ISSN 1471-0498

DEPARTMENT OF ECONOMICS

DISCUSSION PAPER SERIES

THE SUPERIORITY OF BIASED REVIEWERS IN A MODEL OF SIMULTANEOUS SALES

David Gill and Daniel Sgroi

Number 206

September 2004

Manor Road Building, Oxford OX1 3UQ

The Superiority of Biased Reviewers in a Model of Simultaneous Sales [1]

David Gill
Trinity College
University of Oxford

Daniel Sgroi
Churchill College & Dept. of Applied Economics
University of Cambridge

Abstract

This paper considers the impact of reviewers on sales of products of quality unknown to consumers. Sales occur simultaneously after consideration by a reviewer with a known level of bias. Consumers observe the reviewer's decision and a private signal. We find that: (a) with flexible prices and signals that are not too revealing the reviewer most biased against the product is best for profits; (b) with flexible prices and very revealing private signals the reviewer most biased in favor is optimal; (c) with fixed prices then a reviewer biased against, but close to unbiased, is optimal.

Keywords: private information, reviewers, bias, simultaneous sales, marketing
JEL classification: D82, D83, L15

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, pp. 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 the interaction between price and public reviews, and especially the decision of where to send a product for review. This paper attempts to correct this omission.

We focus on the use of reviewers and tests to raise sales for firms with products of quality unknown to consumers, but this is not a signaling paper. A high quality firm cannot differentiate

[1] The paper is a substantial revision of Gill and Sgroi (2003b). Daniel Sgroi can be contacted at daniel.sgroi@econ.cam.ac.uk. David Gill can be contacted at david.gill@economics.ox.ac.uk and would like to thank the Economic and Social Research Council for financial support. 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 seminars in Oxford, Cambridge, Essex, and at the Second World Congress of the Game Theory Society in Marseille for helpful comments and suggestions.

itself from a low quality firm via its actions, as it is costless for the low quality firm to duplicate the high quality firm's actions. As seems reasonable in this context, a firm with even the lowest quality product can still send it to review at the same cost as a firm with the highest quality product. As a result, all of our outcomes will be pooling, and there will be no issues of incentive compatibility or scope for a separating equilibrium. Despite this, we find that a firm will prefer to subject its product to a reviewer with an extreme bias either in favor, or quite counter-intuitively, against the product. To stress, even though there is no gain to be made in demonstrating directly the quality of their product by braving the tough reviewer, firms will still often opt to subject their own product *to the toughest possible reviewer* simply in order to maximize profits by hoping for an endorsement from such a tough reviewer.

First let us consider the context. We consider a firm with either a high or low quality product. Low quality means no utility for the consumer, so a consumer is eager to avoid purchasing a low quality product for any positive price. Information for consumers comes from two sources. They receive a private signal drawn from a distribution guaranteeing an informative though not fully-revealing signal (correct with probability p - the quality of the information). They also observe the decision of a public reviewer or test result. Reviewers are assumed to be better informed than general consumers, because they are allowed to "test out" the product. We only require that reviewers are mildly better informed, letting consumers receive one i.i.d. private signal and reviewers two i.i.d. signals, all from the same distribution. Nevertheless, the extra information they possess means that a reviewer can potentially swamp private opinion. Therefore, a firm must be careful to chose a reviewer type that moves public information in the right direction, given that the types are publicly known. As the reviewers receive multiple signals they will often face mixed evidence of a product's quality. We vary the probability of endorsing on this mixed set of signals, with the most extreme pessimist never endorsing on mixed signals, and the most extreme optimist always endorsing on mixed signals. Consumers know the type of the reviewer or difficulty of test measured by the response to mixed signals. The firm can change price after the result of the review or test is known and consumers can then respond to the price, reviewer's decision and their own signal by purchasing or not. The firm chooses an optimal price given the review but also has the choice of where to send its product for review (or the level of test faced by the product). By allowing both prices and reviews to be chosen we embrace a huge number of real markets where public reviews/tests are possible, from cars, electrical goods, wines, and virtually any product of unknown quality. By also considering situations where prices cannot be changed but reviewers can still be selected we add a variety of cases such as the film industry where ticket prices are not usually a function of perceptions of film quality, and even more generally areas in which price may not be relevant, such as Ph.D. students selecting referees, a high school graduate deciding how tough a university to attend, politicians selecting interviewers, etc.

Now consider: should a car firm put its product through a safety test well above the minimum required? Should a computer game firm send its product to the toughest reviewer in the market? Should a Ph.D. student list well-known extremely tough referees on his C.V.? How tough an auditor or investment bank should a firm launching an I.P.O. choose? We have a trade-off: choose a "yes" man, soft reviewer or easy test and the chances of endorsement are high, but the impact low as consumers will expect endorsement. Choose a pessimistic reviewer or tough test and in all likelihood there will be no endorsement or the test will be failed, but should a product pass, the impact on sales will be positive and strong. We find that in many cases an economic agent with a product of unknown quality will prefer to put its product through the toughest test available. At times it will opt for the lightest test, or easiest reviewer, but when prices can be adjusted it will never go for a non-extreme type. Added to this choice, in the main model the firm needs to consider setting a high price and focusing its sales on those with a strong belief that the product is good, or pricing low to sell to the entire market. In the alternative version of the model the price is fixed and so the firm can only focus on selecting the best test for the product to face. Under such fixed prices the firm will opt instead for the tough test or pessimistic reviewer that is as close as possible to the unbiased or neutral type (which randomizes 50:50 on a mixed set of signals). Much then depends upon whether prices can be changed, but in either case the results are stark and often surprising.

The intuition behind our findings in the flexible price case begins with the fact that all prices are decreasing in the reviewer's probability of endorsement. Endorsement raises beliefs more the tougher the reviewer, while a failure to endorse is not so damaging to beliefs if the reviewer is tougher. In fact, following endorsement prices are strictly convex in reviewer bias, so conditional on endorsement a change in reviewer bias is more powerful where the reviewer is tougher. The probability of an endorsement is linearly increasing in the level of bias of the reviewer and so we find that to maximize profit from endorsement the firm will wish to move towards an extreme level of reviewer bias, either opting for an extreme pessimist to benefit from a steep increase in price while the probability of endorsement falls linearly, or an extreme optimist to benefit from a linear increase in the probability of endorsement, while price does not fall much as the reviewer becomes very optimistic. To maximize sales from a failure to endorse the firm will always opt for the extreme pessimist as both the probability of failing to endorse and prices are decreasing in the probability of endorsement and so in the level of bias. The relevant functions are again convex, and thus, summing over the case of endorsement and failure to endorse, the firm always prefers an extreme. The firm will either choose the extreme pessimist to maximize prices, at the cost of a lower probability of endorsement, or the extreme optimist to maximize the probability of endorsement, at the cost of lower prices. Which extreme is best depends on the quality of information: for p below 0.73, the extreme pessimist is best, while for p above 0.73, the extreme optimist is chosen. We also show, quite intuitively, that opting for no reviewer is worse than selecting even an unbiased reviewer, let alone an extremist.

In the simpler model, when prices cannot be altered we find that the firm will opt for the mildest form of pessimist possible. The intuition is more straightforward. Since the reviewer's signals and consumers' signals are drawn from the same distribution, but the reviewer gets two draws, a pessimist endorsing will swamp private signals, and this is true regardless of the quality of private information. With no endorsement the firm is also better off with the pessimist as his failure to endorse is less damaging. Now by opting for the mildest form of pessimist that is still pessimistic enough to swamp private information on an endorsement, the firm can guarantee the highest expected profits. This is true regardless of the level of bias held by the mildest form of pessimist; for example, if only a very extreme pessimist is available then the firm will opt for that; while if a pessimist arbitrarily close to an unbiased type is available then the firm will instead opt for that. In this case opting for no reviewer is equivalent to selecting the unbiased type.

The results we get are perhaps most counter-intuitive when we consider the case of private information which is informative but not too revealing, so p is in the range between 0.50 and 0.73, and prices can be changed by the firm in response to the reviewer's decision. We get a clear optimal choice for the firm to select a reviewer or test for its product that is as biased against as our model allows. This leaves the probability of endorsement very low, but maintains the twin rewards of high sales if endorsement occurs even with a very high price, and a minimal impact from a failure to endorse since that is exactly what consumers expect from so tough a reviewer or hard a test.

1.1 Related Literature

There is no paper which incorporates the combination of choice of reviewer and choice of price for a firm as we do; however there are various different related literatures.

The paper most closely related to ours is Lerner and Tirole's (2004) recent working paper concerning the role of technology standard setting authorities as certifiers. Similarly to our paper, certifiers are assumed to have an arbitrary bias towards the technology sponsor which determines their endorsement rule. Lerner and Tirole demonstrate a series of interesting results such as the optimality of relying on a single certifier rather than a set or sequence of certifiers. They also discuss some elements of competition between sponsors. However, the model they use has significant differences to ours: their certifiers discover with certainty the quality of the technology they are asked to review; consumers in their model do not receive any private information; and there is no scope for changing price in response to the certification. Therefore, as certifiers cannot overwhelm bad private information or enable a rise in price, Lerner and Tirole do not find any role for certifiers biased *against* the technology. Instead they find that the sponsor prefers the certifier most biased in favor of the new technology on offer, subject to users adopting following an endorsement. This is in stark contrast to our findings, which allow a role for reviewers biased in either direction depending upon the quality of private information and the flexibility of price.

While we consider simultaneous sales after a review, an alternative might consider consumers instead acting in strict sequence, observing the actions of their predecessors. In a sense observing predecessors has a similar feel to observing the public decision of a reviewer or test result. Banerjee (1992) and Bikhchandani, Hirshleifer and Welch (1992) examine a sequential herding model appropriate for this setting. In such a sequential set-up, Gill and Sgroi (2003a) consider the use of biased reviewers, while 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. Gill and Sgroi (2003a) also find that pessimists can be useful, while Sgroi (2002) finds that irrespective of product quality firms would like to use these "guinea pigs".

Caillaud and Tirole (1999) analyze how an ideologically biased political party membership can help party credibility by forcing a centrist leadership to exert effort on improving platform quality (which voters cannot observe directly). A neutral party will rubberstamp political platforms, inducing zero effort, while a biased party may take over formulation of the platform if quality is not high enough, so validation acts as a signal of quality. Our model is somewhat different in that product quality is determined exogenously, so reviewers are not acting as a commitment device, and nor is it meaningful for Caillaud and Tirole to incorporate a notion of price in their work.

Much of this paper allows the firm the choice of reviewer and price, but a number of other papers focus on the price decision alone in a not dissimilar context. Especially relevant are those papers that analyze the use of initial prices to manipulate sales in a learning environment. Various papers such as Caminal and Vives (1996, 1999) and Vettas (1997) focus on low introductory prices to initiate herd-like behavior from customers. In contrast, 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 therefore allows for the possibility of a herd against the sale of a house, but not in favor 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. None of these papers explicitly considers the role of publicly observable reviews, possibly by known biased reviewers, in forming the prior beliefs of consumers. However, the high initial prices in Ottaviani and in Taylor play a qualitatively similar role to the highly critical pessimistic reviewers in this paper: a failure to pass the tough test imposed by either a high initial price or a pessimistic reviewer is not viewed as overly damaging since the test is so unlikely to be passed.

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 favor of this predisposition. For example, Meyer finds that firms who are trying to decide which employee to promote should bias contests in favor 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 favor 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 (who know product quality with certainty and whose aim is to maximize profits) may choose early evaluators that are biased against or in favor of their product depending upon the quality of information available to consumers and reviewers.

Our paper should be contrasted with the literature on experts, in which self-interested experts filter information about the true state of the world (see chapter 10 of Chamley (2004) for a survey). Their self-interest gives rise to incentives to manipulate the messages they send. We, on the other hand, assume that our reviewers have no self-interested motives, apart from taking on biases to appeal to firms. Our work is also different from the 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. Our focus is different as 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, which might make an interesting topic for future research. Finally, our paper is also different from Ottaviani and Pratt (2001), who find that a monopolist may wish to use a public signal of quality such as an outside certifier. Our seller is informed about quality, while in Ottaviani and Pratt both the buyer and seller are uninformed, so a public signal affiliated with the buyer's private information reduces the buyer's informational rents in a second-degree price discrimination setting.

1.2 Overview

The next section develops the model of simultaneous sales and notation used throughout the paper, together with a time-line of the various choices made by firms, consumers and reviewer. Section 3 examines the firm's decision when prices are fixed but a reviewer can be chosen to maximize profits. Section 4 allows the firm the full freedom to choose optimal prices together with an optimal reviewer. Section 5 offers some conclusions. The appendix presents various proofs which are not necessary for understanding the main results in the paper.

2 The Model

Consider a group of $N \in \mathbb{N}_{++}$ consumers who act simultaneously deciding whether to purchase or not purchase some product. Without loss of generality, we normalize the number of customers to 1. The price of purchase is λ, which results in the net gain of $V - \lambda$ where V has prior probability $q = \frac{1}{2}$ of returning 1 or 0, depending on whether the product is of a high or low quality. The firm knows the quality of its product, while 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:*

$$\begin{aligned} \Pr[X_i = H \mid V = 1] &= \Pr[X_i = L \mid V = 0] = p \in \left(\tfrac{1}{2}, 1\right) \\ \Pr[X_i = H \mid V = 0] &= \Pr[X_i = L \mid V = 1] = 1 - p \in \left(0, \tfrac{1}{2}\right) \end{aligned}$$

Consumers update their beliefs using their private information, and purchase if $E[V] > \lambda$. We will assume that where the firm is able to choose price λ, indifferent consumers with $E[V] = \lambda$ purchase. This is without loss of generality, as if at a given price consumers who were indifferent purchased with probability less than one, the firm could always sell to all the indifferent consumers by reducing price by $\epsilon > 0$. However where the price is fixed, so a firm cannot make an ϵ-price reduction, we assume that indifferent consumers flip a coin.[2]

2.1 Reviewers

Before launching its product onto the market, the firm is able to have it publicly reviewed or tested.

We want to think of reviewers as having access to finer information because the firm allows the reviewer to test the product. The simplest way of modeling this is to allow the reviewer to get two signals about the product's quality, instead of just the one received by consumers.

We model the reviewer as making a decision $d \in \{E, F\}$ whether to endorse (E) or fail to endorse (F) 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

[2] Equivalently each consumer may be following a fixed purchase or do not purchase rule, so long as in expectation half of indifferent consumers purchase.

of a process through which even sophisticated reviews quickly get shortened to a binary decision through this process of dissemination.[3] In the case of a test, we can think of E being a pass, while F is a fail.

We consider a continuum of reviewer types who receive two i.i.d. draws from the same signal distribution as consumers. Reviewers endorse if they observe HH, endorse with probability $\phi \in [0,1]$ if they observe HL or LH and fail to endorse if they observe LL. The value of ϕ encapsulates the type of the reviewer. The lower the value of ϕ, the tougher the test that is imposed.[4]

Definition 2 *When $\phi = 1$ the reviewer is an "extreme optimist", who endorses having seen any combination of signals from $\{HH, HL, LH\}$. $\phi = \frac{1}{2}$ corresponds to an "unbiased reviewer" who flips a fair coin on observing a set of mixed signals. $\phi = 0$ corresponds to an "extreme pessimist" who only endorses when HH is observed. Reviewers with $\phi \in [0, \frac{1}{2})$ are termed "pessimistic", and those with $\phi \in (\frac{1}{2}, 1]$ "optimistic".*

We assume that reviewer types are common knowledge, perhaps generated through a known history of endorsement decisions, and that the firm is able to choose the reviewer type. Examples of situations in which the firm is able to pick the type of pre-launch reviewer include advance screenings of movies, the use of beta-testers in the software industry, early tastings in the wine industry, test drives by media outlets in the car industry, and sending free review copies of products to relevant journals, newspapers and magazines. In more general settings, doctoral students can choose the toughness of referees, political parties can selectively leak future policy ideas to the press and so on.

2.2 Pricing and Reviewer Choice

We consider two alternative frameworks, either:

1. The price is pre-set to be $\widehat{\lambda} = \frac{1}{2}$, which is the unconditional expectation of the product's value. This corresponds to the situation where firms have no direct control over price, and cannot or will not adjust price in response to a reviewer's observed decision.

2. The firm has complete control over the price, and chooses to set an optimal price as a function of both the type of reviewer it has chosen and that reviewer's endorsement decision.

[3] Modeling 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."

[4] Reviewers' bias may be general, or may be directed at a specific product or group of products. For example, a film critic may be known to be biased towards romantic comedies.

2.3 Timing Summary

The timing is complex but accords with what we might expect to see in practice in a variety of real-world markets and situations.

- Nature chooses product quality $V \in \{0, 1\}$, which also denotes the utility to consumers gross of price.
- Quality is observed by the firm, but not the consumers who have a 50:50 prior belief.
- The firm selects a publicly known reviewer type from a continuum of types defined by ϕ.
- The reviewer receives two independent signals, either high (H), or low (L), each of which is correct with probability $p > \frac{1}{2}$; and so sees $X_R \in \{HH, LH, HL, LL\}$.
- The reviewer decides whether to endorse:
  - If $X_R = HH$ then the product is endorsed.
  - If $X_R \in \{HL, LH\}$ then the reviewer endorses with known probability ϕ.
  - If $X_R = LL$ there is no endorsement.
- In the fixed price model the price is set at $\widehat{\lambda} = \frac{1}{2}$. In the flexible price model, conditional on the reviewer's type and decision the firm can choose a profit maximizing price.
- The consumers:
  - Receive a single independent signal $X_i \in \{H, L\}$ also correct with probability $p > \frac{1}{2}$;
  - Observe the reviewer's type and endorsement decision;
  - And, in the flexible price model observe the chosen price λ, while in the fixed price model they observe the price of $\widehat{\lambda} = \frac{1}{2}$.
- Consumers then decide whether to purchase or not simultaneously.

2.4 Restriction to High Quality Products

Throughout, for conciseness, we consider just the firm with a good product. By standard signaling considerations, a bad product firm will be forced to copy the choice of the good product firm to avoid immediately revealing its type and so get zero sales, so we do not analyze that case. A separating equilibrium in which the good product firm makes strictly positive profit is not possible, as the bad product firm would copy the choice of the good product firm, and so be believed to be good and make the same profit as the good product firm (selling to all consumers). Thus, we restrict attention to pooling equilibria in which the bad product firm is forced to follow the good product firm's preference. Such equilibria can always be supported by the belief that any firm which deviates from the good product firm's preferences is selling a low quality product. Note that in such pooling equilibria, consumers will be unable to infer anything about product quality from the firm's choice of reviewer or price, and hence the statement in the introduction that this *cannot* reasonably be considered a signaling paper.

3 Fixed Prices

In this section, we consider the case in which the firm is restricted to setting a price $\widehat{\lambda} = \frac{1}{2}$. We first show how consumers' purchasing choices are affected by their post-review beliefs, moving on to show how different reviews affect these beliefs. Finally, we determine the optimal reviewer choice for the firm.

3.1 Partitioning Consumer Decisions

Let us consider consumer decisions for a given price $\widehat{\lambda} = \frac{1}{2}$, after a reviewer has made a public decision. This will allow us to demonstrate how consumers' choices can be easily partitioned according to the impact of the reviewer's decision on beliefs. We can think of the decision by the reviewer as updating the prior from $q = \frac{1}{2}$ (where consumers have no preconceptions) to $q = q^*$ where q^* is the probability that the product quality is high. Alternatively think of q^* directly as a general prior determined explicitly through the review process. We start by making the following Remark, which is used implicitly throughout.

Remark 1

$$\frac{\Pr\left[V=1|X_i\right]}{\Pr\left[V=0|X_i\right]} = \frac{\frac{\Pr[X_i|V=1]\Pr[V=1]}{\Pr[X_i]}}{\frac{\Pr[X_i|V=0]\Pr[V=0]}{\Pr[X_i]}} = \frac{\Pr\left[X_i|V=1\right]q^*}{\Pr\left[X_i|V=0\right](1-q^*)}$$

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 signal they have observed if the product were good to the probability if the product were bad, suitably weighted by the updated prior. Using this Remark, we can now determine how customers will behave for any possible updated prior q^* following the reviewer's decision.

Where $q^* > p$, all consumers will purchase. A consumer i is least likely to purchase if $X_i = L$. Taking this case, we have an odds ratio $\frac{\Pr[V=1|L]}{\Pr[V=0|L]} = \frac{(1-p)q^*}{p(1-q^*)} = \frac{q^*-pq^*}{p-pq^*} > 1$ since $q^* > p$. A symmetrical argument shows that where $q^* < 1-p$, no consumers will purchase. Where $q^* \in (1-p, p)$, following a H signal $\frac{\Pr[V=1|H]}{\Pr[V=0|H]} = \frac{pq^*}{(1-p)(1-q^*)} > 1$ as $q^* > (1-p)$, so the customer purchases. Following a L signal $\frac{\Pr[V=1|L]}{\Pr[V=0|L]} = \frac{(1-p)q^*}{p(1-q^*)} < 1$ as $q^* < p$, so the customer does not purchase. Where $q^* = p$, following a L signal, $\frac{\Pr[V=1|L]}{\Pr[V=0|L]} = \frac{(1-p)q^*}{p(1-q^*)} = 1$ as $q^* = p$, so the consumer is indifferent and flips a coin. A H signal is more positive, so the consumer purchases. By symmetry, where $q^* = 1-p$, following a H signal the consumer flips a coin, while following a L signal the consumer does not purchase. The following Lemma summarizes this information.

Lemma 1 *Consumers will respond to an updated prior as follows: (a) if $q^* > p$ then consumer i will buy; (b) if $q^* = p$, following a H signal the consumer buys, while following a L signal, the customer flips a coin; (c) if $q^* \in (1-p, p)$ then consumer i will buy if and only if $X_i = H$; (d) if $q^* = 1-p$, following a H signal the consumer flips a coin, while following a L signal the consumer will not buy; (e) if $q^* < 1-p$ then consumer i will not buy.*

This lemma is crucial to our understanding of the impact of endorsement decisions by different reviewer types. As we will see, by selecting a particular type of reviewer, the firm can effectively select from a menu of possible ranges of the updated prior.

3.2 Impact of Reviews on Beliefs

Here, we determine how different reviews impact on consumers' beliefs. We begin by finding the updated prior faced by consumers in the event of endorsement (denoted by q_E^*) and a failure to endorse (denoted by q_F^*). If we observe endorsement, then:

$$
\begin{aligned}
q_E^* &\equiv \Pr[V=1 \mid E] \\
&= \frac{\left[p^2+2p(1-p)\phi\right]}{[p^2+2p(1-p)\phi]+\left[(1-p)^2+2(1-p)p\phi\right]} = \frac{p^2+2p(1-p)\phi}{p^2+(1-p)^2+4p(1-p)\phi}
\end{aligned}
$$

While with no endorsement, we have:

$$
\begin{aligned}
q_F^* &\equiv \Pr[V=1 \mid F] \\
&= \frac{\left[(1-p)^2+2p(1-p)(1-\phi)\right]}{\left[(1-p)^2+2p(1-p)(1-\phi)\right]+[p^2+2(1-p)p(1-\phi)]} = \frac{(1-p)^2+2p(1-p)(1-\phi)}{(1-p)^2+p^2+4p(1-p)(1-\phi)}
\end{aligned}
$$

Next we note some properties of these expressions. Firstly, note that q_E^* is decreasing in ϕ.

$$\begin{aligned}\frac{dq_E^*}{d\phi} &= \frac{2p(1-p)\left[p^2+(1-p)^2+4p(1-p)\phi\right]-4p(1-p)\left[p^2+2p(1-p)\phi\right]}{\left[p^2+(1-p)^2+4p(1-p)\phi\right]^2} \\ &= \frac{2p(1-p)(1-2p)}{\left[p^2+(1-p)^2+4p(1-p)\phi\right]^2} < 0\end{aligned}$$

since for $p > \frac{1}{2}$, the denominator is always strictly positive, $2p(1-p) > 0$, but $(1-2p) < 0$. Next note that q_F^* is also decreasing in ϕ.

$$\begin{aligned}\frac{dq_F^*}{d\phi} &= \frac{-2p(1-p)\left[(1-p)^2+p^2+4p(1-p)(1-\phi)\right]+4p(1-p)\left[(1-p)^2+2p(1-p)(1-\phi)\right]}{\left[(1-p)^2+p^2+4p(1-p)(1-\phi)\right]^2} \\ &= \frac{2p(1-p)(1-2p)}{\left[(1-p)^2+p^2+4p(1-p)(1-\phi)\right]^2} < 0\end{aligned}$$

similarly to the q_E^* case. That these updated priors should be decreasing in ϕ is perfectly natural: an endorsement is better news the tougher the reviewer, while a failure to endorse is not such bad news.

With an unbiased reviewer, so setting $\phi = \frac{1}{2}$, we have:

$$q_E^*\left(\phi = \frac{1}{2}\right) = \frac{p^2+p(1-p)}{p^2+(1-p)^2+2p(1-p)} = p \qquad (1)$$

$$q_F^*\left(\phi = \frac{1}{2}\right) = \frac{(1-p)^2+p(1-p)}{(1-p)^2+p^2+2p(1-p)} = 1-p \qquad (2)$$

As q_E^* and q_F^* are both strictly decreasing in ϕ, this immediately implies that $q_E^*\left(\phi < \frac{1}{2}\right) > p$, $q_E^*\left(\phi > \frac{1}{2}\right) < p$, $q_F^*\left(\phi < \frac{1}{2}\right) > 1-p$ and $q_F^*\left(\phi > \frac{1}{2}\right) < 1-p$. Furthermore, $q_E^*(\phi = 1) > \frac{1}{2}$ and $q_F^*(\phi = 0) < \frac{1}{2}$:

$$\begin{aligned}q_E^*(\phi = 1) &= \frac{p^2+2p(1-p)}{p^2+(1-p)^2+4p(1-p)} > \frac{1}{2} \Leftrightarrow p^2 > (1-p)^2 \\ q_F^*(\phi = 0) &= \frac{(1-p)^2+2p(1-p)}{(1-p)^2+p^2+4p(1-p)} < \frac{1}{2} \Leftrightarrow (1-p)^2 < p^2\end{aligned}$$

Using the fact that q_E^* and q_F^* are both strictly decreasing in ϕ, it follows that $q_E^* > 1-p$ and $q_F^* < p$ $\forall\phi$. All this information is summarized in the following Lemma:

Lemma 2 *(a)* $q_E^*\left(\phi < \frac{1}{2}\right) > p$; *(b)* $q_E^*\left(\phi = \frac{1}{2}\right) = p$; *(c)* $q_E^*\left(\phi > \frac{1}{2}\right) \in (1-p, p)$ *and* $q_F^*\left(\phi < \frac{1}{2}\right) \in (1-p, p)$; *(d)* $q_F^*\left(\phi = \frac{1}{2}\right) = 1-p$; *(e)* $q_F^*\left(\phi > \frac{1}{2}\right) < 1-p$.

3.3 Choice of Reviewer

With no choice of price, the firm's decision problem is reduced to the choice of reviewer ϕ drawn from the continuum of types, $\phi \in [0,1]$. We can now determine the expected impact of different reviewer types on profits, starting with pessimists. Using part (a) of Lemmas 1 and 2, all consumers will purchase if any form of pessimist endorses. From part (c) of the same, a pessimist will return expected sales of p, even having failed to endorse (recalling that we have normalized the number of customers to 1).[5] Therefore, profits under a pessimist are given by:

$$\begin{aligned}\Pi\left[\phi < \frac{1}{2}\right] &= \widehat{\lambda}\left\{\Pr[E] + p\Pr[F]\right\} \\ &= \tfrac{1}{2}\left\{\left[p^2 + 2p(1-p)\phi\right] + p\left[(1-p)^2 + 2p(1-p)(1-\phi)\right]\right\} \\ &= \tfrac{1}{2}\left\{\left[2p(1-p)\phi\right](1-p) + p^2 + p(1-p)^2 + 2p^2(1-p)\right\} \end{aligned} \quad (3)$$

As $\frac{d\Pi}{d\phi} > 0$ to maximize $\Pi\left[\phi < \frac{1}{2}\right]$ we select ϕ arbitrarily close to $\frac{1}{2}$ to get profits arbitrarily close to:

$$\tfrac{1}{2}\left[2p(1-p)^2 + p^2 + 2p^2(1-p)\right] = \tfrac{1}{2}\left[p + p(1-p)\right]$$

In the case of an unbiased reviewer, from part (b) of Lemmas 1 and 2 endorsement returns expected sales of $p + \frac{1}{2}(1-p)$, while from part (d) of the same a failure to endorse returns expected sales of $\frac{1}{2}p$. Overall expected profits are therefore

$$\begin{aligned}\Pi\left[\phi = \frac{1}{2}\right] &= \tfrac{1}{2}\left\{\left[p + \tfrac{1}{2}(1-p)\right]\Pr[E] + \tfrac{1}{2}p\Pr[F]\right\} \\ &= \tfrac{1}{2}\left\{\left[p + \tfrac{1}{2}(1-p)\right]p + p(1-p)\right\} = \tfrac{1}{2}p \end{aligned} \quad (4)$$

Note that the unbiased reviewer acts just like a reviewer who receives just one signal and endorses if and only if the signal is high. In either case, $q_E^* = \Pr[E] = p$ and $q_F^* = \Pr[F] = 1-p$.

To complete our comparison, in the case of an optimistic reviewer, from part (c) of Lemmas 1 and 2 endorsement leads to expected sales of p, while from part (e) of the same a failure to endorse results in zero sales. Thus overall profits are

$$\Pi\left[\phi > \frac{1}{2}\right] = \tfrac{1}{2}\left\{p\Pr[E] + [0]\Pr[F]\right\} = \tfrac{1}{2}p\left[p^2 + 2p(1-p)\phi\right] \quad (5)$$

[5] By sales we mean the number of units sold, not revenue. Profits of course equal revenue since there are no costs. In this section profits equal half sales since the price is fixed at one half.

Since $\frac{d\Pi}{d\phi} > 0$ to maximize $\Pi\left[\phi > \frac{1}{2}\right]$ we must select the highest value of ϕ available:

$$\begin{aligned} \max_{\phi \leq 1} \Pi\left[\phi > \frac{1}{2}\right] &= \Pi\left[\phi = 1\right] \\ &= \tfrac{1}{2}p\left[p^2 + 2p\left(1-p\right)\right] = \tfrac{1}{2}p^2(2-p) \end{aligned} \tag{6}$$

As $p + p(1-p) > p$ the optimal pessimist provides strictly higher profits than the unbiased reviewer. In turn, the unbiased reviewer provides strictly higher profits than the optimal optimist, as $p > p^2(2-p) \Leftrightarrow (1-p)^2 > 0$. Thus the pessimistic reviewer arbitrarily close to the unbiased reviewer maximizes profits.

Proposition 1 *For a fixed price $\widehat{\lambda} = \frac{1}{2}$, the optimal choice of reviewer for a good product firm is a pessimist that is arbitrarily close to the unbiased reviewer type, i.e., with ϕ smaller than but arbitrarily close to $\frac{1}{2}$.*

The increase in sales from using the almost unbiased reviewer is $p(1-p)$, which is substantial for most values of p. For p close to $\frac{1}{2}$, the optimal choice of reviewer increases sales by close to 25% of the total stock of consumers. $p(1-p)$ is decreasing in p, but the optimal reviewer increases sales by at least 10% of the total stock of consumers $\forall p < 0.887$, and by 5% $\forall p < 0.947$. The percentage increase in sales, $\frac{p(1-p)}{p} = 1-p$, starts close to 50% for p close to $\frac{1}{2}$, and decreases linearly in p.

By availing itself of the full continuum of types, the firm is able to select a pessimistic reviewer type that is arbitrarily close to the unbiased reviewer, but with just enough pessimism to convince all consumers to disregard their private signals and purchase on endorsement, and to follow their private signals and effectively ignore the pessimist's decision if he fails to endorse. This minimizes the downside of pessimists, by raising the probability of endorsement as much as possible without the reviewer becoming unbiased, while making use of the huge advantage of pessimists: that they can greatly improve sales but can never harm sales when prices are fixed. All forms of optimists have the opposite characteristic: they can harm sales, but cannot improve them, as consumers are expecting to see them endorse. Unlike the optimist, the unbiased reviewer may result in higher sales if an endorsement occurs, but unlike the pessimist, may result in lower sales if there is no endorsement, and so stands somewhere in the middle. The subtle choice of a very mild pessimist, almost identical to the unbiased type, presents the firm with a discontinuous jump in profits of $\frac{1}{2}p(1-p)$.

Figure 1 below summarizes the impact of different reviews on sales, highlighting the role of the updated priors: even an endorsement decision by an optimist will only produce a prior which is as high as that created when a pessimist fails to endorse! In both cases a prior is generated in the $(1-p, p)$ region, within which expected sales are constant as consumers then all purchase if and only if they receive a H signal.[6] The diagram makes clear the preference for an almost unbiased pessimist over a truly unbiased reviewer: the probability of endorsement is almost exactly the same, but sales under either an endorsement or a failure jump up discontinuously.

Figure 1

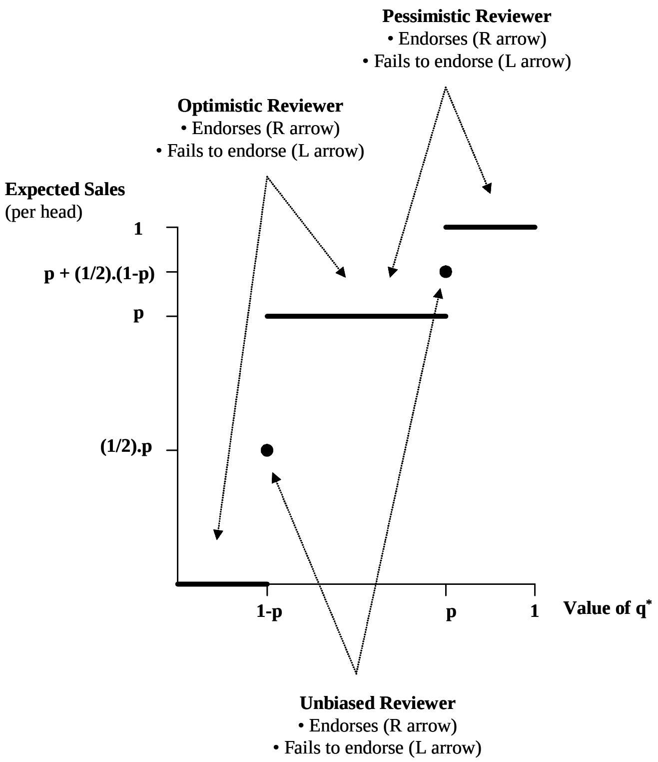

A reasonable question we might ask is whether the result is driven by the fact that the reviewer is close to unbiased (and hence as likely to endorse as a pessimist can be) or the fact that the reviewer is pessimistic. We can investigate this in two ways. Most importantly we move to a flexible price setting in the next section, to generalize the model and to check robustness to an extra dimension of freedom in the firm's decision. There we find that price flexibility always entails a preference for *extreme* bias.

Secondly, we might ask how a firm would respond to a reduced range of choice of reviewer. In particular we can restrict the firm to the unbiased reviewer and biased types that are a discrete

[6] In Gill and Sgroi (2003a), which considers consumers deciding in sequence after the public review, a failure to endorse by a pessimist, which leaves $q^* < 0.5$, is worse than endorsement by an optimist, which leaves $q^* > 0.5$. For simultaneous sales this is not an issue since all that matters is that in both cases $q^* \in (1-p, p)$. It seems that in a sequential sales model, learning from other consumers' actions partially offsets the benefits of a pessimist, though a pessimist remains better overall for expected sales for a good product firm.

distance from the unbiased type. To do this, consider any choice of three types $\phi \in \left\{\phi_P, \frac{1}{2}, \phi_O\right\}$ where $\phi_P \in [0, \frac{1}{2})$ and $\phi_O \in (\frac{1}{2}, 1]$. We know already from earlier in the section that $\Pi[\phi = 1] < \Pi\left[\phi = \frac{1}{2}\right]$, and furthermore that $\Pi[\phi = 1] \geq \Pi[\phi_O]$. Thus, $\Pi[\phi_O] < \Pi\left[\phi = \frac{1}{2}\right]$. Since we also know that $\Pi\left[\phi < \frac{1}{2}\right]$ is strictly increasing in ϕ it must be that $\Pi[\phi_P] \geq \Pi[\phi = 0]$. From (3), $\Pi[\phi = 0] = \frac{1}{2}\left[p^2 + p(1-p)^2 + 2p^2(1-p)\right]$. From (4), $\Pi\left[\phi = \frac{1}{2}\right] = \frac{1}{2}p$. Thus:

$$\begin{aligned} \Pi[\phi = 0] &> \Pi\left[\phi = \frac{1}{2}\right] \Leftrightarrow p + (1-p)^2 + 2p(1-p) > 1 \Leftrightarrow \\ p + 1 - 2p + p^2 + 2p - 2p^2 &> 1 \Leftrightarrow p(1-p) > 0 \end{aligned}$$

Hence $\Pi[\phi = 0] > \Pi\left[\phi = \frac{1}{2}\right]$. Therefore, we have found that $\Pi[\phi_P] \geq \Pi[\phi = 0] > \Pi\left[\phi = \frac{1}{2}\right] > \Pi[\phi_O]$, proving that for any choice of such types, the pessimist is always strictly preferred:

Proposition 2 *For any choice of three reviewer types $\phi \in \left\{\phi_P, \frac{1}{2}, \phi_O\right\}$ such that $\phi_P \in [0, \frac{1}{2})$ and $\phi_O \in (\frac{1}{2}, 1]$, the good product firm strictly prefers the pessimistic type ϕ_P to the unbiased type, who in turn is strictly preferred to the optimistic type ϕ_O.*

So far, we have assumed that the firm must pick a reviewer. If it chose no reviewer at all, it would sell to all consumers with high signals, receiving expected profits of $\frac{1}{2}p$. From (4), we see that this is just as for the choice of the unbiased reviewer.

Proposition 3 *For a fixed price $\widehat{\lambda} = \frac{1}{2}$, the choice of no reviewer is equivalent to the choice of an unbiased reviewer type $\phi = \frac{1}{2}$, and hence is not optimal.*

We saw above that the unbiased reviewer acts just like a one signal reviewer. Thus, consumers' private unbiased private signals add no value,[7] and so in effect consumers see just one signal, explaining the equivalence. As we will see, this equivalence no longer holds once the firm is able to vary price in response to the review.

4 Flexible Prices

With fixed prices we know that a pessimist is favored by the firm, but this pessimist must be chosen as close as possible to the unbiased type. We next consider what happens when the firm is be able to change the price of the product after the reviewer's decision is made public. The

[7] If my private unbiased signal confirms the reviewer's first signal, I act as if I had seen no second signal. If the second signal contradicts the first, I am indifferent and so may as well follow the first signal. This explains the preference for signals biased in favor of initial predispositions, as in Meyer (1991).

optimal price will then be conditional on the reviewer's decision and type and on the product quality.

Let $\mu_{d,X_i} = \Pr[V=1|d, X_i]$ where the reviewer's decision $d \in \{E, F\}$ and the consumer's signal $X_i \in \{H, L\}$. Using Bayes' Rule and the fact that the consumers have 50:50 initial priors over the good's quality:

$$\mu_{d,X_i} = \frac{\Pr[d, X_i|V=1]}{\Pr[d, X_i|V=1] + \Pr[d, X_i|V=0]}$$

After the reviewer's decision, the firm will choose to set price $\lambda = \mu_{d,H}$ to sell to all those who received high private signals or $\lambda = \mu_{d,L}$ to sell to all consumers including those with low private signals. Any other price would be sub-optimal. With $\lambda > \mu_{d,H}$ no consumers will buy, with $\lambda \in (\mu_{d,L}, \mu_{d,H})$ only the high signal consumers will buy at a price lower than $\mu_{d,H}$, and with $\lambda < \mu_{d,L}$ everybody buys at a price lower than $\mu_{d,L}$.

Given we have normalized the number of customers to 1, expected profits as a function of these prices are $\pi_d(\mu_{d,H}) = p\mu_{d,H}$ and $\pi_d(\mu_{d,L}) = \mu_{d,L}$. We define $\pi_d^* = \max\{p\mu_{d,H}, \mu_{d,L}\}$ to be the maximum expected profits achievable given the reviewer's decision, which will be a function of p and ϕ. Then the firm's optimization problem reduces to maximizing *ex ante* expected profits Π:

$$\max_{\phi\in[0,1]} \Pi = [\Pr[E].\pi_E^* + \Pr[F].\pi_F^*]$$

We begin by showing that Π is strictly convex in ϕ, which implies that the firm will always choose an extreme type, choosing either an extreme pessimist with $\phi = 0$ or an extreme optimist with $\phi = 1$. We then move on to show which extreme is best over the range of p. Let

$$\begin{aligned}
\Pi_1 &= \Pr[E]\,p\mu_{E,H} + \Pr[F]\,p\mu_{F,H} \\
\Pi_2 &= \Pr[E]\,\mu_{E,L} + \Pr[F]\,\mu_{F,L} \\
\Pi_3 &= \Pr[E]\,\mu_{E,L} + \Pr[F]\,p\mu_{F,H} \\
\Pi_4 &= \Pr[E]\,p\mu_{E,H} + \Pr[F]\,\mu_{F,L}
\end{aligned}$$

First, we note the following crucial Lemma.

Lemma 3 $\frac{d^2(\Pr[d]p\mu_{d,H})}{d\phi^2} > 0$ *and* $\frac{d^2(\Pr[d]\mu_{d,L})}{d\phi^2} > 0$ *for* $d \in \{E, F\}$, *so* $\frac{d^2\Pi_i}{d\phi^2} > 0$ *for* $i \in \{1, 2, 3, 4\}$.

Proof. See Appendix. ■

Now,

$$\begin{aligned}
\Pi &= [\Pr[E].\pi_E^* + \Pr[F].\pi_F^*] = \Pr[E]\max\left\{p\mu_{E,H}, \mu_{E,L}\right\} + \Pr[F]\max\left\{p\mu_{F,H}, \mu_{F,L}\right\} \\
&= \max\left\{\Pr[E]\,p\mu_{E,H}, \Pr[E]\,\mu_{E,L}\right\} + \max\left\{\Pr[F]\,p\mu_{F,H}, \Pr[F]\,\mu_{F,L}\right\} \\
&= \max_{i\in\{1,2,3,4\}}\{\Pi_i\}
\end{aligned}$$

as to find the maximum of the sum of the terms in Π_i, we just add the maximum of the first term to the maximum of the second term. By Lemma 3, Π_i are all strictly convex in ϕ. Consider any interior value $\widehat{\phi}$. By the convexity of Π_i $\forall i$, moving along the highest of the Π_i at $\widehat{\phi}$ towards either $\phi = 1$ or $\phi = 0$, or possibly both, will always increase profits. Therefore $\widehat{\phi}$ cannot maximize Π, giving the following proposition.[8]

Proposition 4 *With flexible prices, the good product firm always strictly prefers either the extreme pessimist ($\phi = 0$) or the extreme optimist ($\phi = 1$), or both, to any intermediate reviewer with $\phi \in (0, 1)$.*

To provide an intuition for this proposition, first consider that all prices $\{\mu_{E,H}, \mu_{E,L}, \mu_{F,H}, \mu_{F,L}\}$ are decreasing in ϕ. So, an endorsement raises beliefs more the tougher the reviewer, while a failure to endorse is not so damaging to beliefs if the reviewer is tougher. Furthermore, $\mu_{E,H}$ and $\mu_{E,L}$ can be shown to be strictly convex in ϕ, so conditional on endorsement a decrease in ϕ is more powerful where the reviewer is tougher. This can be seen by noting that the relative probability of the reviewer having seen two good signals to having seen mixed signals is $\frac{p^2}{2p(1-p)\phi}$, which is convex in ϕ.[9] On the other hand, $\Pr[E]$ is linearly increasing in ϕ. Thus, to maximize $\Pr[E]p\mu_{E,H}$ or $\Pr[E]\mu_{E,L}$, the firm will wish to choose an extreme ϕ, either setting $\phi = 0$ to benefit from a steep increase in price while the probability of endorsement falls linearly, or setting $\phi = 1$ to benefit from a linear increase in the probability of endorsement, while price does not fall much towards the end.

To maximize $\Pr[F]p\mu_{F,H}$ or $\Pr[F]\mu_{F,L}$, the firm will always want to set $\phi = 0$, as both $\Pr[F]$ and the prices are decreasing in ϕ. The proof of Lemma 3 shows that $\Pr[F]p\mu_{F,H}$ and $\Pr[F]\mu_{F,L}$

[8] Alternatively, we could prove the proposition by using the fact that the maximum of a set of strictly convex functions all defined over the same range must itself be strictly convex, so Π must also be strictly convex, and hence Proposition 4 follows. See Rockafellar (1970).

[9] As $\Pr[HH|E] = \frac{\Pr[E|HH]\Pr[HH]}{\Pr[E]} = \frac{p^2}{\Pr[E]}$ and
$\Pr[\{HL \text{ or } LH\}|E] = \frac{\Pr[E|\{HL \text{ or } LH\}]\Pr[\{HL \text{ or } LH\}]}{\Pr[E]} = \frac{2p(1-p)\phi}{\Pr[E]}$.

are also convex. Thus, summing over the case of endorsement and failure to endorse, the firm always prefers an extreme ϕ. The firm will either choose the extreme pessimist to maximize prices, at the cost of a lower probability of endorsement, or the extreme optimist to maximize the probability of endorsement, at the cost of lower prices.

Next, we show which extreme is preferred for different values of p. We start by finding π_d^* for $\phi = 1$ and $\phi = 0$.

Lemma 4 *(i) Following a failure to endorse, and with* $\phi = 1$, $\mu_{F,L}$ *is superior for* $p \in (\frac{1}{2}, p_1)$ *and* $p\mu_{F,H}$ *is superior for* $p \in [p_1, 1)$, *where* $p_1 \simeq 0.594$.

(ii) Following a failure to endorse, and with $\phi = 0$, $\mu_{F,L}$ *is superior for* $p \in (\frac{1}{2}, p_2)$ *and* $p\mu_{F,H}$ *is superior for* $p \in [p_2, 1)$, *where* $p_2 \simeq 0.607$.

(iii) Following endorsement, and with $\phi = 1$, $\mu_{E,L}$ *is superior for* $p \in (\frac{1}{2}, p_3)$ *and* $p\mu_{E,H}$ *is superior for* $p \in [p_3, 1)$, *where* $p_3 \simeq 0.635$.

(iv) Following endorsement, and with $\phi = 0$, $\mu_{E,L}$ *is superior* $\forall p \in (\frac{1}{2}, 1)$.

Proof. See Appendix. ■

The negative effect of rejection is stronger for higher p, and as p goes up, the proportion of consumers with positive private signals increases. Thus, for $p < p_1$ the firm sells to everybody following a rejection, as the failure to endorse is not too damaging and relatively few consumers receive high private signals, but for $p \geq p_2$, the firm focuses on just those with high signals. In a small range $p \in [p_1, p_2)$, the degree of bias affects the choice: the firm sells to everybody if and only if the reviewer is pessimistic, as rejection then sends a weaker signal. The positive signal arising from an endorsement by an extreme pessimist is strong enough that the firm wants to sell to everybody for all p. Endorsement by an extreme optimist sends a much weaker signal, so for $p \geq p_3$ the proportion of high signal consumers and the price that the firm can charge them are high enough that the firm sells just to them. We can now determine whether the extreme pessimist or optimist is superior for any given p.

Proposition 5 *With flexible prices, the good product firm prefers the extreme pessimist for* $p \in \left(\frac{1}{2}, p_4\right)$, *while the extreme optimist is preferred for* $p \in [p_4, 1)$, *where* $p_4 = -1 + \sqrt{3} \simeq 0.732$. *The preferences are strict, except at* p_2 *and* p_4, *where the firm is indifferent between the two extremes.*

Proof. See Appendix. ■

The extreme pessimist is best for lower values of p while for higher values the extreme optimist is best. This is by no means obvious given the complex interaction of various functions of p, but with the assistance of some diagrams the optimal choice of extremist becomes apparent. In Figure 2, the two continuously increasing lines show $\Pr[E].\pi^*_E$ evaluated over the range of p, the higher line for $\phi = 1$ and the lower line for $\phi = 0$. We can see that $\Pr[E].\pi^*_E$ is always higher with the extreme optimist: the higher probability of endorsement overwhelms the lower price. The two decreasing lines show $\Pr[F].\pi^*_F$, the higher line for $\phi = 0$ and the lower line for $\phi = 1$. In contrast to the endorsement case, $\Pr[F].\pi^*_F$ is always higher under the extreme pessimist, as both prices and $\Pr[F]$ are higher. Note that the kinks in the curves represent values of p at which the firm moves from setting a low price to sell to everybody, to a high price to sell only to those who receive high private signals.

Figure 2

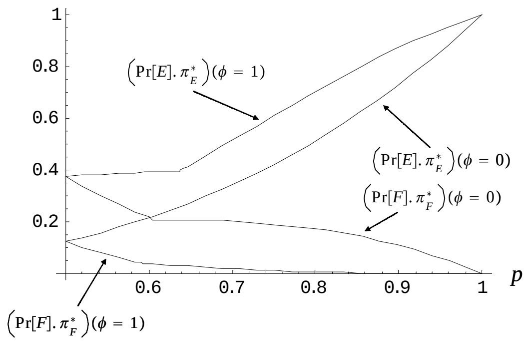

Figure 3

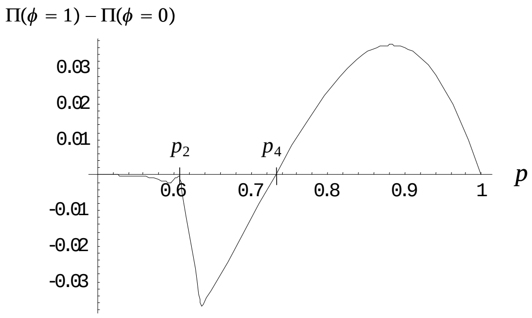

Given p, the firm wants to maximize $\Pr[E].\pi^*_E + \Pr[F].\pi^*_F$ by setting either $\phi = 1$ or $\phi = 0$. Thus it has to trade-off the divergent preferences in the endorsement and failure to endorse cases. For high enough p, the increase in profits in the endorsement case from choosing the optimist outweigh the decrease in the failure case. For low p, the firm prefers the extreme pessimist who allows a large increase in price following endorsement, but who does not reduce price too much following a failure to endorse. For high p, the firm prefers the extreme optimist as the probability of the extreme optimist endorsing is very high with high p, while a catastrophic failure to endorse becomes very unlikely. Figure 3 shows overall profits from the extreme optimist less those from the extreme pessimist over the range of p. As can be seen, choosing the better extreme increases profits of the good product firm by up to almost 4% of the maximum feasible profit.

The firm could of course opt not to have its product reviewed. While in Section 3 we showed that in the fixed price case no review is equivalent to the choice of an unbiased reviewer type, with flexible prices that is not the case. As the firm can change price in response to the unbiased reviewer's decision, we would expect a difference to emerge, and indeed it does: the choice of no reviewer is always worse than an unbiased reviewer in the flexible price case.

Proposition 6 *With flexible prices, the choice of no reviewer is strictly worse than the unbiased reviewer type $\phi = \frac{1}{2}$ and a fortiori strictly worse than the optimal reviewer.*

Proof. See Appendix. ■

The next two figures illustrate the magnitude of the optimal reviewer's superiority in the flexible price case. Choosing the optimal reviewer increases profits dramatically. Figure 4 shows that the optimal reviewer increases profits of the good product by up to 8% over the unbiased reviewer. The optimal reviewer increases profits by up to 15% over no review, as illustrated in Figure 5.

Figure 4: $\frac{\left(\max_{\phi\in[0,1]} \Pi\right) - \Pi\left(\phi=\frac{1}{2}\right)}{\Pi\left(\phi=\frac{1}{2}\right)}$

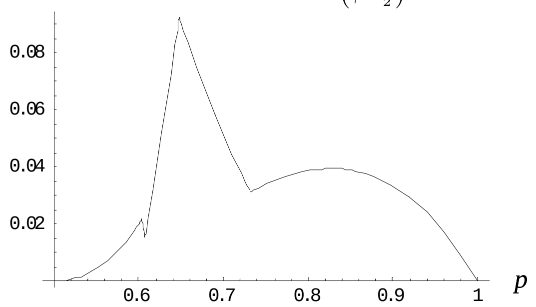

Figure 5: $\frac{\left(\max_{\phi\in[0,1]} \Pi\right) - \Pi(\text{No Reviewer})}{\Pi(\text{No Reviewer})}$

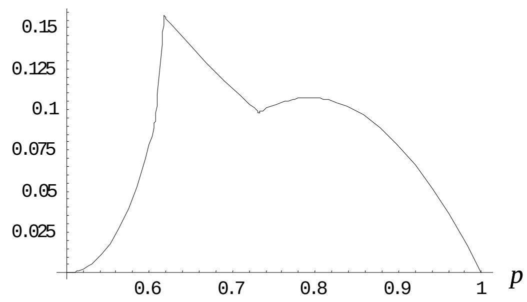

5 Conclusions

The results in this paper provide an integration of two key choices for the producer of any new product whose quality is uncertain: the choice of an initial price, and the early product marketing strategy. Product testing and review is essential in many industries: a glance at the shelves of any news-stand will indicate how many magazines and journals subsist solely to provide reviews of new products ranging from books, films, video games, computers, electrical goods, cars, clothes, food, wines and restaurants. In some cases prices will rise or fall as a result of the decision of reviewers, in others prices are standardized, but often the firm will have a choice of where to send its product for review. We can also think in terms of a test for a new product, which might be the achievement of a safety standard, hallmark or known standard of excellence, and the analysis in this paper is equally applicable when new products face tests rather than reviewers.

The paper has shown that quite apart from any standard signaling arguments, firms with products of quality unknown to consumers will tend towards sending their products to extremists for assessment, whether this be reviewers with known biases or tests that are publicly known to be extremely tough or soft. We have also seen that sending a product to review rather than avoiding a reviewer altogether is sensible. After the test result or reviewer's endorsement decision is known, the firm can then select an optimal price, and consumers will purchase or not based on the price, public review and their own private information. Considering the case where a product is good, quite remarkably when private information and the reviewer's own signals are informative but not too revealing, it is the extreme pessimist which will best raise profits for the firm. By enabling the firm to set a very high price following endorsement, the tough reviewer increases expected profits while keeping the impact of a failure to endorse to a minimum. As private signals become more revealing the extreme optimist becomes the better choice, as the likelihood of an endorsement rises, while the firm is less likely to suffer the hugely damaging impact of an optimist failing to endorse. We have also seen that when prices cannot be changed, the best choice of reviewer must be a pessimist, but one that is as close to the unbiased type as is available. The bad product firm is always forced to copy the good product firm's choice of reviewer to avoid immediately revealing its type.

To give some examples, we might consider a new computer where price will almost certainly reflect which awards or endorsements it may receive. Here it makes sense to evaluate the quality of information about the product, and then send the product to the known toughest or softest reviewer as appropriate. When a new car is marketed in one country, we might ask how far the firm should go in meeting minimum safety requirements by passing the easiest test available, or facing the toughest. Alternatively, we might ask how tough a referee a Ph.D. student should choose to support a job application or how tough an interviewer a politician should pick: here perhaps the best choice is to chose the known tough type who is closest to unbiased, since it is not clear how "prices" might be altered in the light of the interview or reference.

The findings in the paper might explain the survival of reviewers with well-known harsh styles, biases and critical approaches. While we might think firms likely to avoid such reviewers the results in this paper show that while these tough reviewers are likely to not offer any endorsement, the tremendous gains when they do endorse might actually make them popular with firms overall. On the other hand, the proliferation of very soft review journals, which may be supported by the very firms they seek to assess, can also be explained when the quality of information available to consumers and reviewers is very high. Where there is no real price variable, or prices are generally held fixed in the light of reviews, we also have a role for review types closer to unbiased. In each case our choice of optimal reviewer is sharp, and may seem strongly counter-intuitive, but our findings support the survival of extremists in the market-place.

In order to focus on the heart of the issue, the modeling assumptions have been chosen to make the analysis as tractable as possible. However, interesting extensions might look at the case when the firm has generalized beliefs 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. We might also explicitly consider how different types of reviewer evolve, or how consumers learn in more complex settings.

References

Albano, G. L. and Lizzeri, A. (2001) 'Strategic Certification and Provision of Quality', *International Economic Review*, **42**, 267-283.

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.

Caillaud, B. and Tirole, J. (1999) 'Party Governance and Ideological Bias', *European Economic Review*, **43**, 779-789.

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.

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

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. (2003a) 'Product Launches with Biased Reviewers: The Importance of Not Being Earnest', *Cambridge Working Papers in Economics* 0334, Cambridge University.

Gill, D. and Sgroi, D. (2003b) 'The Superiority of Tough Reviewers in a Model of Simultaneous Sales', *Cambridge Working Papers in Economics* 0335, Cambridge University.

Lerner, J. and Tirole, J. (2004) 'A Model of Forum Shopping, with Special Reference to Standard Setting Organizations', *Harvard NOM Research Paper No. 04-31.*

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.

Ottaviani, M. and Pratt, A. (2001) 'The Value of Public Information in Monopoly', *Econometrica*, **69**, 1673-1683.

Robinson, J. (1999), *The Oxford Companion to Wine (2nd ed.),* Oxford University Press.

Rockafellar, R. T. (1970) *Convex Analysis*, Princeton University Press.

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.

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.

Appendix

Proof of Lemma 3. First, note that:

$$\begin{aligned} p\mu_{E,H} &= p.\frac{\left[p^2 + 2p(1-p)\phi\right] p}{\left[p^2 + 2p(1-p)\phi\right] p + \left[(1-p)^2 + 2p(1-p)\phi\right](1-p)} \\ &= \frac{p^2\left[p^2 + 2p(1-p)\phi\right]}{p^3 + (1-p)^3 + 2p(1-p)\phi} \end{aligned} \tag{7}$$

$$\begin{aligned} \mu_{E,L} &= \frac{\left[p^2 + 2p(1-p)\phi\right](1-p)}{\left[p^2 + 2p(1-p)\phi\right](1-p) + \left[(1-p)^2 + 2p(1-p)\phi\right] p} \\ &= \frac{\left[p^2 + 2p(1-p)\phi\right]}{\left[p^2 + 2p(1-p)\phi\right] + \left[(1-p)p + 2p^2\phi\right]} = \frac{\left[p^2 + 2p(1-p)\phi\right]}{2p\phi + p} \end{aligned} \tag{8}$$

$$\begin{aligned} p\mu_{F,H} &= p.\frac{\left[(1-p)^2 + 2p(1-p)(1-\phi)\right] p}{\left[(1-p)^2 + 2p(1-p)(1-\phi)\right] p + \left[p^2 + 2p(1-p)(1-\phi)\right](1-p)} \\ &= \frac{p^2\left[(1-p)^2 + 2p(1-p)(1-\phi)\right]}{p(1-p)\left\{\left[(1-p) + 2p(1-\phi)\right] + \left[p + 2(1-p)(1-\phi)\right]\right\}} \\ &= \frac{p^2\left[(1-p)^2 + 2p(1-p)(1-\phi)\right]}{p(1-p)(3-2\phi)} \end{aligned} \tag{9}$$

$$\begin{aligned} \mu_{F,L} &= \frac{\left[(1-p)^2 + 2p(1-p)(1-\phi)\right](1-p)}{\left[(1-p)^2 + 2p(1-p)(1-\phi)\right](1-p) + \left[p^2 + 2p(1-p)(1-\phi)\right] p} \\ &= \frac{\left[(1-p)^2 + 2p(1-p)(1-\phi)\right](1-p)}{(1-p)^3 + p^3 + 2p(1-p)(1-\phi)} \\ &= \frac{\left[(1-p)^2 + 2p(1-p)(1-\phi)\right](1-p)}{(1-2p+p^2)(1-p) + p^3 + 2p(1-p) - 2p(1-p)\phi} \\ &= \frac{\left[(1-p)^2 + 2p(1-p)(1-\phi)\right](1-p)}{p^2 + (1-p) - 2p(1-p)\phi} \end{aligned} \tag{10}$$

Let $Q = p(1-p)(3-2\phi)$ and $R = p(1-p)(1+2\phi)$. Note that $Q \in (0,1)$ and $R \in (0,1)$ as $p(1-p) \in \left(0, \frac{1}{4}\right)$.[10] Then:

$$\frac{dQ}{d\phi} = \frac{d(1-R)}{d\phi} = -2p(1-p) < 0 \tag{11}$$

$$\frac{dR}{d\phi} = \frac{d(1-Q)}{d\phi} = 2p(1-p) > 0 \tag{12}$$

[10] $p(1-p)$ is continuous over $p > 0$, and strictly decreasing over $p \in \left(\frac{1}{2}, 1\right)$ as $1 - 2p < 0$. As $p(1-p) = \frac{1}{4}$ at $p = \frac{1}{2}$, we conclude that $p(1-p) \in \left(0, \frac{1}{4}\right)$ for $p \in \left(\frac{1}{2}, 1\right)$.

Using (7) and (9):

$$\Pr[E]\, p\mu_{E,H} = \frac{p^2\left[p^2+2p(1-p)\phi\right]^2}{p^3+(1-p)^3+2p(1-p)\phi} \tag{13}$$

$$\Pr[F]\, p\mu_{F,H} = \frac{p^2\left[(1-p)^2+2p(1-p)(1-\phi)\right]^2}{p(1-p)\,(3-2\phi)} \tag{14}$$

But:

$$\begin{aligned}
p^2+2p(1-p)\phi &= p^2-p(1-p)(3-2\phi)+3p(1-p)=(1-p)\,(2p-1)+(1-Q)\,,\\
(1-p)^2+2p(1-p)(1-\phi) &= (1-p)^2+p(1-p)(3-2\phi)-p(1-p)=(1-p)(1-2p)+Q,\\
p^3+(1-p)^3+2p(1-p)\phi &= p^3+(1-p)(1-p)^2-p(1-p)(3-2\phi)+3p(1-p)\\
&= p^3+(1-p)(1+p+p^2)-Q=1-Q
\end{aligned}$$

Therefore:

$$\begin{aligned}
\Pr[E]\, p\mu_{E,H} &= \frac{p^2\left[(1-p)\,(2p-1)+(1-Q)\right]^2}{1-Q}\\
&= \frac{p^2\,(1-p)^2\,(2p-1)^2}{1-Q}+2p^2\,(1-p)\,(2p-1)+p^2\,(1-Q) \qquad (15)\\
\Pr[F]\, p\mu_{F,H} &= \frac{p^2\left[(1-p)(1-2p)+Q\right]^2}{Q}\\
&= \frac{p^2(1-p)^2(1-2p)^2}{Q}+2p^2(1-p)(1-2p)+p^2Q \qquad (16)
\end{aligned}$$

So, given $p\in(\frac{1}{2},1)$ and $Q\in(0,1)$, and using the fact that $\frac{d^2\left(Z^{-1}\right)}{dZ^2}=\frac{d(-Z^{-2})}{dZ}=2Z^{-3}$:

$$\frac{d^2\left(\Pr[E]\, p\mu_{E,H}\right)}{d(1-Q)^2} = \frac{2p^2\,(1-p)^2\,(2p-1)^2}{(1-Q)^3}>0 \tag{17}$$

$$\frac{d^2\left(\Pr[F]\, p\mu_{F,H}\right)}{dQ^2} = \frac{2p^2(1-p)^2(1-2p)^2}{Q^3}>0 \tag{18}$$

Thus, using (12) and (17):

$$\frac{d^2\left(\Pr[E]\, p\mu_{E,H}\right)}{d\phi^2}=\frac{d\left(\frac{d\left(\Pr[E]p\mu_{E,H}\right)}{d(1-Q)}\frac{d(1-Q)}{d\phi}\right)}{d\phi}=\frac{d^2\left(\Pr[E]\, p\mu_{E,H}\right)}{d(1-Q)^2}\left(\frac{d(1-Q)}{d\phi}\right)^2>0$$

Similarly, using (11) and (18):

$$\frac{d^2\left(\Pr[F]\, p\mu_{F,H}\right)}{d\phi^2}=\frac{d^2\left(\Pr[F]\, p\mu_{F,H}\right)}{dQ^2}\left(\frac{dQ}{d\phi}\right)^2>0$$

Using (8) and (10):

$$\begin{aligned}
\Pr[E]\,\mu_{E,L} &= \frac{(1-p)\left[p^2+2p(1-p)\phi\right]^2}{p(1-p)\,(2\phi+1)} &(19)\\
\Pr[F]\,\mu_{F,L} &= \frac{\left[(1-p)^2+2p(1-p)(1-\phi)\right]^2(1-p)}{\left[p^2+(1-p)-2p(1-p)\phi\right]} &(20)
\end{aligned}$$

But:

$$\begin{aligned}
p^2+2p(1-p)\phi &= p^2+p(1-p)(1+2\phi)-p(1-p)=p(2p-1)+R,\\
(1-p)^2+2p(1-p)(1-\phi) &= (1-p)^2-p(1-p)(1+2\phi)+3p(1-p)=p(1-2p)+(1-R)\,,\\
p^2+(1-p)-2p(1-p)\phi &= p^2+(1-p)-p(1-p)(1+2\phi)+p(1-p)=1-R
\end{aligned}$$

Therefore:

$$\begin{aligned}
\Pr[E]\,\mu_{E,L} &= \frac{(1-p)\left[p(2p-1)+R\right]^2}{R}\\
&= \frac{(1-p)p^2(2p-1)^2}{R}+2(1-p)p(2p-1)+(1-p)R &(21)\\
\Pr[F]\,\mu_{F,L} &= \frac{(1-p)\left[p(1-2p)+(1-R)\right]^2}{1-R}\\
&= \frac{(1-p)p^2(1-2p)^2}{1-R}+2(1-p)p(1-2p)+(1-p)(1-R) &(22)
\end{aligned}$$

Similarly to (17) and (18):

$$\begin{aligned}
\frac{d^2\left(\Pr[E]\,\mu_{E,L}\right)}{dR^2} &= \frac{(1-p)p^2(2p-1)^2}{R^3}>0\\
\frac{d^2\left(\Pr[F]\,\mu_{F,L}\right)}{d\,(1-R)^2} &= \frac{(1-p)p^2(1-2p)^2}{(1-R)^3}>0
\end{aligned}$$

So, using (11) and (12):

$$\begin{aligned}
\frac{d^2\left(\Pr[E]\,\mu_{E,L}\right)}{d\phi^2} &= \frac{d^2\left(\Pr[E]\,\mu_{E,L}\right)}{dR^2}\left(\frac{dR}{d\phi}\right)^2>0\\
\frac{d^2\left(\Pr[F]\,\mu_{F,L}\right)}{d\phi^2} &= \frac{d^2\left(\Pr[F]\,\mu_{F,L}\right)}{d\,(1-R)^2}\left(\frac{d\,(1-R)}{d\phi}\right)^2>0
\end{aligned}$$

Given that the sum of strictly convex functions must be strictly convex, we can conclude that $\frac{d^2\Pi_i}{d\phi^2}>0$ for $i\in\{1,2,3,4\}$. ■

Proof of Lemma 4. $p\mu_{d,H} \geq \mu_{d,L} \Leftrightarrow \frac{p\mu_{d,H}}{\mu_{d,L}} \geq 1$. Using (9) and (10):

$$\frac{p\mu_{F,H}}{\mu_{F,L}} = \frac{p\left[p^2 + (1-p) - 2p(1-p)\phi\right]}{(1-p)^2\,(3-2\phi)} \tag{23}$$

Where $\phi = 1$, (23) ≥ 1 iff

$$p^3 + p - p^2 - 2p^2 + 2p^3 \geq 1 - 2p + p^2 \Leftrightarrow 3p^3 - 4p^2 + 3p - 1 \geq 0$$

$\frac{d\left(3p^3-4p^2+3p-1\right)}{dp} = 9p^2 - 8p + 3 = p^2 - 8p(1-p) + 3 > 0$ given $p(1-p) < \frac{1}{4}$ from footnote 10. Solving numerically, we find that $3p^3 - 4p^2 + 3p - 1 = 0$ at $p_1 \simeq 0.594$, proving (i).

Where $\phi = 0$, (23) ≥ 1 iff

$$p^3 + p - p^2 \geq 3 - 6p + 3p^2 \Leftrightarrow p^3 - 4p^2 + 7p - 3 \geq 0$$

$\frac{d\left(p^3-4p^2+7p-3\right)}{dp} = 3p^2 - 8p + 7 = -3p(1-p) - 5p + 7 > 0$ given $p(1-p) < \frac{1}{4}$ from footnote 10. Solving numerically, $p^3 - 4p^2 + 7p - 3 = 0$ at $p_2 \simeq 0.607$, proving (ii).

Using (7) and (8):

$$\frac{p\mu_{E,H}}{\mu_{E,L}} = \frac{p^2\,(2p\phi + p)}{p^3 + (1-p)^3 + 2p(1-p)\phi} \tag{24}$$

Where $\phi = 1$, (24) ≥ 1 iff

$$3p^3 \geq p^3 + (1-p)(1+p^2) \Leftrightarrow 3p^3 - p^2 + p - 1 \geq 0$$

$\frac{d\left(3p^3-p^2+p-1\right)}{dp} = 9p^2 - 2p + 1 = 7p^2 - 2p(1-p) + 1 > 0$ given $p(1-p) < \frac{1}{4}$ from footnote 10. Solving numerically, $3p^3 - p^2 + p - 1 = 0$ at $p_3 \simeq 0.635$, proving (iii).

Where $\phi = 0$, (24) ≥ 1 iff $p^3 \geq p^3 + (1-p)^3$ which can never be true, proving (iv). ■

Proof of Proposition 5. We consider four ranges in turn: (i) $p \in \left(\frac{1}{2}, p_1\right)$; (ii) $[p_1, p_2)$; (iii) $[p_2, p_3)$; (iv) $[p_3, 1)$.

(i) From Lemma 4, $\Pi = \Pi_2$ for $p \in \left(\frac{1}{2}, p_1\right)$. From (21) and (22), and using the fact that $(2p-1)^2 = (1-2p)^2$,

$$\Pi_2 = \Pr[E]\,\mu_{E,L} + \Pr[F]\,\mu_{F,L} = \frac{(1-p)p^2(2p-1)^2}{R(1-R)} + (1-p)$$

Therefore,

$$\begin{aligned}
\Pi_2(\phi=0) &> \Pi_2(\phi=1) \Leftrightarrow \\
R(\phi=0)[1-R(\phi=0)] &< R(\phi=1)[1-R(\phi=1)] \Leftrightarrow \\
[p(1-p)][1-p(1-p)] &< [3p(1-p)][1-3p(1-p)] \Leftrightarrow \\
1-p(1-p) &< 3-9p(1-p) \Leftrightarrow \\
8p(1-p) &< 2 \Leftrightarrow p(1-p) < \frac{1}{4}
\end{aligned}$$

which we know to be true from footnote 10.

(ii) From Lemma 4, $\Pi(\phi=0)=\Pi_2(\phi=0)$ and $\Pi(\phi=1)=\Pi_3(\phi=1)$ for $p \in [p_1, p_2)$. $\Pi_2 = \Pr[E]\mu_{E,L} + \Pr[F]\mu_{F,L}$, so using (19) and (20):

$$\begin{aligned}
\Pi_2(\phi=0) &= \frac{(1-p)\left[p^2\right]^2}{p(1-p)} + \frac{\left[(1-p)^2+2p(1-p)\right]^2(1-p)}{\left[p^2+(1-p)\right]} \\
&= p^3 + \frac{\left(1-p^2\right)^2(1-p)}{p^2+(1-p)} = p^3 + \frac{1-p-2p^2+2p^3+p^4-p^5}{1-p+p^2} = \frac{1-p-2p^2+3p^3}{1-p+p^2}
\end{aligned}$$

$\Pi_3 = \Pr[E]\mu_{E,L} + \Pr[F]p\mu_{F,H}$, so using (19) and (14):

$$\begin{aligned}
\Pi_3(\phi=1) &= \frac{(1-p)\left[p^2+2p(1-p)\right]^2}{3p(1-p)} + \frac{p^2\left[(1-p)^2\right]^2}{p(1-p)} \\
&= \frac{p(2-p)^2+3p(1-p)^3}{3} = \frac{7p-13p^2+10p^3-3p^4}{3} \qquad (25)
\end{aligned}$$

Thus:

$$\begin{aligned}
\Pi_2(\phi=0) &> \Pi_3(\phi=1) \Leftrightarrow \\
3-3p-6p^2+9p^3-\left(1-p+p^2\right)\left(7p-13p^2+10p^3-3p^4\right) &> 0 \Leftrightarrow \\
3-10p+14p^2-21p^3+26p^4-13p^5+3p^6 &> 0 \Leftrightarrow \\
\left(-3+7p-4p^2+p^3\right)\left(-1+p-p^2+3p^3\right) &> 0
\end{aligned}$$

From part (ii) of the proof of Lemma 4, $-3+7p-4p^2+p^3<0$ here. From part (iii) of the same proof, $-1+p-p^2+3p^3<0$ here also.

(iii) From Lemma 4, $\Pi=\Pi_3$ for $p \in [p_2, p_3)$.

$\Pi_3 = \Pr[E]\,\mu_{E,L} + \Pr[F]\,p\mu_{F,H}$, so using (19) and (14):

$$\begin{aligned}\Pi_3(\phi=0) &= p^3 + \frac{p^2\left[(1-p)^2+2p(1-p)\right]^2}{3p(1-p)} = \frac{3p^3+p(1-p)(p+1)^2}{3} \\ &= \frac{p+p^2+2p^3-p^4}{3}\end{aligned} \tag{26}$$

Thus, using (25),

$$\begin{aligned}\Pi_3(\phi=0) &\gtreqless \Pi_3(\phi=1) \Leftrightarrow \\ p+p^2+2p^3-p^4 &\gtreqless 7p-13p^2+10p^3-3p^4 \Leftrightarrow \\ -3+7p-4p^2+p^3 &\gtreqless 0\end{aligned}$$

From part (ii) of the proof of Lemma 4, $-3+7p-4p^2+p^3>0$ here, except at p_2 where it equals zero.

(iv) From Lemma 4, $\Pi(\phi=0)=\Pi_3(\phi=0)$ and $\Pi(\phi=1)=\Pi_1(\phi=1)$ for $[p_3,1)$.

From (15) and (16), and using the facts that $(2p-1)^2=(1-2p)^2$ and $Q(\phi=1)=p(1-p)$,

$$\begin{aligned}\Pi_1 &= \Pr[E]\,p\mu_{E,H} + \Pr[F]\,p\mu_{F,H} = \frac{p^2(1-p)^2(2p-1)^2}{Q(1-Q)} + p^2 \\ \therefore\quad \Pi_1(\phi=1) &= \frac{p(1-p)(2p-1)^2}{1-p+p^2} + p^2 = \frac{p-5p^2+8p^3-4p^4}{1-p+p^2} + p^2\end{aligned}$$

Thus, using (26),

$$\begin{aligned}\Pi_3(\phi=0) &\gtreqless \Pi_1(\phi=1) \Leftrightarrow \\ \left(p+p^2+2p^3-p^4-3p^2\right)\left(1-p+p^2\right) &\gtreqless 3\left(p-5p^2+8p^3-4p^4\right) \Leftrightarrow \\ -2p+12p^2-19p^3+7p^4+3p^5-p^6 &\gtreqless 0 \Leftrightarrow \\ p\left(-2+12p-19p^2+7p^3+3p^4-p^5\right) &\gtreqless 0 \Leftrightarrow \\ p(1-p)(-2+10p-9p^2-2p^3+p^4) &\gtreqless 0 \Leftrightarrow \\ p(1-p)(-1+4p-p^2)(2-2p-p^2) &\gtreqless 0 \Leftrightarrow\end{aligned}$$

Clearly, $p(1-p)>0$. Also, $-1+4p-p^2=3p+p(1-p)-1>0$ for $p>\frac{1}{2}$. As $-2-2p<0$, $2-2p-p^2$ is strictly decreasing for $p>\frac{1}{2}$, and equals zero at $p_4=\frac{2-\sqrt{4-(-8)}}{-2}=-1+\sqrt{3}\simeq 0.732$. Thus, we conclude that for $p\in[p_3,p_4)$, $\Pi_3(\phi=0)>\Pi_1(\phi=1)$, for $p\in(p_4,1)$, $\Pi_3(\phi=0)<\Pi_1(\phi=1)$, and at p_4, $\Pi_3(\phi=0)=\Pi_1(\phi=1)$. ■

Proof of Proposition 6. In a slight variation of earlier notation, let $\mu_{X_i} = \Pr[V = 1|X_i]$ where the consumer's signal is given as $X_i \in \{H, L\}$. Using Bayes' Rule and the fact that the consumers have 50:50 initial priors over the good's quality:

$$\mu_{X_i} = \frac{\Pr[X_i|V=1]}{\Pr[X_i|V=1] + \Pr[X_i|V=0]}$$

Similarly to before, the firm will choose to set price $\lambda = \mu_H$ to sell to all those who received high private signals or $\lambda = \mu_L$ to sell to all consumers including those with low private signals. Once again normalizing the number of customers to one, expected profits as a function of these prices are $\pi(\mu_H) = p\mu_H$ and $\pi(\mu_L) = \mu_L$. We define $\Pi[\text{No Reviewer}] = \max\{p\mu_H, \mu_L\}$ to be the maximum expected profits achievable. Now, we know that $\mu_H = \frac{p}{p+(1-p)} = p$. With expected sales of p at the price μ_H, we have total profit of p^2. Furthermore, $\mu_L = \frac{1-p}{(1-p)+p} = 1-p$. Hence $\Pi[\text{No Reviewer}] = \max\{p^2, 1-p\}$.

We next divide the optimal choice of price into two regions. Clearly, $d(p^2+p-1)/dp > 0$ for $p \in \left(\frac{1}{2}, 1\right)$, and $p^2+p-1 = 0$ at $\widetilde{p} = \frac{1}{2}\left(-1+\sqrt{5}\right) \simeq 0.618$. Therefore for $p \geq \widetilde{p}$, $p^2 \geq 1-p$, and for $p < \widetilde{p}$, $p^2 < 1-p$. Thus, the firm will set a price p when $p \geq \widetilde{p}$, and a price $1-p$ otherwise. Profits are p^2 if $p \in [\widetilde{p}, 1)$ and $1-p$ if $p \in (\frac{1}{2}, \widetilde{p})$.

Suppose that if the firm chose an unbiased reviewer, it had to set price $\lambda = \mu_{d,H}$ for $p \in [\widetilde{p}, 1)$ and $\lambda = \mu_{d,L}$ for $p \in (\frac{1}{2}, \widetilde{p})$. We show that even under such a restriction, profits are higher than with no reviewer, proving that profits under optimal price setting must be higher using an unbiased reviewer. Now, given $\phi = \frac{1}{2}$:

$$\begin{aligned}
\Pr[E|V=1] &= \Pr[F|V=0] &&= p^2 + p(1-p) = p \\
\Pr[F|V=1] &= \Pr[E|V=0] &&= (1-p)^2 + p(1-p) = 1-p
\end{aligned}$$

Thus:

$$\begin{aligned}
\Pr[E].p\mu_{E,H} &= p.p\left[\frac{p^2}{p^2+(1-p)^2}\right] = \frac{p^4}{p^2+(1-p)^2} \\
\Pr[F].p\mu_{F,H} &= (1-p).p\left[\frac{(1-p)p}{(1-p)p+p(1-p)}\right] = \frac{(1-p)p}{2} \\
\Pr[E].\mu_{E,L} &= p.\frac{p(1-p)}{p(1-p)+(1-p)p} = \frac{p}{2} \\
\Pr[F].\mu_{F,L} &= (1-p).\frac{(1-p)^2}{(1-p)^2+p^2} = \frac{(1-p)^3}{(1-p)^2+p^2}
\end{aligned}$$

Therefore, for $p \in (\frac{1}{2}, \widetilde{p})$ the unbiased reviewer is superior to no review if:

$$\begin{aligned}
\Pi\left[\phi = \frac{1}{2} \mid p \in (\tfrac{1}{2}, \widetilde{p})\right] &= \frac{p}{2} + \frac{(1-p)^3}{(1-p)^2 + p^2} > (1-p) \qquad (27) \\
&\Leftrightarrow p(1-p)^2 + p^3 + 2(1-p)^3 - 2(1-p)^3 - 2p^2(1-p) > 0 \\
&\Leftrightarrow 1 - 2p + p^2 + p^2 - 2p + 2p^2 > 0 \Leftrightarrow (1-2p)^2 > 0
\end{aligned}$$

which must hold. For $[\widetilde{p}, 1)$ the unbiased reviewer is superior if:

$$\begin{aligned}
\Pi\left[\phi = \frac{1}{2} \mid p \in [\widetilde{p}, 1)\right] &= \frac{p^4}{p^2 + (1-p)^2} + \frac{(1-p)p}{2} > p^2 \qquad (28) \\
&\Leftrightarrow \frac{z}{2} + \frac{(1-z)^3}{(1-z)^2 + z^2} > (1-z) \qquad (29)
\end{aligned}$$

where $z = (1-p)$. But (27) holds for $p \in \left(0, \frac{1}{2}\right)$, so (29) must hold for $z \in \left(0, \frac{1}{2}\right)$, and hence (28) must hold for $p \in \left(\frac{1}{2}, 1\right)$.

We have shown that the profits attainable with no reviewer are always strictly lower than under an unbiased reviewer and by Proposition 5, certainly lower than under the optimal choice of reviewer. ■