

Product Launches with Biased Reviewers: the Importance of Not Being Earnest

David Gill and Daniel Sgroi

July 2003

CWPE 0334

Not to be quoted without permission

Abstract

The standard simple sequential herding model is altered to allow a firm with a new product to have it reviewed publicly before launch. Reviewers are either inherently pessimistic, optimistic or unbiased. We find the counter-intuitive result that a firm with a good product will prefer a pessimistic reviewer. Although firms with a bad product prefer unbiased reviewers, signalling considerations will force them to copy the choice of the good product firm in order to avoid revealing product type. This asymmetric impact provides a strong explanation for the stylized fact that reviewers are often viewed as being very critical.

Keywords: bias, herding, reviewers, sequential sales

JEL Classification: D82, D83, L15

Product Launches with Biased Reviewers: The Importance of Not Being Earnest

David Gill

*Nuffield College
University of Oxford*

Daniel Sgroi ¹

*Churchill College & Dept. of Applied Economics
University of Cambridge*

1. INTRODUCTION

Reviewers can be incredibly powerful. For example, in the wine trade the American reviewer Robert Parker can make or break a new vintage. According to the Oxford Companion to Wine: “His judgements have had a significant effect on market demand and the commercial future of some producers” (Robinson, 1999, 511-512). Despite the powerful effects that reviewers can exert over new products launched onto the market, the literature has paid little or no attention to this important phenomenon. This paper attempts to redress the omission.

We develop a sequential model of sales in which each consumer has some private signal about the value of a product with unknown quality put on the market by a monopolist. The firm can attempt to tilt sales in its favour through the choice of reviewer type to provide more information before sales begin. We consider three types of reviewer: an unbiased reviewer; a pessimist; and an optimist. The pessimist is a known critic of the firm or type of product and has a higher probability of not endorsing, while the optimist is more likely to endorse than the unbiased reviewer. We find that the pessimist, despite having a much lower probability of endorsement, will be preferred by a firm with a good product. An endorsement by such a pessimist provides an excellent signal of the product’s quality, while consumers expect the reviewer to fail to endorse, so receiving no endorsement will not impact too heavily on the firm’s expected profits. On the other hand, a firm with a bad product would prefer to choose an unbiased reviewer, because the risk that the pessimist fails to endorse is too great for such a firm. However, signalling considerations make it likely that the bad product firm will

¹Daniel Sgroi would like to thank AEA Technology for financial support and can be contacted at daniel.sgroi@econ.cam.ac.uk. David Gill would like to thank the Economic and Social Research Council for financial support and can be contacted at david.gill@nuf.ox.ac.uk. Both authors would like to thank Mark Armstrong, Simon Board, Douglas Gale, Michael Grubb, Paul Klemperer, Meg Meyer, David Myatt, Rebecca Stone and participants at the Gorman Workshop in Oxford for helpful comments and suggestions.

be want to copy the choice of the good product firm to avoid immediately revealing product type through the choice of reviewer. This general preference for pessimistic reviewers can explain why reviewers are keen to build reputations for being overly critical, since doing so makes their endorsements much more valuable to firms.

1.1. Biased Reviewers and Toughness of Tests. Although for concreteness we specifically model biased reviewers in the context of a firm launching a new product, the idea of a pessimistic reviewer acting as a tough test is more general. The basic intuition that a tough test might be useful because the passing of such a test sends out a very strong signal, while failing such a test is not too harmful, could apply to a number of other settings with a similar sequential learning structure. For example, a doctoral student facing a sequence of job market interviews might choose a referee well known to be particularly tough, a firm seeking to launch an initial public offering (IPO) might choose a tough auditor to review its accounts, firms wishing to merge might initially file in a country with a tough competition authority, a firm might seek to have its product approved by a tough kitemarking body, and so on.

1.2. Relation to the Literature. Banerjee (1992) and Bikhchandani, Hirshleifer and Welch (1992) examined a sequential herding model appropriate for this setting, which we develop. We abstract from price-setting to focus on the choice of reviewer, but a number of papers analyze the use of initial prices to manipulate sales in a herding environment. In the fixed-price IPO market, Welch (1992) finds that the issuer will prefer to set a low price to guarantee a favourable herd.² However, where prices are flexible, Taylor (1999) and Ottaviani (1999) find that high initial prices are optimal. In Ottaviani, the firm wishes to set a high initial price (relative to perceived quality) to encourage the transmission of information. If price is too low, everybody buys, so consumers do not learn from each other's decisions, while if an expensive good becomes successful, this conveys strong positive information to later buyers. Taylor, concentrating on the housing market, considers the sale of only a single item. Taylor allows for the possibility of a herd against the sale of a house, not in favour of purchase, arguing that sellers should therefore set a high early price. If the house is not sold quickly, late consumers can then attribute the failure to sell to the product being overpriced rather than being of low quality. The high initial price in these two papers plays a qualitatively similar role to critical pessimistic reviewers here: a failure to pass the tough test imposed by either a high initial price or a pessimistic reviewer is not perceived as overly

²Various other papers have applied sequential herding to asset and financial markets where traders can make inferences based on prices, for example Lee (1998) and Avery and Zemsky (1998).

damaging since the test is so unlikely to be passed. There is also a literature on marketing and advertising which focuses on high early prices to signal quality, see for example Bagwell and Riordan (1991).

In contrast, Caminal and Vives (1996, 1999) and Vettas (1997) find that a low initial price is optimal in order to initiate herd-like behaviour from customers. In Caminal and Vives, past prices are not observable, so consumers infer information about quality differentials from previous market shares. Firms are then tempted to lower prices to manipulate future beliefs about quality arising from this period's market shares. Vettas finds low introductory pricing to be optimal (for a high quality firm) where consumers do not receive private signals but learn from previous purchasing decisions via word-of-mouth communication – the low price, resulting in more initial sales, facilitates this form of information transmission.

None of these papers explicitly considers the role of publicly observable reviews, possibly by known biased reviewers and critics, in forming the prior beliefs of consumers, and all of them assume that price differentiation is a feasible and low cost strategy. In many cases, since reviewers exist at no direct cost to firms and are used freely by consumers (in for example the movie and electronics businesses), it seems important to examine the impact of reviewers and ask whether they might fill a similar role to (or even replace the need for) inter-temporal price discrimination in generating high sales.

Gill and SgROI (2003) do consider reviewers defined in a similar way, but where sales are determined simultaneously after a decision made by a reviewer. Observational learning by consumers from other consumers is therefore not an issue. Though the resulting model is somewhat simpler, they also find that pessimists generally promote higher sales. Since there is no learning from other consumers' actions, reviewers are especially important, and the chance that a pessimist may purchase, generating 100% sales, makes a pessimist a much better choice for the firm, irrespective of product quality. SgROI (2002) examines the use of small groups of consumers who are encouraged to decide early and hence act in a similar way to reviewers, providing additional information for later consumers, who do act in sequence. He finds that irrespective of product quality, firms would like to use these "guinea pigs".

The use of biased methods of evaluation has been considered by a number of authors. Calvert (1985) looks at policy-makers choosing between biased advisors. Sah and Stiglitz (1986) consider the choice between a bureaucratic "hierarchical" structure and a decentralized "polyarchical" one for accepting or rejecting potential projects, where the former imposes a tougher test. Fishman and Hagerty (1990) analyze the level of discretion that an entrepreneur should be allowed to use in reporting information to potential investors, where

greater discretion effectively imposes an easier test. Finally, Meyer (1991) looks at the use of biased contests in deciding which employees to promote.

All these papers are concerned with a decision-maker choosing between a number of alternatives, and all come to a similar conclusion: given an initial predisposition towards one of the alternatives, the decision-maker is best off choosing an evaluation procedure which is biased in favour of this predisposition. For example, Meyer finds that firms who are trying to decide which employee to promote should bias contests in favour of the early leader. Just like the firm in our model, the decision-makers in these papers are able to use bias to alter the information partitions to their advantage. Procedures biased in favour of the predisposition are of value because a recommendation that goes against the predisposition is then strong evidence that the predisposition was wrong, while an evaluator who is neutral or biased against the predisposition is unlikely to change the decision-maker's mind if he advises against the predisposition. In contrast, we find that firms with good products (who know product quality with certainty and whose aim is to maximize sales) should choose early evaluators that are biased against their product.³

Avery and Meyer (1999) also look at evaluators who suffer from varying degrees of bias, but where the levels of bias are only imperfectly known to a decision-maker. The authors look at the impact of the decision-maker tracking evaluations over time on the toughness of different evaluators' reports, and hence on the decision-maker's payoff.

1.3. Overview of the Paper. The next section develops the model of sequential sales and notation used throughout the paper. Section 3 considers a firm offering a high quality product and section 4 a low quality product. Section 5 examines the implications for signalling quality not via the decisions of reviewers but by the choice of reviewer. Section 6 considers more general priors and section 7 offers some conclusions.

2. A SIMPLE MODEL OF SEQUENTIAL HERDING

In sequential herd models the main features are the coarseness of actions spaces relative to signal spaces, the exogenously determined sequential ordering of decisions by consumers, and the externality generated by this information structure. Once a herd begins public information will swamp private information and all later decision-makers will copy the actions of their predecessors, but no new information will be available to later decision-makers so

³In the final section of Meyer (1991), an example is presented in which the firm has an option to promote nobody and chooses to bias the contest *against* the early leader because it wants to be sufficiently confident before promoting anybody. This has more of the flavour of our result.

there will be no scope for later agents to “break” the herd. This section sets up a version of the model first used in the seminal herd paper by Bikhchandani, Hirshleifer and Welch (1992).

Consider a sequence of $N \in \mathbb{N}_{++}$ agents, the ordering of which is exogenous and common knowledge, each deciding whether to purchase (Y) or not purchase (N) some product put on the market by a monopolist. Each agent observes the actions (Y or N) of his predecessors. The cost of purchase is $C = \frac{1}{2}$, and results in the gain of V which has prior probability $\frac{1}{2}$ of returning 1 or 0, depending on whether the product is of a high or low quality. The agents each receive a conditionally independent signal about V defined as $X_i \in \{H, L\}$ for agent i . The signals are informative in the following sense.

Definition 1. *Signals are informative, but not fully-revealing, in the sense that:*

$$\Pr[X_i = H \mid V = 1] = \Pr[X_i = L \mid V = 0] = p \in (0.5, 1)$$

$$\Pr[X_i = H \mid V = 0] = \Pr[X_i = L \mid V = 1] = 1 - p \in (0, 0.5)$$

Define the history to agent n as the set of actions of agents 1 to $n - 1$ so $H_{n-1} \equiv \{A_1, A_2, \dots, A_{n-1}\}$ where $A_i \in \{Y, N\}$. Now define the information set of agent i as $I_i \equiv \{H_{i-1}, X_i\}$. It will be the case that in certain circumstances X_i will be inferable from A_i but this will not always be true. Now $X_1 = H \Leftrightarrow A_1 = Y$ and $X_1 = L \Leftrightarrow A_1 = N$. Agent 2 can infer agent 1’s signal, X_1 , from his action, A_1 , and so has an information set $I_2 = \{X_1, X_2\}$. If $X_2 = H$ and $A_1 = Y \Rightarrow X_1 = H$ then agent 2 adopts so $A_2 = Y$. If $X_2 = H$ and $A_1 = N \Rightarrow X_1 = L$ or if $X_2 = L$ and $A_1 = Y \Rightarrow X_1 = H$ agent 2 will have two conflicting signals and will be indifferent, so we require a tie-breaking rule. We use a simple coin-flipping rule which is known to all agents:⁴

Condition 1. *(Tie-breaking) If I_i includes an equal weighting of H and L signals then $\Pr[A_i = Y] = \Pr[A_i = N] = \frac{1}{2}$. This rule is common knowledge.*

Consider a possible chain of events. The *first agent* will purchase if $X_1 = H$ and reject if $X_1 = L$. The *second agent* can infer the signal of the first agent from his action. He will then purchase if $X_2 = H$ having observed purchase by the first agent. If he observed rejection but received the signal $X_2 = H$ then he will flip a coin following the tie-breaking rule. If he receives $X_2 = L$ and $A_1 = N$ then he too will choose $A_2 = N$. If the first agent

⁴Coin flipping is the standard tie-break rule used in herding models. See for example Bikhchandani, Hirshleifer and Welch (1992). Banerjee (1992) instead uses a “follow your own signal” rule, but does this specifically to minimize the chance of a herd.

purchased then he would be indifferent and so flip a coin. The *third agent* is the first to face the possibility of a herd. If he observed two purchases, so $H_2 = \{Y, Y\}$ then $A_3 = Y$ for all X_3 since he knows that $X_1 = H$ and the second agent's signal is also more likely to be H than L , so the weight of evidence is in favour of purchase *regardless of* X_3 . This initiates a *Y cascade*: the third agent will purchase, the *fourth agent* will also purchase, as will the fifth, etc. Similarly if the third agent observes that both previous choices were rejections then he too will reject, so a *N cascade* is initiated. An informational cascade occurs if an individual's action does not depend upon his private information signal. The individual, having observed the actions of those ahead of him in a sequence, who follows the behaviour of the preceding individual, without regard to his own information, is said to be in a cascade.

Definition 2. Informational Cascades. A *Y cascade* is said to have started by agent n if $A_n = Y$ and $A_{i-1} = Y \Rightarrow A_i = Y$ for all X_i thereafter. A *N cascade* is said to have started if $A_n = N$ and $A_{i-1} = N \Rightarrow A_i = N$ for all X_i thereafter.

We should note that a cascade, once started, will last forever as no further information is revealed by agents' actions.⁵ This is so even if it is based on an action which would not be chosen if all agents' signals were common knowledge. Finally, the possibility of convergence to the incorrect outcome through the loss of information contained in later agents' private signals might be phrased in terms of a discernible negative *herd externality* as suggested by Banerjee (1992).

From the model specifications the (conditional) *ex ante* probabilities of a *Y* cascade, *N* cascade, or no cascade after n agents can be derived. Define $Y(n)$ to be a *Y* cascade which has started by agent n . Similarly define $N(n)$ for a *N* cascade and $No(n)$ for no cascade by agent n . For example, $\Pr[Y(2)]$ is simply the probability that the first two agents both choose *Y*. After an even number of n agents the relevant functions are:

$$(2.1) \quad \Pr[Y(n) \mid V = 1] = \frac{p(p+1)}{2} \frac{1 - (p-p^2)^{\frac{n}{2}}}{1 - (p-p^2)}$$

$$(2.2) \quad \Pr[N(n) \mid V = 1] = \frac{(p-2)(p-1)}{2} \frac{1 - (p-p^2)^{\frac{n}{2}}}{1 - (p-p^2)}$$

These expressions, derived in full in Bikhchandani, Hirshleifer and Welch (1992) using geometric progressions, allow us to make a number of clarifying remarks. It can be shown that

⁵If p were different for each individual, for some individuals p_i might be sufficiently high that they gain nothing from observation. This would therefore confound the learning process and might allow herds to be broken. See Smith and Sorensen (2000) for further details. By restricting $p_i = p \forall i$, we do not consider such confounded learning.

(2.1) is increasing in p and n , but (2.2) registers a high probability even for p much greater than $\frac{1}{2}$. For example, for $p = \frac{3}{4}$, there is a 20% chance of an incorrect herd having started by the 10th agent. Therefore, even when the product is good (which is likely to be reflected by the great majority of the signals), a product still faces the prospect of a possible herd against its purchase, because there is a reasonable chance that a few early incorrect signals start an incorrect herd. This is worrying both for consumers and for a firm with a high quality product. The symmetric case where $V = 0$ would apply when the product is of low quality, and the results provide some hope for the manufacturer of such a product, since there is always the chance of a Y cascade. Of course there is no reason for a firm to stay passive in the face of such potential herds.

2.1. Modelling Reviewers. We want to think of reviewers as having access to finer information because the firm allows the reviewer to test out the product. The simplest way of modelling this is to allow the reviewer to get two signals about the product's quality, instead of just the one that the customers get.

We model the reviewer as choosing whether or not to endorse the firm's product. In reality of course, the reviewer may be able to make a finer distinction than simply endorsing the product or not. However, we want to think of the review as being quickly and easily disseminated throughout the population of potential consumers, e.g., through word of mouth, written reports concerning the review and so on. Thus we are thinking of a process through which even sophisticated reviews quickly get shortened to a binary distinction through this process of dissemination.⁶

Reviewers are either inherently optimistic, pessimistic or neutral about the product prior to receiving any signals. We define the types of reviewers as follows.

Definition 3. *Optimistic reviewers endorse the product iff they get two H signals or a H and a L signal. Pessimistic reviewers endorse the product iff they get two H signals. Unbiased reviewers endorse the product if they get two H signals, don't endorse if they get two L signals and flip a coin if they get a H and a L signal.*

We do not consider explicitly how reviewers come into existence or take on biases. However, the preference of different firms for different types of reviewers may justify the existence of biases. Reviewers may differentiate themselves to appeal to certain firms by building a

⁶Modelling an evaluator as condensing more complex information into a simple binary decision follows for example Calvert (1985) and Sah and Stiglitz (1986). As Calvert puts it: "This feature represents the basic nature of advice, a distillation of complex reality into a simple recommendation."

reputation for being of a certain type. There is a literature on payment structures to certification intermediaries, who play a similar role to reviewers - see for example Lizzeri (1999) and Albano and Lizzeri (2001). In these papers, the intermediary has all the bargaining power as it sets the terms of trade via a price and disclosure rule; instead we are thinking of reviewers as responding to the preferences of the firms regarding their level of bias.⁷

We also assume that reviewer types are common knowledge, perhaps generated through a known history of tough or soft reviews.

2.2. Introducing a Stochastic Number of Customers. As was noted above, Bikhchandani, Hirshleifer and Welch (1992) use geometric progressions to solve for herd probabilities for a fixed finite number of potential customers. Once we introduce potential bias through the choice of reviewer, using this method to calculate and compare profits under different symmetric and asymmetric scenarios quickly becomes excessively unwieldy and complicated. Instead we introduce a stream of customers of uncertain length, which allows us to use a recursive solution method to readily solve for and compare expected profits, even where an asymmetry is introduced by the choice of a biased reviewer.

We assume that the size of the pool of potential customers is not known with certainty to the firm. Instead, after each customer decides whether or not to purchase, there is a probability $(1 - \theta)$ that the pool of potential customers comes to an end, where $\theta \in (0, 1)$.⁸ Thus, the expected number of potential customers $E[N] = 1 + \theta + \theta^2 + \dots = \frac{1}{1-\theta}$, and $\theta = \frac{E[N]-1}{E[N]}$. Note as θ goes up, so does $E[N]$, so $\frac{dE[N]}{d\theta} > 0$. One interpretation of $E[N]$ is that it is a measure of how quickly the firm expects a rival technology or product to be developed which will make its own product obsolete. Note that although θ can range from 0 to 1, for reasonable $E[N]$ values, it will be in the upper part of this range. For example, for $E[N] \geq 10$, $\theta \geq 0.9$.

3. GOOD PRODUCT FIRM

In this section we explicitly focus on the case when $V = 1$ so the firm is selling a high quality product.

⁷As we will show, expected sales of the good product are higher and of the bad product lower following a pessimist's review, which suggests that customers may share a preference with the good product firm for pessimistic reviewers.

⁸The analysis is formally equivalent to introducing a standard discount factor where the firm faces an infinite stream of potential customers.

3.1. Deriving Profits from Different Reviewer Types. We start by developing two Remarks, which are used implicitly throughout the proofs of the propositions that follow. Let q be the customers' prior (either their initial 50:50 prior, or an updated prior based on a commonly observed history).

Remark 1.

$$\frac{\Pr[V = 1|I_i]}{\Pr[V = 0|I_i]} = \frac{\frac{\Pr[I_i|V=1]\Pr[V=1]}{\Pr[I_i]}}{\frac{\Pr[I_i|V=0]\Pr[V=0]}{\Pr[I_i]}} = \frac{\Pr[I_i|V = 1] q}{\Pr[I_i|V = 0] (1 - q)}$$

This Remark says that agents, when applying Bayes' Rule to calculate their beliefs about whether the product is more likely to be of good or bad quality, simply need to calculate the ratio of the probability of the information set they have observed if the product were good to the probability if the product were bad, suitably weighted by the prior. Furthermore:

Remark 2. *When calculating beliefs, agents can cancel and ignore opposing H and L signals.*

Proof. Suppose the agent infers an information set I_i . Now suppose instead of I_i , the agent infers I_i^+ , which we define as the set I_i plus a further two opposing H and L signals. Then, using Remark 1:

$$\begin{aligned} \frac{\Pr[V = 1|I_i^+]}{\Pr[V = 0|I_i^+]} &= \frac{\Pr[I_i^+|V = 1] q}{\Pr[I_i^+|V = 0] (1 - q)} = \\ \frac{\Pr[I_i|V = 1] p(1 - p)q}{\Pr[I_i|V = 0] (1 - p)p(1 - q)} &= \frac{\Pr[I_i|V = 1] q}{\Pr[I_i|V = 0] (1 - q)} = \frac{\Pr[V = 1|I_i]}{\Pr[V = 0|I_i]} \end{aligned}$$

■

Normalizing the value of each sale to 1, we can now calculate expected profits to the firm for various potential updated priors of the customers in the sequence, q , following the reviewer's decision. This will then allow us to easily calculate profits from the different reviewer types.

Where the prior is strongly positive or negative, it outweighs any possible signal the first customer might have, so herds start immediately.

Lemma 1. *Where $q > p$, herds for the product will start immediately, so expected profits are $\frac{1}{1-\theta}$.*

Proof. Following a bad signal, using Remark 1, the first customer purchases as:

$$\frac{\Pr[V = 1|I_i]}{\Pr[V = 0|I_i]} = \frac{(1 - p)q}{p(1 - q)} = \frac{q - pq}{p - pq} > 1$$

Thus, the first customer purchases whatever the signal received. The decision is thus uninformative, so the second also purchases, and so on. ■

A symmetrical argument shows:

Lemma 2. *Where $q < 1 - p$, herds against the product will start immediately, so profits are zero.*

Next we find expected profits where $q = \frac{1}{2}$.

Lemma 3. *Where $q = \frac{1}{2}$, expected profits to the good product firm are given by*

$$\Pi_{q=\frac{1}{2}}^G = \frac{p [2 - (1 - p)\theta^2]}{2 [1 - p(1 - p)\theta^2] (1 - \theta)}$$

Proof. See Appendix. The proof is based on a recursive solution to the appropriate decision tree. ■

Now suppose $q = p$. If the first customer gets a positive signal, he buys. If he gets a negative signal this exactly cancels the positive prior, as using Remark 1, $\frac{\Pr(V=1|I_i)}{\Pr(V=0|I_i)} = \frac{1-p}{p} \cdot \frac{q}{1-q} = 1$, so the customer is indifferent and flips a coin. We can think of the second and subsequent customers as starting a new sequence with updated prior q_2 . If the first customer rejects, later customers infer $X_1 = L$, so $q_2 = \frac{1}{2}$. If the first purchases, then he was more likely to observe H than L , sending a positive signal, thus increasing q_2 above q , so $q_2 > p$.⁹ Thus, using Lemma 1:

Lemma 4. *Where $q = p$, expected profits to the good product firm are given by*

$$\Pi_{q=p}^G = \left[p + \frac{1}{2}(1 - p) \right] \left(\frac{1}{1 - \theta} \right) + \frac{1}{2}(1 - p)\theta\Pi_{q=\frac{1}{2}}^G$$

Suppose instead $q = 1 - p$. This case is the symmetric opposite. If the first customer gets a negative signal, he rejects, while if he gets a positive signal, this exactly cancels the negative prior so he flips a coin. Thus, if the first customer buys, later customers infer $X_1 = H$, so $q_2 = \frac{1}{2}$. If the first rejects, then he was more likely to observe L than H , sending a negative signal, so $q_2 < 1 - p$. Thus, using Lemma 2:

Lemma 5. *Where $q = 1 - p$, expected profits to the good product firm are given by*

$$\Pi_{q=1-p}^G = \frac{1}{2}p \left(1 + \theta\Pi_{q=\frac{1}{2}}^G \right)$$

⁹Formally, $\frac{\Pr[V=1|I_i]}{\Pr[V=0|I_i]} = \frac{p+\frac{1}{2}(1-p)}{\frac{1}{2}p+(1-p)} \frac{q}{1-q} > \frac{q}{1-q}$.

Next, we calculate profits where $q \in (\frac{1}{2}, p)$ or $q \in (1 - p, \frac{1}{2})$. As a first step, letting \hat{q} represent a customer's posterior belief following some updating:

Remark 3. Where $q \in (\frac{1}{2}, p)$, following a H signal $\hat{q} > p$, while a L signal outweighs the positive prior and $\hat{q} \in (1 - p, \frac{1}{2})$.

Proof. See Appendix. ■

The following Lemmas can then be derived:

Lemma 6. Where $q \in (\frac{1}{2}, p)$, expected profits to the good product firm are given by

$$\Pi_{q \in (\frac{1}{2}, p)}^G = \frac{p [1 + (1 - p)\theta (1 - \theta)]}{[1 - p(1 - p)\theta^2] (1 - \theta)}$$

Proof. See Appendix. Again, a recursive decision tree is used. ■

Lemma 7. Where $q \in (1 - p, \frac{1}{2})$, expected profits to the good product firm are given by

$$\Pi_{q \in (1 - p, \frac{1}{2})}^G = \frac{p [p + (1 - p) (1 - \theta)]}{[1 - p(1 - p)\theta^2] (1 - \theta)}$$

Proof. See Appendix. ■

We now have the building blocks which allow us to easily calculate and compare profits to the firm from choosing different reviewer types.

Suppose first the reviewer is unbiased. The probability that the reviewer endorses if the product is good is $p^2 + p(1 - p) = p$, while the probability he fails to endorse is $p(1 - p) + (1 - p)^2 = (1 - p)$. Given the initial 50:50 prior and using Bayes' Rule, then following an endorsement the updated prior is:

$$q = \frac{[p^2 + p(1 - p)] \frac{1}{2}}{[p^2 + p(1 - p)] \frac{1}{2} + [p(1 - p) + (1 - p)^2] \frac{1}{2}} = p$$

Following a rejection, we have:

$$q = \frac{[p(1 - p) + (1 - p)^2] \frac{1}{2}}{[p(1 - p) + (1 - p)^2] \frac{1}{2} + [p^2 + p(1 - p)] \frac{1}{2}} = 1 - p$$

So with probability p the reviewer endorses, so $q = p$ and Lemma 4 applies, while with probability $(1 - p)$ the reviewer fails to endorse so $q = (1 - p)$ and Lemma 5 applies. Thus, the expected profits from choosing an unbiased reviewer in the good product case are

$$\Pi_{ur}^G = p\Pi_{q=p}^G + (1-p)\Pi_{q=1-p}^G$$

Which after some algebra gives:

Proposition 1. *Expected profits to a firm with a good product from choosing an unbiased reviewer are given by*

$$\Pi_{ur}^G = \frac{p[2 + 2p(1-p)\theta(1-\theta) - (1-p)\theta]}{2[1 - p(1-p)\theta^2](1-\theta)}$$

Proof. See Appendix. ■

Next, let's consider the pessimistic reviewer. The probability that the reviewer endorses if the product is good is p^2 , while the probability he fails to endorse is $1 - p^2$. Clearly, if the pessimist endorses he sends out a stronger signal for the product than does the unbiased reviewer, so $q > p$,¹⁰ and hence by Lemma 1 a herd for the product starts immediately. If the pessimist fails to endorse, then this is a negative signal, but weaker than that sent out by the unbiased reviewer who does not endorse, so $q \in (1 - p, \frac{1}{2})$.¹¹ Thus, Lemma 7 applies. Expected profits from choosing a pessimistic reviewer in the good product case are then:

$$\Pi_{pr}^G = p^2 \left(\frac{1}{1-\theta} \right) + (1-p^2) \Pi_{q \in (1-p, \frac{1}{2})}^G$$

After some algebra, this gives:

Proposition 2. *Expected profits to a firm with a good product from choosing a pessimistic reviewer are given by*

$$\Pi_{pr}^G = \frac{p[1 + p(1-p) + p^2(1-p)\theta(1-\theta) - (1-p)\theta]}{[1 - p(1-p)\theta^2](1-\theta)}$$

Proof. See Appendix. ■

Finally, we consider the optimistic reviewer. The probability that the reviewer endorses if the product is good is $p^2 + 2p(1-p)$, while the probability he fails to endorse is $(1-p)^2$. Clearly, this case is the symmetric opposite of the pessimistic reviewer case. If the optimist

¹⁰Formally, $q = \frac{p^2 \frac{1}{2}}{p^2 \frac{1}{2} + (1-p)^2 \frac{1}{2}} > p$ iff $p > p^2 + (1-p)^2$ or $p(1-p) > (1-p)^2$, which is clearly true given $p \in (\frac{1}{2}, 1)$.

¹¹Formally, $q = \frac{(1-p^2) \frac{1}{2}}{(1-p^2) \frac{1}{2} + [p^2 + 2p(1-p)] \frac{1}{2}} > 1-p$ iff $(1+p)(1-p) > [1 + 2p(1-p)](1-p)$ which holds iff $2(1-p) < 1$, which is clearly true as $p > \frac{1}{2}$. Also, $q < \frac{1}{2}$ iff $2(1-p^2) < 1 + 2p - 2p^2$ which holds iff $1 < 2p$.

endorses, this sends a positive signal weaker than the one sent out by the unbiased reviewer, so $q \in (\frac{1}{2}, p)$, and hence Lemma 6 applies. If the optimistic reviewer fails to endorse, this sends out a stronger negative signal against the product than does the unbiased reviewer if he fails to endorse, so $q < 1 - p$, and hence by Lemma 2 a herd against the product starts immediately. Expected profits are thus:

$$\Pi_{or}^G = [p^2 + 2p(1 - p)] \Pi_{q \in (\frac{1}{2}, p)}^G$$

And hence:

Proposition 3. *Expected profits from choosing the optimistic reviewer in the good product case are:*

$$\Pi_{or}^G = [p^2 + 2p(1 - p)] \frac{p[1 + (1 - p)\theta(1 - \theta)]}{[1 - p(1 - p)\theta^2](1 - \theta)}$$

3.2. Comparing Profits from Different Reviewer Types. We are now in a position to compare profits from choosing the different types of reviewers.

Proposition 4. $\Pi_{pr}^G > \Pi_{ur}^G$ for all $p \in (\frac{1}{2}, 1)$ and $\theta \in (0, 1)$, so a firm launching a new product will unambiguously prefer to have it reviewed by a pessimistic reviewer over an unbiased one.

Proof. See Appendix. ■

The pessimistic reviewer is more likely than an unbiased reviewer to fail to endorse the product, thus sending a bad signal about product quality to later customers, and less likely to endorse and hence send a positive signal. However, if the pessimist does endorse, this sends a very strong positive signal about quality, strong enough in fact to initiate a cascade in favour of the product (after the pessimist endorses, $q > p$, starting a herd, while endorsement by the unbiased reviewer gives $\Pi_{q=p}^G$). Furthermore, customers understand that the pessimist is more likely to fail to endorse than an unbiased reviewer, thus dissipating the effect of such a failure on later customers' beliefs (failure to endorse by the pessimist leads to $\Pi_{q \in (1-p, \frac{1}{2})}^G$, but failure by the unbiased reviewer leads to $\Pi_{q=1-p}^G$). Overall, these effects overwhelm the more likely failure to endorse, so the good product firm prefers the pessimist.

The following tables show, for various numbers of expected potential customers, the absolute value of the pessimist (in expected terms), expected sales when the firm chooses an unbiased reviewer, and the percentage increase in sales from choosing a pessimistic reviewer over an unbiased one. We present three tables, one with $p = 0.75$, where signals are

halfway between being fully informative and completely uninformative, one with $p = 0.6$, where signals are quite uninformative, and finally one with $p = 0.9$, where signals are quite informative.

Table I: Value of Pessimistic Reviewer With $p = 0.75$

$E[N]$	Absolute Value of Pessimistic Reviewer	Expected Sales With Unbiased Reviewer	Percentage Increase in Expected Sales
10	0.6	8.0	7.8%
25	1.5	20.1	7.4%
50	2.9	40.3	7.3%
100	5.8	80.7	7.2%
200	11.6	161.5	7.2%
500	28.9	403.8	7.2%
1,000	57.7	807.6	7.1%
10,000	577.0	8076.8	7.1%

Table II: Value of Pessimistic Reviewer With $p = 0.6$

$E[N]$	Absolute Value of Pessimistic Reviewer	Expected Sales With Unbiased Reviewer	Percentage Increase in Expected Sales
10	0.4	6.3	6.1%
25	0.9	15.7	5.4%
50	1.6	31.5	5.2%
100	3.2	63.1	5.1%
200	6.4	126.3	5.1%
500	15.9	315.7	5.0%
1,000	31.6	631.5	5.0%
10,000	315.9	6315.7	5.0%

Table III: Value of Pessimistic Reviewer With $p = 0.9$

$E[N]$	Absolute Value of Pessimistic Reviewer	Expected Sales With Unbiased Reviewer	Percentage Increase in Expected Sales
10	0.4	9.3	4.6%
25	1.0	23.4	4.4%
50	2.0	46.9	4.3%
100	4.0	93.9	4.2%
200	7.9	187.9	4.2%
500	19.8	469.7	4.2%
1,000	39.6	939.5	4.2%
10,000	395.6	9395.6	4.2%

We can make a couple of observations from these tables. First, the advantage from choosing a pessimistic reviewer is quite substantial. With averagely informative signals in Table I, the pessimist adds between 7% and 8% to sales. Second, we see that in each case, the percentage increase rapidly converges to a bound.

Proposition 5. *The percentage increase in sales from having a pessimist choose first converges to a bound of $\frac{(1-p)(2p-1)}{(1+p)}$ as $\theta \rightarrow 1$.*

Proof. See Appendix. ■

For example, this bound equals $\frac{1}{14} = 7.1\%$ in the case of $p = 0.75$.

Finally, we compare pessimistic reviewers to optimistic ones:

Proposition 6. *Pessimistic reviewers are always strictly preferred to optimistic ones by a firm with a good product. However, the preference is slight as choosing the pessimist over the optimist cannot increase expected sales by more than $\frac{1}{4}$.*

Proof. See Appendix. ■

Pessimistic reviewers are valued by good product firms because of the strength of the signal they send out if they endorse the product. However, optimistic reviewers are also valuable: the proposition above shows that they are only slightly less so than pessimistic ones. The strength of the positive signal the optimistic reviewer sends out if he endorses is much weaker, because customers understand how optimistic the reviewer is. However, with a good product, the optimistic reviewer is very likely to endorse and so send out a positive signal. Thus, the good product firm has a general preference for biased reviewers

over unbiased ones, who have neither the advantage of the pessimistic reviewer, who sends out a very strong signal if he endorses, nor of the optimistic reviewer, who is very likely to send out a positive signal. Thus we might expect a polarization of reviewers. Reviewers may tend to take more extreme positions in order to make their reviews more attractive to firms with good products.

4. BAD PRODUCT FIRM

In this section, we consider the optimal choice of reviewer type for a firm which knows that its product is bad. We start by noting the following Remark:

Remark 4. *To calculate expected profits to the firm with a bad product, we need just take the expressions for the good product case and swap each p with $1 - p$ and vice-versa.*

Proof. In the case of a bad product, for any sequence of signals, the reviewers and customers will all act in the same way, as they do not know the product quality. However, the probability of each possible sequence of signals changes. At any given point, the probability of a H signal changes from p to $1 - p$ and the probability of a L signal changes from $1 - p$ to p . Thus, to calculate expected profits to the firm with a bad product, we need just take the expressions for the good product case and swap each p with $1 - p$ and vice-versa. ■

Using this, we can derive the following proposition.

Proposition 7. *A firm with a bad product prefers to have its product reviewed by an unbiased reviewer than by a pessimistic one if (and only if)*

$$p > p^* = \frac{1 + \theta(1 - \theta) - \sqrt{[1 + \theta(1 - \theta)]^2 - 2\theta(1 - \theta)(2 - \theta)}}{2\theta(1 - \theta)}$$

$p^* \in (\frac{1}{2}, 1)$, $\frac{dp^*}{dE[N]} < 0$ and as $\theta \rightarrow 1$, $p^* \rightarrow \frac{1}{2}$.

Proof. See Appendix. ■

The condition is not generally onerous. For example, if $E[N] \geq 10$, so $\theta \geq 0.9$, then $p^* \approx 0.53$, so the condition holds for all but the least informative signals. Of course, as $E[N]$ rises, the condition is weakened further, as $\frac{dp^*}{dE[N]} < 0$.

Again, the pessimistic reviewer is more likely than an unbiased one to fail to endorse the product, thus sending a bad signal about product quality to customers. Since the firm knows its product is bad, it realizes that the pessimistic reviewer is less likely to endorse than in the good product case, and so is less likely to send the strong positive signal inherent in

endorsement, but instead is more likely to send the negative signal inherent in a failure to endorse. Thus, for a firm with a bad product, choosing the pessimistic reviewer is too risky.

As for the good product firm, we can calculate the expected value (here negative) of choosing a pessimistic reviewer over an unbiased one, both in absolute and percentage terms, for different quality signals.

From the tables below, we see that having a pessimist choose first is very costly to the bad product firm. For example, with averagely informative signals ($p = 0.75$), the decrease in expected sales is of the order of 30%. We can also note that, as in the good product case, the percentage decrease rapidly converges to a bound as $E[N]$ goes up. Using Remark 4 and Proposition 5, the bound is given by $\frac{p \cdot (1-2p)}{(2-p)} < 0$.

We can also see that as p increases across the tables, for any $E[N]$ the pessimist becomes more costly to the firm. A higher p means signals are more reliable, so the risk to the bad product firm that the pessimist fails to endorse is amplified.

Table IV: Value of Pessimistic Reviewer With $p = 0.75$

$E[N]$	Absolute Value of Pessimistic Reviewer	Expected Sales With Unbiased Reviewer	Percentage Decrease in Expected Sales
10	-0.5	2.0	-23.9%
25	-1.3	4.9	-27.5%
50	-2.8	9.7	-28.7%
100	-5.7	19.3	-29.4%
200	-11.4	38.5	-29.7%
500	-28.7	96.2	-29.9%
1,000	-57.6	192.4	-29.9%
10,000	-576.8	1923.2	-30.0%

Table V: Value of Pessimistic Reviewer With $p = 0.6$

$E[N]$	Absolute Value of Pessimistic Reviewer	Expected Sales With Unbiased Reviewer	Percentage Decrease in Expected Sales
10	-0.2	3.7	-5.7%
25	-0.7	9.3	-7.4%
50	-1.5	18.5	-8.0%
100	-3.1	36.9	-8.3%
200	-6.2	73.7	-8.4%
500	-15.7	184.3	-8.5%
1,000	-31.5	368.5	-8.6%
10,000	-315.7	3684.3	-8.6%

Table VI: Value of Pessimistic Reviewer With $p = 0.9$

$E[N]$	Absolute Value of Pessimistic Reviewer	Expected Sales With Unbiased Reviewer	Percentage Decrease in Expected Sales
10	-0.3	0.7	-53.4%
25	-0.9	1.6	-60.4%
50	-1.9	3.1	-62.9%
100	-3.9	6.1	-64.2%
200	-7.9	12.1	-64.8%
500	-19.7	30.3	-65.2%
1,000	-39.5	60.5	-65.3%
10,000	-395.6	604.4	-65.4%

Finally, we see that the equivalent to Proposition 6 applies here.

Proposition 8. *Pessimistic reviewers are always strictly preferred to optimistic ones by a firm with a bad product. However, the preference is slight as choosing the pessimist over the optimist cannot increase expected sales by more than $\frac{1}{4}$.*

Proof. The proof follows immediately by applying Remark 4 to (A.2) in the proof of Proposition 6, to find that $\Pi_{pr}^B - \Pi_{or}^B = \Pi_{pr}^G - \Pi_{or}^G$, so the proof of Proposition 6 continues to apply. ■

The optimistic reviewer is even worse for the bad product firm than the pessimistic one, though only marginally so. Optimistic reviewers are bad as if they fail to endorse the product this sends out a very strong unfavourable signal to customers ($q < 1 - p$, so a herd against

the product starts immediately), and for a firm with a bad product, the probability that even an optimistic reviewer fails to endorse the product is reasonably high.

So in conclusion, we see that a firm with a bad product prefers to have its product reviewed by an unbiased reviewer over both pessimistic and optimistic ones. If the firm were forced to choose, it would have a slight preference for pessimistic reviewers over optimistic ones.

5. SIGNALLING AND POOLING

We have not yet considered how customers' priors about the product quality might be affected by the choice of reviewer the firm makes. From Propositions 4 and 7, the good product firm wishes to choose a pessimistic reviewer, but this is costly to the bad product firm, who prefers an unbiased reviewer. Thus, a candidate equilibrium in which the good product firm chooses a pessimist, but the bad product firm chooses an unbiased reviewer, will be undone by customers readjusting their priors in the light of the firm's choice - the choice will fully signal or reveal the product type. The bad product firm then has sales of zero, while the good product firm will sell to all the customers in the sequence, so the bad product firm would want to deviate by copying the choice of the good product firm, and hence be believed to be good. In fact, no separating equilibrium is possible. In any candidate separating equilibrium, the firm's type will be revealed, so the bad product firm would want to deviate by copying the choice of the good product firm.

Pooling equilibria, where both types of firm choose the same reviewer, are possible however. There are three candidate pooling equilibria, which we denote by (p, p) , (o, o) and (u, u) for the choice of a pessimist, an optimist and a unbiased reviewer respectively. Note first that (p, p) payoff dominates (o, o) , from Propositions 6 and 8. Thus, we would expect to see (p, p) played over (o, o) .

Next, we consider the beliefs that could sustain (p, p) and (u, u) as Perfect Bayesian Equilibria. Let w be the consumers' symmetric probability belief that $V = 1$. Because to this point we have only considered an initial 50:50 prior before any reviewer makes a decision, we will restrict $w \in \{0, \frac{1}{2}, 1\}$. Then (p, p) can be sustained iff $w = 0$ where the firm deviates and chooses an unbiased reviewer. If the deviator were thought to be good, the firm would always deviate whatever its type, and if the beliefs remained 50:50, the bad product firm would deviate, as it prefers unbiased reviewers given a 50:50 prior.¹² Similarly, (u, u) can also be sustained iff $w = 0$ where the firm deviates and chooses a pessimistic reviewer. As

¹²In the Perfect Bayesian Equilibrium, following a deviation to an optimistic reviewer, beliefs can remain 50:50, or the deviator can be thought to be bad.

before, the deviator cannot be believed to be good, or both types would deviate. If the beliefs remained 50:50, then of course the good product firm would deviate, as it then prefers the pessimist.¹³

Although both (p, p) and (u, u) can be sustained as Perfect Bayesian Equilibria, we argue that (p, p) is a more natural equilibrium. Formally, we assume a mild consistency requirement on the off-equilibrium path beliefs, requiring that both good and bad firms share the same belief about the value of w , and that consumers understand that firm beliefs are consistent in this way. Suppose we begin in the pooling equilibrium in which both types select an unbiased reviewer. If the bad firm deviates to choosing a pessimist, it must do so on the assumption that w is sufficiently high to return an improved payoff, i.e., $w = 1$. But then the good firm would also want to deviate. Hence, having observed a deviation, it must be the case that consumers believe the deviator is at least as likely to be good as it is to be bad, so $w \in \{\frac{1}{2}, 1\}$, which contradicts the above requirement that $w = 0$ for the (u, u) equilibrium to work. The (p, p) equilibrium is not open to the same objection. The same reasoning process (inverting the good and bad firms) leads us to the conclusion that a deviator must be at least as likely to be bad as to be good, which is consistent with $w = 0$, following a deviation to the unbiased reviewer. Such a deviation could follow from the firm believing that $w = \frac{1}{2}$, resulting in only the bad firm deviating.¹⁴

To summarize consider the following consumer thought process, having observed a deviation from (u, u) : “I’ve seen a deviation, who can it be? Well if it could be the bad type then *a fortiori* it could be the good type, so my beliefs can’t be biased toward the bad type.” Allowing such a thought process allows us to invalidate the pooling equilibrium involving the selection of unbiased reviewers.

In the (p, p) equilibrium, the bad product firm will be forced to copy the choice of the good product firm and choose a pessimist in order to make any sales at all. Then, because both the firms are doing the same thing, the customers’ priors are consistent with the firms’ equilibrium strategies. Note that in comparison to the case where customers do not adjust their priors in the light of the firm’s choice, social welfare increases. Sales of the good product are the same, but sales of the bad product are lower, as from Proposition 7, the bad product firm does worse and sells fewer units by choosing a pessimist.¹⁵

¹³Following a deviation to an optimistic reviewer, the deviator must also be believed to be bad.

¹⁴Neither firm wants to deviate to an optimist, unless they believe that $w = 1$ following the deviation, in which case both types would want to deviate. Thus, the belief that $w = \frac{1}{2}$ following such a deviation is the only belief consistent with our requirement.

¹⁵Note also that reputational concerns might push some bad product firms to not release the product onto the market at all.

6. MORE GENERAL PRIORS

So far, we have limited the analysis to the case where customers have 50:50 priors as to whether or not the product being launched on the market is of good quality. This makes sense if we think that, before receiving their private signal or observing the behaviour of other customers or reviewers, the customers have no prior information whatsoever about product quality. Thus, customers receive all their information through observing private signals and public behaviour. Where the priors are not 50:50, a number of questions arise: where does customers' initial information come from, why are we not modelling explicitly the process by which this information is acquired, and finally why can we not think of all private information acquisition as arising from the private signal? The second question is in a sense directly addressed in this paper in the 50:50 case by explicitly modelling reviewers, as they can be considered to generate priors through their endorsement decision. The assumption of 50:50 priors also simplifies the analysis.

Nonetheless, in this section, we consider more general priors. For conciseness, we consider just the good product firm. We find that for priors not too far from 50:50, our main result continues to hold: good product firms will continue to prefer pessimistic reviewers, appropriately defined. We also find that the more informative are private signals, i.e., the higher is p , the greater the range of priors around 50:50 for which the good product firm prefers pessimistic reviewers. Thus, we can conclude that our main result is not a knife-edge one, depending on priors being exactly 50:50.

Here, we let the general initial prior beliefs be q , while we denote the updated prior following the reviewer's decision by q_1 .

We will restrict $q < p$. If $q > p$, the positive prior outweighs any individual signal, and as a result, using Lemma 1, a herd starts straight away in the absence of any review. In that case, analyzing the firm's choice of reviewer to try to manipulate the herding structure is uninteresting. For simplicity, we do not consider the case where $q = p$. We will (for reasons to be explained shortly) also restrict $q > \frac{1}{2}$, so $q \in (\frac{1}{2}, p)$.

Suppose $q > \frac{1}{2}$. Then an unbiased reviewer will endorse the product if he gets a neutral pair of signals, HL or LH , as the two signals cancel and the reviewer is left with the positive prior. Of course, he will also endorse with two positive signals, while if he gets two negative signals he will fail to endorse, as following just the first negative signal, as shown in Remark 3, the reviewer's posterior $\hat{q} < \frac{1}{2}$. By analogy with the 50:50 prior case, we can now define three types of reviewer:

Definition 4. *Unbiased reviewers endorse the product iff they get two H signals or a H and a L signal. Mildly pessimistic reviewers endorse the product if they get two H signals, don't endorse if they get two L signals and flip a coin if they get a H and a L signal. Strongly pessimistic reviewers endorse the product iff they get two H signals.*

We can see that our three types of reviewers are in some sense analogous to the three types in the 50:50 prior case. Based purely on private information, the unbiased reviewer decides in the same way as the optimistic reviewer in the 50:50 case, the mild pessimist decides in the same way as the unbiased reviewer in the 50:50 case, and the strong pessimist decides in the same way as the pessimist in the 50:50 case. Because of the positive prior, a given decision rule based on private information effectively become more pessimistic relative to total information.

The reason we restrict $q > \frac{1}{2}$ is that where $q < \frac{1}{2}$ we would be in a symmetric situation where instead, we would have unbiased reviewers, mildly optimistic reviewers and strongly optimistic ones, so there could be no clear extension of the result that the good product firm prefers pessimists.¹⁶

We are now in a position to state our main result in this section:

Proposition 9. *Where $q < \frac{1+p}{3}$, the good product firm unambiguously prefers mildly pessimistic reviewers to either strongly pessimistic or unbiased ones.*

Proof. See Appendix. ■

Clearly, as p goes up, so does the range of q above $\frac{1}{2}$ for which the mild pessimist is optimal. In fact, we can easily see that for any p , the range of q for which the proposition holds is exactly one third of the permitted q range (from $\frac{1}{2}$ to p): $\frac{1}{3} \left[p - \frac{1}{2} \right] + \frac{1}{2} = \frac{2p-1+3}{6} = \frac{1+p}{3}$.

As in the 50:50 prior case, the (mild) pessimist is preferred to the unbiased reviewer, because though the pessimist is less likely to endorse the product, an endorsement is a very powerful positive signal, starting an immediate herd in favour of the product, while a failure to endorse is less harmful because the reviewer's pessimism is common knowledge. However, the result only holds where q is not too big relative to p . Otherwise, the strength of the prior is such that even an endorsement by the unbiased reviewer is sufficient to start a herd, who is then preferred as he is more likely to endorse.

¹⁶In any case, the symmetry of the problem means that the results for $q > \frac{1}{2}$ extend to $q < \frac{1}{2}$ with appropriate adjustment.

The mild pessimist is preferred to the strong pessimist, because the mild pessimist is sufficiently pessimistic that an endorsement leads to a herd, and so is preferred to the strong pessimist who is less likely to endorse.

As before, the preference for the (mild) pessimist is quite strong. For example, for $p = 0.75$, the mildly pessimistic reviewer adds about 7% to sales over the unbiased reviewer.

7. CONCLUSION

Our paper is the first to consider sequential sales with reviewers in a herding environment, and by explicitly considering reviewers provides a counter-intuitive justification of the existence of highly critical reviewers as useful to firms launching new products onto the market. More generally, the pessimistic reviewer is like choosing an initial tough public test. It is in a good product firm's best interest to choose a pessimistic reviewer. If the pessimist fails to endorse, his decision will not affect later movers' beliefs about the quality of the product too much, but if he does endorse, he is likely to immediately initiate a cascade in favour of the firm's product. Firms with bad products, on the other hand, will prefer unbiased reviewers to avoid the excessive risk of rejection by a pessimist, although signalling considerations make it likely that the bad product firm will want to copy the choice of the good product firm to avoid immediate revelation of quality through the choice of reviewer.

In order to focus on the heart of the issue, the modelling assumptions have been chosen to make the analysis as tractable as possible. However, interesting extensions might look at generalized beliefs of the firm about the quality of its own product, a more general structure of private information signals to reviewers and customers, or the issue of inter-firm competition. Finally, prices are effectively fixed in our model: there is scope for a unified analysis merging our work on reviewer choice with existing models of inter-temporal price discrimination under sequential decision-making.

REFERENCES

- Albano, G. L. and Lizzeri, A. (2001) 'Strategic Certification and Provision of Quality', *International Economic Review*, **42**, 267-283.
- Avery, C. and Meyer, M. (1999) 'Designing Hiring and Promotion Procedures When Evaluators are Biased', *Mimeo*, Oxford University.
- Avery, C. and Zemsky, P. (1998), 'Multi-Dimensional Uncertainty and Herd Behavior in Financial Markets', *American Economic Review*, **88**, 724-748.
- Bagwell, K. and Riordan, M. H. (1991) 'High and Declining Prices Signal Product Quality', *American Economic Review*, **81**, 224-239.
- Banerjee, A. V. (1992) 'A Simple Model of Herd Behavior', *Quarterly Journal of Economics*, **107**, 797-817.
- Bikhchandani, S., Hirshleifer, D. and Welch, I. (1992) 'A Theory of Fads, Fashion, Custom and Cultural Change as Informational Cascades', *Journal of Political Economy*, **100**, 992-1026.
- Calvert, R. L. (1985) 'The Value of Biased Information: A Rational Choice Model of Political Advice', *Journal of Politics*, **47**, 530-555.
- Caminal, R. and Vives, X. (1996) 'Why Market Shares Matter: An Information-Based Theory', *Rand Journal of Economics*, **27**, 221-239.
- Caminal, R. and Vives, X. (1999) 'Price Dynamics and Consumer Learning', *Journal of Economics and Management Strategy*, **8**, 95-131.
- Fishman, M. J. and Hagerty, K. M. (1990) 'The Optimal Amount of Discretion to Allow in Disclosure', *Quarterly Journal of Economics*, **105**, 427-444.
- Gill, D. and Sgroi, D. (2003), 'The Superiority of Tough Reviewers in a Model of Simultaneous Sales', *Cambridge Working Papers in Economics 0335*, Cambridge University.
- Lee, I. H. (1998), 'Market Crashes and Informational Avalanches', *Review of Economic Studies*, **65**, 741-759.
- Lizzeri, A. (1999) 'Information Revelation and Certification Intermediaries', *Rand Journal of Economics*, **30**, 214-231.
- Meyer, M. A. (1991) 'Learning from Coarse Information: Biased Contests and Career Profiles', *Review of Economic Studies*, **58**, 15-41.
- Ottaviani, M. (1999) 'Monopoly Pricing and Social Learning', *Mimeo*, University College London.
- Robinson, J. (1999), *The Oxford Companion to Wine (2nd ed.)*, Oxford University Press, 511-512.

Sah, R. K. and Stiglitz, J. E. (1986), 'The Architecture of Economic Systems: Hierarchies and Polyarchies', *American Economic Review*, **76**, 716-727.

Sgroi, D. (2002) 'Optimizing Information in the Herd: Guinea Pigs, Profits and Welfare', *Games and Economic Behavior*, **39**, 137-166.

Smith, L. and Sorensen, P. (2000) 'Pathological Outcomes of Observational Learning', *Econometrica*, **68**, 371-398.

Taylor, C. R. (1999) 'Time-on-the-Market as a Sign of Quality', *Review of Economic Studies*, **66**, 555-578.

Vettas, N. (1997) 'On the Informational Role of Quantities: Durable Goods and Consumers' Word-of-Mouth Communication', *International Economic Review*, **38**, 915-944.

Welch, I. (1992), 'Sequential Sales, Learning, and Cascades', *Journal of Finance*, **47**, 695-732.

APPENDIX

1. *Proof of Lemma 3.* From Remark 1 and $q = \frac{1}{2}$, if the ratio of the probability of the observed information set if the product were good to that if the product were bad is greater than 1, an unbiased agent should buy, if it less than 1, he should not, and if it is equal to 1, the unbiased agent flips a coin.

The first agent will follow his signal, i.e., will adopt iff $X_1 = H$. Suppose the first agent adopts. This reveals his signal to be H to the second agent. If the second agent also gets a H signal, he therefore also adopts, but if he gets a L signal the H and L signals cancel, so he is indifferent and flips a coin. If the third agent observes two adopt decisions, a Y herd starts. If he gets $X_3 = L$, his signal and that of the first agent cancel, but because the second agent adopted, he is more likely to have observed $X_2 = H$ than $X_2 = L$. Formally, for the third agent:

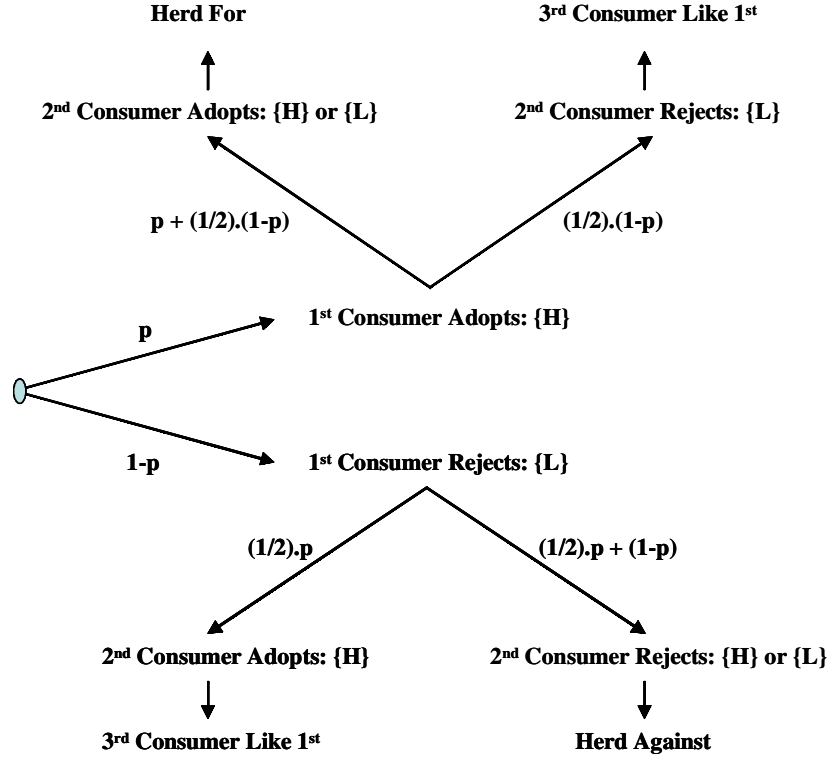
$$\frac{\Pr[V = 1|I_3]}{\Pr[V = 0|I_3]} = \frac{p + \frac{1}{2}(1-p)}{(1-p) + \frac{1}{2}p} > 1$$

So, the third agent adopts if gets a bad signal, and *a fortiori* adopts with a good signal, so a herd for the product has started.

Suppose instead that the first agent adopts, but the second agent rejects. Then the third agent can infer $\{X_1 = H, X_2 = L\}$. These two signals cancel, so the third agent is in exactly the same situation as the first agent before he received a signal.

The case where the first agent rejects is the symmetric opposite. If the second agent also rejects the product, a herd against starts. If the second agent buys, then the third agent is back to exactly the same situation as the first agent. All this can be illustrated in the following decision tree:¹⁷

¹⁷Remember the firm knows its product to be good. The branch probabilities are predicated upon this assumption.



We can use this tree to calculate the expected profit to the firm in this case by finding the payoff down various branches of the tree and multiplying by the probability of the relevant branch. Note that we have a recursive structure, whereby the payoffs from various points further down the tree are equivalent to those from points higher up in the tree, and we solve this recursive structure.

Letting $\Pi_{q=\frac{1}{2}}^G$ equal the discounted value of profits at the beginning of the tree, we get:

$$\begin{aligned}
 \Pi_{q=\frac{1}{2}}^G &= p \left[p + \frac{1}{2}(1-p) \right] \left(\frac{1}{1-\theta} \right) \\
 &\quad + p \left[\frac{1}{2}(1-p) \right] \left(1 + \Pi_{q=\frac{1}{2}}^G \theta^2 \right) \\
 &\quad + (1-p) \left(\frac{1}{2}p \right) (\theta + \Pi_{q=\frac{1}{2}}^G \theta^2) \\
 &\quad + (1-p) \left[\frac{1}{2}p + (1-p) \right] 0
 \end{aligned}$$

So:

$$\begin{aligned}\Pi_{q=\frac{1}{2}}^G &= \frac{1}{2}p(1+p) \left(\frac{1}{1-\theta} \right) \\ &\quad + \frac{1}{2}p(1-p) + \frac{1}{2}p(1-p)\Pi_{q=\frac{1}{2}}^G\theta^2 \\ &\quad + \frac{1}{2}p(1-p)\theta + \frac{1}{2}p(1-p)\Pi_{q=\frac{1}{2}}^G\theta^2\end{aligned}$$

So,

$$\Pi_{q=\frac{1}{2}}^G [1 - p(1-p)\theta^2] (1-\theta) = \frac{1}{2}p [1 + p + (1-p)(1+\theta)(1-\theta)]$$

So,

$$\Pi_{q=\frac{1}{2}}^G = \frac{p [2 - (1-p)\theta^2]}{2 [1 - p(1-p)\theta^2] (1-\theta)}$$

■

2. *Proof of Remark 3.* Following a H signal, $\hat{q} = \frac{pq}{pq+(1-p)(1-q)}$. Thus, $\hat{q} > p$ iff $q > pq + 1 - p - q + pq$, which in turns holds iff $2q(1-p) > 1-p$ or $q > \frac{1}{2}$, which of course we have assumed.

Following a L signal, $\hat{q} = \frac{(1-p)q}{(1-p).q+p(1-q)}$. Thus, by the same reasoning as in the H signal case (switching $1-p$ for p), $\hat{q} > 1-p$ iff $q > \frac{1}{2}$, which we assume. Also, $\frac{\Pr[V=1|I_i]}{\Pr[V=0|I_i]} = \frac{(1-p)q}{p(1-q)} = \frac{q-pq}{p-pq} < 1$, as we assume $q < p$, so $\hat{q} < \frac{1}{2}$. ■

3. *Proof of Lemma 6.* The first customer will purchase iff $X_1 = H$. If the first customer gets a H signal, he clearly adopts as $q > \frac{1}{2}$, while if he gets a L signal he rejects as, from Remark 3, $\hat{q} < \frac{1}{2}$.

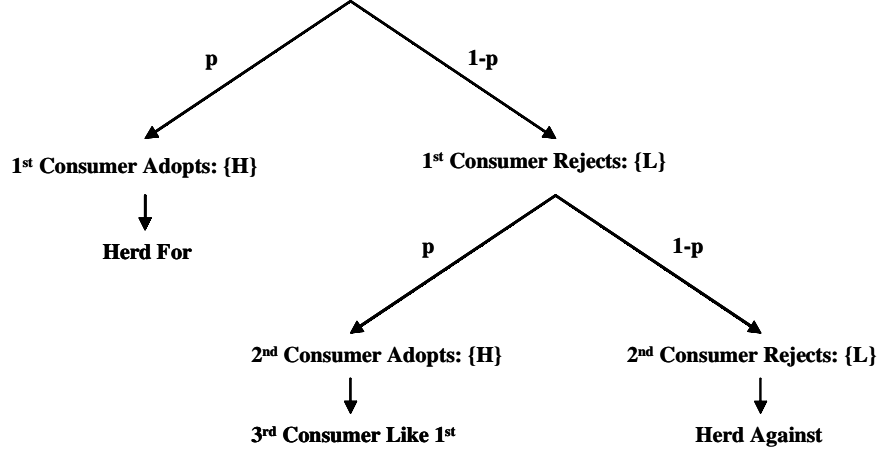
Following adoption by the first customer, a herd for the product starts. The second customer can infer that the first one got a H signal. Thus, from Remark 3, the second customer's updated prior $q_2 > p$, and so by Lemma 1 a herd for the product starts.

If the first agent rejects, then the second customer purchases iff $X_2 = H$. The second customer can infer $X_1 = L$, so if $X_2 = H$, the two signals cancel, and hence the second customer adopts as $q > \frac{1}{2}$. If $X_2 = L$, then the customer has effectively seen two negative signals, and so rejects.

Following rejection by the first customer, and a purchase by the second, the third customer can infer a L and a H signal, which cancel leaving him in exactly the same position as the first customer before he received a signal.

Following rejection by the first two customers, a herd against the product starts. If the third customer gets a H signal, this cancels one of the two inferred L signals, so the customer is left with just one L signal which, by Remark 3, leads to rejection. *A fortiori*, he also rejects if he gets a L signal.

As in the 50:50 prior case, all this information can be summarized in the following decision tree, with branch probabilities given by the fact that we are looking at the good product case.



Thus, for any specific $q \in (\frac{1}{2}, p)$, we can find expected discounted profits to the good product firm as follows:

$$\Pi_{q \in (\frac{1}{2}, p)}^G = p \left(\frac{1}{1 - \theta} \right) + (1 - p)p \left[\theta + \theta^2 \Pi_{q \in (\frac{1}{2}, p)}^G \right]$$

So,

$$\Pi_{q \in (\frac{1}{2}, p)}^G \left[1 - p(1 - p)\theta^2 \right] (1 - \theta) = p + p(1 - p)\theta (1 - \theta)$$

So,

$$\Pi_{q \in (\frac{1}{2}, p)}^G = \frac{p [1 + (1 - p)\theta (1 - \theta)]}{[1 - p(1 - p)\theta^2] (1 - \theta)}$$

But note that this value is independent of q . So the expression for $\Pi_{q \in (\frac{1}{2}, p)}^G$ in fact covers any q in the permitted range. ■

4. *Proof of Lemma 7.* Clearly, the shape of the decision tree, which is determined by customers who do not know the product type, will be the symmetric opposite of the one in the proof of Lemma 6, for the $q \in (\frac{1}{2}, p)$ case. A L signal starts a herd against, just like before a H signal started a herd for, while two H signals start a herd for, just like before two L

signals started a herd against. Again, if the first customer buys but the second does not, the inferred signals cancel.

Thus, for any specific $q \in (1-p, \frac{1}{2})$:

$$\Pi_{q \in (1-p, \frac{1}{2})}^G = p^2 \left(\frac{1}{1-\theta} \right) + p(1-p) \left[1 + \theta^2 \Pi_{q \in (1-p, \frac{1}{2})}^G \right]$$

So,

$$\Pi_{q \in (1-p, \frac{1}{2})}^G \left[1 - p(1-p)\theta^2 \right] (1-\theta) = p^2 + p(1-p)(1-\theta)$$

So,

$$\Pi_{q \in (1-p, \frac{1}{2})}^G = \frac{p[p + (1-p)(1-\theta)]}{[1 - p(1-p)\theta^2](1-\theta)}$$

Again, this profit is independent of the specific q value. ■

5. *Proof of Proposition 1.* We know that:

$$\Pi_{ur}^G = p\Pi_{q=p}^G + (1-p)\Pi_{q=1-p}^G$$

Using Lemma 4, Lemma 5 and Lemma 3, we get:

$$\begin{aligned} \Pi_{ur}^G &= p \left[p + \frac{1}{2}(1-p) \right] \left(\frac{1}{1-\theta} \right) + p\frac{1}{2}(1-p)\theta \frac{p[2 - (1-p)\theta^2]}{2[1 - p(1-p)\theta^2](1-\theta)} \\ &\quad + (1-p)\frac{1}{2}p \left(1 + \theta \frac{p[2 - (1-p)\theta^2]}{2[1 - p(1-p)\theta^2](1-\theta)} \right) \end{aligned}$$

So,

$$\begin{aligned} &\Pi_{ur}^G 2[1 - p(1-p)\theta^2](1-\theta) \\ &= (p^2 + p)[1 - p(1-p)\theta^2] + p^2(1-p)\theta[2 - (1-p)\theta^2] \\ &\quad + p(1-p)[1 - p(1-p)\theta^2](1-\theta) \\ &= p^2 - p^3(1-p)\theta^2 + p - p^2(1-p)\theta^2 \\ &\quad + 2p^2(1-p)\theta - p^2(1-p)^2\theta^3 \\ &\quad + p - p^2 - p(1-p)\theta - p^2(1-p)^2\theta^2 + p^2(1-p)^2\theta^3 \end{aligned}$$

So finally,

$$\begin{aligned}\Pi_{ur}^G &= \frac{2p - 2p^2(1-p)\theta^2 + 2p^2(1-p)\theta - p(1-p)\theta}{2[1 - p(1-p)\theta^2](1-\theta)} \\ &= \frac{p[2 + 2p(1-p)\theta(1-\theta) - (1-p)\theta]}{2[1 - p(1-p)\theta^2](1-\theta)}\end{aligned}$$

■

6. *Proof of Proposition 2.* We know that:

$$\Pi_{pr}^G = p^2 \left(\frac{1}{1-\theta} \right) + (1-p^2) \Pi_{q \in (1-p, \frac{1}{2})}^G$$

Using Lemma 7, we thus get:

$$\Pi_{pr}^G = p^2 \left(\frac{1}{1-\theta} \right) + (1-p^2) \frac{p[p + (1-p) \cdot (1-\theta)]}{[1 - p \cdot (1-p) \cdot \theta^2] \cdot (1-\theta)}$$

So,

$$\begin{aligned}& \Pi_{pr}^G [1 - p(1-p)\theta^2] (1-\theta) \\ &= p^2 [1 - p(1-p)\theta^2] + (1-p^2) p[p + (1-p)(1-\theta)] \\ &= p^2 - p^3(1-p)\theta^2 + p^2 + p(1-p)(1-\theta) - p^4 - p^3(1-p)(1-\theta) \\ &= p^2 - p^3(1-p)\theta^2 + p^2 + p - p^2 - p(1-p)\theta - p^4 - p^3 + p^4 + p^3(1-p)\theta\end{aligned}$$

So we get:

$$\Pi_{pr}^G = \frac{p[1 + p(1-p) + p^2(1-p)\theta(1-\theta) - (1-p)\theta]}{[1 - p(1-p)\theta^2](1-\theta)}$$

■

7. *Proof of Proposition 4.* From Proposition 1 and Proposition 2, $\Pi_{pr}^G - \Pi_{ur}^G$ is equal to:

$$\frac{p[2 + 2p(1-p) + 2p^2(1-p)\theta(1-\theta) - 2(1-p)\theta] - p[2 + 2p(1-p)\theta(1-\theta) - (1-p)\theta]}{2[1 - p(1-p)\theta^2](1-\theta)}$$

So,

$$(A.1) \quad \Pi_{pr}^G - \Pi_{ur}^G = \frac{p(1-p)[2p + 2p^2\theta(1-\theta) - \theta - 2p\theta(1-\theta)]}{2[1 - p(1-p)\theta^2](1-\theta)}$$

The denominator is strictly positive, as is $p(1-p)$, so we just need to show that:

$$2p + 2p^2\theta(1-\theta) - \theta - 2p\theta(1-\theta) > 0$$

So a sufficient condition is that $2p - \theta - 2p\theta(1-\theta) > 0$. This hold iff $2p[1 - \theta(1-\theta)] > \theta$. But $2p > 1$, so a further sufficient condition is that $1 - \theta(1-\theta) > \theta$, or $(\theta - 1)^2 > 0$, which is clearly true. ■

8. *Proof of Proposition 5.* Using (A.1) (in the proof of Proposition 4 above) and Proposition 1:

$$\frac{\Pi_{pr}^G - \Pi_{ur}^G}{\Pi_{ur}^G} = \frac{(1-p)[2p + 2p^2\theta(1-\theta) - \theta - 2p\theta(1-\theta)]}{[2 + 2p(1-p)\theta(1-\theta) - (1-p)\theta]}$$

As $\theta \rightarrow 1$, as claimed, this tends to:

$$\frac{(1-p)(2p-1)}{2 - (1-p)} = \frac{(1-p)(2p-1)}{(1+p)}$$

■

9. *Proof of Proposition 6.* Using Proposition 2 and Proposition 3:

$$\begin{aligned} & [\Pi_{pr}^G - \Pi_{or}^G] \frac{[1 - p(1-p)\theta^2](1-\theta)}{p} \\ &= 1 + p(1-p) + p^2(1-p)\theta(1-\theta) - (1-p)\theta \\ & \quad - p^2 - p^2(1-p)\theta(1-\theta) - 2p(1-p) - 2p(1-p)^2\theta(1-\theta) \\ &= 1 + p - p^2 - \theta + p\theta - p^2 - 2p + 2p^2 - 2p(1-p)^2\theta(1-\theta) \\ &= (1-p)(1-\theta)[1 - 2p(1-p)\theta] \end{aligned}$$

Thus:

$$(A.2) \quad \Pi_{pr}^G - \Pi_{or}^G = \frac{p(1-p)(1-\theta)[1 - 2p(1-p)\theta]}{[1 - p(1-p)\theta^2](1-\theta)}$$

The denominator is strictly positive, as is $p(1-p)(1-\theta)$, so the sign of $\Pi_{pr}^G - \Pi_{or}^G$ and $[1 - 2p(1-p)\theta]$ must be the same. Now, $p(1-p)$ is maximized at $p = \frac{1}{2}$, where $p(1-p) = \frac{1}{4}$. So $2p(1-p)\theta$ must always remain smaller than a half. Thus, $[1 - 2p(1-p)\theta] > 0$, and hence $\Pi_{pr}^G - \Pi_{or}^G > 0$.

Next we need to show that this difference can never exceed $\frac{1}{4}$. For the moment allow $p \in [\frac{1}{2}, 1)$ and $\theta \in [0, 1)$ - before we restricted ourselves to the fully open equivalents. Let

$z = p(1 - p)$, so we can rewrite:

$$\Pi_{pr}^G - \Pi_{or}^G = \frac{z(1 - 2z\theta)}{1 - z\theta^2}$$

$\frac{dz}{dp} = 1 - 2p < 0$ for $p > \frac{1}{2}$. So, given $p \in [\frac{1}{2}, 1)$, $z \in (0, \frac{1}{4}]$. At $p = \frac{1}{2}$, $z = \frac{1}{4}$, and z gradually falls towards zero as p approaches one.

Now for $\theta > 0$, $2z\theta > z\theta^2$ as $2 > \theta$. Thus, $1 - 2z\theta < 1 - z\theta^2$. Both these expressions are strictly positive, so $\frac{z(1 - 2z\theta)}{1 - z\theta^2} < z$. At $\theta = 0$, this expression equals z . Thus, for any $z \in (0, \frac{1}{4}]$, $\Pi_{pr}^G - \Pi_{or}^G$ is maximized at $\theta = 0$. Furthermore, given $\theta = 0$, $\Pi_{pr}^G - \Pi_{or}^G = z$, and so is maximized at the top of z 's range, i.e. at $z = \frac{1}{4}$.

Thus, we have found that for our slightly larger ranges of θ and z , the maximal difference is $\frac{1}{4}$, so of course restricting ourselves again to the fully open ranges the difference can never exceed $\frac{1}{4}$. (In fact, as $\theta \rightarrow 0$ and $z \rightarrow \frac{1}{4}$, the difference gets arbitrarily close to $\frac{1}{4}$.) ■

10. *Proof of Proposition 7.* Using (A.1) in the proof of Proposition 4 above,

$$\Pi_{pr}^G - \Pi_{ur}^G = \frac{p(1 - p) [2p - 2p(1 - p)\theta(1 - \theta) - \theta]}{2 \cdot [1 - p(1 - p)\theta^2] (1 - \theta)}$$

So, using Remark 4, the difference in profits from choosing an unbiased reviewer over a pessimistic one in the bad product case is:

$$\Pi_{ur}^B - \Pi_{pr}^B = \frac{p(1 - p) [\theta - 2(1 - p) + 2p(1 - p)\theta(1 - \theta)]}{2 [1 - p(1 - p)\theta^2] (1 - \theta)}$$

Therefore,

$$\begin{aligned} \Pi_{ur}^B &> \Pi_{pr}^B \\ \Leftrightarrow \theta - 2(1 - p) + 2p(1 - p)\theta(1 - \theta) &> 0 \end{aligned}$$

We can solve to find the values of p^* which satisfy the expression with equality:

$$\begin{aligned} \theta - 2(1 - p) + 2p(1 - p)\theta(1 - \theta) &= 0 \\ \Leftrightarrow \theta - 2 + 2p + p2\theta(1 - \theta) - 2p^2\theta(1 - \theta) &= 0 \\ \Leftrightarrow p^2 [\theta(1 - \theta)] - p [1 + \theta(1 - \theta)] + \frac{2 - \theta}{2} &= 0 \end{aligned}$$

So,

$$p^* = \frac{1 + \theta(1 - \theta) \pm \sqrt{[1 + \theta(1 - \theta)]^2 - 2\theta(1 - \theta)(2 - \theta)}}{2\theta(1 - \theta)}$$

The first thing to note is that we have real roots. This can be shown as follows:

$$\begin{aligned} & [1 + \theta(1 - \theta)]^2 - 2\theta(1 - \theta)(2 - \theta) \\ &= 1 + 2\theta(1 - \theta) + [\theta(1 - \theta)]^2 - 4\theta(1 - \theta) + 2\theta^2(1 - \theta) \\ &= [1 - \theta(1 - \theta)]^2 + 2\theta^2(1 - \theta) > 0 \end{aligned}$$

Then, given $1 + \theta(1 - \theta) > 1$, and noting that $0 < \theta(1 - \theta) \leq \frac{1}{4}$ (as $\theta(1 - \theta)$ is maximized at $\theta = \frac{1}{2}$), we can see that:

$$\frac{1 + \theta(1 - \theta) + \sqrt{[1 + \theta(1 - \theta)]^2 - 2\theta(1 - \theta)(2 - \theta)}}{2\theta(1 - \theta)} > 1$$

So the only potential $p^* \in (\frac{1}{2}, 1)$ is

$$(A.3) \quad p^* = \frac{1 + \theta(1 - \theta) - \sqrt{[1 + \theta(1 - \theta)]^2 - 2\theta(1 - \theta)(2 - \theta)}}{2\theta(1 - \theta)}$$

Now,

$$(A.4) \quad \frac{d[\theta - 2(1 - p) + 2p(1 - p)\theta(1 - \theta)]}{dp} = 2 + 2\theta(1 - \theta)(1 - 2p) > 0$$

as $0 < \theta(1 - \theta) \leq \frac{1}{4}$ (see above) and $1 - 2p > -1$. Thus, given θ , if $p^* \in (\frac{1}{2}, 1)$, then at p^* , $\Pi_{ur}^B = \Pi_{pr}^B$, and $\forall p > p^*$, $\Pi_{ur}^B > \Pi_{pr}^B$.

Now we will see how p^* changes with θ . Note that:

$$(A.5) \quad \frac{d[\theta - 2(1 - p) + 2p(1 - p)\theta(1 - \theta)]}{d\theta} = 1 + 2p(1 - p)(1 - 2\theta) > 0$$

as $0 < p(1 - p) < \frac{1}{4}$ (because $p(1 - p)$ is maximized at $p = \frac{1}{2}$) and $1 - 2\theta > -1$. Take any $\theta \in (0, 1)$ such that $p^* \in (\frac{1}{2}, 1)$, for example $\theta = \frac{1}{2}$ where $p^* = \frac{1.25 - \sqrt{(1.25)^2 - 0.75}}{0.5} \approx 0.70$. If θ increases a little, at the old p^* , $\theta - 2(1 - p) + 2p(1 - p)\theta(1 - \theta) > 0$ by (A.5), so given (A.4), p can fall a little to restore equality at a new lower p^* . Similarly if θ falls a little, p^* needs to increase. Thus $\frac{dp^*}{d\theta} < 0$ at any θ such that $p^* \in (\frac{1}{2}, 1)$.

Furthermore, $p^* \rightarrow \frac{1}{2}$ as $\theta \rightarrow 1$ and $p^* \rightarrow 1$ as $\theta \rightarrow 0$. As $\theta \rightarrow 1$ or $\theta \rightarrow 0$, the numerator and denominator of (A.3) both go to zero. Thus we need to apply l'Hôpital's Rule to find

the limits. Letting $p^* = \frac{f[\theta]}{g[\theta]}$, we have:

$$f'[\theta] = 1 - 2\theta - \frac{1}{2} \frac{\{2[1 + \theta(1 - \theta)](1 - 2\theta) + 2\theta(1 - \theta) + 2\theta(2 - \theta) - 2(1 - \theta)(2 - \theta)\}}{\sqrt{[1 + \theta(1 - \theta)]^2 - 2\theta(1 - \theta)(2 - \theta)}}$$

$$g'[\theta] = 2(1 - 2\theta)$$

As $\theta \rightarrow 1$, $f'[\theta] \rightarrow -1$ and $g'[\theta] \rightarrow -2$, so $p^* \rightarrow \frac{1}{2}$.

As $\theta \rightarrow 0$, $f'[\theta] \rightarrow 2$ and $g'[\theta] \rightarrow 2$, so $p^* \rightarrow 1$.

Finally, we need to show that as θ increases (decreases), p^* does not fall below $\frac{1}{2}$ (rise above 1) and then rise (fall) back again towards its limit. If p^* equalled $\frac{1}{2} - \varepsilon$ (or $1 + \varepsilon$), both (A.4) and (A.5) would continue to hold for small ε . Thus, p^* could not converge back to $\frac{1}{2}$ (or 1), as p^* would eventually start falling (rising) again, so p^* must remain in the range $(\frac{1}{2}, 1)$, and from above $\frac{dp^*}{d\theta} < 0$ everywhere in this range. ■

11. Proof of Proposition 9. We can now calculate the profit from choosing the different types of initial reviewer. We start with the mildly pessimistic reviewer. If the mildly pessimistic reviewer endorses the product, the updated prior for the first customer is, by Bayes' Rule:

$$q_1 = \frac{[p^2 + p(1 - p)]q}{[p^2 + p(1 - p)]q + [(1 - p)^2 + p(1 - p)](1 - q)}$$

$$= \frac{pq}{pq + (1 - p)(1 - q)}$$

Note that this is exactly the same prior as would arise if a single H signal were revealed. Thus, from Remark 3, $q_1 > p$, and hence by Lemma 1 a herd for the product starts. If the mildly pessimistic reviewer fails to endorse the product:

$$q_1 = \frac{[(1 - p)^2 + p(1 - p)]q}{[(1 - p)^2 + p(1 - p)]q + [p^2 + p(1 - p)](1 - q)}$$

$$= \frac{(1 - p)q}{(1 - p)q + p(1 - q)}$$

which is exactly the same updated prior as would arise if a single L signal were revealed. Thus, from Remark 3, the updated prior $q_1 \in (1 - p, \frac{1}{2})$. From that point on, we have a sequence of customers with a prior in $(1 - p, \frac{1}{2})$, so expected profits are $\Pi_{q \in (1 - p, \frac{1}{2})}^G$. Putting all this together, we can conclude that the expected discounted profits to choosing a mildly

pessimistic reviewer are (using $G+$ to indicate a good product firm where $q \in (\frac{1}{2}, p)$):

$$\Pi_{mpr}^{G+} = p \left(\frac{1}{1-\theta} \right) + (1-p) \Pi_{q \in (1-p, \frac{1}{2})}^G$$

Next we calculate the expected discounted profit to the firm from choosing an unbiased reviewer, Π_{ur}^{G+} . If the unbiased reviewer does not endorse, a herd against the product starts. The unbiased reviewer fails to endorse iff he gets two L signals. Following a failure to endorse, these are inferred by the first consumer, who then always rejects just like the third customer in a sequence of customers who has observed two rejections where $q \in (\frac{1}{2}, p)$ (see the proof of Lemma 6).

If the unbiased reviewer endorses, then the updated prior of the first customer $q_1 \in (\frac{1}{2}, p)$. Clearly, given $q > \frac{1}{2}$, $q_1 > \frac{1}{2}$. It remains to be shown that:

$$q_1 = \frac{[p^2 + 2p(1-p)] q}{[p^2 + 2p(1-p)] q + [(1-p)^2 + 2p(1-p)] (1-q)} < p$$

This holds iff:

$$\begin{aligned} (2p - p^2)q &< (2p - p^2)qp + (1 - p^2)(1 - q)p \\ \Leftrightarrow 2pq - p^2q &< 2p^2q - p^3q + p - p^3 - qp + p^3q \\ \Leftrightarrow q &< \frac{p - p^3}{(3p - 3p^2)} = \frac{(1-p)(1+p)}{3(1-p)} = \frac{(1+p)}{3} \end{aligned}$$

which is the condition imposed on q in the Proposition. Thus, from that point on, profits are given by $\Pi_{q \in (\frac{1}{2}, p)}^G$. Therefore, we conclude:

$$\Pi_{ur}^{G+} = [p^2 + 2p(1-p)] \Pi_{q \in (\frac{1}{2}, p)}^G$$

Comparing the mild pessimist to the unbiased reviewer, we find that, using Lemma 6 and Lemma 7:

$$\begin{aligned} \Pi_{mpr}^{G+} - \Pi_{ur}^{G+} &= p \left(\frac{1}{1-\theta} \right) + (1-p) \frac{p^2 + p(1-p)(1-\theta)}{[1 - p(1-p)\theta^2] (1-\theta)} \\ &\quad - [p^2 + 2p(1-p)] \frac{p + p(1-p)\theta(1-\theta)}{[1 - p(1-p)\theta^2] (1-\theta)} \end{aligned}$$

Therefore,

$$\begin{aligned}
& [\Pi_{mpr}^{G+} - \Pi_{ur}^{G+}] [1 - p(1-p)\theta^2] (1-\theta) \\
= & p - p^2(1-p)\theta^2 + p^2(1-p) + p(1-p)^2(1-\theta) \\
& - p^3 - p^3(1-p)\theta(1-\theta) - 2p^2(1-p) - 2p^2(1-p)^2\theta(1-\theta) \\
= & p(1-p) \{ -p\theta^2 + (1-p)(1-\theta) - p^2\theta(1-\theta) - 2p(1-p)\theta(1-\theta) \} \\
& + p + p^2(1-p) - p^3 - 2p^2(1-p) \\
= & p(1-p) \{ -p\theta^2 + 1 - p - \theta + p\theta - p^2\theta + p^2\theta^2 - 2p(1-p)\theta(1-\theta) + 1 \}
\end{aligned}$$

So,

$$\Pi_{mpr}^{G+} - \Pi_{ur}^{G+} = \frac{p(1-p) [2 - p - \theta - p(1-p)\theta(1-\theta)]}{[1 - p(1-p)\theta^2] (1-\theta)}$$

The denominator is strictly positive, as is $p(1-p)$. Thus $\Pi_{mpr}^{G+} > \Pi_{ur}^{G+}$ iff $2 - p - \theta - p(1-p)\theta(1-\theta) > 0$. At $\{p = 1, \theta = 1\}$, $2 - p - \theta - p(1-p)\theta(1-\theta) = 0$. But over $p \in (\frac{1}{2}, 1]$, $\theta \in (0, 1]$

$$\frac{d[2 - p - \theta - p(1-p)\theta(1-\theta)]}{d\theta} = -1 + p(1-p)[2\theta - 1] < 0$$

as the derivative is maximized at $\theta = 1$, where it equals $-1 + p(1-p) < 0$. By a similar reasoning, $\frac{d[2 - p - \theta - p(1-p)\theta(1-\theta)]}{dp} < 0$ over this range. Thus as p and θ fall below $\{1, 1\}$, $2 - p - \theta - p(1-p)\theta(1-\theta)$ rises above zero, and therefore $\Pi_{mpr}^{G+} > \Pi_{ur}^{G+}$ everywhere.

Finally, we calculate expected discounted profits to the firm from choosing a strong pessimist, Π_{spr}^{G+} . If the strong pessimist endorses, this is an even stronger positive signal than an endorsement by the mild pessimist, so again a herd for the product starts. If the strong pessimist fails to endorse, the updated prior of the first customer $q_1 \in (1 - p, \frac{1}{2})$. Clearly, the failure to endorse sends a weaker negative signal than a failure to endorse by the mild pessimist. In the mild pessimist case, after such a failure, $q_1 > 1 - p$, so this must remain the case here. It remains to be shown that $q_1 < \frac{1}{2}$ or equivalently:

$$\begin{aligned}
\frac{\Pr[V = 1|I_i]}{\Pr[V = 0|I_i]} &= \frac{[2p(1-p) + (1-p)^2]q}{[2p(1-p) + p^2](1-q)} \\
&= \frac{(1-p^2)q}{(2p - p^2)(1-q)} < 1
\end{aligned}$$

This holds iff:

$$\begin{aligned} q - p^2 q &< 2p - p^2 - 2pq + p^2 q \Leftrightarrow \\ q &< \frac{2p - p^2}{1 + 2p - 2p^2} \end{aligned}$$

But, this condition must be satisfied as in the Proposition, we have restricted $q < \frac{1+p}{3}$, which is a stronger condition. This last claim is shown as follows:

$$\begin{aligned} \frac{2p - p^2}{1 + 2p - 2p^2} &> \frac{1 + p}{3} \\ \Leftrightarrow 6p - 3p^2 &> 1 + 2p - 2p^2 + p + 2p^2 - 2p^3 \\ \Leftrightarrow 2p^3 - 3p^2 + 3p - 1 &> 0 \Leftrightarrow (2p - 1)(p^2 + 1 - p) > 0 \end{aligned}$$

which is clearly true.

Thus, if the strong pessimist endorses, a herd for the product starts, and if he fails to endorse, $q_1 \in (1 - p, \frac{1}{2})$, so from that point on expected profits are $\Pi_{q \in (1-p, \frac{1}{2})}^G$. Thus we conclude that:

$$\Pi_{spr}^{G+} = p^2 \left(\frac{1}{1 - \theta} \right) + (1 - p^2) \Pi_{q \in (1-p, \frac{1}{2})}^G$$

Compared to Π_{mpr}^{G+} , clearly $\Pi_{mpr}^{G+} > \Pi_{spr}^{G+}$ as the mild pessimist gives a higher chance of a herd and a lower chance of expected profits $\Pi_{q \in (1-p, \frac{1}{2})}^G$. ■