ 's. When $\delta = 0$, the intermediate range shrinks to the tangential point in the borderline case and both prices are optimal at this point. If, given $\bar{p} > 0$, the signals are weak in the sense that the immediate return net of cost from the low price exceeds the expected immediate return net of cost from the high price for all $\lambda \in (0, 1)$, then a sufficiently impatient seller will charge the low price and trigger herding immediately. In contrast, a sufficiently patient seller will charge the low price only for sufficiently high and low λ 's, and demand the high price for all intermediate λ 's. Thus, the results are similar to those of Section 3, with the important modification that there will always be herding at low λ 's and thus herding will arise with probability 1.

It is easy to see that the arguments of Section 4 and 5 carry over to the case where $\bar{p} > 0$. When prices are unobservable to buyers, it cannot be an equilibrium move that a sufficiently patient seller charges $p_1 = p^H(\lambda_1)$, when her equilibrium strategy is a pure Markov strategy. This follows because she would benefit, if she deviates and demands the low instead of the high price at some node of any conjectured equilibrium path that starts with $p_1 = p^H(\lambda_1)$. If buyers are naive, the seller's incentives to give secret discounts persist in the case $\bar{p} > 0$. In fact, since the low price is more attractive relative to the high price when $\bar{p} > 0$, these incentives are even higher than in the case $\bar{p} = 0$.

Consider now the case $c \in (\hat{v}(B), \hat{v}(G))$, i.e., $\bar{p} \in (-1, 0)$. In this case, one of three alternatives is optimal for the seller in any period t : (i) demand the high price $p^H(\lambda_t)$, (ii) charge the low price $p^L(\lambda_t)$, or (iii) exit the market. We examine first the situation where prices are observable to buyers. If either (ii) or (iii) is optimal in some period t , the same decision is optimal in all later periods $\tau > t$. Moreover, for sufficiently small λ 's the seller's optimal decision is to exit the market. Thus, for the model to be interesting the prior λ_1 has to be sufficiently high to make (iii) suboptimal at $t = 1$. Since the seller has to incur the cost c only if she is able to sell the object, an increase in the cost c makes alternatives (i) and (iii) more attractive relative to alternative (ii). A sufficient (but not necessary) condition for (i) to be optimal is $p^H(\lambda_t) > -\bar{p} \geq p^L(\lambda_t)$. Thus, the seller charges $p^L(\lambda_t)$ and triggers herding only if $p^L(\lambda_t) > -\bar{p}$. If the cost c is close to $\hat{v}(G)$ and $p^L(\lambda_1) > -\bar{p} = c - \hat{v}(B)$, the seller will charge $p^L(\lambda_1)$ in $t = 1$ and trigger herding immediately. Since $p^L(\lambda) + \bar{p}$ is strictly convex (and increasing) and $\varphi(\lambda) [p^H(\lambda) + \bar{p}]$ is linear (and increasing) in λ , and since $\Delta(0) = \alpha\bar{p} < 0$ and $\Delta(1) = (1 - \alpha)(1 + \bar{p}) > 0$, these two curves intersect exactly once. Consequently, the seller's optimal policy is to charge the low price and trigger herding whenever λ_t exceeds a critical value (that depends on δ). Depending on the parameters,

the seller either exits or demands the high price whenever λ_t is below that critical value. In the latter case there exists a second critical value of λ , at which the seller exits. Thus, herding may, but need not occur. In particular, as in the case of Proposition 1, herding will never arise at small λ 's. However, learning will stop when $\lambda_t > 0$ becomes sufficiently small because the seller will exit. These arguments show that with the modification that the seller will exit whenever $\lambda_t > 0$ becomes sufficiently small, the analysis of Section 3 carries over to the case $\bar{p} \in (-1, 0)$.²³

The same is true for the analysis of Section 4. If the seller is sufficiently patient, a pure equilibrium strategy that is Markov cannot prescribe her to charge $p_1 = p^H(\lambda_1)$. This follows because she would benefit, if she deviates and demands the low instead of the high price at some node of any conjectured equilibrium path that starts with $p_1 = p^H(\lambda_1)$.²⁴ Finally, the arguments that underlie the results of Section 5 persist for a cost $c \in (\hat{v}(B), \hat{v}(G))$. Rather than using the exit option, a sufficiently patient seller will grant a secret discount and raise the buyers' beliefs about the good state. Consequently, the probability which buyers assign to the good state will almost surely reach in finite time a value such that the seller triggers herding. We conclude from this discussion that the simplifying assumption $c = \hat{v}(B)$ about the seller's cost is not responsible for the basic results of our analysis.

Another simplifying assumption of our model is that there are only two signal realizations and that the probability that the signal is correct does not depend on the state. The "symmetry" of the signal with respect to the two states is helpful for the analysis but not crucial for our results. Consider the case of $K > 2$ signal realizations instead of 2, where K is finite and no signal realization provides perfect information. In this case the seller is confronted with K types of buyers. In each period t , the optimal price will be equal to the updated expected value that one specific (path-dependent) type k_t of the K types assigns to the object, and buyer t will purchase the object if and only if he is of type k_t or "higher" (i.e., has a higher updated expected value than type k_t). Unfortunately the seller's optimization problem is not analytically tractable when there are K signal realizations. However, the intuition for the three general insights listed above carries over to this more general case.

²³Ottaviani (1999) studies a case where $\bar{p} \in (-1, 0)$. He normalizes, in our notation, $v(B) = -1$, $v(G) = 1$, and $c = 0$, which implies the assumption that $c = \frac{1}{2}[\hat{v}(B) + \hat{v}(G)]$. Because of this assumption on cost, the seller's optimal decision at the prior $\lambda_1 = \frac{1}{2}$ is to stay in the market and demand the high price. This holds regardless of the value of α and δ , respectively. However, without this assumption there need not exist a prior such that the high price is optimal for the seller. If, given α and δ , the cost c is sufficiently close to $\hat{v}(G)$, the high price is never optimal. Rather the low price is uniquely optimal for all priors where the seller does not exit the market, and thus the seller will never demand the high price.

²⁴The proof of Lemma 10 has to be modified because of the seller's exit option. But since after exit the seller's profits are zero, this can easily be done.

Consider the second finding. Whenever prices are unobservable the seller has an incentive to manipulate information aggregation by deviating to a price which is below the price she is expected to demand. The fact that the seller may “cheat” and cannot commit not to cheat, will force her, in equilibrium, to trigger herding when she would not do so in the case of observable prices. Because of this, herding will be more common when prices are unobservable. Similarly, the intuition for the other two findings carries over to the case of K signals.

Finally, the assumption that there are only two states of nature does not drive the results either. What matters is the function that maps signal realizations into updated expected values of the object, not the number of states. As long as the seller and rational buyers make identical inferences from the history, the structure of the model and therefore the basic results remain the same.

We conclude that the basic insights of our analysis are fairly robust for the case of a single seller.

Appendix

A.1. Proof of Propositions 1, 2, and 3

The proof of Propositions 1, 2 and 3 proceeds in three steps. First, we introduce a function $F(\lambda)$ that has the property that whenever $F(\lambda)$ is positive, the high price $p^H(\lambda)$ is uniquely optimal. Second, we provide conditions that imply that the low price $p^L(\lambda)$ is uniquely optimal whenever $F(\lambda)$ is negative. Third, we show that the set of λ 's where $F(\lambda) > 0$ consists of all attainable λ 's that lie in a connected interval, which may be empty. The third point together with the first two points implies that the set of λ 's where the high price is optimal also lies in a connected interval, which may be empty.

We proceed as follows. First we derive some results that are relevant for the proofs of all three Propositions. Then we continue with a lemma that we need for the proof of Proposition 1 and conclude the proof of Proposition 1. Next we prove a lemma that we need for the proof of Propositions 2 and 3. Finally we prove these two propositions.

Let $F(\lambda)$ denote the difference in the seller's expected discounted return between (i) charging the high price $p^H(\lambda)$ now and the low price $p^L(\lambda^-)$ or $p^L(\lambda^+)$, respectively (whatever the updated probability of the good state may be), from the next period onwards, and (ii) charging the low price now and forever.

That is

$$F(\lambda) \equiv \left\{ \alpha\lambda + \delta \left[\varphi(\lambda) \frac{p^L(\lambda^+)}{1-\delta} + [1 - \varphi(\lambda)] \frac{p^L(\lambda^-)}{1-\delta} \right] \right\} - \frac{p^L(\lambda)}{1-\delta}.$$

If $F(\lambda) > 0$, then $p^H(\lambda)$ is the uniquely optimal price. The converse does not hold. If $F(\lambda) < 0$ either price may be optimal. However, we show in Lemma 5 that if $p^H(\lambda)$ is optimal for some λ where $F(\lambda) < 0$, then it cannot be that $p^L(\lambda^+)$ and $p^L(\lambda^-)$ are both optimal at λ^+ and λ^- , respectively. Thus, whenever $p^H(\hat{\lambda})$ is optimal for some $\hat{\lambda}$ where $F(\hat{\lambda}) < 0$, or is uniquely optimal for some $\hat{\lambda}$ where $F(\hat{\lambda}) \leq 0$, it can be so only because the high price $p^H(\hat{\lambda})$ is “anchored” at some $\hat{\hat{\lambda}}$ where $F(\hat{\hat{\lambda}}) > 0$. That is, there must be a sequence of signals and corresponding updated λ 's that lead from $\hat{\lambda}$ to $\hat{\hat{\lambda}}$ such that at each of the updated λ 's the associated high price $p^H(\lambda)$ is optimal.

Given λ_1 , let λ^l be an arbitrary element in the set $\Lambda(\lambda_1)$ of attainable λ 's. Let $\varphi_l \equiv \varphi(\lambda^l) \equiv \Pr(s = g | \lambda^l)$. We now proceed with a series of lemmas, followed by the proof of Proposition 1.

Lemma 4. *Given any prior $\lambda_1 \in (0, 1)$, if the seller is sufficiently patient, then the high price is uniquely optimal and herding does not occur immediately, i.e., for each $\lambda_1 \in (0, 1)$, $\exists \bar{\delta} < 1$ such that $V(\lambda_1) > \frac{1}{1-\delta}p^L(\lambda_1)$ for all $\delta \in (\bar{\delta}, 1)$.*

Proof: All we have to show is that for each $\lambda_1 \in (0, 1)$, $F(\lambda_1) > 0$ if δ is sufficiently large. Rearranging $F(\lambda_1)$ gives

$$F(\lambda_1) = \alpha\lambda_1 - p^L(\lambda_1) + \frac{\delta}{1-\delta}[\varphi(\lambda_1)p^L(\lambda_1^+) + (1-\varphi(\lambda_1))p^L(\lambda_1^-) - p^L(\lambda_1)].$$

Since p^L is strictly convex, $[\varphi(\lambda_1)p^L(\lambda_1^+) + (1-\varphi(\lambda_1))p^L(\lambda_1^-) - p^L(\lambda_1)] > 0$. For $\delta \rightarrow 1$, the last term on the right hand side of the equality becomes arbitrarily large, whereas $\alpha\lambda_1 - p^L(\lambda_1)$ is independent of δ . Hence $F(\lambda_1) > 0$ for δ sufficiently large. \blacksquare

Lemma 5. *Let $\{\lambda^l, \lambda^{l+1}, \dots, \lambda^{l+K}\}$, $K \geq 2$, where $\lambda^{l+k} \equiv \Pr(\omega = G \mid \lambda^l, k \text{ signals } s = g)$, be a set of λ 's such that $F(\lambda^{l+k}) \leq 0$ for all $k \in \{0, \dots, K\}$. If $V(\lambda^{l+k}) > \frac{p^L(\lambda^{l+k})}{1-\delta}$ for all $k \in \{1, \dots, K-1\}$, then either $V(\lambda^l) > \frac{p^L(\lambda^l)}{1-\delta}$ or $V(\lambda^{l+K}) > \frac{p^L(\lambda^{l+K})}{1-\delta}$ (or both).*

Proof: The proof is by contradiction. Assume $V(\lambda^l) = \frac{p^L(\lambda^l)}{1-\delta}$ and $V(\lambda^{l+K}) = \frac{p^L(\lambda^{l+K})}{1-\delta}$. We simplify the notation by $V_{l+k} \equiv V(\lambda^{l+k})$, $k \in \{0, 1, \dots, K\}$. Recall that $p^L(\lambda^{l+k}) = \lambda^{l+k-1}$. We have

$$V_{l+k} = \begin{cases} \lambda^{l-1} + \delta V_l & \text{for } k = 0 \\ \alpha\lambda^{l+k} + \delta(1 - \varphi_{l+k})V_{l+k-1} + \delta\varphi_{l+k}V_{l+k+1} & \text{for } k \in \{1, \dots, K-1\} \\ \lambda^{l+k-1} + \delta V_{l+k} & \text{for } k = K \end{cases} \quad (6.1)$$

We define

$$V \equiv \begin{bmatrix} V_l \\ \cdot \\ \cdot \\ \cdot \\ V_{l+K} \end{bmatrix}, \quad L \equiv \begin{bmatrix} \lambda^{l-1} \\ \alpha\lambda^{l+1} \\ \alpha\lambda^{l+2} \\ \cdot \\ \cdot \\ \cdot \\ \alpha\lambda^{l+K-1} \\ \lambda^{l+K-1} \end{bmatrix},$$

$$A \equiv \begin{bmatrix} 1 & 0 & 0 & 0 & \cdot & \cdot & \cdot & 0 & 0 & 0 \\ 1 - \varphi_{l+1} & 0 & \varphi_{l+1} & 0 & \cdot & \cdot & \cdot & 0 & 0 & 0 \\ 0 & 1 - \varphi_{l+2} & 0 & \varphi_{l+2} & \cdot & \cdot & \cdot & 0 & 0 & 0 \\ - & - & - & - & - & - & - & - & - & - \\ 0 & 0 & 0 & 0 & \cdot & \cdot & \cdot & 1 - \varphi_{l+K-1} & 0 & \varphi_{l+K-1} \\ 0 & 0 & 0 & 0 & \cdot & \cdot & \cdot & 0 & 0 & 1 \end{bmatrix}.$$

With these definitions and (6.1) we get

$$V = L + \delta AV.$$

Let I denote the identity matrix. Notice that A is a semipositive square matrix and that each row sum is 1 (thus A has a Frobenius root of 1). This and $\delta \in (0, 1)$ imply that the inverse $(I - \delta A)^{-1}$ exists and is semipositive (see, e.g., Takayama 1974, Theorem 4.D.2, p. 392). Therefore,

$$V = (I - \delta A)^{-1} L. \quad (6.2)$$

We define $P_{l+k} \equiv p^L(\lambda^{l+k}) = \lambda^{l+k-1}$, $k \in \{0, 1, \dots, K\}$, and

$$P \equiv \begin{bmatrix} P_l \\ P_{l+1} \\ \cdot \\ \cdot \\ \cdot \\ P_{l+K} \end{bmatrix} = \begin{bmatrix} \lambda^{l-1} \\ \lambda^l \\ \cdot \\ \cdot \\ \cdot \\ \lambda^{l+K-1} \end{bmatrix}.$$

By assumption, $V \geq \frac{1}{1-\delta}P$ and at least one inequality is strict because by assumption $V(\lambda^{l+k}) > \frac{p^L(\lambda^{l+k})}{1-\delta}$ for all $k \in \{1, \dots, K-1\}$ and $K \geq 2$. Thus,

$$V \not\geq \frac{1}{1-\delta}P. \quad (6.3)$$

Note that

$$AP = \begin{bmatrix} \lambda^{l-1} \\ (1 - \varphi_{l+1}) \lambda^{l-1} + \varphi_{l+1} \lambda^{l+1} \\ \cdot \\ \cdot \\ \cdot \\ (1 - \varphi_{l+k}) \lambda^{l+k-2} + \varphi_{l+k} \lambda^{l+k} \\ \cdot \\ \cdot \\ \cdot \\ (1 - \varphi_{l+K-1}) \lambda^{l+K-3} + \varphi_{l+K-1} \lambda^{l+K-1} \\ \lambda^{l+K-1} \end{bmatrix}$$

and that the assumption $F(\lambda^{l+k}) \leq 0$ for all $k \in \{0, 1, \dots, K\}$ implies

$$\frac{\delta}{1-\delta} AP + L \leq \frac{1}{1-\delta} P,$$

or

$$(1 - \delta) L \leq (I - \delta A) P.$$

Since $(I - \delta A)^{-1}$ is semipositive,

$$(I - \delta A)^{-1} L \leq \frac{1}{1-\delta} P.$$

Together with (6.2) this gives

$$V = (I - \delta A)^{-1} L \leq \frac{1}{1-\delta} P$$

which contradicts (6.3). This contradiction proves the Lemma. \blacksquare

Corollary 1. Let $\{\lambda^l, \dots, \lambda^{l+K}\}$, $K \geq 2$, where $\lambda^{l+k} \equiv \Pr(\omega = G \mid \lambda^l, k \text{ signals } s = g)$, be a set of λ 's such that $F(\lambda^{l+k}) \leq 0$ for all $k \in \{0, \dots, K\}$. If $V(\lambda^l) = \frac{p^L(\lambda^l)}{1-\delta}$ and $V(\lambda^{l+K}) = \frac{p^L(\lambda^{l+K})}{1-\delta}$, then $V(\lambda^{l+k}) = \frac{p^L(\lambda^{l+k})}{1-\delta}$ for all $k \in \{1, \dots, K-1\}$.

Lemma 6. Let $\{\lambda^l, \lambda^{l+1}, \lambda^{l+2}\}$, where $\lambda^{l+k} \equiv \Pr(\omega = G \mid \lambda^l, k \text{ signals } s = g)$, be a set of λ 's such that $F(\lambda^{l+1}) < 0$. If $V(\lambda^{l+k}) = \frac{p^L(\lambda^{l+k})}{1-\delta}$ for all $k \in \{0, 1, 2\}$, then $p^H(\lambda^{l+1})$ cannot be optimal, i.e., the price $p(\lambda^{l+1}) = p^L(\lambda^{l+1})$ is uniquely optimal.

Proof: By assumption $F(\lambda^{l+1}) < 0$, and $V(\lambda^{l+k}) = \frac{p^L(\lambda^{l+k})}{1-\delta}$ for $k \in \{0, 2\}$. The expected payoff from the price $p^H(\lambda^{l+1})$ equals $\alpha \lambda^{l+1} + \delta(1 - \varphi_{l+1}) V_l + \delta \varphi_{l+1} V_{l+2} = F(\lambda^{l+1}) + \frac{p^L(\lambda^{l+1})}{1-\delta} < \frac{p^L(\lambda^{l+1})}{1-\delta}$, since $F(\lambda^{l+1}) < 0$. Therefore, $p^H(\lambda^{l+1})$ is not optimal. \blacksquare

Lemma 7. *If $\alpha^2 > (1 - \alpha)$, there exists a $\lambda' \in [\bar{\lambda}_\alpha, 1)$ such that*

$$F(\lambda) \begin{cases} > 0 & \text{for all } \lambda \in (0, \lambda') \\ = 0 & \text{for } \lambda = \lambda' \\ < 0 & \text{for all } \lambda \in (\lambda', 1). \end{cases}$$

The number λ' is strictly increasing in δ , $\lambda' = \bar{\lambda}_\alpha$ for $\delta = 0$, and $\lambda' \rightarrow 1$ for $\delta \rightarrow 1$.

If $\alpha^2 = (1 - \alpha)$, there exists a $\delta^ \in (0, 1)$ such that for all $\delta \in [0, \delta^*]$, $F(\lambda) < 0$ for all $\lambda \in (0, 1)$. For each $\delta \in (\delta^*, 1)$ there exists a $\lambda' \in (0, 1)$ such that*

$$F(\lambda) \begin{cases} > 0 & \text{for all } \lambda \in (0, \lambda') \\ = 0 & \text{for } \lambda = \lambda' \\ < 0 & \text{for all } \lambda \in (\lambda', 1). \end{cases}$$

The number λ' is strictly increasing in δ , $\lambda' \rightarrow 0$ for $\delta \rightarrow \delta^$, and $\lambda' \rightarrow 1$ for $\delta \rightarrow 1$.*

If $\alpha^2 < (1 - \alpha)$, there exists a $\delta^ \in (0, 1)$ such that for all $\delta \in [0, \delta^*]$, $F(\lambda) < 0$ for all $\lambda \in (0, 1)$. For each $\delta \in [\delta^*, 1)$, there exist a $\lambda'' \in (0, 1)$ and a $\lambda' \in [\lambda'', 1)$ such that*

$$F(\lambda) \begin{cases} > 0 & \text{for all } \lambda \in (\lambda'', \lambda') \\ = 0 & \text{for } \lambda \in \{\lambda'', \lambda'\} \\ < 0 & \text{for all } \lambda \in (0, \lambda'') \cup (\lambda', 1). \end{cases}$$

Whereas λ'' is strictly decreasing in $\delta \in [\delta^, 1)$, λ' is strictly increasing in $\delta \in [\delta^*, 1)$. For all $\delta \in (\delta^*, 1)$, $\lambda'' < \lambda'$; and for $\delta = \delta^*$, $\lambda'' = \lambda'$. For $\delta \rightarrow 1$, $\lambda'' \rightarrow 0$ and $\lambda' \rightarrow 1$.*

Proof: With the definition $G(\lambda) \equiv \varphi(\lambda)p^L(\lambda^+) + [1 - \varphi(\lambda)]p^L(\lambda^-) - p^L(\lambda)$ we get $F(\lambda) = \alpha\lambda - p^L(\lambda) + \frac{\delta}{1-\delta}G(\lambda)$. Since $p^L(\lambda^+) = \lambda$, $p^L(\lambda) = \lambda^-$, and $\lambda^- = \alpha\lambda^- + (1-\alpha)\lambda^- = \varphi(\lambda^-)\lambda + [1 - \varphi(\lambda^-)]p^L(\lambda^-)$,

$$\begin{aligned} G(\lambda) &= [\varphi(\lambda) - \varphi(\lambda^-)]\lambda + [\varphi(\lambda^-) - \varphi(\lambda)]p^L(\lambda^-) = [\varphi(\lambda) - \varphi(\lambda^-)][\lambda - p^L(\lambda^-)] \\ &= [\alpha(\lambda - \lambda^-) + (1 - \alpha)(\lambda^- - \lambda)][\lambda - p^L(\lambda^-)] = (2\alpha - 1)(\lambda - \lambda^-)[\lambda - p^L(\lambda^-)]. \end{aligned}$$

Moreover,

$$\lambda - \lambda^- = \frac{2\alpha - 1}{(1 - \alpha)\lambda + \alpha(1 - \lambda)}\lambda(1 - \lambda)$$

and

$$\lambda - p^L(\lambda^-) = \frac{2\alpha - 1}{(1 - \alpha)^2\lambda + \alpha^2(1 - \lambda)}\lambda(1 - \lambda)$$

and therefore

$$G(\lambda) = (2\alpha - 1)^3 \frac{\lambda^2(1 - \lambda)^2}{[(1 - \alpha)\lambda + \alpha(1 - \lambda)][(1 - \alpha)^2\lambda + \alpha^2(1 - \lambda)]}.$$

Since $p^L(\lambda) = \lambda^-$, $F(\lambda) = 0$ if and only if $\frac{\delta}{1-\delta}G(\lambda) = \lambda^- - \alpha\lambda$. Note that

$$\lambda^- - \alpha\lambda = \frac{(1-\alpha)(1-\alpha\lambda) - \alpha^2(1-\lambda)}{(1-\alpha)\lambda + \alpha(1-\lambda)}\lambda.$$

Hence, $F(\lambda) = 0$ if and only if

$$\begin{aligned} & (2\alpha - 1)^3 \frac{\delta}{1-\delta} \frac{\lambda^2(1-\lambda)^2}{(1-\alpha)^2\lambda + \alpha^2(1-\lambda)} \\ &= [(1-\alpha)(1-\alpha\lambda) - \alpha^2(1-\lambda)]\lambda. \end{aligned} \quad (6.4)$$

One solution to (6.4) is $\lambda = 0$, but we are looking for solutions $\lambda \in (0, 1)$. Thus we can multiply both sides of (6.4) by $\frac{(1-\alpha)^2\lambda + \alpha^2(1-\lambda)}{\lambda(1-\lambda)}$. Rearranging gives

$$\begin{aligned} (2\alpha - 1)^3 \frac{\delta}{1-\delta} \lambda(1-\lambda) &= \alpha^2(1-\alpha)(1-\alpha\lambda) - \alpha^4(1-\lambda) + \\ & (1-\alpha)^3 \frac{(1-\alpha\lambda)\lambda}{1-\lambda} - \alpha^2(1-\alpha)^2\lambda \end{aligned} \quad (6.5)$$

Next, we show that the left-hand side of (6.5) is strictly concave in λ and the right-hand side is strictly convex in λ , which in turn implies that there can be at most two different λ 's that satisfy (6.5). Let the left-hand side be denoted by $h(\lambda) \equiv (2\alpha - 1)^3 \frac{\delta}{1-\delta} \lambda(1-\lambda)$. Since $\alpha > 1/2$, $h(\lambda)'' = -2(2\alpha - 1)^3 \frac{\delta}{1-\delta} < 0$. Thus, $h(\lambda)$ is strictly concave. In addition, $h(0) = h(1) = 0$. We denote the right-hand side of (6.5) by $k(\lambda) \equiv \alpha^2(1-\alpha)(1-\alpha\lambda) - \alpha^4(1-\lambda) + (1-\alpha)^3 \frac{(1-\alpha\lambda)\lambda}{1-\lambda} - \alpha^2(1-\alpha)^2\lambda$. The function $k(\lambda)$ is strictly convex: all the linear terms drop out after differentiating twice and we get $k''(\lambda) = \frac{2(1-\alpha)^4}{(1-\lambda)^3} > 0$. Moreover, $k(0) = \alpha^2(1-\alpha) - \alpha^4$ and $k(\lambda) \rightarrow \infty$ for $\lambda \rightarrow 1$.

Consider first the case $\alpha^2 > (1-\alpha)$. In this case, $k(0) < \alpha^2\alpha^2 - \alpha^4 = 0 = h(0)$. Given the properties of $h(\lambda)$ and $k(\lambda)$, this implies that there is exactly one $\lambda \in (0, 1)$, denoted by λ' , that satisfies (6.5) because $k(\lambda)$ is convex, $k(\lambda) \rightarrow \infty$ for $\lambda \rightarrow 1$, and $h(\lambda)$ is strictly concave. Moreover, $G(\lambda) > 0$ because $p^L(\lambda)$ is concave. Consequently, for $\delta > 0$, $F(\lambda) = 0$ implies $p^L(\lambda) - \alpha\lambda > 0$ and thus $\lambda' \in (\bar{\lambda}_\alpha, 1)$. For $\delta = 0$, $F(\lambda) = 0$ implies $p^L(\lambda) - \alpha\lambda = 0$ and therefore $\lambda' = \bar{\lambda}_\alpha$. If δ increases, $h(\lambda)$ increases for every $\lambda \in (0, 1)$, whereas $k(\lambda)$ is unaffected. Thus, λ' is strictly increasing in δ . In addition, for every $\varepsilon > 0$ there exists a (sufficiently large) $\delta < 1$ such that $h(1-\varepsilon) > k(1-\varepsilon)$, and consequently $\lambda' > 1-\varepsilon$. Because of this, $\lambda' \rightarrow 1$ for $\delta \rightarrow 1$. For $\lambda \in (0, \lambda')$, $h(\lambda) > k(\lambda)$ and therefore $F(\lambda) > 0$. Similarly, $F(\lambda) < 0$ for all $\lambda \in (\lambda', 1)$.

Next, consider the case $\alpha^2 = (1-\alpha)$. In this case, $k(0) = h(0) = 0$. Moreover, $h'(0) = (2\alpha - 1)^3 \frac{\delta}{1-\delta}$ and $k'(0) = -\alpha^3(1-\alpha) + \alpha^4 + (1-\alpha)^3 - \alpha^2(1-\alpha)^2 = \alpha^6$. Therefore, there

exists a unique $\delta^* \in (0, 1)$ such that $h'(0) = k'(0)$. For all $\delta \in [0, \delta^*]$ no $\lambda \in (0, 1)$ solves (6.5), whereas for all $\delta \in (\delta^*, 1)$ the case is analogous to the case where $\alpha^2 > (1 - \alpha)$ except that $\bar{\lambda}_\alpha = 0$.

Finally, consider the case $\alpha^2 < (1 - \alpha)$. In this case, $k(0) > \alpha^2\alpha^2 - \alpha^4 = 0 = h(0)$ and either (6.5) has no solution i.e., $h(\lambda) < k(\lambda)$ for all $\lambda \in (0, 1)$, or there are two solutions, $\lambda'' \in (0, 1)$ and $\lambda' \in [\lambda'', 1)$, that solve (6.5). In the case $\lambda' = \lambda''$ the two solutions are identical. For all $\lambda \in (\lambda'', \lambda')$, $h(\lambda) > k(\lambda)$ and thus $F(\lambda) > 0$. Similarly, $F(\lambda) < 0$ for all $\lambda \in (0, \lambda'') \cup (\lambda', 1)$. If λ''_δ and λ'_δ solve (6.5) for some fixed $\bar{\delta} \in (0, 1)$, an increase in δ will increase $h(\lambda''_\delta)$ as well as $h(\lambda'_\delta)$, whereas $k(\lambda''_\delta)$ and $k(\lambda'_\delta)$ remain unchanged. Moreover, $h(\lambda) > k(\lambda)$ for all $\lambda \in (\lambda''_\delta, \lambda'_\delta)$ if $\delta \in (\bar{\delta}, 1)$. Consequently, λ'' decreases in δ and λ' increases in δ . If λ''_δ and $\lambda'_\delta > \lambda''_\delta$ solve (6.5) for some fixed $\bar{\delta} \in (0, 1)$, a decrease in δ implies that λ'' increases whereas λ' as well as $\lambda' - \lambda''$ decreases. At some δ , denoted by δ^* , $\lambda' = \lambda''$. From this the rest of the lemma follows. \blacksquare

Proposition 1. *Assume that prices are observable. If $\alpha^2 > 1 - \alpha$, there exists a critical probability $\mu^* \geq \bar{\lambda}_\alpha$, $\mu^* \in \Lambda(\lambda_1)$, such that it is uniquely optimal for the seller to demand*

$$p_t = \begin{cases} p^H(\lambda_t) & \text{whenever } \lambda_t < \mu^* \\ p^L(\lambda_t) & \text{whenever } \lambda_t > \mu^* \end{cases}$$

For $\lambda_t = \mu^*$, $p^L(\lambda_t)$ is optimal, but $p^H(\lambda_t)$ may be optimal as well. Moreover, $\mu^* > \bar{\lambda}_\alpha$ for $\delta > 0$, $\mu^* = \min_{\lambda \in \Lambda(\lambda_1) \cap [\bar{\lambda}_\alpha, 1]} \lambda$ for $\delta = 0$, and $\mu^* \rightarrow 1$ for $\delta \rightarrow 1$.

Proof : We know that for any $\lambda \in \Lambda(\lambda_1)$, $F(\lambda) > 0$ implies that $p^H(\lambda)$ is uniquely optimal. We know from Lemma 3 that given δ , $V(\lambda) = \frac{1}{1-\delta}p^L(\lambda)$, if λ is sufficiently close to 1. Define μ^* as the smallest $\mu \in \Lambda(\lambda_1)$ such that for all $\lambda \geq \mu$, $\lambda \in \Lambda(\lambda_1)$, it holds that $V(\lambda) = \frac{1}{1-\delta}p^L(\lambda)$. Since $V(\lambda) > \frac{1}{1-\delta}p^L(\lambda)$ for all $\lambda < \bar{\lambda}_\alpha$, $\lambda \in \Lambda(\lambda_1)$, μ^* exists and $\mu^* \geq \bar{\lambda}_\alpha$. We show by contradiction that $V(\lambda) > \frac{1}{1-\delta}p^L(\lambda)$ for all $\lambda < \mu^*$, $\lambda \in \Lambda(\lambda_1)$. Assume that for some $\lambda < \mu^*$, $\lambda \in \Lambda(\lambda_1)$, it holds that $V(\lambda) = \frac{1}{1-\delta}p^L(\lambda)$, and let λ^l be the largest such $\lambda < \mu^*$. This implies $F(\lambda^l) \leq 0$ (otherwise $p^H(\lambda^l)$ would be uniquely optimal). Hence by Lemma 7, $\lambda' \leq \lambda^l$ and $F(\lambda) < 0$ for all $\lambda > \lambda^l$, $\lambda \in \Lambda(\lambda_1)$. Since by construction $\lambda^l < \mu^*$ and $V(\lambda) > \frac{1}{1-\delta}p^L(\lambda)$ for all $\lambda \in \Lambda(\lambda_1) \cap (\lambda^l, \mu^*)$, and since the definition of μ^* implies $V(\lambda^-) > \frac{1}{1-\delta}p^L(\lambda^-)$ for $\lambda = \mu^*$, there exists a set $\{\lambda^l, \dots, \lambda^{l+K}\}$, $K \geq 2$, that satisfies the assumptions of Lemma 5 and, in addition, $V(\lambda^K) = \frac{p^L(\lambda^K)}{1-\delta}$. Because of this, Lemma 5 implies $V(\lambda^l) > \frac{p^L(\lambda^l)}{1-\delta}$, whereas by construction $V(\lambda^l) = \frac{p^L(\lambda^l)}{1-\delta}$. This contradiction proves that $V(\lambda) > \frac{1}{1-\delta}p^L(\lambda)$ for all $\lambda < \mu^*$, $\lambda \in \Lambda(\lambda_1)$. It follows that $p^H(\lambda)$ is uniquely optimal for all $\lambda < \mu^*$, $\lambda \in \Lambda(\lambda_1)$. Moreover, because of Lemma 7 and $\mu^* \geq \lambda'$ (which follows from $F(\mu^*) \leq 0$), Lemma 6 implies that only $p^L(\lambda)$ is optimal for all $\lambda > \mu^*$, $\lambda \in \Lambda(\lambda_1)$. The rest of Proposition 1 follows from part 1 of Lemma 7. \blacksquare

We define

$$V \equiv \begin{bmatrix} V_l \\ V_{l-1} \\ \cdot \\ \cdot \\ \cdot \\ \cdot \end{bmatrix}, \quad L \equiv \begin{bmatrix} \lambda^{l-1} \\ \alpha\lambda^{l-1} \\ \alpha\lambda^{l-2} \\ \cdot \\ \cdot \\ \cdot \end{bmatrix}, \quad P \equiv \begin{bmatrix} P_l \\ P_{l-1} \\ \cdot \\ \cdot \\ \cdot \\ \cdot \end{bmatrix} \equiv \begin{bmatrix} \lambda^{l-1} \\ \lambda^{l-2} \\ \cdot \\ \cdot \\ \cdot \\ \cdot \end{bmatrix}.$$

With these definitions, the matrix C defined above, and (6.6) we get

$$V = L + \delta CV,$$

or

$$L = (I - \delta C) V. \quad (6.7)$$

The assumption $F(\lambda^{l-k}) \leq 0$ for all $k \in \{0, 1, \dots\}$ implies

$$\frac{\delta}{1-\delta} CP + L \leq \frac{1}{1-\delta} P,$$

or

$$L \leq \frac{1}{1-\delta} (I - \delta C) P. \quad (6.8)$$

>From (6.7) and (6.8) we get

$$(I - \delta C) V \leq \frac{1}{1-\delta} (I - \delta C) P. \quad (6.9)$$

Define the infinite vector $x = (x_1, x_2, \dots) \gg 0$ by $x_1 = \frac{1}{1-\delta}$, $x_k = 1$ for $k \in \{2, 3, \dots\}$. Then the infinite vector $z \equiv x(I - \delta C)$ has the elements $z_1 = 1 - \delta\varphi_{l-1} \in (0, 1]$, $z_2 = 1 - \delta\varphi_{l-2} \in (0, 1]$, $z_k = 1 - \delta(1 - \varphi_{l-k+1} + \varphi_{l-k-1}) \in (0, 1]$ for $k \in \{3, 4, \dots\}$, where $1 - \delta(1 - \varphi_{l-k+1} + \varphi_{l-k-1}) \in (0, 1]$ follows from $\varphi_{l-k+1} - \varphi_{l-k-1} \in (0, 1)$ and $\delta \in (0, 1]$.

Because

$$\begin{aligned} \frac{\lambda^{l-k-1}}{\lambda^{l-k}} &= \frac{1-\alpha}{(1-\alpha)\lambda^{l-k} + \alpha(1-\lambda^{l-k})} = 1 - \frac{(1-\alpha)\lambda^{l-k} + \alpha(1-\lambda^{l-k}) - (1-\alpha)}{(1-\alpha)\lambda^{l-k} + \alpha(1-\lambda^{l-k})} = \\ &= 1 - \frac{(2\alpha-1)(1-\lambda^{l-k})}{(1-\alpha)\lambda^{l-k} + \alpha(1-\lambda^{l-k})} < 1 - \frac{(2\alpha-1)(1-\lambda^l)}{\alpha} \in (0, 1), \end{aligned}$$

the infinite sum $\sum_{k=0}^{\infty} \lambda^{l-k-1}$ converges, i.e., $\sum_{k=0}^{\infty} \lambda^{l-k-1} < \infty$. Since $z_k \in (0, 1]$ for all $k \in \{1, 2, \dots\}$, $0 < zP = \sum_{k=0}^{\infty} z_k \lambda^{l-k-1} \leq \sum_{k=0}^{\infty} \lambda^{l-k-1} < \infty$. Multiplying both sides of (6.9) by $x \gg 0$ gives

$$zV \leq \frac{1}{1-\delta} zP, \quad (6.10)$$

where, as shown above, $zP < \infty$. However, the proof's assumption that $V(\lambda^{l-k}) > \frac{p^L(\lambda^{l-k})}{1-\delta}$ for all $k \in \{1, 2, \dots\}$, together with $V(\lambda^l) = \frac{p^L(\lambda^l)}{1-\delta}$ and $z \gg 0$, implies

$$zV > \frac{1}{1-\delta} zP.$$

This contradiction proves the lemma. ■

Proposition 2. *Assume that prices are observable. If $\alpha^2 = 1 - \alpha$, there exists a discount factor $\delta^* \in [0, 1)$ such that for all $\delta \in [0, \delta^*]$ the uniquely optimal prices are given by $p_t = p^L(\lambda_1)$ for all $t \in \{1, 2, \dots\}$. For each $\delta \in (\delta^*, 1)$ there exists a critical probability $\mu^* \in \Lambda(\lambda_1)$ such that it is uniquely optimal for the seller to demand*

$$p_t = \begin{cases} p^H(\lambda_t) & \text{whenever } \lambda_t < \mu^* \\ p^L(\lambda_t) & \text{whenever } \lambda_t > \mu^* \end{cases}$$

For $\lambda_t = \mu^*$, $p^L(\lambda_t)$ is optimal, but $p^H(\lambda_t)$ may be optimal as well. Finally, $\mu^* \rightarrow 0$ for $\delta \rightarrow \delta^*$, and $\mu^* \rightarrow 1$ for $\delta \rightarrow 1$.

Proof: From Lemma 3 we know that given δ , the low price $p^L(\lambda)$ is uniquely optimal whenever λ is sufficiently close to 1. By Lemma 7, there exists a $\delta^* \in (0, 1)$ such that for all $\delta \in [0, \delta^*]$, $F(\lambda) < 0$ for all $\lambda \in \Lambda(\lambda_1)$, and thus Lemma 8 implies that $p^L(\lambda)$ is optimal for all $\lambda \in \Lambda(\lambda_1)$. Since for any $\varepsilon > 0$, $(0, \varepsilon) \cap \Lambda(\lambda_1) \neq \emptyset$, the rest of the proof of Proposition 2 is analogous to the proof of Proposition 1 (note that for $\alpha^2 = (1 - \alpha)$, $\bar{\lambda}_\alpha = 0$). ■

Proposition 3. *Assume that prices are observable. If $\alpha^2 < 1 - \alpha$, there exist discount factors $\delta^{**} \in (0, 1)$ and $\delta^{***} \in [\delta^{**}, 1)$ such that*

- for all $\delta \in [0, \delta^{**})$, the uniquely optimal prices are given by $p_t = p^L(\lambda_1)$ for all $t \in \{1, 2, \dots\}$;
- for $\delta \in [\delta^{**}, \delta^{***}]$, $p_t = p^L(\lambda_t)$ is optimal for all $\lambda_t \in \Lambda(\lambda_1)$, but $p_t = p^H(\lambda_t)$ is optimal as well for at least one $\lambda_t \in \Lambda(\lambda_1)$;
- for each $\delta \in (\delta^{***}, 1)$ there exist critical probabilities $\mu^* \in \Lambda(\lambda_1)$ and $\mu^{**} \in \Lambda(\lambda_1) \cup \{0\}$, $0 \leq \mu^{**} < \mu^* < 1$, such that it is uniquely optimal for the seller to demand

$$p_t = \begin{cases} p^H(\lambda_t) & \text{whenever } \lambda_t \in (\mu^{**}, \mu^*) \\ p^L(\lambda_t) & \text{whenever } \lambda_t \in (0, \mu^{**}) \cup (\mu^*, 1); \end{cases}$$

for $\lambda_t \in \{\mu^{**}, \mu^*\}$, $p^L(\lambda_t)$ is optimal, but $p^H(\lambda_t)$ may be optimal as well; for $\delta \rightarrow 1$, $\mu^{**} \rightarrow 0$ and $\mu^* \rightarrow 1$.

Proof: From Lemma 3 we know that given δ , the low price $p^L(\lambda)$ is uniquely optimal whenever λ is sufficiently close to 1. By Lemma 7, there exists a $\delta^* \in (0, 1)$ such that for all $\delta \in [0, \delta^*)$, $F(\lambda) < 0$ for all $\lambda \in \Lambda(\lambda_1)$. Hence Lemma 8 implies that $p^L(\lambda)$ is optimal for all $\delta \in (0, \delta^*)$. Although $F(\lambda) > 0$ for some $\lambda \in (0, 1)$ if $\delta > \delta^*$, these λ 's may not be elements of $\Lambda(\lambda_1)$ and thus it may still be the case that $F(\lambda) < 0$ for all $\lambda \in \Lambda(\lambda_1)$. Therefore, define $\delta^{**} \geq \delta^* > 0$ as the supremum of δ in the set $\{\delta \mid F(\lambda) < 0 \text{ for all } \lambda \in \Lambda(\lambda_1)\}$. Together with Lemma 6 the definition of δ^{**} implies that for all $\delta \in (0, \delta^{**})$, $p^L(\lambda)$ is uniquely optimal for all $\lambda \in \Lambda(\lambda_1)$.

Consider now the case $\delta \in (\delta^{**}, 1)$. Recall that $F(\lambda) > 0$ implies that $p^H(\lambda)$ is uniquely optimal. Hence from Lemma 7 we know that if $\delta \in (\delta^{**}, 1)$ is sufficiently large (and therefore λ'' and λ' of Lemma 7 are sufficiently close to 0 and 1, respectively), there exists a $\lambda \in \Lambda(\lambda_1)$ such that $V(\lambda) > \frac{1}{1-\delta}p^L(\lambda)$. Let $\delta^{***} \geq \delta^{**}$ denote the supremum of δ in the set $\{\delta \mid V(\lambda) = \frac{1}{1-\delta}p^L(\lambda) \text{ for all } \lambda \in \Lambda(\lambda_1)\}$. Continuity implies $V(\lambda) = \frac{1}{1-\delta}p^L(\lambda)$ for all $\lambda \in \Lambda(\lambda_1)$, if $\delta = \delta^{***}$. For any $\delta \in (\delta^{***}, 1)$ define μ^{**} as the largest $\mu \in \Lambda(\lambda_1)$ such that for all $\lambda \leq \mu$, $\lambda \in \Lambda(\lambda_1)$, it holds that $V(\lambda) = \frac{1}{1-\delta}p^L(\lambda)$, if such a μ exists; otherwise define $\mu^{**} \equiv 0$. Define μ^* as the smallest $\mu \in \Lambda(\lambda_1)$ such that for all $\lambda \geq \mu$, $\lambda \in \Lambda(\lambda_1)$, it holds that $V(\lambda) = \frac{1}{1-\delta}p^L(\lambda)$. The definition of δ^{***} implies $\mu^{**} < \mu^*$. Since $F(\mu^{**}) \leq 0$ if $\mu^{**} > 0$, and $F(\mu^*) \leq 0$, Lemma 7 implies $F(\lambda) < 0$ for all $\lambda \in \Lambda(\lambda_1)$ that satisfy either $\lambda < \mu^{**}$ or $\lambda > \mu^*$. If $\mu^{**} > 0$, the part of Proposition 3 that relates to $\delta \in (\delta^{***}, 1)$ follows from Corollary 1 and Lemma 8.

Next consider the case $\mu^{**} = 0$. Corollary 1 still implies that $p^L(\lambda)$ is optimal for $\lambda \in (\mu^*, 1)$. We show by contradiction that $p^H(\lambda)$ is uniquely optimal for all $\lambda \in (0, \mu^*)$. Assume that $p^L(\lambda^0)$ is optimal for some $\lambda^0 \in (0, \mu^*) \cap \Lambda(\lambda_1)$. This implies $F(\lambda^0) \leq 0$ and thus $\lambda^0 \leq \lambda''$ because of Corollary 1, Lemma 7, and the definition of μ^* . Therefore, $F(\lambda) \leq 0$ for all $\lambda \in (0, \lambda^0)$. Hence by Lemma 8, $p^L(\lambda)$ is optimal for all $\lambda \in (0, \lambda^0) \cap \Lambda(\lambda_1)$. This implies $\mu^{**} \geq \lambda^0 > 0$ and thus contradicts $\mu^{**} = 0$. Therefore, $p^H(\lambda)$ is uniquely optimal for all $\lambda \in (0, \mu^*) \cap \Lambda(\lambda_1)$.

Finally, consider the case $\delta \in [\delta^{**}, \delta^{***}]$. The definition of δ^{***} implies that for all $\delta \in [\delta^{**}, \delta^{***}]$, $V(\lambda) = \frac{1}{1-\delta}p^L(\lambda)$ for all $\lambda \in \Lambda(\lambda_1)$ and therefore $p^L(\lambda)$ is optimal for all $\lambda \in \Lambda(\lambda_1)$. Moreover, the definitions of δ^{**} and δ^{***} , respectively, imply that $F(\lambda) = 0$ for at least one $\lambda \in \Lambda(\lambda_1)$ if $\delta \in [\delta^{**}, \delta^{***}]$, and therefore $p^H(\lambda)$ is also optimal for at least one $\lambda \in \Lambda(\lambda_1)$. ■

A.2. Proof of Propositions 4 and 5

Propositions 4 and 5 state that there exists a PBE in pure strategies such that $p_t = p^L(\lambda_1)$ for all $t \in \{1, 2, \dots\}$. The respective equilibrium beliefs of the buyers are given by (4.1),

and these beliefs are consistent with buyers' information at all nodes that can be reached either along the equilibrium path or after deviations of the seller. Given these beliefs, the optimal strategy of any buyer $t \in \{1, 2, \dots\}$ depends only on μ_t and s_t . Consequently, for the seller all relevant aspects of the history are captured by λ_t and μ_t in each period t , and thus the seller has a best response to the buyers' strategies and beliefs such that for each t the seller's move depends only on λ_t and μ_t , i.e., the seller's best response is Markov. Because of this, we can employ the value function, which we denote by $W(\lambda, \mu)$, in order to analyze the seller's optimization problem. That is, to any $(\lambda, \mu) \in \Lambda(\lambda_1) \times \Lambda(\lambda_1)$, $W(\lambda, \mu)$ assigns the maximum expected payoff that the seller can achieve by playing a pure Markov strategy, given buyers' strategies and beliefs. We have to show that, given buyers' strategies and beliefs, the seller's strategy to charge $p_t = p^L(\lambda_1)$ for all $t \in \{1, 2, \dots\}$ maximizes her expected payoff over all pure strategies that are Markov (and thus over *all* pure strategies, since the seller's best response is Markov). For any prior μ^l and any positive integer k we define $\mu^{l+k} \equiv \Pr(\omega = G \mid \mu^l, k \text{ signals } s = g)$ and $\mu^{l-k} \equiv \Pr(\omega = G \mid \mu^l, k \text{ signals } s = b)$, analogously to λ^{l+k} and λ^{l-k} , respectively. First we prove the following Lemma.

Lemma 9. *Let $\alpha^2 \leq 1 - \alpha$. Assume that each buyer t believes that the seller has charged the low price $p^L(\mu_\tau)$ to each previous buyer $\tau < t$ who has purchased the object, and the high price $p^H(\mu_\tau)$ to each previous buyer $\tau < t$ who has not purchased the object. If for some $\lambda_t = \mu_t = \mu^l$ it holds that $W(\mu^l, \mu^l) - W(\mu^{l-1}, \mu^{l-1}) \geq \frac{1}{1-\delta} [p^L(\mu^l) - p^L(\mu^{l-1})]$, then $p^L(\mu^l)$ is optimal for the seller at $(\lambda, \mu) = (\mu^l, \mu^l)$.*

Proof: The proof is by contradiction. Assume that the high price $p^H(\mu^l)$ is uniquely optimal at $(\lambda_t, \mu_t) = (\mu^l, \mu^l)$. Then

$$W(\mu^l, \mu^l) = \varphi(\mu^l) p^H(\mu^l) + \delta \varphi(\mu^l) W(\mu^{l+1}, \mu^l) + \delta [1 - \varphi(\mu^l)] W(\mu^{l-1}, \mu^{l-1})$$

and $W(\mu^l, \mu^l) \geq p^L(\mu^l) + \delta W(\mu^l, \mu^l)$. Since $\alpha^2 \leq 1 - \alpha$, $\varphi(\mu^l) p^H(\mu^l) < p^L(\mu^l)$. Thus, $p^H(\mu^l)$ can only be optimal, if $\delta > 0$ and

$$\varphi(\mu^l) [W(\mu^{l+1}, \mu^l) - W(\mu^l, \mu^l)] > [1 - \varphi(\mu^l)] [W(\mu^l, \mu^l) - W(\mu^{l-1}, \mu^{l-1})].$$

Note that $W(\mu^l, \mu^l) \leq W(\lambda, \mu^l)$ for $\lambda \geq \mu^l$. Furthermore, $W(\mu^l, \mu^l) > \frac{p^L(\mu^l)}{1-\delta}$ by assumption, hence for $\lambda \geq \mu^l$ it follows that $W(\lambda, \mu^l) > \frac{p^L(\mu^l)}{1-\delta}$. In addition, we know that $\varphi(\lambda) < \alpha$ for all λ , and combined this implies that $W(\lambda, \mu^l) < \frac{1}{1-\delta} \alpha p^H(\mu^l)$ for all λ . Together with the lemma's assumption that

$$W(\mu^l, \mu^l) - W(\mu^{l-1}, \mu^{l-1}) \geq \frac{1}{1-\delta} [p^L(\mu^l) - p^L(\mu^{l-1})]$$

this gives

$$\varphi(\mu^l) [\alpha p^H(\mu^l) - p^L(\mu^l)] > [1 - \varphi(\mu^l)] [p^L(\mu^l) - p^L(\mu^{l-1})]$$

or

$$\alpha \frac{\varphi(\mu^l) p^H(\mu^l)}{p^L(\mu^l)} - \varphi(\mu^l) > [1 - \varphi(\mu^l)] \left[1 - \frac{p^L(\mu^{l-1})}{p^L(\mu^l)} \right]. \quad (6.11)$$

Since

$$1 - \frac{p^L(\mu^{l-1})}{p^L(\mu^l)} = 1 - \frac{\mu^{l-2}}{\mu^{l-1}} = 1 - \frac{1 - \alpha}{1 - \varphi(\mu^{l-1})} = \frac{\alpha - \varphi(\mu^{l-1})}{1 - \varphi(\mu^{l-1})},$$

and since $\varphi(\mu^l) p^H(\mu^l) < p^L(\mu^l)$ for $\alpha^2 \leq 1 - \alpha$, (6.11) implies

$$\alpha - \varphi(\mu^l) > \frac{1 - \varphi(\mu^l)}{1 - \varphi(\mu^{l-1})} [\alpha - \varphi(\mu^{l-1})]$$

or $\left[1 + \frac{\varphi(\mu^l) - \varphi(\mu^{l-1})}{1 - \varphi(\mu^l)} \right] [\alpha - \varphi(\mu^l)] > \alpha - \varphi(\mu^{l-1})$, hence

$$\frac{\alpha - \varphi(\mu^l)}{1 - \varphi(\mu^l)} [\varphi(\mu^l) - \varphi(\mu^{l-1})] > \varphi(\mu^l) - \varphi(\mu^{l-1}).$$

Since $\varphi(\mu^l) - \varphi(\mu^{l-1}) > 0$, it must be that $\frac{\alpha - \varphi(\mu^l)}{1 - \varphi(\mu^l)} > 1$, which contradicts $\alpha < 1$. This contradiction proves the lemma. \blacksquare

Proposition 4. *Assume that previous prices are unobservable for buyers. If $\alpha^2 \leq 1 - \alpha$ or $\alpha^2 > 1 - \alpha$ and $\lambda_1 \in [\bar{\lambda}_\alpha, 1)$, there exists a PBE such that $p_t = p^L(\lambda_1)$ for all $t \in \{1, 2, \dots\}$. Thus, there is a PBE where the seller always charges the price $p^L(\lambda_1)$ and herding arises immediately at $t = 1$.*

Proof: We show that there exists a PBE where the seller charges the price $p_t = p^L(\mu_t)$ whenever $\mu_t = \lambda_t$. All buyers $j \in \{\tau + 1, \dots\}$ believe that in each period $t \in \{1, \dots, \tau\}$ the price \hat{p}_t of (4.1) was demanded. Each buyer t buys the object if and only if $p_t \leq \Pr(\omega = G \mid \mu_t, s_t)$. Obviously the buyers' strategies are optimal after every history of actions. If $p_t = p^L(\mu_t)$ whenever $\mu_t = \lambda_t$, then the buyers' beliefs are consistent with Bayesian updating and the seller's equilibrium strategy and observed actions. Thus, we have to show that given the buyers' strategies and beliefs it is optimal for the seller to charge $p_t = p^L(\mu_t)$ whenever $\mu_t = \lambda_t$. We do so by proving that there is no $\lambda_1 \in (0, 1)$ such that it is optimal for the seller to deviate from $p^L(\lambda_1)$ and instead charge $p_1 = p^H(\lambda_1)$. Let $\mu^l \equiv \lambda_1$.

Consider first the case where $\alpha^2 \leq 1 - \alpha$. Recall that, in this case $\varphi(\mu) p^H(\mu) < p^L(\mu)$ for all $\mu \in (0, 1)$. We prove the result by contradiction. Assume that $p_1 = p^L(\mu^l)$ is not optimal for some $\mu^l \in (0, 1)$. It follows from Lemma 9 that

$$\frac{1}{1 - \delta} [p^L(\mu^l) - p^L(\mu^{l-1})] > W(\mu^l, \mu^l) - W(\mu^{l-1}, \mu^{l-1}) > \frac{1}{1 - \delta} p^L(\mu^l) - W(\mu^{l-1}, \mu^{l-1}),$$

and thus $\gamma \equiv W(\mu^{l-1}, \mu^{l-1}) - \frac{1}{1-\delta} p^L(\mu^{l-1}) > 0$, i.e., $p^L(\mu^{l-1})$ is not optimal at $(\lambda_t, \mu_t) = (\mu^{l-1}, \mu^{l-1})$. Because of Lemma 9 this implies

$$W(\mu^{l-2}, \mu^{l-2}) - \frac{1}{1-\delta} p^L(\mu^{l-2}) > W(\mu^{l-1}, \mu^{l-1}) - \frac{1}{1-\delta} p^L(\mu^{l-1}) = \gamma > 0.$$

Repeating the argument gives

$$W(\mu^{l-k}, \mu^{l-k}) - \frac{1}{1-\delta} p^L(\mu^{l-k}) > \gamma > 0 \text{ for all } k \in \{0, 1, \dots\}.$$

On the other hand, $W(\mu^{l-k}, \mu^{l-k}) < \frac{1}{1-\delta} p^H(\mu^{l-k})$ and thus $\lim_{k \rightarrow \infty} W(\mu^{l-k}, \mu^{l-k}) = 0$. Therefore,

$$\begin{aligned} 0 < \gamma &\leq \lim_{k \rightarrow \infty} \left[W(\mu^{l-k}, \mu^{l-k}) - \frac{1}{1-\delta} p^L(\mu^{l-k}) \right] \\ &= \lim_{k \rightarrow \infty} W(\mu^{l-k}, \mu^{l-k}) - \frac{1}{1-\delta} \lim_{k \rightarrow \infty} p^L(\mu^{l-k}) = 0 \end{aligned}$$

This contradiction implies that it is a best response for the seller to charge the price $p_t = p^L(\mu_t)$ whenever $\mu_t = \lambda_t$, and thus proves Proposition 4 for $\alpha^2 \leq 1 - \alpha$.

Next, we show that immediate herding also is an equilibrium outcome when $\alpha^2 > 1 - \alpha$ and $\lambda_1 \in [\bar{\lambda}_\alpha, 1)$, where $\bar{\lambda}_\alpha \equiv \frac{\alpha^3 - (1-\alpha)^2}{\alpha^3 - (1-\alpha)^3}$. We show by contradiction that the seller has no incentive to deviate from $p_1 = p^L(\mu^l)$, where $\mu^l \equiv \lambda_1$. Suppose $p_1 = p^L(\mu^l)$ is not optimal for some $\mu^l \in (0, 1)$. Then,

$$p^L(\mu^l) + \delta \frac{p^L(\mu^l)}{1-\delta} < \varphi(\mu^l) p^H(\mu^l) + \delta \varphi(\mu^l) W(\mu^{l+1}, \mu^l) + \delta [1 - \varphi(\mu^l)] W(\mu^{l-1}, \mu^{l-1}).$$

Straightforward calculation shows that $\mu^l \geq \bar{\lambda}_\alpha$ is equivalent to $p^L(\mu^l) \geq \alpha p^H(\mu^l)$. Thus, $p^L(\mu^l) \geq \alpha p^H(\mu^l) > \varphi(\mu^l) p^H(\mu^l)$. If $p^H(\mu^l)$ is optimal at $(\lambda, \mu) = (\mu^l, \mu^l)$, then $p^H(\mu^l)$ is optimal at $(\lambda, \mu) = (\mu^{l+k}, \mu^l)$ for all $k \in \{1, 2, \dots\}$, because $W(\mu^{l+k}, \mu^l) \geq W(\mu^l, \mu^l)$. Therefore, $\frac{\alpha p^H(\mu^l)}{1-\delta} > W(\mu^{l+1}, \mu^l) \geq W(\mu^{l-1}, \mu^{l-1})$. Thus, the deviation from the low price is optimal only if $p^L(\mu^l) < \alpha p^H(\mu^l)$, which contradicts $\lambda_1 \geq \bar{\lambda}_\alpha$. ■

Proposition 5. *Assume that previous prices are unobservable for buyers. If $\alpha^2 > 1 - \alpha$ and $\lambda_1 \in (0, \bar{\lambda}_\alpha)$, there exists a $\bar{\delta} \in (0, 1)$ such that for all $\delta > \bar{\delta}$ there exists a PBE where $p_t = p^L(\lambda_1)$ for all $t \in \{1, 2, \dots\}$. Thus, there is a PBE where the seller always charges the price $p^L(\lambda_1)$ and herding arises immediately at $t = 1$, provided the seller is sufficiently patient.*

Proof: All buyers $j \in \{\tau + 1, \dots\}$ believe that in each period $t \in \{1, \dots, \tau\}$ the price \hat{p}_t of (4.1) was demanded. Each buyer t buys the object if and only if $p_t \leq \Pr(\omega = G \mid \mu_t, s_t)$. Obviously the buyers' strategies are optimal after every history of actions. Furthermore, if

$p_t = p^L(\mu_t)$ whenever $\mu_t = \lambda_t$, then the buyers' beliefs are consistent with Bayesian updating and the seller's equilibrium strategy and observed actions. Thus, we have to show that given the buyers' strategies and beliefs it is optimal for the seller to charge $p_t = p^L(\mu_t)$ whenever $\mu_t = \lambda_t$, which implies that $p_1 = p^L(\lambda_1)$ for all $\lambda_1 \in (0, 1)$.

We need to show that for all $\lambda_1 < \bar{\lambda}_\alpha$ there exists a $\bar{\delta} \in (0, 1)$ such that for all $\delta > \bar{\delta}$ there exists a PBE where $p_t = p^L(\lambda_1)$ for all $t \in \{1, 2, \dots\}$. We prove this by contradiction. Notice that (because $\varphi(\lambda)$ increases in λ) if for some pair (λ, μ) the high price $p^H(\mu)$ is optimal, this price must also be optimal for any pair (λ', μ) where $\lambda' > \lambda$.

Let $\mu^l \equiv \lambda_1$. We distinguish between three cases that exhaust all possibilities of potentially profitable deviations by the seller: (i) The seller demands the high price $p^H(\mu^l)$ until for the first time a buyer refuses to buy, and the low price $p^L(\mu^{l-1})$ thereafter; (ii) the seller demands the high price $p^H(\mu^l)$ until for the first time a buyer refuses to buy, then demands the high price $p^H(\mu^{l-1})$ until for the second time a buyer refuses to buy, and demands the low price $p^L(\mu^{l-2})$ thereafter; (iii) relative to the respective buyer's belief μ^{l-k} , $k \in \{0, 1, 2\}$, the seller demands the high price $p^H(\mu^{l-k})$ until three buyers have refused to buy, and some price $p_t \leq p^H(\mu^{l-3})$ thereafter. All cases that are not covered by (i) and (ii) are included in (iii), because buyer's beliefs imply $\mu_{t+1} \leq \mu_t$ for all t .

For any probability $\psi = \Pr(\omega = G)$ of the good state and any $k \in \{1, 2, \dots\}$ let

$$\pi_k(\psi) \equiv \Pr[s_1 = \dots = s_{k-1} = g, s_k = b \mid \psi] = (1 - \alpha)\alpha^{k-1}\psi + \alpha(1 - \alpha)^{k-1}(1 - \psi)$$

denote the probability conditional on ψ that out of k buyers the first $k - 1$ observe the good signal and the last one observes the bad signal. For any $\psi \in [0, 1]$ we have $\sum_{k=1}^{\infty} \pi_k(\psi) = 1$ and

$$\begin{aligned} \sum_{k=1}^{\infty} \pi_k(\psi) \delta^{k-1} &= (1 - \alpha) \psi \sum_{k=1}^{\infty} (\alpha\delta)^{k-1} + \alpha(1 - \psi) \sum_{k=1}^{\infty} [(1 - \alpha)\delta]^{k-1} \\ &= \frac{(1 - \alpha)\psi}{1 - \alpha\delta} + \frac{\alpha(1 - \psi)}{1 - (1 - \alpha)\delta}. \end{aligned}$$

Moreover, $\alpha > 1/2$ implies $\frac{1-\alpha}{1-\alpha\delta} \leq \frac{\alpha}{1-(1-\alpha)\delta}$ and thus $\frac{1-\alpha}{1-\alpha\delta} \leq \frac{(1-\alpha)\psi}{1-\alpha\delta} + \frac{\alpha(1-\psi)}{1-(1-\alpha)\delta} \leq \frac{\alpha}{1-(1-\alpha)\delta}$ for every $\psi \in [0, 1]$.

Case (i): Let the expected payoff from the deviating strategy that is described under (i) above be denoted as $\hat{U}^{(i)}$. Since $\sum_{j=0}^{t-2} \delta^j p^H(\mu^l) = \frac{1-\delta^{t-1}}{1-\delta} p^H(\mu^l)$,

$$\hat{U}^{(i)} = \sum_{t=1}^{\infty} \pi_t(\mu^l) \left[\frac{1 - \delta^{t-1}}{1 - \delta} p^H(\mu^l) + \delta^t \frac{p^L(\mu^{l-1})}{1 - \delta} \right].$$

The expected payoff from the equilibrium strategy is $\frac{1}{1-\delta} p^L(\mu^l)$. From $\sum_{t=1}^{\infty} \pi_t = 1$ and $\sum_{t=1}^{\infty} \pi_t(\mu^l) \delta^{t-1} = \frac{(1-\alpha)\mu^l}{1-\alpha\delta} + \frac{\alpha(1-\mu^l)}{1-(1-\alpha)\delta}$ we get

$$\hat{U}^{(i)} = \frac{p^H(\mu^l)}{1-\delta} \left[1 - \frac{(1-\alpha)\mu^l}{1-\alpha\delta} - \frac{\alpha(1-\mu^l)}{1-(1-\alpha)\delta} \right] + \delta \frac{p^L(\mu^{l-1})}{1-\delta} \left[\frac{(1-\alpha)\mu^l}{1-\alpha\delta} + \frac{\alpha(1-\mu^l)}{1-(1-\alpha)\delta} \right].$$

Since $\frac{1-\alpha}{1-\alpha\delta} \leq \frac{(1-\alpha)\mu^l}{1-\alpha\delta} + \frac{\alpha(1-\mu^l)}{1-(1-\alpha)\delta} \leq \frac{\alpha}{1-(1-\alpha)\delta}$, $\hat{U}^{(i)} > \frac{1}{1-\delta} p^L(\mu^l)$ implies

$$\frac{\alpha(1-\delta)}{1-\alpha\delta} p^H(\mu^l) + \frac{\alpha\delta}{1-(1-\alpha)\delta} p^L(\mu^{l-1}) > p^L(\mu^l).$$

Because $\frac{p^H(\mu^l)}{p^L(\mu^l)} < \frac{\alpha^2}{(1-\alpha)^2}$, dividing both sides of the inequality by $p^L(\mu^l)$ gives

$$(1-\delta) \frac{\alpha^3}{(1-\alpha\delta)(1-\alpha)^2} + \frac{\alpha\delta}{1-(1-\alpha)\delta} \frac{p^L(\mu^{l-1})}{p^L(\mu^l)} > 1.$$

Since for $\delta \rightarrow 1$ the left hand side converges to $\frac{p^L(\mu^{l-1})}{p^L(\mu^l)} < 1$, this inequality cannot hold if δ is sufficiently large. This proves that the deviating strategy does not lead to a larger expected payoff than always charging the low price $p^L(\mu^l)$, provided δ is sufficiently large.

Case (ii): Let the expected payoff from the deviating strategy that is described under (ii) above be denoted as $\hat{U}^{(ii)}$. Then,

$$\hat{U}^{(ii)} = \sum_{t=1}^{\infty} \pi_t(\mu^l) \left[\frac{1-\delta^{t-1}}{1-\delta} p^H(\mu^l) + \delta^t \hat{U}(t, \mu^{l-1}) \right]$$

where

$$\hat{U}(t, \mu^{l-1}) \equiv \sum_{\tau=1}^{\infty} \pi_{\tau}(\mu^{l+t-2}) \left[\frac{1-\delta^{\tau-1}}{1-\delta} p^H(\mu^{l-1}) + \delta^{\tau} \frac{p^L(\mu^{l-2})}{1-\delta} \right]$$

is the seller's expected payoff at $t+1$, if her first failure to sell at the high price had occurred at t . After this has happened at some t , the seller's strategy resembles the strategy in case (i), and thus $\hat{U}(t, \mu^{l-1})$ resembles $\hat{U}^{(i)}$.

Analogously to case (i), we can calculate

$$\begin{aligned} & \sum_{t=1}^{\infty} \pi_t(\mu^l) \delta^t \hat{U}(t, \mu^{l-1}) \\ = & \sum_{t=1}^{\infty} \pi_t(\mu^l) \delta^t \left\{ \frac{p^H(\mu^{l-1})}{1-\delta} \left[1 - \frac{(1-\alpha)\mu^{l+t-2}}{1-\alpha\delta} - \frac{\alpha(1-\mu^{l+t-2})}{1-(1-\alpha)\delta} \right] + \right. \\ & \left. \delta \frac{p^L(\mu^{l-2})}{1-\delta} \left[\frac{(1-\alpha)\mu^{l+t-2}}{1-\alpha\delta} + \frac{\alpha(1-\mu^{l+t-2})}{1-(1-\alpha)\delta} \right] \right\} \\ \leq & \frac{1}{1-\delta} \sum_{t=1}^{\infty} \pi_t(\mu^l) \delta^t \left[\frac{\alpha(1-\delta)}{1-\alpha\delta} p^H(\mu^{l-1}) + \frac{\alpha\delta}{1-(1-\alpha)\delta} p^L(\mu^{l-2}) \right] \\ \leq & \frac{1}{1-\delta} \frac{\alpha\delta}{1-(1-\alpha)\delta} \left[\frac{\alpha(1-\delta)}{1-\alpha\delta} p^H(\mu^{l-1}) + \frac{\alpha\delta}{1-(1-\alpha)\delta} p^L(\mu^{l-2}) \right] \end{aligned}$$

where the last inequality follows from $\sum_{t=1}^{\infty} \pi_t(\mu^l) \delta^{t-1} = \frac{(1-\alpha)\mu^l}{1-\alpha\delta} + \frac{\alpha(1-\mu^l)}{1-(1-\alpha)\delta} \leq \frac{\alpha}{1-(1-\alpha)\delta}$. Together with $\sum_{t=1}^{\infty} \pi_t(\mu^l) (1-\delta^{t-1}) p^H(\mu^l) = p^H(\mu^l) \left[1 - \frac{(1-\alpha)\mu^l}{1-\alpha\delta} - \frac{\alpha(1-\mu^l)}{1-(1-\alpha)\delta}\right] \leq p^H(\mu^l) \frac{\alpha(1-\delta)}{1-\alpha\delta}$ these calculations show that $\hat{U}^{(ii)}(\mu^l) > \frac{1}{1-\delta} p^L(\mu^l)$ implies

$$\frac{\alpha(1-\delta)}{1-\alpha\delta} p^H(\mu^l) + \frac{\alpha\delta}{1-(1-\alpha)\delta} \left[\frac{\alpha(1-\delta)}{1-\alpha\delta} p^H(\mu^{l-1}) + \frac{\alpha\delta}{1-(1-\alpha)\delta} p^L(\mu^{l-2}) \right] > p^L(\mu^l).$$

Since $p^H(\mu^l) > p^H(\mu^{l-1})$, it follows that

$$\left[1 + \frac{\alpha\delta}{1-(1-\alpha)\delta} \right] \frac{\alpha(1-\delta)}{(1-\alpha\delta)} p^H(\mu^l) + \left[\frac{\alpha\delta}{1-(1-\alpha)\delta} \right]^2 p^L(\mu^{l-2}) > p^L(\mu^l).$$

Following the same approach as in case (i), this implies

$$(1-\delta) \frac{\alpha^3}{(1-\alpha\delta)(1-\alpha)^2} \left[1 + \frac{\alpha\delta}{1-(1-\alpha)\delta} \right] + \left[\frac{\alpha\delta}{1-(1-\alpha)\delta} \right]^2 \frac{p^L(\mu^{l-2})}{p^L(\mu^l)} > 1,$$

which cannot hold if δ is sufficiently large. Consequently, $\hat{U}^{(ii)} < \frac{1}{1-\delta} p^L(\mu^l)$ for sufficiently large δ .

Case (iii): This case includes all strategies where the seller charges the high price until three buyers have revealed a bad signal. Let $\hat{U}^{(iii)}$ denote the maximum expected payoff that the seller can achieve by deviations of this type. The proof that for a sufficiently patient seller none of these deviating strategies results in a higher expected payoff than $\frac{1}{1-\delta} p^L(\mu^l)$, is again similar to the proofs for the cases (i) and (ii). The only difference is the continuation value term since we don't specify whether the seller charges the high or the low price at nodes (λ, μ^{l-3}) . However, since $W(\lambda, \mu^{l-3}) \geq W(\lambda, \mu^{l-K})$ for all $K \geq 3$, we need only focus on an upper bound for $W(\lambda, \mu^{l-3})$. Let t , τ , and k , respectively, denote the (random) period when for the first, the second, and the third time, respectively, a buyer has revealed a bad signal. Thus at $T \equiv t + \tau + k + 1$ buyers' beliefs become μ^{l-3} and the optimal price is bounded from above by $p^H(\mu^{l-3}) = \mu^{l-2} = p^L(\mu^{l-1})$. It follows that for any realization T the seller's expected payoff at T is bounded by $\frac{1}{1-\delta} p^L(\mu^{l-1})$. Consequently,

$$\begin{aligned} \hat{U}^{(iii)} \leq & \sum_{t=1}^{\infty} \pi_t(\mu^l) \left\{ \frac{1-\delta^{t-1}}{1-\delta} p^H(\mu^l) + \delta^t \sum_{\tau=1}^{\infty} \pi_{\tau}(\mu^{l+t-2}) \left[\frac{1-\delta^{\tau-1}}{1-\delta} p^H(\mu^{l-1}) \right. \right. \\ & \left. \left. + \delta^{\tau} \sum_{k=1}^{\infty} \pi_k(\mu^{l+t+\tau-4}) \left(\frac{1-\delta^{k-1}}{1-\delta} p^H(\mu^{l-2}) + \delta^k \frac{p^L(\mu^{l-1})}{1-\delta} \right) \right] \right\}, \end{aligned}$$

where the second (third) sum relates to the continuation value after buyer t (buyers t and $t + \tau$) has (have) declined to buy the object. Hence, $\hat{U}^{(iii)} > \frac{1}{1-\delta} p^L(\mu^l)$ implies

$$p^H(\mu^l) \sum_{t=1}^{\infty} \pi_t(\mu^l) \left\{ (1-\delta^{t-1}) + \delta^t \sum_{\tau=1}^{\infty} \pi_{\tau}(\mu^{l+t-2}) [(1-\delta^{\tau-1}) + \right.$$

$$\left. \delta^\tau \sum_{k=1}^{\infty} \pi_k(\mu^{l+t+\tau-4}) (1 - \delta^{k-1}) \right\} + \sum_{t=1}^{\infty} \pi_t(\mu^l) \delta^t \sum_{\tau=1}^{\infty} \pi_\tau(\mu^{l+t-2}) \delta^\tau \sum_{k=1}^{\infty} \pi_k(\mu^{l+t+\tau-4}) \delta^k p^L(\mu^{l-1}) > p^L(\mu^l).$$

Following the same approach as earlier, this inequality and some calculation shows that $\hat{U}^{(iii)} > \frac{1}{1-\delta} p^L(\mu^l)$ implies

$$\frac{(1-\delta)\alpha^3}{(1-\alpha\delta)(1-\alpha)^2} \left[1 + \frac{\alpha\delta}{1-(1-\alpha)\delta} + \left(\frac{\alpha\delta}{1-(1-\alpha)\delta} \right)^2 \right] + \left[\frac{\alpha\delta}{1-(1-\alpha)\delta} \right]^3 \frac{p^L(\mu^{l-1})}{p^L(\mu^l)} > 1.$$

We conclude again that for sufficiently large δ it is not possible that $\hat{U}^{(iii)} > \frac{1}{1-\delta} p^L(\mu^l)$.

It follows that for sufficiently large values of δ it is a best response for the seller to charge the low price $p^L(\mu^l)$ in each period. Hence there exists a $\bar{\delta} \in (0, 1)$ such that for all $\delta > \bar{\delta}$ there exists a PBE where $p_t = p^L(\lambda_1)$ for all $t \in \{1, 2, \dots\}$. Moreover, since $\frac{p^L(\mu^{l-2})}{p^L(\mu^{l-1})} < \frac{p^L(\mu^{l-1})}{p^L(\mu^l)} < \frac{1-\alpha}{(1-\alpha)\bar{\lambda}_\alpha + \alpha(1-\bar{\lambda}_\alpha)} < 1$, $\bar{\delta}$ can be chosen independently of $\lambda_1 \in (0, \bar{\lambda}_\alpha)$. ■

A.3. Proof of Lemma 10

Lemma 10. *If either (i) $\alpha\lambda_1 \leq \lambda_1^-$, i.e., signals are weak or borderline, or signals are strong and $\lambda_1 \in [\bar{\lambda}_\alpha, 1)$, or (ii) $\delta > \hat{\delta}$ for some sufficiently large $\hat{\delta} \in (0, 1)$, then there does not exist a PBE where for all signal realizations the price is $p_t = p^H(\lambda_t)$ for all $t \in \{1, 2, \dots\}$, that is, where the seller always demands the high price.*

Proof: We first show that the Lemma holds when $\alpha\lambda_1 \leq \lambda_1^-$. The proof is by contradiction. Assume it is a PBE to charge always the high price. Then the seller's expected equilibrium payoff is

$$\begin{aligned} U &= E \left[\sum_{t=1}^{\infty} \delta^{t-1} p^H(\lambda_t) a_t \mid \lambda_1 \right] = E \left\{ \sum_{t=1}^{\infty} \delta^{t-1} E[p^H(\lambda_t) a_t \mid \lambda_t] \mid \lambda_1 \right\} \\ &= E \left[\sum_{t=1}^{\infty} \delta^{t-1} \alpha \lambda_t \mid \lambda_1 \right] = \alpha \lambda_1 + \delta \frac{\alpha \lambda_1}{1-\delta}, \end{aligned}$$

where the last equality follows because $\{\lambda_t\}_{t=1}^{\infty}$ is a martingale and thus $E(\lambda_t \mid \lambda_1) = \lambda_1$.

Consider the following deviation of the seller. The seller demands the low price $p_1 = p^L(\lambda_1) = \lambda_1^-$ instead of the high price $p^H(\lambda_1)$ in period 1, and after that the seller always demands the high price, i.e., $p_t = p^H(\mu_t)$ for all $t \in \{2, 3, \dots\}$, where μ_t is the inference of buyer t from the history of actions. Since buyer $t \in \{2, 3, \dots\}$ cannot observe the seller's deviation, he will update according to the seller's equilibrium strategy, $p_t = p^H(\lambda_t)$ for all

$t \in \{1, 2, \dots\}$, which in particular implies that he believes that $p_1 = p^H(\lambda_1)$. To distinguish the stochastic process of the seller's beliefs that is generated by the deviating strategy from the respective stochastic process $\{\lambda_t\}_{t=2}^\infty$ that is generated by the equilibrium strategy, we denote the former by $\{\hat{\lambda}_t\}_{t=2}^\infty$. Thus, when the seller deviates her beliefs at $t \geq 2$ are given by $\hat{\lambda}_t$, and the respective inference of buyer t is given by $\mu_t = \hat{\lambda}_t^+$, $t \in \{2, 3, \dots\}$. Denoting the seller's expected payoff from the deviating strategy by \hat{U} we get

$$\begin{aligned}
\hat{U} &= \lambda_1^- + E \left[\sum_{t=2}^{\infty} \delta^{t-1} p^H(\hat{\lambda}_t^+) a_t \mid \lambda_1 \right] \\
&= \lambda_1^- + E \left\{ \sum_{t=2}^{\infty} \delta^{t-1} E \left[p^H(\hat{\lambda}_t^+) a_t \mid \hat{\lambda}_t \right] \mid \lambda_1 \right\} \\
&= \lambda_1^- + E \left[\sum_{t=2}^{\infty} \delta^{t-1} \varphi(\hat{\lambda}_t) p^H(\hat{\lambda}_t^+) \mid \lambda_1 \right] \\
&= \lambda_1^- + E \left[\sum_{t=2}^{\infty} \delta^{t-1} \varphi(\hat{\lambda}_t) p^H(\hat{\lambda}_t) \frac{p^H(\hat{\lambda}_t^+)}{p^H(\hat{\lambda}_t)} \mid \lambda_1 \right] \\
&= \lambda_1^- + E \left[\sum_{t=2}^{\infty} \delta^{t-1} \frac{\alpha}{\alpha \hat{\lambda}_t^+ + (1-\alpha)(1-\hat{\lambda}_t^+)} \alpha \hat{\lambda}_t \mid \lambda_1 \right] \\
&= \lambda_1^- + E \left[\sum_{t=2}^{\infty} \delta^{t-1} \gamma_t \alpha \hat{\lambda}_t \mid \lambda_1 \right],
\end{aligned}$$

where γ_t is defined by $\gamma_t \equiv \frac{\alpha}{\alpha \hat{\lambda}_t^+ + (1-\alpha)(1-\hat{\lambda}_t^+)} > 1$, $t \in \{2, 3, \dots\}$. Since $\gamma_t > 1$ for all t and since $\{\hat{\lambda}_t\}_{t=2}^\infty$ is a martingale where $\hat{\lambda}_2 = \lambda_1$,

$$\hat{U} = \lambda_1^- + E \left[\sum_{t=2}^{\infty} \delta^{t-1} \gamma_t \alpha \hat{\lambda}_t \mid \lambda_1 \right] > \lambda_1^- + E \left[\sum_{t=2}^{\infty} \delta^{t-1} \alpha \hat{\lambda}_t \mid \lambda_1 \right] = \lambda_1^- + \delta \frac{\alpha \lambda_1}{1-\delta}.$$

It must hold that $U \geq \hat{U}$, and thus $\alpha \lambda_1 > \lambda_1^-$. This proves that the strategy $p_t = p^H(\lambda_t)$ for all $t \in \{1, 2, \dots\}$ cannot be an equilibrium strategy whenever $\alpha \lambda_1 \leq \lambda_1^-$.

Consider now the case $\alpha \lambda_1 > \lambda_1^-$, i.e., $\alpha^2 > 1 - \alpha$ and $\lambda_1 \in (0, \bar{\lambda}_\alpha)$. We proceed in two steps. First, we show that if λ_1 is sufficiently small, a sufficiently patient seller will immediately deviate from the hypothetical equilibrium strategy and demand the low price $p^L(\lambda_1)$ instead of the high price $p^H(\lambda_1)$. Second, we show that this implies that for any prior $\lambda_1 \in (0, \bar{\lambda}_\alpha)$ a sufficiently patient seller will deviate from the hypothetical equilibrium strategy at some node that is reached with positive probability.

For the first step we consider the deviation where the seller demands the low price $p^L(\lambda_1)$ in period $t = 1$ and the high price $p^H(\hat{\lambda}_t^+)$ thereafter. If $s_1 = g$, the deviating seller demands

the same price that she would have demanded along the equilibrium path and receives the same expected revenue in period $t = 2$ and in all subsequent periods. If $s_1 = b$, the deviation to $p^L(\lambda_1)$ in period 1 results in a higher expected revenue in each period $t \in \{2, 3, \dots\}$. Thus, if the expected gain in period $t = 2$ alone compensates for the expected loss in period $t = 1$, then the seller will deviate. For $s_1 = b$ and $s_2 = b$, the seller's revenue in period 2 is zero, since in this case there is no sale. The remaining case is $s_1 = b$, $s_2 = g$, which ex ante has probability $[1 - \varphi(\lambda_1)]\varphi(\lambda_1^-) = \alpha(1 - \alpha)$. In this case, the seller receives $p^H(\lambda_1^-) = \lambda_1$ in period $t = 2$ along the equilibrium path and $p^H(\lambda_1^+)$ if she has deviated in period 1. Therefore, a necessary condition for the seller not to deviate is

$$\alpha\lambda_1 - \lambda_1^- - \alpha(1 - \alpha)\delta[p^H(\lambda_1^+) - \lambda_1] \geq 0.$$

The necessary condition that no seller, however patient, deviates is that

$$\begin{aligned} 0 &\leq \alpha\lambda_1 - \lambda_1^- - \alpha(1 - \alpha)[p^H(\lambda_1^+) - \lambda_1] \\ &= \lambda_1 - (1 - \alpha)^2\lambda_1 - \lambda_1^- - \alpha(1 - \alpha)p^H(\lambda_1^+) \\ &= \lambda_1 - (1 - \alpha)^2\lambda_1 - \frac{1 - \alpha}{(1 - \alpha)\lambda_1 + \alpha(1 - \lambda_1)}\lambda_1 - \frac{\alpha^3(1 - \alpha)}{\alpha^2\lambda_1 + (1 - \alpha)^2(1 - \lambda_1)}\lambda_1, \end{aligned}$$

which implies

$$(1 - \alpha)\frac{1}{(1 - \alpha)\lambda_1 + \alpha(1 - \lambda_1)} + \alpha\frac{\alpha^2(1 - \alpha)}{\alpha^2\lambda_1 + (1 - \alpha)^2(1 - \lambda_1)} + (1 - \alpha)^2 \leq 1. \quad (6.12)$$

Since $\frac{1}{(1 - \alpha)\lambda_1 + \alpha(1 - \lambda_1)} > \frac{1}{\alpha} > 1$ for all λ_1 and $\frac{\alpha^2(1 - \alpha)}{\alpha^2\lambda_1 + (1 - \alpha)^2(1 - \lambda_1)} \geq 1$ for sufficiently small λ_1 (because this ratio converges to $\frac{\alpha^2}{(1 - \alpha)} > 1$ for $\lambda_1 \rightarrow 0$), the inequality (6.12) is violated if λ_1 is sufficiently small. For $\lambda_1 \rightarrow 0$ the left hand side of (6.12) becomes $\frac{1 - \alpha}{\alpha} + \frac{\alpha^3}{(1 - \alpha)}$. If we define $\hat{\delta}$ by $\frac{1 - \alpha}{\alpha} + \hat{\delta}\frac{\alpha^3}{(1 - \alpha)} = 1$, then $\frac{1 - \alpha}{\alpha} + \delta\frac{\alpha^3}{(1 - \alpha)} > 1$ for all $\delta > \hat{\delta}$, and thus for each $\delta > \hat{\delta}$ there exists a sufficiently small λ_1 such that the seller will deviate from the hypothetical equilibrium strategy.

Finally, consider a prior λ_1 where (6.12) is satisfied. Let $\lambda^l \equiv \lambda_1$. The previous analysis implies that for each $\delta > \hat{\delta}$ there exists a λ^{l-K} that is “ K steps below λ_1 ”, i.e., $\lambda^{l-K} \equiv \Pr(\omega = G \mid \lambda_1, K \text{ signals } s = b)$, such that the inequality (6.12) is violated when λ_1 is replaced by λ^{l-K} , i.e.,

$$\frac{1 - \alpha}{(1 - \alpha)\lambda^{l-K} + \alpha(1 - \lambda^{l-K})} + \frac{\alpha^3(1 - \alpha)}{\alpha^2\lambda^{l-K} + (1 - \alpha)^2(1 - \lambda^{l-K})} + (1 - \alpha)^2 > 1. \quad (6.13)$$

The realization $s_t = b$ for all $t \in \{1, \dots, K\}$ has positive probability and gives $\lambda_{K+1} = \lambda^{l-K}$ along the equilibrium path. If in $t = K + 1$ the seller deviates and demands $p^L(\lambda_{K+1})$

instead of $p^H(\lambda_{K+1})$, buyer $K+1$ observes the deviation, but will nevertheless purchase the object for the price $p^L(\lambda_{K+1})$ because the only rational inference from the public history $h_K = (0, \dots, 0)$ is $p_t = p^H(\lambda_t)$ and $s_t = b$ for all $t \in \{1, \dots, K\}$ and thus $\lambda_{K+1} = \lambda^{l-K}$ even though $p_{K+1} = p^L(\lambda_{K+1})$ differs from the equilibrium price $p^H(\lambda_{K+1})$. That is, buyer $K+1$ will rationally believe that the seller deviates in period $t = K+1$ for the first time. Buyers $\tau \geq K+2$ are unable to observe the seller's deviation in period $t = K+1$ and thus will buy the object at the respective high price if and only if $s_\tau = g$. Consequently, the previous analysis with λ_1 being replaced by $\lambda_{K+1} = \lambda^{l-K}$ applies, and because of (6.13) the seller will deviate. \blacksquare

A.4. Proof of Propositions 6 and 7

Proposition 6. *Assume that previous prices are unobservable for buyers. If either $\alpha^2 \leq 1 - \alpha$ or $\alpha^2 > 1 - \alpha$ and $\lambda_1 \in [\bar{\lambda}_\alpha, 1)$, then any pure strategy PBE where the seller's equilibrium strategy is Markov has immediate herding, i.e., $p_t = p^L(\lambda_1)$ for all $t \in \{1, 2, \dots\}$, as equilibrium outcome. That is, under these conditions immediate herding is the unique equilibrium outcome.*

Proof: Existence follows from Proposition 4. Notice that $\alpha\lambda_1 \leq \lambda_1^-$, since $\alpha^2 \leq 1 - \alpha$ or $\alpha^2 > 1 - \alpha$ and $\lambda_1 \geq \bar{\lambda}_\alpha \geq \bar{\lambda}_\alpha$. The proof of uniqueness is by contradiction. Assume that there is a pure strategy PBE where the seller's equilibrium strategy is Markov and $p_1 = p^H(\lambda_1)$. The seller's equilibrium strategy prescribes for each $\lambda_t \in \Lambda(\lambda_1)$ a price $p^*(\lambda_t)$, and by assumption $p^*(\lambda_1) = p^H(\lambda_1)$. Lemma 10 and the assumption that the seller's equilibrium strategy is Markov imply that along the equilibrium path herding must occur at some attainable $\lambda_L < \lambda_1$ or at some attainable $\lambda_H > \lambda_1$, or both. Assume first that there is a pure strategy PBE where there is herding only at some $\lambda_L < \lambda_1$, i.e., the equilibrium price $p^*(\lambda_t)$ is given by

$$p^*(\lambda_t) = \begin{cases} p^H(\lambda_t) & \text{for } \lambda_t > \lambda_L \\ p^L(\lambda_t) & \text{for } \lambda_t = \lambda_L \end{cases}$$

and $\lambda_t \geq \lambda_L$ for all t . For any t , let $\chi_t \equiv \Pr(\lambda_t = \lambda_L \mid \lambda_1) < 1$ denote the probability that $\lambda_t = \lambda_L$ and thus $p_t = p^L(\lambda_L)$. The seller's expected revenue in period t conditional on λ_t is $\alpha\lambda_t$ if $\lambda_t > \lambda_L$, and $p^L(\lambda_L)$ if $\lambda_t = \lambda_L$. Since the stochastic process $\{\lambda_t\}_{t=1}^\infty$ is a martingale, the seller's expected revenue in any period t conditional on λ_1 is

$$\begin{aligned} E[p^*(\lambda_t) \mid \lambda_1] &= \chi_t p^L(\lambda_L) + (1 - \chi_t) E[\alpha\lambda_t \mid \lambda_1, \lambda_t > \lambda_L] \\ &= \chi_t p^L(\lambda_L) + E[\alpha\lambda_t \mid \lambda_1] - \chi_t E[\alpha\lambda_t \mid \lambda_1, \lambda_t = \lambda_L] \end{aligned}$$

$$\begin{aligned}
&= \chi_t [p^L(\lambda_L) - \alpha\lambda_L] + \alpha\lambda_1 \\
&= \chi_t [p^L(\lambda_L) - \alpha\lambda_L] - [p^L(\lambda_1) - \alpha\lambda_1] + p^L(\lambda_1) \\
&< p^L(\lambda_1)
\end{aligned}$$

where the last inequality follows from the facts that $\lambda_L < \lambda_1$, and that the assumptions of the proposition imply, $p^L(\lambda_1) - \alpha\lambda_1 > 0$, and $p^L(\lambda_L) - \alpha\lambda_L < p^L(\lambda_1) - \alpha\lambda_1$. It follows that the seller's expected payoff from her equilibrium strategy is less than $\frac{1}{1-\delta}p^L(\lambda_1)$. But the seller can guarantee herself a payoff of $\frac{1}{1-\delta}p^L(\lambda_1)$ by charging $p^L(\lambda_1)$ in each period. The reason is that $\mu_t \geq \lambda_1$ for any buyer t who observes purchases from all previous buyers. This proves that in a PBE that satisfies the assumptions of the proposition there must be herding at some attainable $\lambda \geq \lambda_1$. Let λ_H denote the lowest attainable $\lambda \geq \lambda_1$ where herding occurs. If the PBE prescribes the seller to demand the high price for all $\lambda_t < \lambda_1$, i.e., if for $\lambda_t < \lambda_1$ there is no herding, define $\lambda_L \equiv 0$; otherwise let λ_L denote the highest attainable $\lambda < \lambda_1$ for which the PBE prescribes the low price and thus herding. Assume that there exists a PBE that satisfies our assumptions and has $\lambda_H > \lambda_1$. Define the integer K by $\Pr(\omega = G \mid \lambda_1, K \text{ signals } s = g) = \lambda_H$, i.e., λ_H is “ K steps above λ_1 .” Consider the seller's deviation to charge $p^L(\lambda_1)$ in the first K periods, i.e., $t \in \{1, \dots, K\}$ and $p^L(\lambda_H)$ thereafter. Since for any buyer $t \in \{1, \dots, K\}$ the minimum estimated probability for the good state is λ_1 and since buyers $t \in \{K + 1, \dots\}$ cannot observe the seller's deviation, $a_t = 1$ for all t . Moreover, starting with $t = K + 1$ the seller gets the maximum per period revenue $p^L(\lambda_H) = \lambda_H^-$. The expected payoff that is associated with this deviating strategy is $\widehat{U} = \sum_{t=1}^K \delta^{t-1} p^L(\lambda_1) + \delta^K \frac{1}{1-\delta} p^L(\lambda_H)$. If the seller follows the equilibrium strategy, her expected revenue in period $t \in \{1, \dots, K\}$ conditional on λ_1 is either $E[p^*(\lambda_t) a_t \mid \lambda_1, \lambda_L < \lambda_t < \lambda_H] = E[\alpha\lambda_t \mid \lambda_1, \lambda_L < \lambda_t < \lambda_H]$ or $p^L(\lambda_L)$, since it takes at least $K + 1$ periods to get herding at λ_H and thus either $\lambda_L < \lambda_t < \lambda_H$ or $\lambda_t = \lambda_L$. Repeating the argument employed above we get $E[p^*(\lambda_t) a_t \mid \lambda_1] < p^L(\lambda_1)$ for $t \in \{1, \dots, K\}$. For the expected equilibrium payoff U^* this implies $U^* < \sum_{t=1}^K \delta^{t-1} p^L(\lambda_1) + \delta^K \frac{1}{1-\delta} p^L(\lambda_H) = \widehat{U}$. This contradicts the assumption that there is a PBE that satisfies our assumptions and has $\lambda_H > \lambda_1$. Consequently, $\lambda_H = \lambda_1$, i.e., under the assumptions of the proposition $p_t = p^L(\lambda_1)$ for all $t \in \{1, 2, \dots\}$ is the only equilibrium outcome. \blacksquare

Proposition 7. *Assume that previous prices are unobservable for buyers. If $\alpha^2 > 1 - \alpha$ and $\lambda_1 \in (0, \bar{\lambda}_\alpha)$, there exists a $\bar{\delta} \in (0, 1)$ such that for each $\delta > \bar{\delta}$ any pure strategy PBE where the seller's equilibrium strategy is Markov has immediate herding, i.e., $p_t = p^L(\lambda_1)$ for all $t \in \{1, 2, \dots\}$, as equilibrium outcome. That is, under these conditions immediate herding is the unique equilibrium outcome.*

Proof: Existence follows from Proposition 5 for all $\delta > \bar{\delta}$. For uniqueness consider first the case where $\bar{\lambda}_\alpha \leq \lambda_1 < \bar{\bar{\lambda}}_\alpha$. In this case, $\alpha\lambda_1 \leq \lambda_1^-$ and hence the proposition follows from a proof identical to that provided for Proposition 6.

Next consider the case where $\lambda_1 < \bar{\lambda}_\alpha$. The proof is by contradiction. Assume there is a PBE with a different equilibrium outcome. Denote the respective equilibrium prices by $p^*(\lambda_t)$, $\lambda_t \in \Lambda(\lambda_1)$. By assumption, $p^*(\lambda_1) = p^H(\lambda_1)$. Let λ_L denote the largest attainable $\lambda < \lambda_1$ such that $p^*(\lambda) = p^L(\lambda)$, if such a λ exists, and define $\lambda_L \equiv 0$ otherwise. Similarly let λ_H denote the smallest attainable $\lambda > \lambda_1$ such that $p^*(\lambda) = p^L(\lambda)$, if such a λ exists, and define $\lambda_H \equiv 1$ otherwise. That is, if $\lambda_L > 0$ there is herding at λ_L below λ_1 , and if $\lambda_H < 1$ there is herding at λ_H above λ_1 . Along the equilibrium path, $\lambda_L \leq \lambda_t \leq \lambda_H$ for all $t \in \{1, 2, \dots\}$. The PBE generates a particular martingale $\{\lambda_t\}_{t=1}^\infty$ which we denote by $\{\tilde{\lambda}_t\}_{t=1}^\infty$. We consider the deviation where the seller demands $p^L(\lambda_1)$ instead of $p^*(\lambda_1) = p^H(\lambda_1)$ in period 1 and “mimics” the PBE from period 2 onwards by demanding the price $\hat{p}(\hat{\lambda}_t) = p^*(\hat{\lambda}_t^+) \in \{p^L(\hat{\lambda}_t^+), p^H(\hat{\lambda}_t^+)\}$, where $\hat{\lambda}_t$ denotes the seller’s estimation $\Pr(\omega = G | \lambda_1; H_{t-1})$ along the deviation. Buyer $t \in \{2, 3, \dots\}$ cannot observe the seller’s deviation at $t = 1$ and therefore estimates the probability of the good state to be $\hat{\lambda}_t^+$ for each $\hat{\lambda}_t \in \Lambda(\lambda_1)$.

Define (in slight abuse of notation) the function $a(s_t)$, $s_t \in \{b, g\}$, by $a(s_t) = 0$ for $s_t = b$ and $a(s_t) = 1$ for $s_t = g$, and let $d(\hat{\lambda}_t, \tilde{\lambda}_t, s_t)$ denote the difference in the seller’s revenue in period $t \in \{2, 3, \dots\}$, if she deviates rather than follows the equilibrium strategy. That is,

$$d(\hat{\lambda}_t, \tilde{\lambda}_t, s_t) \equiv \begin{cases} \left[p^H(\hat{\lambda}_t^+) - p^H(\tilde{\lambda}_t) \right] a(s_t) & \text{if } \hat{p}(\hat{\lambda}_t) = p^H(\hat{\lambda}_t^+) \text{ and } p^*(\tilde{\lambda}_t) = p^H(\tilde{\lambda}_t) \\ p^H(\hat{\lambda}_t^+) a(s_t) - p^L(\tilde{\lambda}_t) & \text{if } \hat{p}(\hat{\lambda}_t) = p^H(\hat{\lambda}_t^+) \text{ and } p^*(\tilde{\lambda}_t) = p^L(\tilde{\lambda}_t) \\ p^L(\hat{\lambda}_t^+) - p^H(\tilde{\lambda}_t) a(s_t) & \text{if } \hat{p}(\hat{\lambda}_t) = p^L(\hat{\lambda}_t^+) \text{ and } p^*(\tilde{\lambda}_t) = p^H(\tilde{\lambda}_t) \\ p^L(\hat{\lambda}_t^+) - p^L(\tilde{\lambda}_t) & \text{if } \hat{p}(\hat{\lambda}_t) = p^L(\hat{\lambda}_t^+) \text{ and } p^*(\tilde{\lambda}_t) = p^L(\tilde{\lambda}_t) \end{cases}$$

In period 1 the deviation generates the revenue $p^L(\lambda_1)$, whereas $p^*(\lambda_1) = p^H(\lambda_1)$ generates the expected revenue $\alpha\lambda_1 > p^L(\lambda_1)$. Consider the difference in period revenues for $t \in \{2, 3, \dots\}$. If $s_1 = g$, $\tilde{\lambda}_t = \hat{\lambda}_t^+$ and thus the buyers’ beliefs and the seller’s prices along the equilibrium path are identical to the buyer’s beliefs and seller’s prices along the deviating path for all $t \in \{2, 3, \dots\}$. Consequently, the seller’s revenue is identical along the equilibrium and the deviating path, i.e., $d(\hat{\lambda}_t, \tilde{\lambda}_t, s_t) = 0$ for all $t \in \{2, 3, \dots\}$ for each realization of signals $(s_1, s_2, \dots) = (g, s_2, s_3, \dots)$.

Consider now the case $s_1 = b$. In this case, $\tilde{\lambda}_t = \hat{\lambda}_t^-$ for all $t \in \{2, 3, \dots\}$ as long as $\lambda_L < \hat{\lambda}_t^+ < \lambda_H$ and $\lambda_L < \tilde{\lambda}_t < \lambda_H$, i.e., provided neither at the equilibrium path nor at the

deviating path herding has occurred. Notice that $\tilde{\lambda}_t = \lambda_H$ implies that $\lambda_L < \tilde{\lambda}_t \leq \hat{\lambda}_t^+$ for all $\tau \in \{1, \dots, t\}$ and $\hat{\lambda}_t^+ = \lambda_H$ because whenever $\tilde{\lambda}_t = \lambda_H$, $\hat{\lambda}_\tau^+ = \lambda_H$ at some $\tau < t$, i.e., herding at λ_H along the equilibrium path implies herding at λ_H along the deviating path. Similarly, herding at λ_L along the deviating path implies herding at λ_L along the equilibrium path because whenever $\hat{\lambda}_t^+ = \lambda_L$ it must be the case that $\tilde{\lambda}_\tau = \hat{\lambda}_\tau^- = \lambda_L$ in some earlier period $\tau < t$. These arguments imply that in the case $s_1 = b$,

$$d\left(\hat{\lambda}_t, \tilde{\lambda}_t, s_t\right) = \begin{cases} \left[p^H\left(\hat{\lambda}_t^+\right) - p^H\left(\hat{\lambda}_t^-\right)\right] a\left(s_t\right) & \text{if } \lambda_L < \hat{\lambda}_t^+ < \lambda_H \text{ and } \lambda_L < \tilde{\lambda}_t < \lambda_H \\ 0 & \text{if } \hat{\lambda}_t^+ = \lambda_L \text{ or } \tilde{\lambda}_t = \lambda_H \\ p^H\left(\hat{\lambda}_t^+\right) a\left(s_t\right) - p^L\left(\lambda_L\right) & \text{if } \lambda_H > \hat{\lambda}_t^+ > \tilde{\lambda}_t = \lambda_L \\ p^L\left(\lambda_H\right) - p^H\left(\tilde{\lambda}_t\right) a\left(s_t\right) & \text{if } \lambda_H = \hat{\lambda}_t^+ > \tilde{\lambda}_t > \lambda_L \\ p^L\left(\lambda_H\right) - p^L\left(\lambda_L\right) & \text{if } \hat{\lambda}_t^+ = \lambda_H \text{ and } \tilde{\lambda}_t = \lambda_L \end{cases}$$

Using the notation $\hat{\lambda}_t^{++} \equiv p^H\left(\hat{\lambda}_t^+\right)$, we get $E\left[p^H\left(\hat{\lambda}_t^+\right) a\left(s_t\right) \mid \hat{\lambda}_t, s_1 = b\right] = \varphi\left(\hat{\lambda}_t^-\right) p^H\left(\hat{\lambda}_t^+\right) = \varphi\left(\hat{\lambda}_t^-\right) \hat{\lambda}_t \frac{\hat{\lambda}_t^{++}}{\hat{\lambda}_t} = \alpha \hat{\lambda}_t^- \frac{\hat{\lambda}_t^{++}}{\hat{\lambda}_t}$, and thus

$$E\left[d\left(\hat{\lambda}_t, \tilde{\lambda}_t, s_t\right) \mid \hat{\lambda}_t, \tilde{\lambda}_t, s_1 = b\right] = \begin{cases} \left(\frac{\hat{\lambda}_t^{++}}{\hat{\lambda}_t} - 1\right) \alpha \hat{\lambda}_t^- & \text{if } \lambda_L < \hat{\lambda}_t^+ < \lambda_H \text{ and } \lambda_L < \tilde{\lambda}_t < \lambda_H \\ 0 & \text{if } \hat{\lambda}_t^+ = \lambda_L \text{ or } \tilde{\lambda}_t = \lambda_H \\ \alpha \hat{\lambda}_t^- \frac{\hat{\lambda}_t^{++}}{\hat{\lambda}_t} - \lambda_L^- & \text{if } \lambda_H > \hat{\lambda}_t^+ > \tilde{\lambda}_t = \lambda_L \\ \lambda_H^- - \alpha \tilde{\lambda}_t & \text{if } \lambda_H = \hat{\lambda}_t^+ > \tilde{\lambda}_t > \lambda_L \\ \lambda_H^- - \lambda_L^- & \text{if } \hat{\lambda}_t^+ = \lambda_H \text{ and } \tilde{\lambda}_t = \lambda_L \end{cases}$$

for all $t \in \{2, 3, \dots\}$. Next, we show that unless $\hat{\lambda}_t^+ = \lambda_L$ or $\tilde{\lambda}_t = \lambda_H$ the terms on the right hand side are strictly positive. First, $\left(\frac{\hat{\lambda}_t^{++}}{\hat{\lambda}_t} - 1\right) \alpha \hat{\lambda}_t^- > 0$. Second, because $\lambda_L^+ \leq \lambda_1$ and $\alpha \lambda > \lambda^-$ for all $\lambda \leq \lambda_1$, $\alpha \lambda_L > \lambda_L^-$ and $\alpha \lambda_L^+ > \lambda_L$. Therefore, if $\hat{\lambda}_t > \lambda_L$, $\alpha \hat{\lambda}_t^- \frac{\hat{\lambda}_t^{++}}{\hat{\lambda}_t} - \lambda_L^- > \alpha \hat{\lambda}_t^- - \lambda_L^- \geq \alpha \lambda_L - \lambda_L^- > 0$, whereas if $\hat{\lambda}_t = \lambda_L$, $\alpha \hat{\lambda}_t^- \frac{\hat{\lambda}_t^{++}}{\hat{\lambda}_t} - \lambda_L^- > \alpha \lambda_L^- \frac{\lambda_L^{++}}{\alpha \lambda_L^+} - \lambda_L^- = \left(\frac{\lambda_L^{++}}{\lambda_L^+} - 1\right) \lambda_L^- > 0$. Consequently, $\alpha \hat{\lambda}_t^- \frac{\hat{\lambda}_t^{++}}{\hat{\lambda}_t} - \lambda_L^- > 0$ for $\hat{\lambda}_t^+ > \lambda_L$. Third, $\lambda_H^- - \alpha \tilde{\lambda}_t \geq (1 - \alpha) \lambda_H^- > 0$ for $\lambda_H > \tilde{\lambda}_t$. Together with the analysis of the case $s_1 = g$ this implies

$$E\left[d\left(\hat{\lambda}_t, \tilde{\lambda}_t, s_t\right) \mid \lambda_1\right] \geq 0 \quad \text{for all } t \in \{2, 3, \dots\}. \quad (6.14)$$

Thus, for $t \geq 2$ the expected future return from the deviation exceeds the respective return from the equilibrium strategy. Whenever this future gain is sufficient to compensate for the loss from deviating in the first period the seller will deviate. We distinguish between several cases and show that in each case the deviation is beneficial for a sufficiently patient seller.

First, consider the case $\lambda_H = \lambda_1^+$. That is, along the equilibrium path the seller triggers herding as soon as $\tilde{\lambda}_t = \lambda_1^+$. In that case the deviating seller gets the maximal revenue $p^L(\lambda_1^+) = \lambda_1$ for sure in each period $t \in \{2, 3, \dots\}$. If the seller follows the equilibrium strategy her revenue in period t depends on $\tilde{\lambda}_t$. In particular, if $\tilde{\lambda}_t \leq \lambda_1$ for all t , the seller receives an expected revenue that is bounded from above by $\alpha\lambda_1 < \lambda_1$ in each period $t \in \{2, 3, \dots\}$. Let $\pi(\lambda_1) \equiv \Pr(\tilde{\lambda}_t \leq \lambda_1 \text{ for all } t \mid \lambda_1) > 0$ denote the probability of the event $\{\tilde{\lambda}_t \leq \lambda_1 \text{ for all } t\}$. Since $\pi(\lambda_1)$ decreases in λ_1 , $\bar{\pi} \equiv \pi(\bar{\lambda}_\alpha) < \pi(\lambda_1)$ for all $\lambda_1 \in (0, \bar{\lambda}_\alpha)$. Hence an equilibrium where $\lambda_1^+ = \lambda_H$ can only be sustained if

$$\begin{aligned} \alpha\lambda_1 - \lambda_1^- &\geq \frac{\delta}{1-\delta}\lambda_1 - (1-\bar{\pi})\frac{\delta}{1-\delta}\lambda_1 - \bar{\pi}\frac{\delta}{1-\delta}\alpha\lambda_1 \\ &= \frac{\delta}{1-\delta}\bar{\pi}(1-\alpha)\lambda_1 \end{aligned}$$

and thus

$$\alpha - \frac{1-\alpha}{\alpha} > \alpha - \frac{1-\alpha}{(1-\alpha)\lambda_1 + \alpha(1-\lambda_1)} > \frac{\delta}{1-\delta}\bar{\pi}(1-\alpha).$$

There exists a $\hat{\delta}_1 \in (0, 1)$ that is independent of λ_1 such that this inequality is violated for all $\delta > \hat{\delta}_1$.

Next, consider the case $\lambda_L = \lambda_1^-$ and $\lambda_H > \lambda_1^+$. That is, along the equilibrium path the seller triggers herding as soon as $\tilde{\lambda}_t = \lambda_1^-$, but not when $\tilde{\lambda}_t = \lambda_1^+$. If $s_1 = b$, the seller's revenue along the equilibrium path is $p^L(\lambda_1^-)$ in period $t = 2$, whereas the expected revenue of the deviating seller in period $t = 2$ is $\varphi(\lambda_1^-)p^H(\lambda_1^+) = \alpha\frac{\lambda_1^-}{\lambda_1}p^H(\lambda_1^+)$. Hence an equilibrium where $\lambda_1^- = \lambda_L$ can only be sustained if $\alpha\lambda_1 - \lambda_1^- \geq \delta[1 - \varphi(\lambda_1)]\left[\alpha\frac{\lambda_1^-}{\lambda_1}p^H(\lambda_1^+) - p^L(\lambda_1^-)\right]$. Since $\lambda_1^- > p^L(\lambda_1^-)$ and $\lambda_1^- = \frac{(1-\alpha)\lambda_1}{1-\varphi(\lambda_1)}$ this implies $\alpha\lambda_1 > \delta\alpha(1-\alpha)p^H(\lambda_1^+)$ and thus

$$1 > \delta(1-\alpha)\frac{p^H(\lambda_1^+)}{\lambda_1} = \delta\frac{(1-\alpha)\alpha^2}{\alpha^2\lambda_1 + (1-\alpha)^2(1-\lambda_1)}.$$

Because $\frac{(1-\alpha)\alpha^2}{\alpha^2\lambda_1 + (1-\alpha)^2(1-\lambda_1)} \rightarrow \frac{\alpha^2}{1-\alpha} > 1$ for $\lambda_1 \rightarrow 0$, there exists a $\bar{\lambda}_1 > 0$ and an $\varepsilon > 0$ such that $\frac{(1-\alpha)\alpha^2}{\alpha^2\lambda_1 + (1-\alpha)^2(1-\lambda_1)} > 1 + \varepsilon$ for all $\lambda_1 \in (0, \bar{\lambda}_1)$. Consequently, for all $\lambda_1 \in (0, \bar{\lambda}_1)$ it holds that $\delta\frac{(1-\alpha)\alpha^2}{\alpha^2\lambda_1 + (1-\alpha)^2(1-\lambda_1)} > 1$ for all $\delta > \hat{\delta}_2 \equiv \frac{1}{1+\varepsilon} \in (0, 1)$. Thus, if the seller is sufficiently patient, $p^*(\lambda_1) = p^H(\lambda_1)$ implies that $\lambda_1 \geq \bar{\lambda}_1$ for some $\bar{\lambda}_1 > 0$. We continue the proof by showing that if the seller is sufficiently patient, $\lambda_1 \geq \bar{\lambda}_1$ leads to a contradiction as well.

By assumption, along the equilibrium path the seller triggers herding as soon as $\tilde{\lambda}_t = \lambda_1^-$. Let λ_1^{--} denote the probability “one step below λ_1^- ” (or, “two steps below λ_1^- ”), i.e., $\lambda_1^{--} \equiv \Pr(\omega = G \mid \lambda_1^-, s = b) = p^L(\lambda_1^-)$. If $s_1 = b$, the seller's revenue along the equilibrium path is $p^L(\lambda_1^-)$ in each period $t \in \{2, 3, \dots\}$, whereas the expected revenue of the deviating seller in period t conditional on $\hat{\lambda}_t^+$ is at least $\varphi(\lambda_1^{--})p^H(\lambda_1)$ as long as $\hat{\lambda}_t^+ > \lambda_1^-$ and $p^L(\lambda_1^-)$

as soon as $\hat{\lambda}_t^+ = \lambda_1^-$. Note that $\varphi(\lambda_1^-) p^H(\lambda_1) = \alpha \lambda_1^- \frac{\lambda_1^+}{\lambda_1^-} > \frac{\lambda_1^+}{\lambda_1^-} p^L(\lambda_1^-) > p^L(\lambda_1^-)$ because $\alpha \lambda_1 > \lambda_1^-$ and $\lambda_1^- = p^L(\lambda_1^-)$. Let $\chi(\lambda_1) \equiv \Pr(\hat{\lambda}_t^+ > \lambda_1^- \text{ for all } t \mid \lambda_1, s_1 = b) > 0$ denote the probability that conditional on $s_1 = b$ herding at $\lambda_L = \lambda_1^-$ will not occur along the deviating path. Given λ_1 , this probability is lowest when $\lambda_H = 1$, and this lower bound increases in λ_1 . Therefore, there exists a $\bar{\chi} > 0$ such that $\chi(\lambda_1) > \bar{\chi}$ for all $\lambda_1 \geq \bar{\lambda}_1$. Thus, conditional on $s_1 = b$ the expected revenue of the deviating seller in each period $t \in \{2, 3, \dots\}$ is at least $\frac{\lambda_1^+}{\lambda_1^-} p^L(\lambda_1^-)$ with probability $\bar{\chi} > 0$ and at least $p^L(\lambda_1^-)$ with the complementary probability $1 - \bar{\chi}$. Hence if $\lambda_1 \geq \bar{\lambda}_1$ an equilibrium where $p^*(\lambda_1) = p^H(\lambda_1)$ and $\lambda_1^- = \lambda_L$ can only be sustained if

$$\begin{aligned} \alpha \lambda_1 - \lambda_1^- &\geq [1 - \varphi(\lambda_1)] \left[\bar{\chi} \frac{\delta}{1 - \delta} \frac{\lambda_1^+}{\lambda_1^-} p^L(\lambda_1^-) + (1 - \bar{\chi}) \frac{\delta}{1 - \delta} p^L(\lambda_1^-) - \frac{\delta}{1 - \delta} p^L(\lambda_1^-) \right] \\ &= \frac{\delta}{1 - \delta} \bar{\chi} [1 - \varphi(\lambda_1)] \left(\frac{\lambda_1^+}{\lambda_1^-} - 1 \right) p^L(\lambda_1^-) > 0. \end{aligned}$$

There exists a $\hat{\delta}_3 \in (0, 1)$ that is independent of $\lambda_1 \in [\bar{\lambda}_1, \bar{\lambda}_\alpha)$ such that this inequality is violated for all $\delta > \hat{\delta}_3$.

Consider now the case where $0 < \lambda_L < \lambda_1^-$. That is, along the equilibrium path the seller triggers herding at some $\lambda_L = \Pr(\omega = G \mid \lambda_1, K \text{ signals } s = b)$ that is “ K steps below λ_1 ,” where $K \geq 2$. The event that the first $K - 1$ buyers all receive bad signals has positive probability. If the seller follows the equilibrium strategy and that event realizes, the first $K - 1$ buyers decline to buy the object and the updated probability of the good state is $\lambda_L^+ < \lambda_H^-$. Therefore, the seller is in the same situation as in the case where $\lambda_L = \lambda_1^-$ and $\lambda_H > \lambda_1^+$, and thus will deviate whenever $\delta > \max(\hat{\delta}_2, \hat{\delta}_3)$.

The remaining case is the one where $\lambda_L = 0$ and $\lambda_H > \lambda_1^+$. Because of (6.14) the argument of the proof of Lemma 10 that led to the inequalities (6.12) and (6.13), respectively, applies here as well. Consequently, there exists a $\hat{\delta}_4 \in (0, 1)$ that is independent of λ_1 such that the seller will deviate whenever $\delta > \hat{\delta}_4$.

Define $\bar{\delta} \equiv \max(\hat{\delta}_1, \hat{\delta}_2, \hat{\delta}_3, \hat{\delta}_4)$. We have shown that the assumption that there exists a pure strategy PBE where the seller’s equilibrium strategy is Markov and where $p^*(\lambda_1) = p^H(\lambda_1)$, leads to a contradiction whenever $\delta > \bar{\delta}$, where $\bar{\delta}$ is independent of λ_1 . Consequently, $p^*(\lambda_1) = p^L(\lambda_1)$ in any such PBE and the proposition follows. \blacksquare

A.5. Proof of Proposition 8

Given buyers’ beliefs, in each period t all aspects of the history that are relevant for the seller, are captured by λ_t and μ_t . Therefore, the seller has a best response to the buyers’ strategies and beliefs such that for each t the seller’s move depends only on λ_t and μ_t , i.e., a

best response that is Markov. Because of this, we can employ the value function, which we denote by $W(\lambda, \mu)$, in order to analyze the seller's optimization problem. That is, to any $(\lambda, \mu) \in \Lambda(\lambda_1) \times \Lambda(\lambda_1)$, $W(\lambda, \mu)$ assigns the maximum expected payoff that the seller can achieve by playing a pure Markov strategy, given buyers' strategies and beliefs. Notice that for all t , it must be that $\lambda_t \leq \mu_t$.

Proposition 8. *Assume that the seller may grant secret discounts and buyers are naïve. Let μ^* denote the critical probability of Proposition 1, 2, and 3, respectively, and $\mu^{*+} \equiv \Pr(\omega = G \mid \mu^*, s = g)$. If (i) $\alpha^2 \leq 1 - \alpha$ or (ii) $\alpha^2 > 1 - \alpha$ and $\lambda_1 \in [\bar{\lambda}_\alpha, 1)$, the seller immediately posts and actually charges the low price $p^L(\lambda_1)$ and triggers herding whenever this is uniquely optimal in the observable prices case; otherwise she posts the high price $p^H(\mu_t)$ and grants a secret discount $d_t = p^H(\mu_t) - p^L(\mu_t)$ for the first T periods $t \in \{1, \dots, T\}$, where T is a finite, deterministic integer. In the latter situation, $\mu_{T+1} = \mu^*$ if $p^L(\mu^*)$ is uniquely optimal at μ^* when prices are observable, and $\mu_{T+1} = \mu^{*+}$ if $p^H(\mu^*)$ is also optimal at μ^* when prices are observable. In period $T + 1$ the seller posts and charges $p^L(\mu^*)$ or $p^L(\mu^{*+})$, respectively, and triggers herding.*

Proof: We first consider the case where $\lambda_1 \in (0, \mu^{**}) \cup (\mu^*, 1)$, where μ^{**} is the critical probability defined in Proposition 3. In this case the low price $p^L(\lambda_1)$ is uniquely optimal in the observable prices case. Given the assumption on the buyers' beliefs, the seller will post and charge the low price $p^L(\lambda_1)$, and thus herding is triggered immediately. The same is true if $\lambda_1 = \mu^{**}$ or $\lambda_1 = \mu^*$ and the low price $p^L(\lambda_1)$ is uniquely optimal in the observable prices case.

Next we consider the case where $\lambda_1 \in (\mu^{**}, \mu^*)$. Assume first that for $\lambda_t \in \{\mu^{**}, \mu^*\}$ the low price $p^L(\lambda_t)$ is uniquely optimal in the observable prices case. Since $\lambda_1 \in (\mu^{**}, \mu^*)$, the seller posts $p^H(\lambda_1)$ as this would be the seller's uniquely optimal price in the observable prices case. Note that if $\mu_t \geq \mu^*$ for some $t \in \{2, 3, \dots\}$, then the seller will post and charge the low price, and $W(\lambda_t, \mu_t) = \frac{p^L(\mu_t)}{1-\delta}$ whenever $\mu_t \geq \mu^*$ (since $\lambda_t \leq \mu_t$). For the following inductive argument let $\mu^l = \mu^*$, and define $\mu^{l-k} \equiv \Pr(\omega = G \mid \mu^l, k \text{ signals } s = b)$ for any positive integer k . Note that $\lambda_1 < \mu^*$ implies $\mu^{l-1} \geq \lambda_1$. First we show that for any $\lambda_t \leq \mu^{l-1}$ it holds that at $(\lambda_t, \mu_t) = (\lambda_t, \mu^{l-1})$ the seller's expected payoff from posting $p^H(\mu^{l-1})$ and charging $p^L(\mu^{l-1})$ exceeds that of posting and charging $p^H(\mu^{l-1})$, i.e.,

$$p^L(\mu^{l-1}) + \delta W(\lambda_t, \mu^l) > \varphi(\lambda_t) p^H(\mu^{l-1}) + \delta [\varphi(\lambda_t) W(\lambda_t^+, \mu^l) + (1 - \varphi(\lambda_t)) W(\lambda_t^-, \mu^{l-2})].$$

The inequality holds because if $\lambda_t \leq \mu_t$ and either $\alpha^2 \leq (1 - \alpha)$ or $\alpha^2 > (1 - \alpha)$ and $\mu_t \geq \bar{\lambda}_\alpha$, the immediate return from the low price exceeds that from the high price, i.e., $p^L(\mu^{l-1}) - \varphi(\lambda_t) p^H(\mu^{l-1}) \geq p^L(\mu^{l-1}) - \alpha \mu^{l-1} > 0$, and because $W(\lambda_t^-, \mu^{l-2}) \leq W(\lambda_t, \mu^l) =$

$W(\lambda_t^+, \mu^l) = \frac{p^L(\mu^l)}{1-\delta}$. Hence at $(\lambda_t, \mu_t) = (\lambda_t, \mu^{l-1})$ it is optimal to post the high price and charge the low price, and $W(\lambda_t, \mu^{l-1}) = p^L(\mu^{l-1}) + \frac{\delta}{1-\delta}p^L(\mu^l)$, for all $\lambda_t \leq \mu^{l-1}$.

With $W(\lambda_t, \mu^{l-1})$ independent of λ_t , we can repeat the argument for any (λ_t, μ^{l-k}) and thus derive that the seller charges the low price at λ_1 . Hence the seller will post $p^H(\mu_t)$ and charge $p^L(\mu_t)$, buyer t will purchase the object, and $\mu_{t+1} = \mu_t^+ > \mu_t$. This process continues until at some finite T for the first time μ^* is reached. By assumption $p^L(\mu^*)$ is uniquely optimal in the observable prices case for $\lambda = \mu^*$, hence the seller posts and charges $p^L(\mu^*)$.

Finally, consider the case that for $\lambda_t = \mu^{**}$ or for $\lambda_t = \mu^*$ the high price $p^H(\lambda_t)$ is (also) optimal in the observable prices case. The previous analysis shows that in these cases the seller will post the high price and charge the low price at $\lambda_1 = \mu^{**}$ and $\lambda_t = \mu^*$, respectively. ■

A.6. Proof of Proposition 9

Let μ^* be the critical probability of Proposition 1. We spell out the proof of Proposition 9 under the assumption that at μ^* only the low price $p^L(\mu^*)$ is optimal when prices are observable. However, if at μ^* the high price $p^H(\mu^*)$ is (also) optimal when prices are observable, the argument is identical except that in the following μ^* must be replaced by μ^{*+} .

Let μ^l be the maximal $\mu \in \Lambda(\lambda_1)$ that is below $\bar{\lambda}_\alpha$, i.e., $\mu^l \equiv \max_{\mu \in \Lambda(\lambda_1) \cap (0, \bar{\lambda}_\alpha)} \mu$. We define $\mu^{l+k} \equiv \Pr(\omega = G \mid \mu^l, k \text{ signals } s = g)$ and $\mu^{l-k} \equiv \Pr(\omega = G \mid \mu^l, k \text{ signals } s = b)$ for any positive integer k . Note that, given δ and λ_1 , there is a finite positive integer, say $N_\delta \geq 1$, such that $\mu^{l+N_\delta} = \mu^*$. Proposition 8 implies that the seller will post $p^H(\mu^{l+k})$ but give a secret discount and actually charge only $p^L(\mu^{l+k})$ for all $\mu^{l+k} < \mu^*$, $k \in \{1, \dots, N_\delta - 1\}$. Moreover, $\alpha\mu^{l+1-h} - p^L(\mu^{l+1-h}) > 0$ and $\alpha\mu^{l+1+h} - p^L(\mu^{l+1+h}) < 0$ for all $h \in \{1, 2, \dots\}$, and $\alpha\mu^{l+1} - p^L(\mu^{l+1}) \leq 0$.

Given any $\mu_\tau = \mu^{l-k}$, $k \in \{0, 1, \dots\}$, and any $\delta \in [0, 1)$, the seller can reach μ^* in $k + N_\delta$ periods, if she posts $p^H(\mu_t)$ but actually charges only $p^L(\mu_t)$ in each period $t \in \{\tau, \tau + 1, \dots, \tau + k + N_\delta - 1\}$. In period $\tau + k + N_\delta$ the buyer updates $\mu_{\tau+k+N_\delta} = \mu^*$ from the perceived history, and the seller posts and charges $p^L(\mu^*)$ and triggers herding. In period τ the seller's discounted expected payoff (for the time that starts with period τ) from this strategy is

$$w(\mu^{l-k}) = \sum_{j=0}^{k+N_\delta-1} \delta^j p^L(\mu^{l-k+j}) + \delta^{k+N_\delta} \frac{p^L(\mu^*)}{1-\delta}. \quad (6.15)$$

Next we prove the following Lemma.

Lemma 11: To each $\pi \in (0, 1)$ there exists a $\delta_\pi \in (0, 1)$ such that for all $\lambda_1 \in (0, 1)$ the following holds: for all $k \in \{0, 1, \dots\}$,

$$\begin{aligned} & \pi \delta^k \sum_{j=1}^{N_\delta-1} \delta^j p^L(\mu^{l+j}) + \frac{\pi}{1-\delta} [\delta^{k+N_\delta} p^L(\mu^*) - \delta^k \alpha \mu^l] \\ & > \delta^k \sum_{h=0}^{\infty} \delta^{-h} [\alpha \mu^{l-h} - p^L(\mu^{l-h})] \end{aligned} \quad (6.16)$$

for all $\delta \in (\delta_\pi, 1)$, where δ_π decreases in π .

Proof: The proof is by induction. For $k = 0$ inequality (6.16) becomes

$$\pi \sum_{j=1}^{N_\delta-1} \delta^j [p^L(\mu^{l+j}) - \alpha \mu^l] + \frac{\pi}{1-\delta} \delta^{N_\delta} [p^L(\mu^*) - \alpha \mu^l] > \sum_{h=0}^{\infty} \delta^{-h} [\alpha \mu^{l-h} - p^L(\mu^{l-h})],$$

where the parameters μ^l , μ^* , and N_δ depend on λ_1 . Define $\bar{\mu}^l \equiv \Pr(\omega = G \mid \bar{\lambda}_\alpha, s = b)$ and $\bar{N}_\delta \equiv \min_{\mu^l \in [\bar{\mu}^l, \bar{\lambda}_\alpha]} N_\delta$. Notice that to each $\lambda_1 \in (0, 1)$ there corresponds a $\mu^l \in [\bar{\mu}^l, \bar{\lambda}_\alpha)$. Define $\beta \equiv \bar{\mu}^l / \bar{\lambda}_\alpha < 1$ and consider the infinite sum $\sum_{h=0}^{\infty} \delta^{-h} [\alpha \mu^{l-h} - p^L(\mu^{l-h})]$. Since

$$\frac{\mu^{l-h-1}}{\mu^{l-h}} = \frac{1-\alpha}{(1-\alpha)\mu^{l-h} + \alpha(1-\mu^{l-h})} \leq \frac{\mu^{l-1}}{\mu^l} \leq \frac{\bar{\mu}^l}{\bar{\lambda}_\alpha} = \beta$$

for all $h \in \{1, 2, \dots\}$, $\mu^{l-h} \leq \beta^h \mu^l < \beta^h \bar{\lambda}_\alpha$ for all $h \in \{1, 2, \dots\}$ for all $\lambda_1 \in (0, 1)$. Therefore, for $\delta > \beta$ we get

$$0 < \sum_{h=0}^{\infty} \delta^{-h} [\alpha \mu^{l-h} - p^L(\mu^{l-h})] \leq \alpha \sum_{h=0}^{\infty} \delta^{-h} \mu^{l-h} < \alpha \bar{\lambda}_\alpha \frac{\delta}{\delta - \beta}.$$

Because of Proposition 1, $\mu^* \rightarrow 1$ and thus $\bar{N}_\delta \rightarrow \infty$ for $\delta \rightarrow 1$. Since $p^L(\bar{\mu}^{l+j}) - \alpha \bar{\lambda}_\alpha \geq p^L(\bar{\mu}^{l+2}) - \alpha \bar{\lambda}_\alpha > 0$ for all $j \geq 2$, this implies $\sum_{j=2}^{\bar{N}_\delta-1} \delta^j [p^L(\bar{\mu}^{l+j}) - \alpha \bar{\lambda}_\alpha] \rightarrow \infty$ for $\delta \rightarrow 1$. Hence there exists a sufficiently large $\delta_\pi \in (0, 1)$ such that $\pi \sum_{j=2}^{\bar{N}_\delta-1} \delta^j [p^L(\bar{\mu}^{l+j}) - \alpha \bar{\lambda}_\alpha] > \alpha \bar{\lambda}_\alpha \delta / (\delta - \beta)$ for all $\delta \in (\delta_\pi, 1)$. Consequently, if $\delta \in (\delta_\pi, 1)$,

$$\begin{aligned} & \pi \sum_{j=1}^{N_\delta-1} \delta^j [p^L(\mu^{l+j}) - \alpha \mu^l] + \frac{\pi}{1-\delta} \delta^{N_\delta} [p^L(\mu^*) - \alpha \mu^l] > \pi \sum_{j=2}^{\bar{N}_\delta-1} \delta^j [p^L(\bar{\mu}^{l+j}) - \alpha \bar{\lambda}_\alpha] \\ & > \alpha \bar{\lambda}_\alpha \frac{\delta}{\delta - \beta} > \sum_{h=0}^{\infty} \delta^{-h} [\alpha \mu^{l-h} - p^L(\mu^{l-h})], \end{aligned}$$

holds for all $\lambda_1 \in (0, 1)$. This proves that for all $\lambda_1 \in (0, 1)$ the inequality (6.16) is satisfied for $k = 0$ provided $\delta \in (\delta_\pi, 1)$ for some sufficiently large $\delta_\pi \in (0, 1)$. Moreover, δ_π decreases in π .

Assume that inequality (6.16) holds for some k , given some fixed $\delta \in (\delta_\pi, 1)$. Given this δ , we get for $k + 1$:

$$\begin{aligned}
& \pi \delta^{k+1} \sum_{j=1}^{N_\delta-1} \delta^j p^L(\mu^{l+j}) + \frac{\pi}{1-\delta} \left[\delta^{(k+1)+N_\delta} p^L(\mu^*) - \delta^{k+1} \alpha \mu^l \right] \\
= & \delta \left\{ \pi \delta^k \sum_{j=1}^{N_\delta-1} \delta^j p^L(\mu^{l+j}) + \frac{\pi}{1-\delta} \left[\delta^{k+N_\delta} p^L(\mu^*) - \delta^k \alpha \mu^l \right] \right\} \\
> & \delta^{k+1} \sum_{h=0}^{\infty} \delta^{-h} \left[\alpha \mu^{l-h} - p^L(\mu^{l-h}) \right],
\end{aligned}$$

i.e., inequality (6.16) holds for $k + 1$ as well and the lemma follows. \blacksquare

Proposition 9. *Assume that the seller may grant secret discounts and buyers are naïve. Let μ^* denote the critical probability of Proposition 1 and $\mu^{*+} \equiv \Pr(\omega = G \mid \mu^*, s = g)$. If $\alpha^2 > 1 - \alpha$ and $\lambda_1 \in (0, \bar{\lambda}_\alpha)$, there exists a discount factor $\bar{\delta} \in (0, 1)$ such that for all $\delta \in (\bar{\delta}, 1)$ the seller posts (and charges) the price $p^L(\mu^*)$ in finite time with probability 1 if $p^L(\mu^*)$ is uniquely optimal at μ^* when prices are observable, and posts (and charges) $p^L(\mu^{*+})$ in finite time with probability 1 if $p^H(\mu^*)$ is also optimal at μ^* when prices are observable. Thus, herding arises in finite time with probability 1, provided the seller is sufficiently patient.*

Proof: Consider an equilibrium where, given $\lambda_1 \in (0, 1)$ and $\delta \in [0, 1)$, the probability that herding occurs in finite time is less than 1. In this case there must be a set of realizations of signals (s_1, s_2, \dots) that has positive probability, where herding will not occur and where the seller's updated probability that herding will not occur converges to 1. Therefore, given any $\pi \in (0, 1)$ there must be a node (λ_T, μ_T) that is reached with positive probability such that the seller's updated probability that herding will not occur is larger than some given $\pi \in (0, 1)$. Given μ_T there is a nonnegative k such that $\mu_T = \mu^{l-k}$. At such a node $(\lambda_T, \mu_T) = (\lambda_T, \mu^{l-k})$ the seller's expected payoff $W(\lambda_T, \mu^{l-k})$ is bounded by

$$\begin{aligned}
W(\lambda_T, \mu^{l-k}) & < \sum_{j=0}^k \delta^j \alpha \mu^{l-k+j} + (1-\pi) \sum_{j=k+1}^{k+N_\delta-1} \delta^j p^L(\mu^{l-k+j}) + \\
& (1-\pi) \delta^{k+N_\delta} \frac{p^L(\mu^*)}{1-\delta} + \pi \delta^k \frac{\alpha \mu^l}{1-\delta}.
\end{aligned}$$

Because at $\mu_T = \mu^{l-k}$ the seller can always guarantee herself the payoff $w(\mu^{l-k})$ specified by (6.15), it must hold that $w(\mu^{l-k}) \leq W(\lambda_T, \mu^{l-k})$ and thus

$$\begin{aligned}
& \sum_{j=0}^{k+N_\delta-1} \delta^j p^L(\mu^{l-k+j}) + \delta^{k+N_\delta} \frac{p^L(\mu^*)}{1-\delta} \\
< & \sum_{j=0}^k \delta^j \alpha \mu^{l-k+j} + (1-\pi) \sum_{j=k+1}^{k+N_\delta-1} \delta^j p^L(\mu^{l-k+j}) + \\
& (1-\pi) \delta^{k+N_\delta} \frac{p^L(\mu^*)}{1-\delta} + \pi \delta^k \frac{\alpha \mu^l}{1-\delta}.
\end{aligned} \tag{6.17}$$

However, rearranging (6.16) and using the fact that $\sum_{j=0}^k \delta^j \alpha \mu^{l-k+j} = \delta^k \sum_{h=0}^k \delta^{-h} \alpha \mu^{l-h}$ implies

$$\sum_{j=0}^k \delta^j \alpha \mu^{l-k+j} + \delta^k \sum_{h=k+1}^{\infty} \delta^{-h} \alpha \mu^{l-h} = \delta^k \sum_{h=0}^{\infty} \delta^{-h} \alpha \mu^{l-h};$$

the fact that

$$\sum_{j=0}^{k+N_\delta-1} \delta^j p^L(\mu^{l-k+j}) = \delta^k \sum_{h=0}^k \delta^{-h} p^L(\mu^{l-h}) + \sum_{j=k+1}^{k+N_\delta-1} \delta^j p^L(\mu^{l-k+j});$$

and the fact that $\sum_{j=k+1}^{k+N_\delta-1} \delta^j p^L(\mu^{l-k+j}) = \delta^k \sum_{j=1}^{N_\delta-1} \delta^j p^L(\mu^{l+j})$, shows that

$$\begin{aligned}
& \sum_{j=0}^{k+N_\delta-1} \delta^j p^L(\mu^{l-k+j}) + \delta^{k+N_\delta} \frac{p^L(\mu^*)}{1-\delta} \\
> & \sum_{j=0}^k \delta^j \alpha \mu^{l-k+j} + (1-\pi) \sum_{j=k+1}^{k+N_\delta-1} \delta^j p^L(\mu^{l-k+j}) + \\
& (1-\pi) \delta^{k+N_\delta} \frac{p^L(\mu^*)}{1-\delta} + \pi \delta^k \frac{\alpha \mu^l}{1-\delta} + \delta^k \sum_{h=k+1}^{\infty} \delta^{-h} [\alpha \mu^{l-h} - p^L(\mu^{l-h})]
\end{aligned}$$

for all $\delta \in (\delta_\pi, 1)$ because of Lemma 11. Since the last term on the right hand side is positive, (6.17) cannot hold for $\delta \in (\delta_\pi, 1)$. Because δ_π decreases in π , the proposition follows with $\bar{\delta} \equiv \inf_{\pi \in (0,1)} \delta_\pi$. ■

References

- C. Avery and P. Zemsky, “Multidimensional Uncertainty and Herd Behavior in Financial Markets,” *American Economic Review* 88 (1998), 724-748.
- A. Banerjee, “A Simple Model of Herd Behavior,” *Quarterly Journal Economics* 107 (1992), 797-818.
- D. Bergemann and J. Välimäki, “Experimentation in Markets,” *Review of Economic Studies* 67 (2000), 231-234.
- S. Bikhchandani, D. Hirshleifer, and I. Welch, “A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades,” *Journal of Political Economy* 100 (1992), 992–1026.
- S. Bikhchandani, D. Hirshleifer and I. Welch, “Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades,” *Journal of Economic Perspectives* 12 (1998) 151-170.
- A. Devenow and I. Welch, “Rational Herding in Financial Economics,” *European Economic Review* 40 (1996), 603-615.
- D. Gale, “What Have We Learned from Social Learning?” *European Economic Review* 40 (1996), 617-28.
- A. Mas-Colell, M. Whinston, and J. Green, *Microeconomic Theory*, Oxford University Press, (1995).
- Z. Neeman and G. O. Orosel, “Herding and the Winner’s Curse in Markets with Sequential Bids,” *Journal of Economic Theory* 85 (1999), 91-121.
- M. Ottaviani, “Monopoly Pricing with Social Learning,” unpublished manuscript, 1999.
- L. Smith and P. Sørensen, “Pathological Outcomes of Observational Herding,” *Econometrica* 68 (2000), 371-398.
- A. Takayama, “Mathematical Economics,” Dryden Press, Hinsdale, Ill., 1974.
- I. Welch, “Sequential Sales, Learning, and Cascades,” *Journal of Finance* 47 (1992), 695–732.
- X. Vives, “Social Learning and Rational Expectations,” *European Economic Review* 40 (1996), 589-601.