The Optimal Choice of Pre-launch Reviewer:
How Best to Transmit Information using Tests and Conditional Pricing

David Gill and Daniel Sgroi

No 877

WARWICK ECONOMIC RESEARCH PAPERS

DEPARTMENT OF ECONOMICS

THE UNIVERSITY OF WARWICK

The Optimal Choice of Pre-launch Reviewer: How Best to Transmit Information using Tests and Conditional Pricing

David Gill
University of Southampton

Daniel Sgroi
University of Warwick

Abstract

A principal who knows her type can face public testing to help attract endorsements from agents. Tests are pass/fail and have an innate toughness (bias) corresponding to a trade-off between the higher probability of passing a softer test and the greater impact on agents' beliefs from passing a tougher test. Conditional on the test result, the principal also selects the price of endorsement. The principal always wants to be tested, and chooses the toughest or softest test available depending upon the precision of the agents' and tests' information. Applications abound in industrial organization, political economy and labor economics.

Keywords: tests, reviewers, Bayesian learning, information transmission, bias, marketing

JEL classification: D82, D83, L15

This version: 19 November 2008

1 Introduction

This paper investigates the use of public tests to transmit information about type. A principal wishes to attract endorsements from a group of agents. The principal knows her own type (high or low), but is unable to verifiably disclose her type to the agents, each of whom receives some private information about the principal's type. If she so chooses, the principal can subject herself to a public test in an attempt to convince agents that she is a high type. Tests return a binary "pass" or "fail" decision and are characterized by a publicly known level of toughness (or bias) that corresponds to the probability of passing. By choosing to be tested, the principal can influence the learning process of the agents. In particular, in selecting the toughness of the test to be faced, the principal trades off the higher probability of passing a softer test against the greater impact on agents' beliefs from passing a tougher test. Conditional on the chosen test toughness and test decision, the principal also selects the price that agents must pay to endorse her. We want to discover both what sort of test a principal might choose and how the test interacts with pricing. In short, we want to analyze the best way for any principal to use public tests and prices to maximize the proceeds from endorsements.[1]

By keeping the context abstract we hope to capture a host of relevant problems. Applications abound, for example: a firm launching a new product can choose who to send it to from a range of pre-launch reviewers with known biases;[2] an issuer of stocks or bonds can attempt to receive certification from investment banks and rating agencies which differ in their reputation and independence; a technology sponsor can choose between standard setting organizations of varying

[1] Note that we are not analyzing the problem of optimal *test expertise*: the principal has to choose between different tests which all base their decisions on signals of the same quality. Also note that we do not analyze what test agents would prefer: in all the applications outlined in the next paragraph, it is the principal who chooses the test type.

[2] A computer software or game producer might consider approaching a magazine or online site with different standards for a "preview". Sites and magazines typically have a known toughness. For example, general video games websites will often list not only the results of reviews and previews, but also give some indication of the reviewer's toughness, while "official" games magazines owned by the same company that produces a game might be considered soft. Early success for software in a preview can have huge implications for the sales of the product, especially in the pre-order market. When Microsoft launched Windows Vista in November 2006, initially all information about the new operating system came through official Microsoft sources, followed by an exclusive preview by Paul Thurrot, the well-know pro-Microsoft blogger and editor of Windows IT Pro Magazine. To give an opposing example, Microsoft in collaboration with Gearbox Software developed a PC version of their hit Xbox game "Halo" in 2003. Gamespy was one of the few review sites that did not award the original Xbox version of Halo near perfect marks, and went on record saying they believed the original game "wasn't quite as perfect as other critics made it out to be". Despite their tough reputation, Gearbox Software *invited* Gamespy to preview the PC version of the game in advance of its public release.

toughness ranging from fully independent to largely captive;[3] a University research spin-off looking for funds can select from a pool of referees with different academic and professional reputations to review the project; a film-maker can choose to premiere his movie at a prestigious film festival such as Cannes, where competition for prizes is fierce, or at a smaller festival with less competition; a recent Ph.D. on the job market has to decide whether to risk selecting a well-respected but tough professor as a referee or going for a softer option; a prospective student has to choose how tough a degree programme to attempt; and a politician with a policy to "sell" to potential supporters can seek support from a variety of think-tanks or policy institutes with known policy biases.

In each case, the principal selects from a range of tests, reviewers or accreditors of varying toughness. The interpretation of the "price" of endorsement varies across the applications. Linking to the examples above, "price" can take the form of a standard market price, the size of an ownership stake, contractual terms offered to exhibitors of movies, a salary, or the degree of compromise required for a policy to be approved.

In the context of our model, a low type principal can costlessly duplicate the actions of a high type principal. As a result, all of our equilibria will be pooling so the principal cannot use her choice of test to signal type directly. Nevertheless, tests play a crucial role in information transmission. We find that the principal always chooses to be tested, so tests complement the choice of price, and that the ability to condition price on the test result convexifies the principal's payoffs: the principal always chooses an extreme test, selecting either the toughest or softest test available.

Where the private signals received by the agents are of low precision and the agents' prior belief about the principal's type is not too high, which might correspond to a new type of product or innovative idea, the principal will choose the toughest test available, maximizing the impact of passing the test. In that case, the principal accepts a higher risk of failing the test in order to launch her product, idea or policy with a bang if she passes. If on the other hand agents' signals are of high enough precision, perhaps because the product or idea is well-known, the softest test is chosen (except where, in the good state, the information received by the tests is heavily negatively biased), maximizing the probability of passing but dampening the impact of a pass on beliefs and

[3] The descriptions of bias in technology standard setters and raters of stock or bond issues are from Lerner and Tirole (2006), p. 1091. To give an example of varying toughness in product certification, in Europe Red Book's *BRE certification* claims to be a more rigorous alternative to the much milder assessment provided by *CE marking*.

hence on prices.

Remarkably, the literature has paid almost no attention to the use of public tests to transmit information, especially when used in combination with price; however there are various related literatures. A number of articles analyze the use of initial prices to manipulate sales in a learning environment. For example, Taylor (1999) and Bose et al. (2007, 2008) find that high initial prices, whose effects are similar to the choice of a tough test, can be optimal. In Bose et al., the firm (which unlike the principal in our model does not know its own quality) 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 (the analogue of passing a tough test), this conveys strong positive information to later buyers. Taylor, concentrating on the housing market, finds a high price to be optimal as a failure to sell a house early (the analogue of failing a tough test) can then be attributed to overpricing rather than low quality. By contrast, in Caminal and Vives (1996, 1999), in which early prices are unobservable to later consumers, and in Welch (1992), in which prices cannot be conditioned on the history of purchases, low introductory prices are optimal.

Lerner and Tirole's (2006) paper focuses on the role of technology standard setting authorities as certifiers.[4] Similarly to our tests, the certifiers have an arbitrary bias towards the technology sponsor which determines their decision rule. Lerner and Tirole's model has significant differences to ours: the sponsor is not perfectly informed about the quality of its technology; the chosen certifier discovers with certainty the quality of the technology it is asked to review; consumers do not receive any private information; and there is no incentive to set price in response to the certification (as the certifier's endorsement rule is sensitive to any anticipated price response to its decision). Therefore, as certifiers cannot counter 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 tests that are soft (biased towards) or tough (biased against) depending upon model parameters.

In a setting where both the buyer and seller are uninformed about quality, Ottaviani and Prat

[4] Chiao, Lerner and Tirole (2007) empirically test Lerner and Tirole's model, while Farhi, Lerner and Tirole (2005) extend the model to a dynamic setting.

(2001) find that a monopolist may wish to use a public signal of quality such as an (unbiased) outside certifier. In their model, a public signal affiliated with the buyer's private information reduces the buyer's informational rents in a second-degree price discrimination setting.

Our paper should also 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 testers have no self-interested motives, apart from taking on differential toughness levels. Our work is also distinct from the literature on payment structures to certification intermediaries. For example in Lizzeri (1999) and Albano and Lizzeri (2001) the question is how intermediaries affect the quality chosen by the firm, while in our model type is fixed. They show that disclosure may turn out to be incomplete, but do not allow any bias.

We assume flexible prices, but in some settings prices cannot be varied conditional on the test result. Gill and Sgroi (2008) considers a sequential endorsement model with public tests, but without prices and for a very specific signal structure. As the length of the sequence of agents tends to one, that model can be re-interpreted as a simplified analogue of the simultaneous endorsement model in this paper, but with fixed prices. A key result in the paper is that a tough test is always preferred. When prices cannot be adjusted in light of the test, passing a tough test swamps bad private information, while a failure to pass a tough test is not enough to damage good private information, so tough tests are valuable. Where available, the principal will choose a tough test that is very close to unbiased to maximize the chance of passing. In Sgroi (2002) multiple public decisions made by consumers at the start of a product's life-cycle act in a similar way to a public test, and one of the key results is that a firm (or social planner) should optimally select to use such early decision makers to boost profits (or welfare). While this provides some support for the result that a public test should always be used, the paper has neither prices nor a notion of toughness, so can offer no guide as to the optimal toughness of a public test.

1.1 Overview

Section 2 describes the model. Section 3 examines the principal's problem, ruling out separating equilibria and then deriving the principal's payoff function. Section 4 proves that the principal will always prefer to be tested. Section 5 shows that the principal always selects the toughest or softest test available. Section 6 performs a pairwise comparison between the two candidates for the optimal test type. Section 7 concludes. Omitted proofs are provided in the appendix.

2 The Model

Nature draws the type $V \in \{0, 1\}$ of a principal with $q \equiv \Pr[V = 1] \in (0, 1)$. The principal discovers whether she is a "good" type ($V = 1$) or "bad" type ($V = 0$), but cannot verifiably disclose this information. $N \in \mathbb{N}_{++}$ agents simultaneously decide whether or not to endorse the principal after each receiving a private signal about V. Endorsement is a general concept which could, for example, encompass adopting some new technology, funding a research project, subscribing to an I.P.O., purchasing a product, making a job offer etc. The principal chooses the price of endorsement $\lambda \in \mathbb{R}_+$, which gives each agent a value to endorsing of $V - \lambda$ while not endorsing returns zero. Before choosing the price, if she wishes the principal can subject herself to a public test, which she either passes or fails. If the principal chooses to be tested, she can choose the type of test, which becomes public knowledge, and she can condition the price on the result, while the agents observe the test result and price before making their endorsement decision. The principal's aim is to maximize revenue, which equals the price of endorsement multiplied by the number of agents who choose to endorse. We outline the test technology, followed by a formal description of the agents' and principal's decisions, below.

2.1 The Test

The test draws one private signal Z from the set $\mathbb{Z} \equiv \{H, U, L\}$ with representative member z. H is a positive ("high") signal about V, L is a negative ("low") signal, and U is uninformative. Let $p_V^z \equiv \Pr[Z = z|V]$. Clearly, $\sum_{\mathbb{Z}} p_V^z = 1$. To make U uninformative we restrict $p_1^U = p_0^U$. To make H positive, we restrict $p_1^H > p_0^H$ and to make L negative, we restrict $p_1^L < p_0^L$. This implies that:

$$p_1^H - p_0^H = p_0^L - p_1^L > 0 \tag{1}$$

We assume that $p_V^z > 0$, which implies that no signal is fully informative. It is plausible that the test receives better quality information than any individual agent, but we do not need to impose this as a formal assumption.

The test publicly returns a binary decision $d \in \mathbb{D} \equiv \{d_1, d_2\}$, with $\phi_d^z \equiv \Pr[d|Z=z]$ and $\sum_{\mathbb{D}} \phi_d^z = 1$, so $\Pr[d|V] = \sum_{\mathbb{Z}} p_V^z \phi_d^z$ and $\sum_{\mathbb{D}} \Pr[d|V] = 1$. Modelling a test (or an evaluator) as condensing more complex information into a simple binary decision is a common assumption in the literature.[5] Notice also that, conditional on Z, the test decision is independent of the publicly known prior q.

Suppose that an agent, who understands the test technology, conditions his belief on the test result together with generic private information I, which is independent of Z (conditionally on V) with $\Pr[V=1|I] \in (0,1)$. When $\sum_{\mathbb{Z}} \phi_d^z > 0$, so $\Pr[d|V] > 0$, decision d is good news about V if $\phi_d^H > \phi_d^L$, bad news if $\phi_d^H < \phi_d^L$ and provides no news if $\phi_d^H = \phi_d^L$, in the sense that:[6]

$$\Pr[V=1|I,d] \gtreqless \Pr[V=1|I] \Leftrightarrow \phi_d^H \gtreqless \phi_d^L$$

When $\sum_{\mathbb{Z}} \phi_d^z = 0$, so $\Pr[d|V] = 0$, decision d is never observed and so cannot provide any news about the principal's type.

Thus, *w.l.o.g.* we label the decisions P ("pass") and F ("fail") such that $\phi_P^H \geq \phi_P^L$, which of course implies that $\phi_F^H \leq \phi_F^L$ as $\sum_{\mathbb{D}} \phi_d^z = 1$, and formally we model the choice not to be tested by a choice of $\phi_P^H = \phi_P^L$. Note that in the language of Milgrom (1981), P is "more favorable" news than F about V as the Monotone Likelihood Ratio Property is satisfied.

The test type $\boldsymbol{\phi} \equiv \left\{\phi_P^H, \phi_P^U, \phi_P^L\right\} \in [0,1]^3$ chosen by the principal is common knowledge, perhaps generated through a known history of pass or fail decisions. Depending on the application, the choice of test type might consist of a choice between different reviewers, referees, accreditation bodies and so on.

[5] For example see Calvert (1985), Sah and Stiglitz (1986), Farhi et al. (2005), Lerner and Tirole (2006), Chiao et al. (2007), Gill and Sgroi (2008) and Demange (2008). As Calvert (1985, p. 534) puts it: “This feature represents the basic nature of advice, a distillation of complex reality into a simple recommendation.” The coarseness of the binary report relative to the information received by the test is a key driver of our results: see footnote 9 for further discussion.

[6] $\frac{\Pr[V=1|I,d]}{\Pr[V=1|I]} \gtreqless 1 \Leftrightarrow \frac{\Pr[d|V=1]}{\Pr[d|V=1]\Pr[V=1|I]+\Pr[d|V=0](1-\Pr[V=1|I])} \gtreqless 1 \Leftrightarrow \sum_{\mathbb{Z}} p_1^z \phi_d^z \gtreqless \sum_{\mathbb{Z}} p_0^z \phi_d^z \Leftrightarrow \left(p_1^H - p_0^H\right) \phi_d^H \gtreqless \left(p_0^L - p_1^L\right) \phi_d^L \Leftrightarrow \phi_d^H \gtreqless \phi_d^L$. The first part follows using Bayes' Rule together with $\Pr[I,d|V] = \Pr[d|V]\Pr[I|V]$ and dividing top and bottom by $\Pr[I]$, the second last part follows from $p_1^U = p_0^U$ and the last part from (1).

2.2 Agents

A typical agent i chooses an action $A_i \in \{Y, N\}$, where Y denotes "endorsement" and N "no endorsement", and is a risk neutral expected utility maximizer with utility $u(A_i, V) \in \mathbb{R}$ of the form $u(Y, V) = V - \lambda$ and $u(N, V) = 0$. Each agent draws one private signal X_i from the finite set $\mathbb{X} \equiv \{0, 1, ..., M\}$ with representative member m and $M \geq 0$. The draws are *i.i.d.* conditional on V and are conditionally independent of Z. Let $p_V^m \equiv \Pr[X_i = m|V]$ where $p_V^m \in (0, 1]$ and $\sum_{\mathbb{X}} p_V^m = 1$. The assumption that $p_V^m > 0$ implies that no signal is fully informative and also rules out signals which are never drawn.

Let the posterior belief having observed a private signal $X_i = m$, a choice of test $\boldsymbol{\phi}$, a test result d and a price λ be

$$\mu \equiv \Pr[V = 1|X_i = m, \boldsymbol{\phi}, d, \lambda]$$

Each agent endorses iff his $\mu \geq \lambda$. The endorse at indifference rule is *w.l.o.g.* as if at a given λ indifferent agents endorsed with probability less than one, the principal could shave price by a small $\epsilon > 0$ and get all the indifferent agents to endorse. Throughout, *w.l.o.g.,* we normalize the number of agents to 1.

We let μ^m be an agent's belief that the principal is of a good type having observed only $X_i = m$, so

$$\mu^m \equiv \Pr[V = 1|X_i = m] = \frac{p_1^m q}{p_1^m q + p_0^m (1 - q)} \in (0, 1) \tag{2}$$

Notice that $\mu^m \gtreqless q \Leftrightarrow p_1^m \gtreqless p_0^m$, so when $p_0^m = p_1^m \;\; \forall m$, the model includes the situation where all signals are uninformative (which holds trivially when $M = 0$).

2.3 The Principal

Having discovered her type V, the principal selects the test type $\boldsymbol{\phi}$, with $\phi_P^z \in [0, 1]$. As outlined above, choosing not to be tested is modelled as choosing an uninformative test with $\phi_P^H = \phi_P^L$. After the test result, the principal selects a price λ (which can be conditioned on V, d and the chosen test type) to maximize expected revenue R. Initially, the principal selects the test type $\boldsymbol{\phi}$ to maximize expected revenue given the post-test pricing rule. We restrict the principal to pure strategies.

3 The Principal's Problem

This section details the principal's problem. Before deriving the revenue function, we first point out that separation by choice of test type is not possible and then justify restricting attention to the good type of principal.

3.1 Restriction to the Choice of Good Type of Principal

Nature's draw of the principal's type results in a game of imperfect information, so we apply perfect Bayesian equilibrium (PBE) as the equilibrium concept. Our first result is that all PBEs must be pooling. Suppose an equilibrium existed in which the bad type of principal had a different strategy from the good type. Given the restriction to pure strategies, by following her strategy the bad type would immediately reveal herself to be bad at any node where the two strategies indicated different actions. She would then receive no endorsements, and so would want to deviate by duplicating the good type's action at every such node, thereby costlessly making the agents believe her to be good.

This establishes that in equilibrium the principal's choice of ϕ and λ does not signal type, so agents will be unable to infer anything from the choice of test or price *per se*. Instead the agents will have to rely on the outcome of the test for information.

Lemma 1 *All PBEs are pooling, so in equilibrium the principal's choice of test and price is uninformative.*

By ruling out signalling via the choice of test and price, we focus attention on the role of the test itself in transmitting information to the agents.[7] Furthermore, because we have assumed that the principal's type is non-verifiable, our model is not a game of persuasion à la Milgrom (1981) so the principal is unable to reveal information directly to try to separate from other types.

The lemma implies that in equilibrium an agent's posterior belief μ that the principal is of a good type is given by

$$\mu_d^m \equiv \frac{\Pr[d|V=1,\phi]\,\mu^m}{\Pr[d|V=1,\phi]\,\mu^m + \Pr[d|V=0,\phi]\,(1-\mu^m)} \tag{3}$$

[7] More generally, price choices can play a signalling role. See, e.g., Judd and Riordan (1994) and Mahenc (2004).

as when an agent uses Bayes' Rule to calculate his posterior belief conditional on d and $X_i = m$ he can first calculate his belief conditional on his private signal, yielding μ^m, and then use Bayes' Rule to adjust for the impact of the test result, using μ^m as the prior.

We have an equilibrium selection issue because any pooled strategies can form a pooling equilibrium supported by the belief that a deviator is of a bad type. Conditional on playing a pooling equilibrium, let Ω represent the good type of principal's set of optimal $\boldsymbol{\phi}$ and conditional λ pairs. We assume that only pooling PBEs with strategies in Ω are played. This can be justified by assuming that starting from pooled strategies outside of Ω, observing deviations consistent with a strategy in Ω weakly increases the agents' beliefs that the principal is of a good type. Then a good type will have a strict incentive to deviate at the first node she reaches where the pooled strategy differs from one of the strategies in Ω (a bad type might also want to do so).

Under this assumption, in equilibrium a bad principal is forced to follow a good principal's choice of test type and conditional pricing rule. Thus, *throughout we restrict attention to the good type of principal's choice of test and price conditional on pooling, which always form a PBE* (i.e., we characterize Ω).

3.2 Deriving the Revenue Function

After the test's decision d, the principal chooses an optimal price λ from the set $\left\{\mu_d^0, \mu_d^1, ..., \mu_d^M\right\}$. By setting $\lambda = \mu_d^m$, the principal receives endorsements from all those agents who received private signals at least as strong as $X_i = m$, i.e., from those agents who received $X_i \in \left\{k : \mu^k \geq \mu^m\right\}$,[8] and so faces a standard price-quantity trade-off. The principal will never choose a price outside this set, as if $\lambda \notin \left\{\mu_d^0, \mu_d^1, ..., \mu_d^M\right\}$ the principal could raise price to $\min\limits_{m \in \mathbb{X}} \{\mu_d^m : \mu_d^m \geq \lambda\}$ without losing any endorsements (unless $\lambda > \max\limits_{m \in \mathbb{X}} \mu_d^m$ in which case the principal receives no endorsements).

Let r_d represent expected revenue after a decision d, given a choice of price μ_d^m. We have normalized the number of agents to 1, and, as explained in the previous subsection, we focus on the good type of principal's choice, conditional on pooling. Thus $r_d = \mu_d^m \sum\limits_{k:\mu^k \geq \mu^m} p_1^k$, so the

[8] If $\sum_{\mathbb{Z}} \phi_d^z > 0$, $\Pr[d|V, \boldsymbol{\phi}] > 0$ so from (3) and letting m_i represent a particular m, $\mu_d^{m_i} \gtreqless \mu_d^{m_j} \Leftrightarrow \mu^{m_i} \gtreqless \mu^{m_j}$. The paragraph after (4) discusses the case where $\sum_{\mathbb{Z}} \phi_d^z = 0$.

maximal revenue achievable r_d^*, which is a function of $\boldsymbol{\phi}$, is given by:

$$r_d^* = \max_{m \in \mathbb{X}} \mu_d^m \sum_{k:\mu^k \geq \mu^m} p_1^k$$

The principal's optimal choice of test therefore reduces to:

$$\max_{\boldsymbol{\phi} \in [0,1]^3} R = \textstyle\sum_{\mathbb{D}} \Pr\left[d|V=1,\boldsymbol{\phi}\right] r_d^* \tag{4}$$

We next deal with a technical issue, which is that when $\sum_{\mathbb{Z}} \phi_d^z = 0$, $\Pr\left[d|V=1,\boldsymbol{\phi}\right] = \Pr\left[d|V=0,\boldsymbol{\phi}\right] = 0$ so from (3) μ_d^m is not defined and therefore neither is r_d^*. However, in such a case the belief is irrelevant as $\Pr\left[d|V=1,\boldsymbol{\phi}\right] = 0$. Hence we can define $\mu_d^m \equiv 0$, so $r_d^* = 0$ and the expression for R remains valid.

We can re-write expected revenue as follows:

$$R = \textstyle\sum_{\mathbb{D}} \max_{m \in \mathbb{X}} \Pr\left[d|V=1,\boldsymbol{\phi}\right] \mu_d^m \sum_{k:\mu^k \geq \mu^m} p_1^k \tag{5}$$

Note that R is continuous in ϕ_P^z for all $z \in \mathbb{Z}$. As the maximum of continuous functions is continuous, we simply need to show the continuity of $\Pr\left[d|V=1,\boldsymbol{\phi}\right] \mu_d^m$ in ϕ_d^z, noting that when $d = F$, using $\phi_P^z = 1 - \phi_F^z$, continuity in ϕ_F^z implies continuity in ϕ_P^z. Now $\Pr\left[d|V,\boldsymbol{\phi}\right] = \sum_{\mathbb{Z}} p_V^z \phi_d^z$ is continuous in ϕ_d^z for all $z \in \mathbb{Z}$. Thus at any $\boldsymbol{\phi}$ such that $\sum_{\mathbb{Z}} \phi_d^z > 0$, so $\Pr\left[d|V,\boldsymbol{\phi}\right] > 0$, we can establish the required continuity using (3). At $\boldsymbol{\phi}$ such that $\sum_{\mathbb{Z}} \phi_d^z = 0$, so $\Pr\left[d|V,\boldsymbol{\phi}\right] = 0$, continuity can be shown applying L'Hôpital's Rule: letting $\overline{z}$ represent a specific member of $\mathbb{Z}$, when $\sum_{\mathbb{Z}\setminus\overline{z}} \phi_d^z = 0$, $\lim_{\phi_d^{\overline{z}} \downarrow 0} \Pr\left[d|V=1,\boldsymbol{\phi}\right] \mu_d^m = \lim_{\phi_d^{\overline{z}} \downarrow 0} \frac{2\Pr[d|V=1,\boldsymbol{\phi}] p_1^{\overline{z}} \mu^m}{p_1^{\overline{z}} \mu^m + p_0^{\overline{z}} (1-\mu^m)} = 0$ as required.

4 Choosing to be Tested

In this section, we show that the principal will always choose to be tested, as *any* test is better than no test. From Section 2.1 not being tested is equivalent to choosing $\phi_P^H = \phi_P^L$. If the principal chooses $\phi_P^H = \phi_P^L$, $\Pr\left[d|V=1,\boldsymbol{\phi}\right] = \Pr\left[d|V=0,\boldsymbol{\phi}\right]$ given $p_1^U = p_0^U$. Thus, when $\Pr\left[d|V=1,\boldsymbol{\phi}\right] > 0$, from (3) $\mu_d^m = \mu^m$, i.e., the test result does not change the agents' beliefs. Therefore, using (4) and $\sum_{\mathbb{D}} \Pr[d|V, \boldsymbol{\phi}] = 1$, expected revenue from choosing $\phi_P^H = \phi_P^L$ is given by:

$$R\left(\phi_P^H = \phi_P^L\right) = \max_{m \in \mathbb{X}} \mu^m \sum_{k:\mu^k \geq \mu^m} p_1^k$$

Let $\widehat{m}$ be an m which maximizes $\mu^m \sum_{k:\mu^k \geq \mu^m} p_1^k$, so $R\left(\phi_P^H = \phi_P^L\right) = \mu^{\widehat{m}} \sum_{k:\mu^k \geq \mu^{\widehat{m}}} p_1^k$. Suppose that the principal is restricted to choosing a particular test type $\overline{\boldsymbol{\phi}}$ with $\phi_P^H \neq \phi_P^L$, and, conditional on the test decision d, a price $\overline{\lambda} \equiv \mu_d^{\widehat{m}}$ to target agents with signals at least as strong as $X_i = \widehat{m}$ for endorsement. Letting $\overline{R}$ represent the resulting expected revenue, and remembering that from Section 3.1 we are considering the choice of the good type of principal:

$$\overline{R} = \sum_{\mathbb{D}} \Pr\left[d|V = 1, \overline{\boldsymbol{\phi}}\right] \mu_d^{\widehat{m}} \sum_{k:\mu^k \geq \mu^{\widehat{m}}} p_1^k$$

The following lemma implies that $\sum_{\mathbb{D}} \Pr\left[d|V = 1, \overline{\boldsymbol{\phi}}\right] \mu_d^{\widehat{m}} > \mu^{\widehat{m}}$, so $\overline{R} > R\left(\phi_P^H = \phi_P^L\right) = \mu^{\widehat{m}} \sum_{k:\mu^k \geq \mu^{\widehat{m}}} p_1^k$.

Lemma 2 $\forall m$, $\sum_{\mathbb{D}} \Pr\left[d|V = 1, \overline{\boldsymbol{\phi}}\right] \mu_d^m > \mu^m$

Proof. See appendix. ■

Clearly, if the principal chooses $\overline{\boldsymbol{\phi}}$ but is unrestricted in her conditional pricing rule, her expected revenue will be at least $\overline{R}$. As $\overline{R} > R\left(\phi_P^H = \phi_P^L\right)$, this implies that any test with $\phi_P^H \neq \phi_P^L$ gives strictly higher expected revenue than choosing $\phi_P^H = \phi_P^L$, i.e., than choosing not to be tested, giving the following result.

Theorem 1 *The principal strictly prefers any test with $\phi_P^H \neq \phi_P^L$ to choosing $\phi_P^H = \phi_P^L$. Thus the principal always wants to be tested.*

The proof works by restricting the principal to choosing a conditional price to target the same set of agents as would be targeted in the absence of a test and showing that even under this restriction on pricing the principal prefers any given test to not being tested. *A fortiori* she prefers to be tested in the absence of this restriction.

To get some intuition for this result, suppose that the principal did not know her type with certainty, but instead shared belief $\mu^{\widehat{m}}$ with the lowest belief agents in the set she is forced to target under the pricing restriction. Then

$$\sum_{\mathbb{D}} \Pr\left[d|\overline{\phi}\right] \mu_d^{\widehat{m}} = \sum_{\mathbb{D}} \frac{\left(\Pr\left[d|V=1,\overline{\phi}\right]\mu^{\widehat{m}} + \Pr\left[d|V=0,\overline{\phi}\right]\left(1-\mu^{\widehat{m}}\right)\right)\Pr\left[d|V=1,\overline{\phi}\right]\mu^{\widehat{m}}}{\Pr\left[d|V=1,\overline{\phi}\right]\mu^{\widehat{m}} + \Pr\left[d|V=0,\overline{\phi}\right]\left(1-\mu^{\widehat{m}}\right)} = \mu^{\widehat{m}}$$

which means that

$$\Pr\left[P|\overline{\phi}\right]\left(\mu_P^{\widehat{m}} - \mu^{\widehat{m}}\right) = \Pr\left[F|\overline{\phi}\right]\left(\mu^{\widehat{m}} - \mu_F^{\widehat{m}}\right) \tag{6}$$

Thus under the conditional pricing restriction, testing would not affect expected revenues at all as the price rise following a pass would be exactly compensated by the price reduction following a fail. This occurs because the shift in beliefs following a decision exactly reflects the shared perceived probability of that decision.

However, we are considering the choice of a principal who knows her type to be good, and so knows that a pass is more likely than is believed to be the case by the agents. Thus, the left-hand side of (6) is now strictly greater than the right-hand side, so even when the restriction on the conditional pricing rule is imposed the principal strictly prefers any test to not being tested. Of course, unrestricted pricing will increase revenues further.

In conclusion, the ability to change price in response to the test decision, together with the principal's knowledge about her own type, ensures that the principal strictly prefers any given test to not being tested at all.

5 Optimal Choice of Test Type

Having established that the principal wishes to be tested, we next consider the principal's optimal choice of test type, remembering from Section 3.1 that we are considering the good type's choice.

First we define toughness. In the previous section we saw that when $\phi_P^H = \phi_P^L$, $\mu_P^m = \mu_F^m = \mu^m$ so the test does not change the agents' beliefs. When $\phi_P^H \neq \phi_P^L$, the test becomes informative: $\mu_P^m > \mu^m > \mu_F^m$, i.e., a pass raises agents' beliefs while a fail lowers them. Furthermore, it is straightforward to show that $\frac{\partial \mu_P^m}{\partial \phi_P^U} < 0$ and $\frac{\partial \mu_F^m}{\partial \phi_P^U} < 0$, i.e., as the test becomes more likely to return a pass on receiving an uninformative signal, the positive impact on agents' beliefs of passing the test becomes weaker while the negative impact of failing the test becomes stronger. On the other hand, $\frac{\partial \Pr[P|V,\phi]}{\partial \phi_P^U} > 0$ so the test is more likely to return a pass when the probability of passing on an uninformative signal is higher.

Thus, for fixed $\phi_P^H \neq \phi_P^L$, the principal faces a crucial trade-off: the higher ϕ_P^U the more likely the test is to be passed, but the weaker the positive impact of passing on beliefs (and hence on prices) and the stronger the negative impact of failing. Therefore we can think of tests with lower ϕ_P^U as being tougher while we think of those with higher ϕ_P^U as being softer. The polar cases of $\phi_P^U = 0$ and $\phi_P^U = 1$ provide the extremes.

Definition 1 *For fixed $\phi_P^H \neq \phi_P^L$, we define test "toughness" to be decreasing in the test's probability of returning a pass on receiving an uninformative signal $\left(\phi_P^U\right)$ and test "softness" to be increasing in this probability.*

One way of thinking about toughness is as representing the chosen test's innate bias towards or against the principal. A high value of ϕ_P^U indicates that the test is happy to pass the principal with high probability without good cause, a low value of ϕ_P^U similarly indicates a bias against the principal, and $\phi_P^U = \frac{1}{2}$ represents an unbiased test.

There is no corresponding notion of toughness in the case of ϕ_P^H or ϕ_P^L. A higher ϕ_P^H raises beliefs more after a pass but lowers them more after a fail, as does a lower ϕ_P^L. We will show later in this section that the principal, who knows her type to be good, always chooses to set $\phi_P^H = 1$ and $\phi_P^L = 0$, thus maximizing μ_P^m and minimizing μ_F^m for any given ϕ_P^U and hence ensuring that information is transmitted to the agents as clearly as possible.

Next we turn to a lemma outlining the convexity of $\Pr\left[d|V=1,\boldsymbol{\phi}\right]\mu_d^m$, which will be useful throughout the rest of this section.

Lemma 3 *For $\phi_P^H \neq \phi_P^L$, and for all $m \in \mathbb{X}$:*

(i) $\Pr\left[d|V=1,\boldsymbol{\phi}\right]\mu_d^m$ *is weakly convex in ϕ_P^z for all $z \in \mathbb{Z}$;*

(ii) $\Pr\left[d|V=1,\boldsymbol{\phi}\right]\mu_d^m$ *is strictly convex in ϕ_P^U;*

(iii) When $\phi_P^U < 1$, $\Pr\left[F|V=1,\boldsymbol{\phi}\right]\mu_F^m$ *is strictly convex in (a) ϕ_P^H and (b) ϕ_P^L;*

(iv) When $\phi_P^U > 0$, $\Pr\left[P|V=1,\boldsymbol{\phi}\right]\mu_P^m$ *is strictly convex in (a) ϕ_P^H and (b) ϕ_P^L.*

Proof. See appendix. ■

Lemma 3 allows us to establish that when $\phi_P^H \neq \phi_P^L$ expected revenue R is strictly convex in ϕ_P^z for all $z \in \mathbb{Z}$. The maximum of a set of convex functions is itself convex, while the maximum of a set of strictly convex functions is strictly convex (see Rockafellar, 1970). Either part (iii) or part (iv) must apply (or both), so using $\sum\limits_{k:\mu^k\geq\mu^m} p_1^k > 0$ it follows that $\max\limits_{m\in\mathbb{X}} \Pr\left[d|V=1,\boldsymbol{\phi}\right]\mu_d^m \sum\limits_{k:\mu^k\geq\mu^m} p_1^k$ is always strictly convex in ϕ_P^H and ϕ_P^L for at least one element of $\mathbb{D}$, while from part (i) the expression is always weakly convex. From part (ii) it is always strictly convex in ϕ_P^U. Thus, using the fact that summing a weakly convex and strictly convex function returns a strictly convex function, (5) itself must always be strictly convex.

This strict convexity implies that for fixed $\phi_P^H \neq \phi_P^L$, the good type of principal maximizes revenue at an extreme choice of test, selecting either the toughest $\left(\phi_P^U = 0\right)$ or softest $\left(\phi_P^U = 1\right)$ test available. Together with Theorem 1, the strict convexity implies that for fixed ϕ_P^U, the principal maximizes revenue by selecting $\left\{\phi_P^H = 1, \phi_P^L = 0\right\}$, so the test always passes on receiving a high signal and fails on receiving a low signal. The following lemma summarizes.

Lemma 4 *(i) When $\phi_P^H \neq \phi_P^L$, expected revenue R is strictly convex in ϕ_P^z for all $z \in \mathbb{Z}$;*

(ii) For fixed $\phi_P^H \neq \phi_P^L$, R is maximized at, and only at, an extreme choice of $\phi_P^U = 0$ and/or $\phi_P^U = 1$, i.e., either the toughest or softest test possible is optimal;

(iii) For a fixed ϕ_P^U, R is maximized at, and only at, a choice of $\left\{\phi_P^H = 1, \phi_P^L = 0\right\}$, i.e., the optimal test always returns a pass on receiving a high signal and a fail on receiving a low signal.

Proof. See appendix. ■

At first sight, the result in part (ii) might appear counter-intuitive: it is not immediately clear why the trade-off between beliefs (and hence prices) and the probability of passing should always be resolved at an extreme. To provide an intuition note first that all beliefs and hence prices μ_d^m are decreasing in ϕ_P^U: a pass raises beliefs more the tougher the test, while a fail is not so damaging to beliefs if the test is tougher. Furthermore, μ_P^m can be shown to be strictly convex in ϕ_P^U, so conditional on passing an increase in toughness (decrease in ϕ_P^U) is more powerful where the test is tougher. This convexity can be seen by using (3), (1) and $p_1^U = p_0^U$ to re-write

$$\begin{aligned}\mu_P^m &= \frac{\Pr[P|V=1,\boldsymbol{\phi}]\mu^m}{\Pr[P|V=1,\boldsymbol{\phi}]-(\Pr[P|V=1,\boldsymbol{\phi}]-\Pr[P|V=0,\boldsymbol{\phi}])(1-\mu^m)} \\ &= \frac{\Pr[P|V=1,\boldsymbol{\phi}]\mu^m}{\Pr[P|V=1,\boldsymbol{\phi}]-\left(p_1^H-p_0^H\right)\left(\phi_P^H-\phi_P^L\right)(1-\mu^m)}\end{aligned}$$

and noting that $\Pr\left[P|V=1,\boldsymbol{\phi}\right]$ is linear in ϕ_P^U. The convexity of μ_P^m implies that $\max\limits_{m\in\mathbb{X}} \mu_P^m \sum\limits_{k:\mu^k\geq\mu^m} p_1^k$ is also strictly convex. Overall, $\Pr\left[P|V=1,\boldsymbol{\phi}\right] \max\limits_{m\in\mathbb{X}} \mu_P^m \sum\limits_{k:\mu^k\geq\mu^m} p_1^k$ is strictly convex in ϕ_P^U. To maximize this, the good type of principal would choose an extreme ϕ_P^U, either setting $\phi_P^U = 0$ to benefit from a steep increase in prices (given μ_P^m is decreasing and convex in ϕ_P^U) while the probability of passing falls linearly as ϕ_P^U goes towards zero, or setting $\phi_P^U = 1$ to benefit from a linear increase in the probability of passing while prices do not fall much as ϕ_P^U goes towards one.

To maximize $\Pr[F|V=1,\boldsymbol{\phi}] \max_{m\in\mathbb{X}} \mu_F^m \sum_{k:\mu^k\geq\mu^m} p_1^k$, the principal would want to set $\phi_P^U = 0$, as both $\Pr[F|V=1,\boldsymbol{\phi}]$ and the prices μ_F^m are decreasing in ϕ_P^U. As in the pass case, this expression is strictly convex in ϕ_P^U. Summing over the pass and fail cases, (4) is then strictly convex, so the good type of principal will choose either the extreme tough test, to maximizes prices at the cost of a lower probability of passing, or the extreme soft test, to maximize the probability of passing at the cost of lower prices.

Regarding part (iii), a good type of principal wants information to be transmitted to the agents as clearly as possible, as she is certain of her type, whereas agents always think that there is some chance that the principal is a bad type, whatever private signals they receive. Thus, she wants the test to return a pass on a high signal and a fail on a low signal, i.e., she wants to set $\phi_P^H = 1$ and $\phi_P^L = 0$ to maximize μ_P^m and minimize μ_F^m, thus making the pass and fail decisions as informative as possible. If not, the agents will always think there is some chance that a pass followed a low signal and vice-versa, adding noise to the information transmission process. At the opposite limit, when ϕ_P^H and ϕ_P^L tend towards one another we saw in Section 4 that the test becomes completely uninformative, whatever the choice of ϕ_P^U.[9]

From Lemma 4, expected revenue R can be raised if we do not have $\phi_P^H = 1$, $\phi_P^L = 0$ and $\phi_P^U \in \{0,1\}$, and by the continuity of R in ϕ_P^z for all $z \in \mathbb{Z}$ established in Section 3.2 a $\boldsymbol{\phi}$ which maximizes R must exist. Thus we have fully characterized the optimal choice of test for the good type of principal.

[9] This discussion clarifies why a simpler model in which the test receives a binary signal (i.e., $p_1^U = p_0^U = 0$) is of little interest. The principal would simply select $\phi_P^H = 1$ and $\phi_P^L = 0$, thus choosing a test which perfectly reveals its signal. The key trade-off faced by the principal, namely the choice of test toughness on receiving an uninformative signal, would be lost. Thus the coarse nature of the test's report is crucial to the analysis. A formal proof that $\phi_P^H = 1$ and $\phi_P^L = 0$ remains optimal when $p_1^U = p_0^U = 0$ would proceed as in the main text, but noting that from the proof of Lemma 3(iii) & (iv), $\Pr[d|V=1,\boldsymbol{\phi}]\,\mu_P^m$ is strictly convex in ϕ_P^H for at least one element of $\mathbb{D}$ as $\phi_P^L > 0$ or $\phi_F^L > 0$ (or both), and similarly for convexity in ϕ_P^L, and also noting that Theorem 1 continues to apply when $p_1^U = p_0^U = 0$.

Theorem 2 *The principal always selects* $\phi_P^H = 1$, $\phi_P^L = 0$ *and* $\phi_P^U \in \{0, 1\}$, *i.e., the principal chooses a test which is as tough or as soft as possible on receiving an uninformative signal, and which always returns a pass on receiving a high signal and a fail on receiving a low signal.*

In conclusion, we have discovered that the principal will always select the toughest or softest test available. Moreover, applying the strict convexity of R from Lemma 4(i) allows us to show that for any finite set of ϕ_P^U values of size $S > 2$ such that $\phi_{P1}^U > \phi_{P2}^U > ... > \phi_{PS}^U$, and for $\phi_P^H \neq \phi_P^L$, ϕ_{P1}^U or ϕ_{PS}^U would be strictly preferred by the principal to any intermediate test with $\phi_{Ps}^U \notin \left\{\phi_{P1}^U, \phi_{PS}^U\right\}$. (Given the choice of ϕ_{Ps}^U, Lemma 4(iii) continues to show that $\left\{\phi_P^H = 1, \phi_P^L = 0\right\}$ is optimal.) Hence, restricting the choice range from a continuum of test types to a finite set leaves extreme tests preferred.

The next section examines the trade-off faced by the principal in more detail, determining when each of the two candidates for the optimum is best. We will see that the prior and the probabilities of receiving different signals for both the test and agents have a profound impact on the optimal choice of test toughness.

6 Pairwise Comparison

We have just seen that the principal always chooses $\phi_P^H = 1$ and $\phi_P^L = 0$ together with either the toughest ($\phi_P^U = 0$) or softest ($\phi_P^U = 1$) test possible. In this section we analyze the principal's choice between the toughest and softest tests, maintaining $\phi_P^H = 1$ and $\phi_P^L = 0$ throughout. We first derive analytical results when agents' private information is sufficiently informative or sufficiently uninformative. We then provide some numerical examples to illustrate the complexity of the principal's choice more generally and to provide some feel for the importance of choosing the optimal test.

6.1 Analytical Results

As we are going to vary the p_V^m's, we select an arbitrary but constant numbering of the M agent signals. Let $R\left(\phi_P^U\right)$ denote the good type of principal's revenue when she chooses $\boldsymbol{\phi} = \left\{1, \phi_P^U, 0\right\}$. Suppose that we restrict the principal to targeting agents with signals at least as strong as $X_i = m$

for endorsement whatever the decision d by setting price $\lambda = \mu_d^m$. Referring back to (3), this price nonetheless varies in the decision and chosen test type. Let

$$R\left(\phi_P^U, m\right) \equiv \sum_{\mathbb{D}} \Pr\left[d|V=1, \boldsymbol{\phi} = \left\{1, \phi_P^U, 0\right\}\right] \mu_d^m \sum_{k:\mu^k \geq \mu^m} p_1^k$$

denote revenue under this restriction. The following lemma tells us whether the toughest or softest test is optimal under such a restriction.

Lemma 5 $R(1,m) \gtreqless R(0,m) \Leftrightarrow 2(\mu^m - 1)\left(p_1^H - p_0^H\right) + p_1^H \gtreqless p_1^L$

Proof. See appendix. ■

Let γ represent the vector $\left(p_1^0,\ p_1^1, ...,\ p_1^M,\ p_0^0,\ p_0^1, ...,\ p_0^M\right)$, which outlines the probabilities of the agents receiving the different private signals. Let $p_V^m(\gamma)$ and $\mu^m(\gamma)$ denote the values of p_V^m and μ^m at a particular γ. Let Γ be the set of allowable γ, i.e., those for which all $p_V^m \in (0,1]$.

Suppose that at a given γ, denoted by $\overline{\gamma}$, the agents' signals are all fully informative. This implies that a subset of the signals $\mathbb{X}^+ \neq \emptyset$ are "positive" with $p_1^m(\overline{\gamma}) > 0$ and $p_0^m(\overline{\gamma}) = 0$ so $\mu^m(\overline{\gamma}) = 1$ and a subset $\mathbb{X}^- \neq \emptyset$ are "negative" with $p_1^m(\overline{\gamma}) = 0$ and $p_0^m(\overline{\gamma}) > 0$ so $\mu^m(\overline{\gamma}) = 0$, with $\mathbb{X}^+ \cup \mathbb{X}^- = \mathbb{X}$. Let $\overline{\Gamma}$ be the set of fully informative γ's. Our model does not permit $\gamma \in \overline{\Gamma}$, i.e., $\Gamma \cap \overline{\Gamma} = \emptyset$, but we can look at what happens as $\gamma \in \Gamma$ approaches $\overline{\gamma} \in \overline{\Gamma}$.[10]

Suppose instead that at a given γ, denoted by $\underline{\gamma}$, the agents' signals are all fully uninformative (so the agents receive no private information). This would imply that $\forall m$, $p_1^m(\underline{\gamma}) = p_0^m(\underline{\gamma}) > 0$ so $\mu^m(\underline{\gamma}) = q$. Let $\underline{\Gamma}$ be the set of fully uninformative γ. Our model permits $\gamma \in \underline{\Gamma}$ as a special case, i.e., $\underline{\Gamma} \subset \Gamma$.

Remember that from Section 3.1 we are considering the choice of the good type of principal. The next theorem outlines the good type's preference between the toughest and softest tests when agent signals are sufficiently informative or sufficiently uninformative.

[10] Permitting fully informative agent signals would make the choice of test irrelevant: a good type of principal would know that $X_i \in \mathbb{X}^+$ must be received. As $\mu^m(\overline{\gamma}) = 1\ \forall m \in \mathbb{X}^+$, $R = 1$ whatever the test type. A bad type of principal could do nothing about agents receiving $X_i \in \mathbb{X}^-$, so it would receive no revenue.

Theorem 3

(i) For $\gamma \in \Gamma$ sufficiently close to (but distinct from) any arbitrary $\overline{\gamma}$, so agent signals are sufficiently informative, the principal strictly prefers the softest test when $p_1^H > p_1^L$ and strictly prefers the toughest test when $p_1^H \leq p_1^L$.

(ii) For $\gamma \in \Gamma \backslash \underline{\Gamma}$ sufficiently close to (but distinct from) any arbitrary $\underline{\gamma}$, so agent signals are sufficiently uninformative, the principal strictly prefers the softest test when the prior $q > \overline{q} = \frac{1}{2} + \frac{p_1^L - p_0^H}{2(p_1^H - p_0^H)}$ and strictly prefers the toughest test when $q \leq \overline{q}$.

(iii) For $\gamma \in \underline{\Gamma}$, so agent signals are fully uninformative, the preference is the same as in (ii), except that the principal is indifferent between the softest and toughest tests when $q = \overline{q}$.

Proof. See appendix. ∎

If we impose a weak symmetry assumption on the structure of the signals received by the test, these results can be considerably simplified. Setting $p_1^H = p_0^L$ means assuming that the probability of the test receiving the positive signal H when $V = 1$ is the same as the probability of receiving the negative signal L when $V = 0$. Since $p_1^U = p_0^U$, this immediately implies $p_1^L = p_0^H$, so the probability of the test receiving the negative signal when $V = 1$ is the same as the probability of receiving the positive signal when $V = 0$. Under such symmetry, $p_1^H > p_1^L$ as $p_1^H > p_0^H$ always given H is a high signal, and clearly $\overline{q} = \frac{1}{2}$, giving the following:

Corollary 1 (to Theorem 3)

Assuming symmetry of the test's signals, in the sense that $p_1^H = p_0^L$:

(i) For sufficiently informative agent signals, the principal strictly prefers the softest test.

(ii) For sufficiently uninformative agent signals, the principal strictly prefers the softest test when the prior $q > \overline{q} = \frac{1}{2}$ and strictly prefers the toughest test when $q \leq \overline{q} = \frac{1}{2}$.

(iii) For fully uninformative agent signals, the preference is the same as in (ii), except that the principal is indifferent between the softest and toughest tests when $q = \overline{q} = \frac{1}{2}$.

When agent signals are sufficiently informative, the proof of Theorem 3 shows that the principal targets all agents with positive signals for endorsement. Part (i) of Corollary 1 reveals that the principal then prefers the softest test when the signals received by the test are symmetric. The good type of principal expects agents to start with strong positive beliefs about her type, so there is little upside from risking the toughest test. Part (i) of the Theorem states that this preference for the softest test generalizes to asymmetric test signals, unless these signals are so asymmetric that p_1^L rises above p_1^H, in which case the toughest test is preferred. Remember that L is a low signal so $p_1^L < p_0^L$. Under symmetry $p_1^H = p_0^L$, so $p_1^L < p_1^H$. For this inequality to be reversed, the signals need to become so asymmetric that in the good state the low signal is more likely to be received by the test than the high signal. In this peculiar scenario, the good type of principal is quite likely to fail the test, so avoiding the worst case outcome of failing the softest test becomes paramount, despite the fact that the high p_1^L and low p_1^H mean the agents place less weight on the test's decision.

When agent signals are sufficiently or fully uninformative, the proof of Theorem 3 shows that the principal targets every agent for endorsement. Furthermore, because the agents place little or no weight on their private information, agents start with beliefs close to the prior q. Parts (ii) and (iii) of Theorem 3 tell us that the principal prefers the toughest test when the prior is below a threshold $\overline{q}$ and the softest test when the prior is above this threshold. The advantage of the toughest test is the strong impact of a pass on beliefs, which is more valuable when these beliefs start low and so have more scope to rise. When beliefs start higher, the principal plumps for the softest test which is likely to be passed by the principal who knows her type to be good. The astute reader might wonder about the negative impact of failing such a test: the second numerical example below illustrates that this risk shrinks the softest test's relative advantage when compared to the toughest test's advantage for low q.

When the signals received by the test are symmetric, parts (ii) and (iii) of Corollary 1 inform us that the threshold $\overline{q} = \frac{1}{2}$. When signals become asymmetric, this threshold shifts around. Starting from a position of symmetry with $p_1^L = p_0^H$, we can introduce asymmetry by raising or lowering p_0^L while shifting p_0^H by an equal magnitude but in the opposite direction and keeping other parameters constant. Alternatively, we can move p_1^L and p_1^H in opposite directions. Let us first consider what happens when we change p_0^L by $\triangle_0$ and p_0^H by $-\triangle_0$. The threshold then shifts to become $\overline{q} = \frac{1}{2} + \frac{\triangle_0}{2\left(p_1^H - \left(p_0^H - \triangle_0\right)\right)}$ (using the initial value of p_0^H). Remembering that $p_1^H - p_0^H > 0$, $\frac{\partial \overline{q}}{\partial \triangle_0} = \frac{p_1^H - p_0^H}{2\left(p_1^H - \left(p_0^H - \triangle_0\right)\right)^2} > 0$. As a result, any $\triangle_0 > 0$, which corresponds to any size of increase in p_0^L and decrease in p_0^H, causes $\overline{q}$ to rise. For the good type of principal the probability of passing either test stays the same, while after the change the agents put more weight on the test result (as the signal received by the test has become more accurate in the bad state). The boost from passing the toughest test thus becomes stronger, while failing the softest test becomes more damaging, leaving the principal to prefer the toughest test over a greater range of q, i.e., $\overline{q}$ rises. When $\triangle_0 < 0$, so p_0^L falls and p_0^H rises, the effects are reversed and $\overline{q}$ falls.

Again starting from symmetry, when we change p_1^L by $\triangle_1$ and p_1^H by $-\triangle_1$ the threshold becomes $\overline{q} = \frac{1}{2} + \frac{\triangle_1}{2((p_1^H - \triangle_1) - p_0^H)}$ (using the initial value of p_1^H), so $\frac{\partial \overline{q}}{\partial \triangle_1} = \frac{p_1^H - p_0^H}{2((p_1^H - \triangle_1) - p_0^H)^2} > 0$. As a result, any $\triangle_1 > 0$, which corresponds to any size of increase in p_1^L and decrease in p_1^H, causes $\overline{q}$ to rise. After the change, the agents put less weight on the test result (as the signal received by the test has become less accurate in the good state). Analogously to the case where $\triangle_0 < 0$, we might expect this to make the softest test more attractive. However, there is a now a countervailing effect which happens to dominate: the good type of principal understands that her probability of failing either test has risen, making the softer test less attractive due to the damage from failing such a test.[11] Overall, the principal prefers the toughest test over a greater range, i.e., $\overline{q}$ rises. When $\triangle_1 < 0$, so p_1^L falls and p_1^H rises, the effects are reversed and $\overline{q}$ falls.

6.2 Numerical Examples

Providing a general description of the principal's choice between the toughest and softest tests when the agents' private information is of intermediate quality remains intractable (as illustrated in footnote 13). Instead, this section describes the principal's problem using numerical methods in a simplified environment. We should make clear, however, that a principal in her specific environment can discover her optimal strategy: for *given parameter values* the principal's choice of prices and hence test is straightforward to calculate using (5).

In order to illustrate the principal's price choice rule and revenue functions in numerical plots, we need to restrict the number of parameters to two. To do this, we first assume that the agents draw their signal from a binary set ($M = 1$) and that the signals are symmetric, so $p_1^1 = p_0^0 = p_A$ and $p_0^1 = p_1^0 = 1 - p_A$ for $p_A \in \left[\frac{1}{2}, 1\right)$ which measures the precision or informativeness of the agents' private information. Clearly, $m = 1$ is a "high" signal as $\mu^1 \geq q$ while $m = 0$ is a "low"

[11] If we change p_0^L, p_0^H, p_1^L and p_1^H simultaneously, setting $\triangle_1 > 0$ and $\triangle_0 = -\triangle_1 < 0$ to retain symmetry, the increase in the probability of failing from $\triangle_1 > 0$ pushes towards the toughest test, while the push towards the softest test from the agents putting less weight on the test result now arises from both the $\triangle_1$ and $\triangle_0$ changes. The effects exactly cancel, so $\overline{q}$ remains at $\frac{1}{2}$, as outlined in parts (ii) and (iii) of Corollary 1. The argument is reversed if $\triangle_1 < 0$ and $\triangle_0 = -\triangle_1 > 0$.

signal as $\mu^0 \le q$. We further assume that the test's signal structure has the following symmetric form: $p_1^H = p_0^L = (p_T)^2$, $p_1^U = p_0^U = 2p_T(1-p_T)$ and $p_0^H = p_1^L = (1-p_T)^2$ for $p_T \in \left(\frac{1}{2}, 1\right)$. One interpretation of this structure is that the test receives two conditionally independent draws, each of precision p_T, from a binary signal set, with H corresponding to two high draws, L two low draws and U a high and low draw (which cancel out). Whatever the interpretation, the precision of the test's signal structure increases in p_T in the sense that p_1^H rises while p_1^U and p_1^L fall. In our first example, we also set $q = \frac{1}{2}$ so the agents' prior is uninformative. In the second example, we set $p_A = \frac{1}{2}$, so agents receive no private information, but we allow a general prior $q \in (0,1)$.

6.2.1 Example 1

In this example, p_A and p_T are the model's parameters and the prior $q = \frac{1}{2}$. We begin by providing a numerical analysis of the principal's choice of price. Conditional on the chosen test type $\phi_P^U \in \{0,1\}$ and the test decision d, the principal chooses between setting a low price $\lambda = \mu_d^0$ to target all the agents for endorsement and a higher price $\lambda = \mu_d^1$ to target only those agents who observed the high private signal. For any choice of softest or toughest test and a test decision, there will a unique threshold $\overline{p}_A\left(p_T, \phi_P^U \in \{0,1\}, d\right) \in \left(\frac{1}{2}, 1\right)$ such that the principal chooses the low price for $p_A < \overline{p}_A$ and the high price for $p_A > \overline{p}_A$. From the proof of Theorem 3 we know that for sufficiently low p_A the principal selects the low price while for sufficiently high p_A she selects the high price. Furthermore, as p_A rises the high price becomes relatively more attractive: more agents receive the high signal (as we are considering the choice of the good type of principal), while μ_d^1 rises and μ_d^0 falls as the private signal is becoming more informative. Thus a unique threshold must exist.[12]

We can write $\overline{p}_A$ as an implicit function (by setting $\mu_d^1 p_A - \mu_d^0 = 0$) and Figure 1 plots this function numerically for the four possible cases: the gray dotted line when the softest test is failed, the gray full line when the toughest test is failed, the black dotted line when the softest test is passed and finally the black full line when the toughest test is passed.

[12] Such unique thresholds will continue to exist for any agent and test signal structure, and prior q, so long as private signals remain binary.

Figure 1: $\overline{p}_A$ as a function of p_T

Gray dotted line: $\overline{p}_A\left(p_T, 1, F\right)$; Gray full line: $\overline{p}_A\left(p_T, 0, F\right)$

Black dotted line: $\overline{p}_A\left(p_T, 1, P\right)$; Black full line: $\overline{p}_A\left(p_T, 0, P\right)$

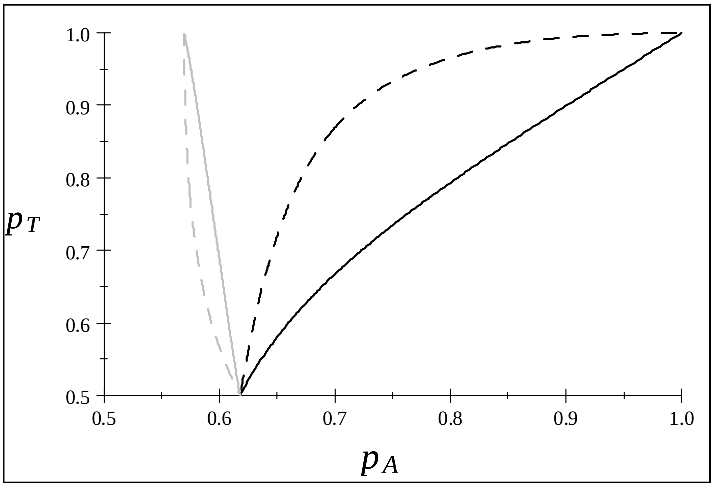

Moving horizontally across Figure 1 for any p_T, we observe that the order of the thresholds is preserved: the better the news from the test result across the four cases, the higher the $\overline{p}_A$ threshold at which the principal switches to targeting just the agents with high private signals. Moving vertically, we also see that better news from the test raises the threshold: as p_T rises (so signals received by the test become more precise) passes become better news, so the two right-hand curves show a positive slope, while fails become worse news, so the two left-hand curves exhibit a negative slope. The effect is intuitive: the beliefs of the agents receiving low signals start lower, and so tend to move up relatively more as the news from the test result improves, explaining the greater relative attraction of targeting all the agents for endorsement. Interestingly in the region between the two full lines, having chosen the toughest test the principal targets all agents following a pass, but after a fail targets only those agents who are more optimistic, having received a high private signal. The same occurs for the softest test between the two dotted lines.

This numerical analysis tells us which price the principal sets given ϕ_P^U and d. Providing a general description of the principal's choice of test remains intractable.[13] However, we can illustrate the principal's choice of test using the following numerical plots. When interpreting these, remember that the *maximum possible revenue is 1*, corresponding to the principal being

[13] For example, if we attempt to determine the choice of test in the region between the two black lines in Figure 1, in which the high price is chosen unless the toughest test is passed, factoring the difference in revenues returns a multivariate polynomial in the numerator of order seven in p_T and order four in p_A.

endorsed by all agents at a price of 1 when all the agents are completely convinced that the principal is of a good type.

Figure 2: $R(1) - R(0)$

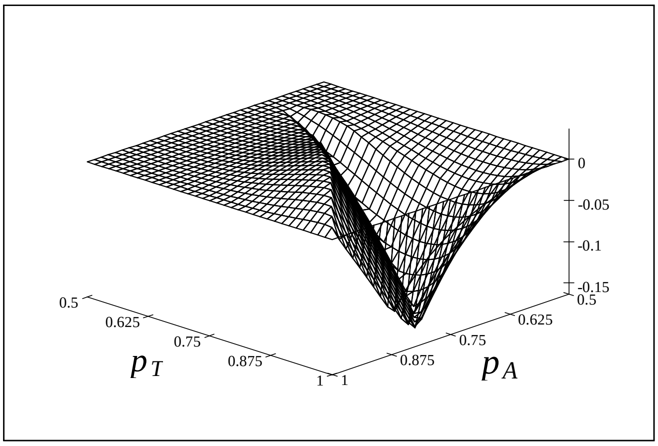

Figure 3: $\max\{R(1), R(0)\} - R(No\ Test)$

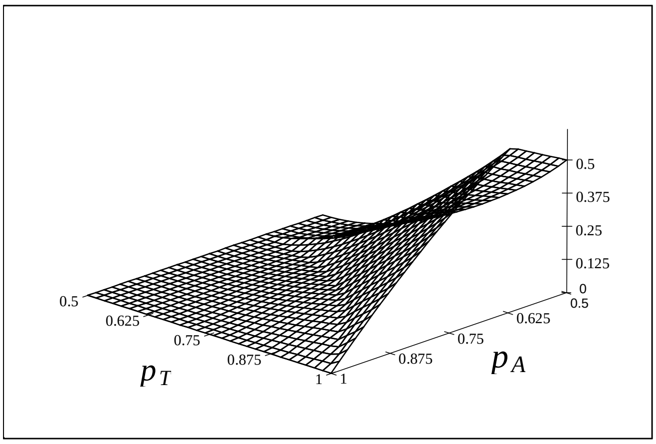

Figure 4: $R(1) - R(0)$ for $p_T = 0.8$

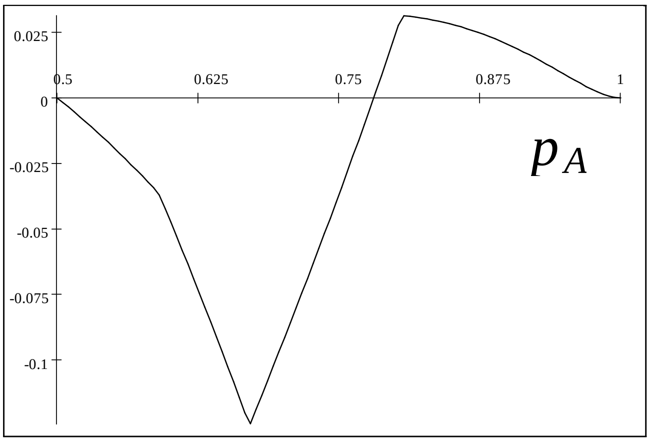

Figure 5: $R(1) - R(0)$ for $p_T = 0.55$

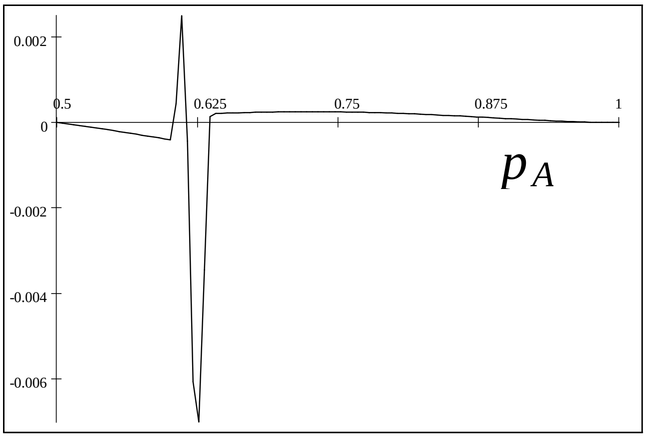

Figure 2 shows the difference in revenue between choosing the softest test and choosing the toughest test. First, note that the choice of test matters a lot: depending on the parameters the difference in revenue can reach over 15% of the maximum possible revenue.

Applying Corollary 1 to the example considered here informs us that the toughest test is preferred for sufficiently low precision of agents' private information p_A while the softest test is preferred for sufficiently high p_A. Figure 2 of course confirms this, but also shows that over most of the range of p_T (the precision of the signals received by the test) the toughest test is preferred until p_A becomes sufficiently large, after which the softest test is optimal, so the result from Corollary 1 extends in a natural way to intermediate p_A. Figure 4 exemplifies this when $p_T = 0.8$. However, for very low p_T the principal's preference becomes more complicated, switching from tough to soft to tough and then back to soft again as p_A rises. This complexity, which is indicative of the analytical intractability of the principal's choice problem in general, is illustrated in Figure 5 for $p_T = 0.55$. Note that each of the turning points in Figure 5 matches one of the $\overline{p}_A$ thresholds from Figure 1 at which the principal's pricing rule switches from the low to the high price in one of the four cases.

Figure 2 also shows that when the toughest test is preferred, choosing the better test tends to be more important than when the softest test is optimal. When p_A is high, most agents receive the high signal, and such a signal is very informative. Thus the softest test tends to be preferred as there is no need to risk a tough test to push beliefs up strongly in the event of a pass. The advantage of the softest test is that passing is very likely; however, the scope for beliefs to fall after failing a soft test shrinks the advantage relative to the toughest test. On the other hand, the toughest test tends to be a strong favorite when it is preferred due the potential for pushing low starting beliefs up a lot on passing, while beliefs don't fall much after a fail.

Figure 3 illustrates the difference in revenue between choosing the optimal test and not being tested at all ($R(No\ Test) = \max\left\{\mu^1 p_A, \mu^0\right\}$). We know from Theorem 1 that this difference must be positive. The figure shows that, unsurprisingly, as the precision p_T of the test's signal structure goes up, choosing to be tested becomes more important. Figure 2 shows that correctly choosing between the toughest and softest tests also tends to become more important as p_T rises, except that for p_T very close to 1 the choice starts to matter less and less.[14] This effect arises because in this example, as p_T tends to 1 the probability of the test receiving the uninformative signal goes to zero, so the difference between soft and tough tests becomes insignificant.

6.2.2 Example 2

Our first example above gives no feel for the importance of the prior. Here we consider a general prior, but set $p_A = \frac{1}{2}$ so agents receive no private information and the model's parameters are p_T and q. The principal no longer needs to select between a high and low price, explaining why applying Corollary 1(iii) provides an analytical description of the principal's choice of test for

[14] Figure 2 does not show $R(1)-R(0)$ clearly for p_T very close to 1 due to the limited number of points we can plot while retaining overall plot clarity. Our model does not permit $p_T = 1$: the choice of test would then become irrelevant as there would be no uninformative signal. However, our analytical expressions are continuous for $p_T \in \left[\frac{1}{2}, 1\right]$, so we can infer that the choice of test matters less and less as p_T tends to one.

this example. When interpreting the following plots, remember again that the maximum possible revenue is 1.

Figure 6: $R(1) - R(0)$

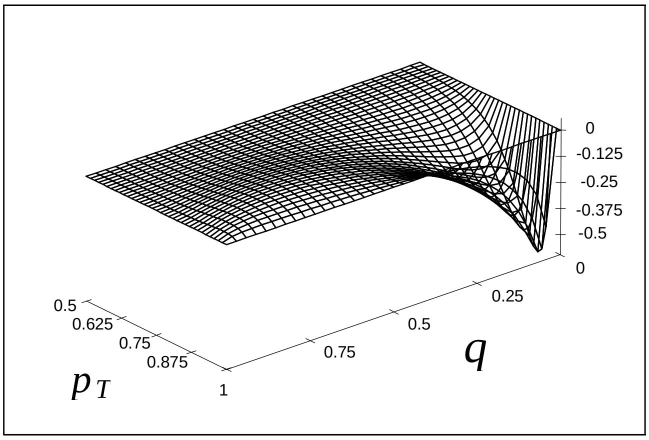

Figure 7: $\max\{R(1), R(0)\} - R(No\ Test)$

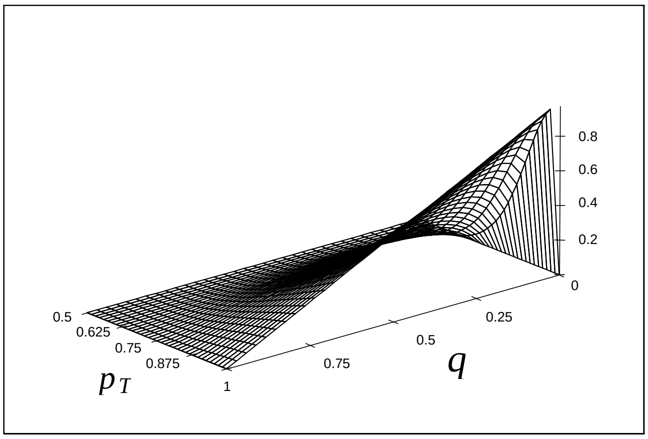

Figure 6 confirms the result from Corollary 1(iii): the toughest test is preferred for $q < \frac{1}{2}$ and the softest test is preferred for $q > \frac{1}{2}$. The difference in revenue between choosing the toughest and softest tests is bigger when q is quite low and p_T is quite high (exceeding 50% of the maximum revenue in some cases): the scope for raising beliefs and hence prices is high and the test receives good quality signals, so choosing a tough test becomes very attractive for the good type of principal. When q is high, the softest test is best, but the scope for beliefs to fall after failing such a test shrinks the advantage relative to the toughest test.

Figure 7 illustrates the difference in revenue between choosing the optimal test and not being tested at all. As dictated by Theorem 1, this difference is always positive. Similarly to Example 1, the importance of being tested goes up with p_T, and, consistent with the discussion of Figure 6, the importance of the test is higher when q is low.

7 Conclusion

The results in this paper provide an integration of two key choices for any principal (firm, job market candidate, politician) hoping to convince agents (consumers, employers, supporters) that she is a good type: the choice of a price for her services, and the use of public testing as an early marketing strategy.

Quite apart from any standard signalling arguments, we have found that a principal whose type is unknown to agents will tend towards facing extreme tests for assessment, whether this be reviewers with known biases, or tests or referees that are publicly known to be extremely tough or soft. We have also seen that avoiding the testing process is not optimal, so testing is a complement to optimal pricing. After the test result is known, the principal can then select an optimal price, and agents will endorse (purchase, hire or support) based on the price, test result, test type and their own private information. Perhaps most remarkably, in many cases it is the toughest test that will be best for the principal in expectation. This would appear to agents as though the principal is opting to put herself through a "baptism of fire", hoping for a powerful pass from a tough test. By enabling the principal to set a very high price following such a pass, the tough test increases expected revenue while keeping the impact of a fail to a minimum. Alternatively, when the principal subjects herself to the softest test, which entails a high probability of a pass, this would appear much like garnering the support of a "yes-man" or subjecting herself to scrutiny from a body strongly biased in favor of her policies or products.

We examine the conditions under which the toughest or softest test will be optimal for the principal and find that much depends upon the informativeness of the signals received by the agents and tests, and on the size of the common prior belief that the principal is a good type. Loosely, if the private signals received by the agents are relatively uninformative and the agents start with a prior which is unfavorable to the good type, which might correspond to an innovative product or idea, the principal prefers the toughest test to launch the product or idea with a bang on passing. If, however, the agents receive relatively informative private signals, instead the principal selects the softest test (as long as in the good state the tests do not receive low signals more frequently than high signals). This might correspond to a well-known product or idea, where the principal does not need to risk a tough test and instead chooses the softest test to maximize the probability of passing.

A priori we might think principals are likely to want to avoid tough tests, reviewers who are biased against them, tough accreditation standards, difficult academic qualifications or known tough referees. However, the results in this paper show that despite tough tests being difficult to pass, the tremendous gains when a pass is obtained might be enough to make them popular with principals. On the other hand, we can also better understand the popularity of soft tests, friendly referees, soft review journals or easy reviewers as a function of the quality of the diverse sources of information available.

Appendix

Proof of Lemma 2. From (3)

$$\sum_{\mathbb{D}} \Pr\left[d|V=1,\overline{\boldsymbol{\phi}}\right]\mu_d^m = \sum_{\mathbb{D}} \frac{\left(\Pr\left[d|V=1,\overline{\boldsymbol{\phi}}\right]\right)^2\mu^m}{\Pr\left[d|V=1,\overline{\boldsymbol{\phi}}\right]\mu^m+\Pr\left[d|V=0,\overline{\boldsymbol{\phi}}\right](1-\mu^m)}$$

Let $a \equiv \Pr\left[P|V=1,\overline{\boldsymbol{\phi}}\right]$, so $1-a = \Pr\left[F|V=1,\overline{\boldsymbol{\phi}}\right]$, and $b \equiv \Pr\left[P|V=0,\overline{\boldsymbol{\phi}}\right]$, so $1-b = \Pr\left[F|V=0,\overline{\boldsymbol{\phi}}\right]$. As $\phi_P^H \neq \phi_P^L$, $\Pr\left[d|V,\overline{\boldsymbol{\phi}}\right] > 0$, so both denominators in the expression above are strictly positive. Thus we need to show that:

$$\begin{aligned} &a^2\left((1-a)\mu^m + (1-b)\left(1-\mu^m\right)\right) + (1-a)^2\left(a\mu^m + b\left(1-\mu^m\right)\right) > \\ &a(1-a)\left(\mu^m\right)^2 + b\left(1-b\right)\left(1-\mu^m\right)^2 + \left(a(1-b)+b\left(1-a\right)\right)\mu^m\left(1-\mu^m\right) \end{aligned}$$

Re-arranging, this collapses to $(a-b)^2\left(1-\mu^m\right)^2 > 0$. From (2), $\mu^m \in (0,1)$. Using $p_1^U = p_0^U$

$$a-b = \phi_P^H\left(p_1^H - p_0^H\right) - \phi_P^L\left(p_0^L - p_1^L\right)$$

so, using $\phi_P^H \neq \phi_P^L$ together with (1), $a-b \neq 0$. ■

Proof of Lemma 3. Let $\overline{z}$ represent a specific member of $\mathbb{Z}$. Then, using (3),

$$\Pr\left[d|V=1,\boldsymbol{\phi}\right]\mu_d^m = \frac{\left(p_1^{\overline{z}}\phi_d^{\overline{z}}+\sum_{\mathbb{Z}\setminus\overline{z}} p_1^z\phi_d^z\right)^2\mu^m}{\left(p_1^{\overline{z}}\phi_d^{\overline{z}}+\sum_{\mathbb{Z}\setminus\overline{z}} p_1^z\phi_d^z\right)\mu^m+\left(p_0^{\overline{z}}\phi_d^{\overline{z}}+\sum_{\mathbb{Z}\setminus\overline{z}} p_0^z\phi_d^z\right)(1-\mu^m)}$$

The denominator is strictly positive as $\phi_P^H \neq \phi_P^L$ implies $\phi_d^H \neq \phi_d^L$ given $\sum_{\mathbb{D}}\phi_d^z = 1$, so $\sum_{\mathbb{Z}} p_V^z\phi_d^z > 0$ given $p_V^z > 0$. We can show that:

$$\frac{\partial^2 \Pr[d|V=1,\boldsymbol{\phi}]\mu_d^m}{\partial\left(\phi_d^{\overline{z}}\right)^2} = \frac{2\mu^m(1-\mu^m)^2\left(\sum_{\mathbb{Z}\setminus\overline{z}}\phi_d^z\left(p_0^{\overline{z}}p_1^z - p_0^z p_1^{\overline{z}}\right)\right)^2}{\left(\left(p_1^{\overline{z}}\phi_d^{\overline{z}}+\sum_{\mathbb{Z}\setminus\overline{z}} p_1^z\phi_d^z\right)\mu^m+\left(p_0^{\overline{z}}\phi_d^{\overline{z}}+\sum_{\mathbb{Z}\setminus\overline{z}} p_0^z\phi_d^z\right)(1-\mu^m)\right)^3} \geq 0 \tag{7}$$

Now, $\frac{\partial \Pr[d|V=1,\boldsymbol{\phi}]\mu_d^m}{\partial\left(1-\phi_d^{\overline{z}}\right)} = \frac{\partial \Pr[d|V=1,\boldsymbol{\phi}]\mu_d^m}{\partial\phi_d^{\overline{z}}}.(-1)$ so

$$\frac{\partial^2 \Pr[d|V=1,\boldsymbol{\phi}]\mu_d^m}{\partial\left(1-\phi_d^{\overline{z}}\right)^2} = \frac{\partial^2 \Pr[d|V=1,\boldsymbol{\phi}]\mu_d^m}{\partial\left(\phi_d^{\overline{z}}\right)^2} \tag{8}$$

Furthermore, using (7) and noting that the denominator is strictly positive from above and $\mu^m \in (0,1)$ from (2):

$$\frac{\partial^2 \Pr[d|V=1,\boldsymbol{\phi}]\mu_d^m}{\partial\left(\phi_d^{\overline{z}}\right)^2} > 0 \quad \text{for} \quad \sum_{\mathbb{Z}\backslash\overline{z}} \phi_d^z \left(p_0^{\overline{z}} p_1^z - p_0^z p_1^{\overline{z}}\right) \neq 0 \tag{9}$$

(i) follows immediately from (7) for $d = P$, and follows from (7) for $d = F$ using $\phi_P^z = 1 - \phi_F^z$ so convexity in ϕ_F^z implies convexity in ϕ_P^z by (8).

(ii) From (9), $\Pr[d|V=1,\boldsymbol{\phi}]\,\mu_d^m$ is strictly convex in ϕ_d^U if $\phi_d^H\left(p_0^U p_1^H - p_0^H p_1^U\right) \neq \phi_d^L\left(p_0^L p_1^U - p_0^U p_1^L\right)$. Using $p_1^U = p_0^U > 0$ and (1), this reduces to $\phi_d^H \neq \phi_d^L$, which holds from above. Thus the result holds for $d = P$, and also holds for $d = F$ using $\phi_P^z = 1 - \phi_F^z$ so convexity in ϕ_F^U implies convexity in ϕ_P^U by (8).

(iii) & (iv)

From (9) $\Pr[d|V=1,\boldsymbol{\phi}]\,\mu_d^m$ is strictly convex in ϕ_d^H if $\phi_d^U\left(p_0^H p_1^U - p_0^U p_1^H\right) \neq \phi_d^L\left(p_0^L p_1^H - p_0^H p_1^L\right)$. Using $p_1^U = p_0^U > 0$ and (1), the left-hand side is strictly negative while the right-hand side is weakly positive when $\phi_d^U > 0$. Thus (iv)(a) holds, and (iii)(a) holds using $\phi_P^z = 1 - \phi_F^z$ so $\phi_F^U > 0$ and convexity in ϕ_F^H implies convexity in ϕ_P^H by (8).

From (9) $\Pr[d|V=1,\boldsymbol{\phi}]\,\mu_d^m$ is strictly convex in ϕ_d^L if $\phi_d^U\left(p_0^L p_1^U - p_0^U p_1^L\right) \neq \phi_d^H\left(p_0^H p_1^L - p_0^L p_1^H\right)$. Using $p_1^U = p_0^U > 0$ and (1), the left-hand side is strictly positive while the right-hand side is weakly negative when $\phi_d^U > 0$. Thus (iv)(b) holds, and (iii)(b) holds using $\phi_P^z = 1 - \phi_F^z$ so $\phi_F^U > 0$ and convexity in ϕ_F^L implies convexity in ϕ_P^L by (8). ■

Proof of Lemma 4. (i) is proved in the text.

(ii) By the continuity of R in ϕ_P^U, established in Section 3.2, a maximum must exist. Suppose that at a maximum, $\phi_P^U \notin \{0,1\}$. By strict convexity, a strictly higher R can be achieved by appropriately choosing to raise or lower ϕ_P^U, giving a contradiction.

(iii) By the continuity of R in ϕ_P^H and ϕ_P^L, established in Section 3.2, a maximum must exist. Suppose that at a maximum, it is not the case that $\left\{\phi_P^H = 1, \phi_P^L = 0\right\}$. From Theorem 1, $\phi_P^H \neq \phi_P^L$ at this maximum, so from (i) R is strictly convex in ϕ_P^H and ϕ_P^L. Remember that in Section 2.1, *w.l.o.g.* we labelled the decisions such that $\phi_P^H \geq \phi_P^L$. Thus $\left\{\phi_P^H = 0, \phi_P^L = 1\right\}$ is ruled out, while $\phi_P^H \neq \phi_P^L$ rules out $\left\{\phi_P^H = 1, \phi_P^L = 1\right\}$ or $\left\{\phi_P^H = 0, \phi_P^L = 0\right\}$, so at this maximum $\phi_P^z \in (0,1)$ for $z = H$ or $z = L$ or both. By strict convexity, a strictly higher R can be achieved by appropriately choosing to raise or lower such a ϕ_P^z, giving a contradiction. ■

Proof of Lemma 5. From (1), $p_0^L = p_1^H - p_0^H + p_1^L$, and using $p_1^U = p_0^U$ we can see that $p_0^U = 1 - p_1^H - p_1^L$. Thus, using (3) and dividing through by $\sum_{k:\mu^k \geq \mu^m} p_1^k > 0$:

$$\begin{aligned} \frac{R(1,m)}{\sum_{k:\mu^k \geq \mu^m} p_1^k} &= \frac{\left(p_1^H + p_1^U\right)^2 \mu^m}{\left(p_1^H + p_1^U\right)\mu^m + \left(p_0^H + p_0^U\right)(1-\mu^m)} + \frac{\left(p_1^L\right)^2 \mu^m}{p_1^L \mu^m + p_0^L (1-\mu^m)} \qquad (10) \\ &= \frac{\left(1-p_1^L\right)^2 \mu^m}{\left(1-p_1^L\right)\mu^m + \left(p_0^H + 1 - p_1^H - p_1^L\right)(1-\mu^m)} + \frac{\left(p_1^L\right)^2 \mu^m}{p_1^L \mu^m + \left(p_1^H - p_0^H + p_1^L\right)(1-\mu^m)} \\ &= \frac{\left(1-p_1^L\right)^2 \mu^m}{-(1-\mu^m)\left(p_1^H - p_0^H\right) + 1 - p_1^L} + \frac{\left(p_1^L\right)^2 \mu^m}{(1-\mu^m)\left(p_1^H - p_0^H\right) + p_1^L} \end{aligned}$$

$$\begin{aligned} \frac{R(0,m)}{\sum_{k:\mu^k \geq \mu^m} p_1^k} &= \frac{\left(p_1^H\right)^2 \mu^m}{p_1^H \mu^m + p_0^H (1-\mu^m)} + \frac{\left(p_1^U + p_1^L\right)^2 \mu^m}{\left(p_1^U + p_1^L\right)\mu^m + \left(p_0^U + p_0^L\right)(1-\mu^m)} \qquad (11) \\ &= \frac{\left(p_1^H\right)^2 \mu^m}{p_1^H \mu^m + p_0^H (1-\mu^m)} + \frac{\left(1-p_1^H\right)^2 \mu^m}{\left(1-p_1^H\right)\mu^m + \left(1-p_0^H\right)(1-\mu^m)} \\ &= \frac{\left(p_1^H\right)^2 \mu^m}{-(1-\mu^m)\left(p_1^H - p_0^H\right) + p_1^H} + \frac{\left(1-p_1^H\right)^2 \mu^m}{(1-\mu^m)\left(p_1^H - p_0^H\right) + 1 - p_1^H} \end{aligned}$$

Let $w \equiv (1-\mu^m)\left(p_1^H - p_0^H\right)$. We know that $w > 0$ as $\mu^m \in (0,1)$ from (2) and $p_1^H > p_0^H$.

Let $t \equiv \left(-w + 1 - p_1^L\right)\left(w + p_1^L\right)\left(-w + p_1^H\right)\left(w + 1 - p_1^H\right)$. We know that $t > 0$ as each of the four parts is strictly positive, being a re-arrangement of one of the four denominators in expressions (10) and (11), which in turn are all strictly positive. Then

$$\begin{aligned} \frac{(R(1,m) - R(0,m))}{\sum_{k:\mu^k \geq \mu^m} p_1^k} \frac{t}{\mu^m} &= \left(\left(1 - 2p_1^L\right)\left(w + p_1^L\right) + \left(p_1^L\right)^2\right)\left(-w + p_1^H\right)\left(w + 1 - p_1^H\right) \\ &\quad - \left(\left(1 - 2p_1^H\right)\left(-w + p_1^H\right) + \left(p_1^H\right)^2\right)\left(-w + 1 - p_1^L\right)\left(w + p_1^L\right) \end{aligned}$$

which simplifies to $-w^2\left(1 - p_1^H - p_1^L\right)\left(2w - p_1^H + p_1^L\right)$, and the result then follows, using $1 - p_1^H - p_1^L > 0$ as $p_1^U > 0$. ■

Proof of Theorem 3. (i) As $\gamma \in \Gamma \to \overline{\gamma}$, $p_V^m \to p_V^m(\overline{\gamma})$, $\mu^m \to \mu^m(\overline{\gamma}) = 1$ for $m \in \mathbb{X}^+$ and $\mu^m \to \mu^m(\overline{\gamma}) = 0$ for $m \in \mathbb{X}^-$.

For $\gamma \in \Gamma$ sufficiently close to $\overline{\gamma}$, the good type of principal targets all those agents whose signals are positive at $\overline{\gamma}$ for endorsement by setting price $\lambda = \min_{m\in\mathbb{X}^+} \mu_d^m$, conditioned on d. By targeting a smaller group, in expectation the principal loses the endorsement of a proportion of agents at least as big as $\min_{m\in\mathbb{X}^+} p_1^m$. As $\gamma \to \overline{\gamma}$, $\min_{m\in\mathbb{X}^+} p_1^m \to \min_{m\in\mathbb{X}^+} p_1^m(\overline{\gamma}) > 0$. On the other hand, the gain in price from targeting a smaller group vanishes as $\gamma \to \overline{\gamma}$: from (3) $\mu_d^m \to 1 \ \ \forall m \in \mathbb{X}^+$ as $\mu^m \to 1$ and, given $\phi_P^H = \phi_F^L = 1$ here, $\Pr[d|V, \boldsymbol{\phi}] > 0$. The principal will not target a bigger group: in the limit price would have to go to zero as $\forall m \in \mathbb{X}^-$, $\mu^m \to 0$ so $\mu_d^m \to 0$.

From (3), and letting m_i represent a particular m, $\mu_d^{m_i} \gtreqless \mu_d^{m_j} \Leftrightarrow \mu^{m_i} \gtreqless \mu^{m_j}$. Thus at a given γ the same $m \in \mathbb{X}^+$ (or set of m's) minimizes μ_d^m whatever d or ϕ_P^U. Let $\widetilde{m}(\gamma)$ be a minimizing m. Then for $\gamma \in \Gamma$ sufficiently close to $\overline{\gamma}$, from above the principal sets $\lambda = \mu_d^{\widetilde{m}(\gamma)}$, so $R(\phi_P^U) = R(\phi_P^U, \widetilde{m}(\gamma))$. Thus we can apply Lemma 5 to show that $R(1) \gtreqless R(0) \Leftrightarrow 2\left(\mu^{\widetilde{m}(\gamma)} - 1\right)\left(p_1^H - p_0^H\right) + p_1^H \gtreqless p_1^L$. As $\gamma \to \overline{\gamma}$, $\mu^{\widetilde{m}(\gamma)} \to 1$ given $\mu^m \to \mu^m(\overline{\gamma}) = 1$ for $m \in \mathbb{X}^+$. Thus for $\gamma \in \Gamma$ sufficiently close to $\overline{\gamma}$, $R(1) > R(0)$ when $p_1^H > p_1^L$, and, given $\mu^{\widetilde{m}(\gamma)} < 1$ (from (2)) and $p_1^H > p_0^H$, $R(1) < R(0)$ when $p_1^H \leq p_1^L$.

(ii) As $\gamma \in \Gamma \backslash \underline{\Gamma} \to \underline{\gamma}$, $p_V^m \to p_V^m(\underline{\gamma})$ and $\mu^m \to \mu^m(\underline{\gamma}) = q$ for all m.

For $\gamma \in \Gamma \backslash \underline{\Gamma}$ sufficiently close to $\underline{\gamma}$, the good type of principal targets all agents for endorsement by setting price $\lambda = \min_{m\in\mathbb{X}} \mu_d^m$, conditioned on d. By targeting a smaller group, in expectation the principal loses the endorsement of a proportion of agents at least as big as $\min_{m\in\mathbb{X}} p_1^m$. As $\gamma \to \underline{\gamma}$, $\min_{m\in\mathbb{X}} p_1^m \to \min_{m\in\mathbb{X}} p_1^m(\underline{\gamma}) > 0$. On the other hand the gain in price from targeting a smaller group vanishes as $\gamma \to \underline{\gamma}$: from (3) $\mu_d^m \to \frac{\Pr[d|V=1,\boldsymbol{\phi}]q}{\Pr[d|V=1,\boldsymbol{\phi}]q + \Pr[d|V=0,\boldsymbol{\phi}](1-q)} \ \ \forall m \in \mathbb{X}$ as $\mu^m \to q$.

An argument paralleling that in (i) then shows that for $\gamma \in \Gamma \backslash \underline{\Gamma}$ sufficiently close to $\underline{\gamma}$, $R(1) \gtreqless R(0) \Leftrightarrow 2\left(\mu^{\widetilde{m}(\gamma)} - 1\right)\left(p_1^H - p_0^H\right) + p_1^H \gtreqless p_1^L$, where $\widetilde{m}(\gamma)$ is now the $m \in \mathbb{X}$ which minimizes μ_d^m. Thus

$$
\begin{aligned}
R(1) \ &\gtreqless \ R(0) \Leftrightarrow 2\mu^{\widetilde{m}(\gamma)}\left(p_1^H - p_0^H\right) \gtreqless \left(p_1^H - p_0^H\right) + \left(p_1^L - p_0^H\right) \\
&\Leftrightarrow \ \mu^{\widetilde{m}(\gamma)} \gtreqless \tfrac{1}{2} + \tfrac{p_1^L - p_0^H}{2\left(p_1^H - p_0^H\right)} \equiv \overline{q}
\end{aligned}
$$

As $\gamma \to \underline{\gamma}$, $\mu^{\widetilde{m}(\gamma)} \to q$ given $\mu^m \to \mu^m(\underline{\gamma}) = q$ for all $m \in \mathbb{X}$. Thus for $\gamma \in \Gamma \backslash \underline{\Gamma}$ sufficiently close to $\underline{\gamma}$, $R(1) > R(0)$ when $q > \overline{q}$ and, given $\widetilde{m}(\gamma)$ is the worst signal about type so $\mu^{\widetilde{m}(\gamma)} < q$,

$R(1) < R(0)$ when $q \leq \overline{q}$.

(iii) When $\gamma \in \underline{\Gamma}$, all signals are equally uninformative. Thus, the principal targets all agents for endorsement by setting $\lambda = \mu_d^m = \frac{\Pr[d|V=1,\phi]q}{\Pr[d|V=1,\phi]q+\Pr[d|V=0,\phi](1-q)}$ and, using $\mu^m = q$, we can invoke Lemma 5 to show that $R(1) \gtreqless R(0) \Leftrightarrow 2(q-1)\left(p_1^H - p_0^H\right) + p_1^H \gtreqless p_1^L \Leftrightarrow q \gtreqless \overline{q}$. ■

References

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

Bose, S., Orosel, G., Ottaviani, M. and Vesterlund, L. (2007) 'Dynamic Monopoly Pricing with Herding', *RAND Journal of Economics*, 37, 910-928.

Bose, S., Orosel, G., Ottaviani, M. and Vesterlund, L. (2008) 'Monopoly Pricing in the Binary Herding Model', *Economic Theory*, 37, 203-241.

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.

Chiao, B., Lerner, J. and Tirole, J. (2007) 'The Rules of Standard Setting Organizations: An Empirical Analysis', *RAND Journal of Economics*, 38, 905-930.

Demange, G. (2008) 'Sharing Information in Web Communities', *Paris School of Economics Working Papers* 2008-34.

Farhi, E., Lerner, J. and Tirole J. (2005) 'Certifying New Technologies', *Journal of the European Economic Association*, 3, 734-744.

Gill, D. and Sgroi, D. (2008) 'Sequential Decisions with Tests', *Games and Economic Behavior*, 63, 663-678.

Judd, K. L. and Riordan, M. (1994) 'Price and Quality in a New Product Monopoly', *Review of Economic Studies*, 61, 773-789.

Lerner, J. and Tirole, J. (2006) 'A Model of Forum Shopping', *American Economic Review*, 96, 1091-1113.

Lizzeri, A. (1999) 'Information Revelation and Certification Intermediaries', *RAND Journal of Economics*, 30, 214-231.

Mahenc, P. (2004) 'Influence of Informed Buyers in Markets Susceptible to the Lemons Problem', *American Journal of Agricultural Economics*, 86, 649-659.

Milgrom, P. R. (1981), 'Good News and Bad News: Representation Theorems and Applications', *Bell Journal of Economics*, 12, 380-391.

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

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.

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