

Sequential Decisions with Tests

David Gill

*Trinity College and Dept. of Economics
University of Oxford*

Daniel SgROI ¹

*Faculty of Economics and Churchill College
University of Cambridge*

¹For financial support, Daniel SgROI would like to thank AEA Technology plc and David Gill would like to thank the Economic and Social Research Council. Both authors would like to thank Simon Board, Simon Cowan, Douglas Gale, Paul Klemperer, Meg Meyer, David Myatt, Pierre Regibeau and participants at presentations in Birkbeck, Cambridge, Essex, Oxford, the Second World Congress of the Game Theory Society in Marseille and the Workshop on Informational Herding Behavior in Copenhagen for helpful comments and suggestions.

ABSTRACT

We consider a principal-agent problem where the principal wishes to be endorsed by a sequence of agents, but cannot truthfully reveal type. In the standard "herding" model, the agents learn from each other's decisions, which can lead to cascades on a given decision when later agents' private information is swamped. We augment the standard model to allow the principal to subject herself to a test designed to provide public information about her type. She must decide how tough a test to attempt from a continuum of test types, which involves trading off the higher probability of passing an easier test against the greater impact from passing a tougher test. We find that the principal will always choose to be tested, and will prefer a tough test to a neutral or easy one.

Keywords: Bayesian updating, endorsements, herding, sequential decision-making, tests

1. INTRODUCTION

When a firm develops a new product, it will often consider where to send it for pre-launch review or accreditation. A job applicant facing a sequence of interviews must select suitable referees. Film studios decide who to invite to pre-launch viewings and premieres. A politician may wish to leak a new policy to the media, and has a choice of which media outlet to target. A firm launching an initial public offering can choose from auditors of different reputations to review its accounts. The list goes on. In each case a principal ultimately wants to sell products, win votes, or generally be endorsed by a group of agents, and has the option of being publicly tested before seeking endorsement. Despite the powerful effects that success or failure of such a test may have on the performance of the principal, the literature has paid little attention to the role of such public testing, especially the decision that a principal must make when a variety of tests are available which differ in their degree of toughness. The present paper attempts to correct for this omission.

The principal is assumed to be either good or bad for agents, who need to estimate the relative likelihood of the two types before making their endorsement decision. The agents decide in sequence and are granted three sources of information. As in the standard “herding” model, they receive private information, which perhaps relates to prior experience of the principal or her product or policies, and can observe each other’s endorsement decisions in an attempt to learn something about other agents’ private information. The sequential nature of decision-making can allow cascades on a given decision to develop when later agents’ private information is swamped by the information revealed by the decisions of earlier agents. We introduce a third source of information by allowing the principal to subject herself to a test designed to provide public information about her type before any agent has made his endorsement decision. She must decide how tough a test to attempt from a continuum of test types, which involves trading off the higher probability of passing an easier test against the greater impact from passing a tougher test.

As seems reasonable in this context, a bad type of principal can costlessly duplicate the choice of test chosen by a good principal. 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. Therefore the choice of a tough test does not have the advantage of signaling the type of principal. Without a role for signaling by choice of test, we might assume that since an easy test is by definition the most likely to be passed, it must naturally be the first choice for any principal. However tough tests have two innate advantages. They generate a stronger impact on agents in the event of a pass, and they are less damaging in the event of a fail. We find that the principal will always choose to be tested, and perhaps surprisingly will prefer a tough test to a neutral or easy one despite the lack of an immediate signaling advantage through the choice of test.

1.1. Related Literature. The paper most closely related to ours is Lerner and Tirole’s (2004) recent (and independently developed) paper concerning the role of technology standard setting authorities as certifiers.² Along with Lerner and Tirole, our paper is the first to consider a principal seeking an initial public endorsement who can choose from amongst a continuum of tests of varying bias or toughness. Similarly to our tests, Lerner and Tirole’s certifiers have an arbitrary bias towards the technology sponsor which determines their endorsement rule. The model has significant differences to ours. Crucially, Lerner and Tirole do not allow consumers to either receive any private information or observe other agents’ actions. Furthermore, the sponsor is not perfectly informed about quality and the chosen certifier discovers with certainty the quality of the technology it is asked to review. As certifiers cannot overwhelm bad private information or trigger cascades, Lerner and Tirole do not find any role for certifiers biased against the technology. Instead they find the somewhat unsurprising result that the sponsor prefers the certifier most biased in favor of the new technology on offer, subject to users adopting following a positive decision by the certifier.

In a model with sequential sales, SgROI (2002) examines the use of small groups of consumers who are encouraged to decide early, acting as “guinea pigs” and providing additional information for later consumers. SgROI’s paper is similar to ours in that in both cases the principal is able to manipulate herd probabilities to her advantage. However, the method of manipulation and basic trade-off in SgROI are quite different to the ones analyzed here. In SgROI, there is no notion of bias or toughness. Instead, the principal is able to manipulate the structure of agents’ decision-making, trading off a *delay* in the start of the herd against a greater chance that the herd ends up moving in a favorable direction. In this paper, in contrast, the sequence of agent decisions is assumed fixed, while the principal chooses a tough public test to try to start an *immediate* herd on passing, trading this off against a higher chance of failing the test.

Taylor (1999) and Bose et al. (2005) find that high prices can be optimal in a similar way to tough tests in this paper. In Bose et al., if an expensive good becomes successful, 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 can then be attributed to overpricing rather than low quality.

Our work should be contrasted with the literature on self-interested experts, who manipulate the messages they send about the true state of the world (see chapter 10 of Chamley (2004) for a survey). Our tests are purely mechanical: the level of toughness is fixed and commonly known. Our work is also different from the literature on payment structures to certification intermediaries, where the question is how intermediaries affect the quality chosen by the firm (e.g., Albano and Lizzeri, 2001). In our

²Chiao et al. (2005) empirically test Lerner and Tirole’s model, while Farhi et al. (2005) extend the model to a dynamic setting.

model, quality is fixed. Finally, our analysis is different from Ottaviani and Prat (2001), who find that a monopolist may wish to use a public signal of quality such as an outside certifier. Our principal is informed about her type, while in Ottaviani and Prat 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.

2. THE MODEL

A risk-neutral principal, whose objective is to maximize the number of endorsements from a sequence of agents, may be good or bad for the agents. The principal knows her type, but the agents do not, and there is no easy means of truthful revelation. A stream of agents of uncertain length act in an exogenously ordered sequence, choosing an action $A_i \in \{Y, N\}$ where Y denotes an endorsement and N a rejection. This "endorsement" is a general concept which could, for example, encompass adopting some new technology, voting for a candidate in an election, purchasing a product, watching a movie, making a job offer etc.³ After each agent decides, the sequence of agents comes to an end with probability $1 - \theta \in (0, 1)$, so the expected number of agents is $\frac{1}{1-\theta}$.⁴ Introducing this uncertainty, as opposed to a sequence of determinate length as in the standard Bikhchandani et al. (1992) herding model, allows us to use a recursive solution method, greatly simplifying the analysis in the presence of the asymmetries introduced by the choice of biased tests. The payoff to an agent from Y is V which has prior probability $\frac{1}{2}$ of returning 1 or -1 , depending on whether the principal is a good or bad type, while the payoff from N is 0, leaving agents indifferent before additional information is obtained. The agents each receive a conditionally independent signal about V defined as $X_i \in \{H, L\}$ for agent i . The signals are informative but not fully revealing in the sense that:

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

Before facing the stream of agents, the principal can opt to be publicly tested. We want to think of the test as by its very nature accessing finer information about the principal's type than a single typical private signal. The simplest way of modeling this is to allow the test to involve the draw of two i.i.d. signals from the same distribution as agents. The test generates a binary decision $d \in \{P, F\}$ whether to

³In many specific cases we might need to add more content to the model such as prices in an industrial organization context, a voting rule in a political economy context, etc. We wish to leave this open, but the addition of such features is straightforward. See Gill and SgROI (2004) for a specific application to purchasing decisions with flexible prices, though in a simultaneous sales context.

⁴Equivalently, the sequence is infinite and the principal discounts at a rate θ . If the principal is attempting to sell a product, push a new technology or advance a particular policy then $1 - \theta$ may represent a measure of how quickly the principal expects a rival product, technology or policy to be developed which will make her own obsolete.

pass (P) or fail (F) the principal. Modeling an evaluation as condensing more complex information into a simple binary decision follows for example Calvert (1985) and Sah and Stiglitz (1986).⁵ We consider a continuum of test types: tests are passed with a signal draw of HH , pass with probability $\phi \in [0, 1]$ on a draw of HL or LH and fail on a draw of LL . The value of ϕ encapsulates the type of the test. The lower the value of ϕ , the tougher the test.

Definition 1. *Tests with $\phi \in [0, \frac{1}{2})$ are termed “tough”, those with $\phi \in (\frac{1}{2}, 1]$ “easy”, and those with $\phi = \frac{1}{2}$, which pass on a coin flip on observing a set of mixed signals, “neutral”.*

The toughness of a particular test is common knowledge, perhaps generated through a known history of pass or fail decisions, and the principal is able to choose the test type. We have left the notion of test fairly abstract, but depending on the application, the choice of test might consist of a choice between different reviewers, interviewers, referees, accreditation bodies and so on.

Agent i will observe the choice of test, the test result, the actions of his predecessors $H_{i-1} \equiv \{A_1, \dots, A_{i-1}\}$ and his private signal, and will endorse if $E[V \mid \phi, d, H_{i-1}, X_i] > 0$. Where $E[V] = 0$, the agent flips a fair coin.⁶ The sequential ordering allows agents to learn from each other’s decisions as well as from the result of the test, combining this information with their own private signals. Potentially, more information may be transmitted to agents later in the sequence. However, agents may fall into an informational cascade in the sense of Bikhchandani et al. (1992), where public information swamps private information. Once a given agent in the sequence rationally disregards his own private information, nothing further is revealed to later agents who will then also all disregard their private information and copy the choice of their predecessor agent. Of course there is no reason for a principal to stay passive in the face of such potential cascades, and by selecting a suitable test the principal can hope to raise the chance of a cascade in her favor and diminish the chance of a cascade going against her.

Throughout, for conciseness, we consider just the good type of principal. By standard signaling considerations, a bad type of principal will be forced to copy the choice of the good type to avoid immediately revealing type and so receive no endorsements. A separating equilibrium is not possible, as the bad type would copy the choice of the good principal, and so be believed to be good and obtain the same outcome as the good type (receiving endorsements from all agents). Thus, we restrict attention to pooling equilibria in which the bad type of principal is forced to follow the good type’s preference.

⁵As Calvert puts it (p. 534): “This feature represents the basic nature of advice, a distillation of complex reality into a simple recommendation.”

⁶Coin flipping is the standard tie-break rule used in herding models. See for example Bikhchandani et al. (1992). Equivalently each agent may be following a fixed selection rule, so long as in expectation half of indifferent agents endorse the principal and half do not. Banerjee (1992) instead uses a “follow your own signal” rule, but does this specifically to minimize the chance of a herd.

Such equilibria can always be supported by the belief that any principal who deviates from the good type's preferences must be a bad type. Note that in such pooling equilibria, agents will be unable to infer anything about the type of the principal from the principal's *choice of test*, and so will have to rely on observing the *outcome of the test* to provide additional information about the principal's type.

All omitted proofs can be found in the Appendix.

3. AGENT LEARNING

Observing the choice of test, the test result and the actions of his predecessors, agent i will be able to update his prior belief that the principal is good from $\frac{1}{2}$ to $q_i \equiv \Pr[V = 1 \mid \phi, d, H_{i-1}]$. We start by deriving two remarks, which are used implicitly throughout.

Remark 1. *Given his private signal X_i and updated prior q_i , agent i Bayes updates as follows:*

$$\frac{\Pr[V = 1 \mid X_i]}{\Pr[V = -1 \mid X_i]} = \frac{\frac{\Pr[X_i \mid V=1] \Pr[V=1]}{\Pr[X_i]}}{\frac{\Pr[X_i \mid V=-1] \Pr[V=-1]}{\Pr[X_i]}} = \frac{\Pr[X_i \mid V = 1] q_i}{\Pr[X_i \mid V = -1] (1 - q_i)}$$

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

We can now determine agent i 's endorsement decision.

Lemma 1. *The i^{th} agent will respond to an updated prior as follows: (a) if $q_i > p$ then i will endorse; (b) if $q_i = p$, following a H signal i endorses, while following a L signal, he flips a coin; (c) if $q_i \in (1 - p, p)$ then i will endorse if and only if $X_i = H$; (d) if $q_i = 1 - p$, following a H signal i flips a coin, while following a L signal he rejects; (e) if $q_i < 1 - p$ then i rejects.*

Where q_i is strongly positive or negative, it outweighs any possible private signal agent i might receive, leading to a cascade. Where $q_i > p$, from Lemma 1 agent i endorses whatever the signal received. The decision is thus uninformative, so agent $i + 1$ also endorses, and so on. A symmetrical argument applies where $q_i < 1 - p$.

Lemma 2. *Where $q_i > p$, we have a cascade on endorsement, i.e., agent i and all subsequent agents endorse. Where $q_i < 1 - p$, we have a cascade on rejection, i.e., agent i and all subsequent agents fail to endorse.*

Next, we determine how different test results impact on the first agent's beliefs. The updated prior faced by the first agent in the event of a pass (denoted by $q_1 = q_P$) and a fail (denoted by $q_1 = q_F$) will be as follows:

$$q_P \equiv \Pr[V = 1 \mid \text{Pass}] = \frac{[p^2 + 2p(1-p)\phi]}{[p^2 + 2p(1-p)\phi] + [(1-p)^2 + 2(1-p)p\phi]} = \frac{p^2 + 2p(1-p)\phi}{p^2 + (1-p)^2 + 4p(1-p)\phi}$$

$$q_F \equiv \Pr[V = 1 \mid \text{Fail}] = \frac{[(1-p)^2 + 2p(1-p)(1-\phi)]}{[(1-p)^2 + 2p(1-p)(1-\phi)] + [p^2 + 2(1-p)p(1-\phi)]} = \frac{(1-p)^2 + 2p(1-p)(1-\phi)}{(1-p)^2 + p^2 + 4p(1-p)(1-\phi)}$$

Note that:

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

That the updated priors should be decreasing in ϕ is perfectly natural: a pass is better news the tougher the test, while a fail is not such bad news. Using this, we can derive:

Lemma 3. (a) $q_P(\phi < \frac{1}{2}) > p$; (b) $q_P(\phi = \frac{1}{2}) = p$; (c) $q_P(\phi > \frac{1}{2}) \in (\frac{1}{2}, p)$; (d) $q_F(\phi < \frac{1}{2}) \in (1-p, \frac{1}{2})$; (e) $q_F(\phi = \frac{1}{2}) = 1-p$; (f) $q_F(\phi > \frac{1}{2}) < 1-p$.

Together with Lemma 1, Lemma 3 shows how the test determines the first agent's decision.

4. CHOICE OF TEST

In the previous section, we calculated the effect of different test results on q_1 ; now we calculate the expected number of endorsements, π , for the good type of principal $\forall q_1$. From Lemma 2, $q_1 > p$ leads to an immediate cascade on endorsement, so $\pi_{q_1 > p} = \frac{1}{1-\theta}$, while $q_1 < (1-p)$ leads to an immediate cascade on rejection, so $\pi_{q_1 < 1-p} = 0$. The following lemma uses a recursive solution to the appropriate decision trees. Note that in cases (ii) and (iii) π is independent of the specific q_1 value.

Lemma 4. For a good principal,

$$(4.1) \quad (i) \quad \pi_{q_1 = \frac{1}{2}} = \frac{p[2-(1-p)\theta^2]}{2[1-p(1-p)\theta^2](1-\theta)}$$

$$(4.2) \quad (ii) \quad \pi_{q_1 \in (\frac{1}{2}, p)} = \frac{p[1+(1-p)\theta(1-\theta)]}{[1-p(1-p)\theta^2](1-\theta)}$$

$$(4.3) \quad (iii) \quad \pi_{q_1 \in (1-p, \frac{1}{2})} = \frac{p[p+(1-p)(1-\theta)]}{[1-p(1-p)\theta^2](1-\theta)}$$

Now suppose $q_1 = p$. From Lemma 1 if the first agent gets a positive signal he endorses, while if he gets a negative signal this exactly cancels the positive prior so he is indifferent and flips a coin. We can think of the second and subsequent agents as starting a new sequence with updated prior q_2 . If the first agent rejects, later agents infer $X_1 = L$, so $q_2 = \frac{1}{2}$. If the first endorses, then he is more likely to have observed H than L , sending a positive signal, thus increasing q_2 above q_1 , so $q_2 > p$.⁷ Thus,

$$(4.4) \quad \pi_{q_1 = p} = [p + \frac{1}{2}(1-p)] \left(\frac{1}{1-\theta} \right) + \frac{1}{2}(1-p)\theta\pi_{q_1 = \frac{1}{2}}$$

Suppose instead $q_1 = 1-p$. This case is the symmetric opposite. If the first agent gets a L signal, he rejects, while if he gets a H signal, this exactly cancels the negative prior so he flips a coin. Thus, if

⁷Formally, $\frac{q_2}{1-q_2} = \frac{\Pr[V=1|A_1=Y]}{\Pr[V=-1|A_1=Y]} = \frac{p+\frac{1}{2}(1-p)}{(1-p)+\frac{1}{2}p} \frac{q_1}{1-q_1} > \frac{q_1}{1-q_1}$.

the first agent endorses, later agents infer $X_1 = H$, so $q_2 = \frac{1}{2}$. If the first rejects, then he is more likely to have observed L than H , sending a negative signal, so $q_2 < 1 - p$. Thus,

$$(4.5) \quad \pi_{q_1=1-p} = \frac{1}{2}p \left(1 + \theta \pi_{q_1=\frac{1}{2}} \right)$$

Of course, $\pi_{q_1>p} > \pi_{q_1=p} > \pi_{q_1 \in (\frac{1}{2}, p)} > \pi_{q_1=\frac{1}{2}} > \pi_{q_1 \in (1-p, \frac{1}{2})} > \pi_{q_1=1-p} > \pi_{q_1 < 1-p}$.

We now have the building blocks which allow us to easily calculate and compare the expected number of endorsements for the principal from choosing different test types, $\Pi[\phi]$. Suppose first that the test is neutral. Then, using Lemma 3 and the fact that $\Pr[P] = p^2 + 2p(1-p)\phi$ while $\Pr[F] = (1-p)^2 + 2p(1-p)(1-\phi)$:

$$(4.6) \quad \Pi \left[\phi = \frac{1}{2} \right] = \Pr[P] \pi_{q_1=p} + \Pr[F] \pi_{q_1=1-p} = p \pi_{q_1=p} + (1-p) \pi_{q_1=1-p}$$

With a tough test, using Lemmas 2 and 3, a pass starts an immediate cascade on endorsement with $q_1 > p$, while a fail leaves $q_1 \in (1-p, \frac{1}{2})$. Thus:

$$(4.7) \quad \Pi \left[\phi \in [0, \frac{1}{2}) \right] = \{p^2 + 2p(1-p)\phi\} \pi_{q_1>p} + \{(1-p)^2 + 2p(1-p)(1-\phi)\} \pi_{q_1 \in (1-p, \frac{1}{2})}$$

Note that $\frac{d(\Pi[\phi \in [0, \frac{1}{2})])}{d\phi} > 0$ as $\pi_{q_1>p}$ and $\pi_{q_1 \in (1-p, \frac{1}{2})}$ are constant in ϕ , $\pi_{q_1>p} > \pi_{q_1 \in (1-p, \frac{1}{2})}$ and $\frac{d\Pr[P]}{d\phi} = -\frac{d\Pr[F]}{d\phi} > 0$.

Finally, we consider an easy test. This case is the symmetric opposite of the tough test case. Using Lemmas 2 and 3, a fail starts an immediate cascade on rejection as $q_1 < 1-p$, while a pass sends a positive signal weaker than the one sent out when a neutral test is passed, so $q_1 \in (\frac{1}{2}, p)$, giving:

$$(4.8) \quad \Pi \left[\phi \in (\frac{1}{2}, 1] \right] = \{p^2 + 2p(1-p)\phi\} \pi_{q_1 \in (\frac{1}{2}, p)}$$

Comparing the different expected number of endorsements, we find the following proposition.

Proposition 1. *For any choice of three test types $\phi \in \{\phi_T, \frac{1}{2}, \phi_E\}$ such that $\phi_T \in [0, \frac{1}{2})$ and $\phi_E \in (\frac{1}{2}, 1]$, the good type of principal strictly prefers the tough test ϕ_T to the neutral test $\phi = \frac{1}{2}$ and to the easy test ϕ_E .*

Despite the fact that tough tests are less likely to be passed, the good principal prefers any tough test to any neutral or easy one. She prefers tough tests because of the strong impact on agents' decisions from a pass, which leads to a cascade on endorsement, while a fail in a tough test is not too costly as the test is known to be tough, diluting the impact of failure on agents' beliefs.

Next we define the concept of ϵ -optimization.

Definition 2. To ϵ -optimize Π over a range of ϕ values, $\Omega \subseteq [0, 1]$, means to select a $\phi \in \Omega$ such that $\Pi > \sup_{\phi \in \Omega} \Pi - \epsilon$.

The following lemma follows from the fact that $\frac{d(\Pi[\phi \in [0, \frac{1}{2}]])}{d\phi} > 0$, so $\sup_{\phi < \frac{1}{2}} \Pi = p\pi_{q_1 > p} + (1-p)\pi_{q_1 \in (1-p, \frac{1}{2})}$.

Lemma 5. For any $\epsilon > 0$ the principal can ϵ -optimize over $\phi \in [0, \frac{1}{2})$ by selecting ϕ sufficiently close to $\frac{1}{2}$, thus achieving $\Pi > \sup_{\phi < \frac{1}{2}} \Pi - \epsilon$.

Informally, the lemma says that the optimal tough test is one whose toughness is arbitrarily mild (i.e., with ϕ strictly below but arbitrarily close to $\frac{1}{2}$).

Proposition 1 tells us that any tough type of test beats the neutral test or any easy type of test, which implies that $\sup_{\phi \in [0, 1]} \Pi = \sup_{\phi < \frac{1}{2}} \Pi$. Together with Lemma 5 this gives:

Proposition 2. For any $\epsilon > 0$ the principal can ϵ -optimize over $\phi \in [0, 1]$ by selecting ϕ sufficiently close to but below $\frac{1}{2}$, thus achieving $\Pi > \sup_{\phi \in [0, 1]} \Pi - \epsilon$.

Informally, the optimal test type is a tough test which has arbitrarily mild toughness. A pass in *any* tough type of test leads to an immediate cascade on endorsement, so the principal prefers a tough test that is as close as possible to the neutral test type so as to maximize the probability of a pass.⁸

Finally, we find that the neutral test is strictly preferred to no test at all. The signal arising from the neutral test allows some information transmission through the sequence of agents which, in expectation, is valuable to the good type of principal.

Proposition 3. The choice of taking no test is strictly worse than taking the neutral type of test $\phi = \frac{1}{2}$ and a fortiori strictly worse than a tough test type.

⁸For analytical tractability, we have chosen a simple binary signal structure. Though beyond the scope of this paper, one might wonder to what extent our results carry over to a richer signal space. Our model crucially depends on the possibility of cascades, which are manipulated by the principal. As is well understood in the herding literature (see Smith and Sorensen, 2000), cascades require that the strength of private signals be bounded. With unbounded signals, any public belief can be overwhelmed by a strong enough signal, so herds can always be broken and beliefs must eventually converge to the truth. Thus, our results would not extend to Gaussian continuous signals (see Chamley, 2004, p. 27). However, we conjecture that our main intuition that tough tests are best would hold under a richer discrete bounded signal space, as the principal should still want to risk a tough test to try to get an early favorable cascade started. The principal could then choose a test just tougher than that required to make an agent with the worst possible signal indifferent conditional on the test being passed. However, if such extreme signals are sufficiently unlikely, we conjecture that the principal will instead choose a somewhat weaker tough test to give a high, but not certain, chance of an immediate favorable cascade in the event of a pass. The fact that a pass by a test ϵ tougher than a fair one outweighs the worst possible signal is specific to the simple binary signal case. More generally the principal will still be ϵ -optimizing in that for whichever signal is targeted as the endorse threshold conditional on a pass, the test needs to be ϵ tougher than the one giving indifference with that signal. With a continuum of bounded signals, the property of ϵ -optimization will disappear, unless the principal targets the worst possible signal.

The following two figures illustrate the size of the increase in the expected number of endorsements from opting to face a tough test. Figure 1 shows the percentage increase in the expected number of endorsements from taking the toughest test over a neutral test, which peaks at 25%, while Figure 2 shows the percentage increase from using the tough test type which is almost neutral, which peaks at 50%.

Figure 1: $\frac{\Pi[\phi=0]-\Pi[\phi=\frac{1}{2}]}{\Pi[\phi=\frac{1}{2}]}$





Figure 2: $\frac{\sup_{\phi < \frac{1}{2}} \Pi - \Pi[\phi=\frac{1}{2}]}{\Pi[\phi=\frac{1}{2}]}$



5. CONCLUSION

Within the model presented in which test toughness is tightly defined, as are the available sources of information for agents, we find that a principal of type unknown to a sequence of agents should seek to face a public test if this is possible. Furthermore, that principal should seek out the mildest form of tough test available. If there are only very tough tests, then those are the ones that should be selected. Tests that are relatively easy are not optimal, as they provide too damaging a signal in the event of a fail, and too little gain in the event of a pass. Converting these results into practical normative advice, job applicants might consider resisting the temptation to approach relatively soft referees. Firms should avoid “yes men” reviewers for their products. A politician should consider opting to select where to be interviewed, or where to leak new policies, based on the simple premise of first ruling out optimistic or positively biased journalists and then selecting the mildest of those who are intrinsically biased against the politician and his policies. At a descriptive level, we have an explanation for the existence of tough tests, biased newspapers, tough referees, etc. that have a well known harsh or overly critical style, and yet are regularly chosen. When selecting a test to take, a reviewer to observe your product, a referee to provide a letter of recommendation or an interviewer to face, the old Roman proverb is perhaps the best summary of our findings: fortune favors the brave.

APPENDIX

Proof of Remark 2. Suppose the agent starts with a prior z and observes or infers an information set I_i . Now suppose that instead of I_i , the agent infers I_i^+ , which we define as the set I_i plus a further H and then a L signal. Then, using Remark 1 (replacing q_i with z), $\frac{\Pr[V=1|I_i^+]}{\Pr[V=-1|I_i^+]} = \frac{\Pr[I_i^+|V=1]z}{\Pr[I_i^+|V=-1](1-z)} = \frac{\Pr[I_i|V=1]p(1-p)z}{\Pr[I_i|V=-1](1-p)p(1-z)} = \frac{\Pr[V=1|I_i]}{\Pr[V=-1|I_i]}$. \square

Proof of Lemma 1. Where $q_i > p$, the i^{th} agent will endorse. Agent i is least likely to endorse if $X_i = L$. Taking this case, we have an odds ratio $\frac{\Pr[V=1|L]}{\Pr[V=-1|L]} = \frac{(1-p)q_i}{p(1-q_i)} = \frac{q_i - pq_i}{p - pq_i} > 1$ since $q_i > p$. A symmetrical argument shows that where $q_i < 1 - p$, agent i will reject. Where $q_i \in (1 - p, p)$, following a H signal $\frac{\Pr[V=1|H]}{\Pr[V=-1|H]} = \frac{pq_i}{(1-p)(1-q_i)} > 1$ as $q_i > (1 - p)$, so i endorses. Following a L signal $\frac{\Pr[V=1|L]}{\Pr[V=-1|L]} = \frac{(1-p)q_i}{p(1-q_i)} < 1$ as $q_i < p$, so i rejects. Where $q_i = p$, following a L signal, $\frac{\Pr[V=1|L]}{\Pr[V=-1|L]} = \frac{(1-p)q_i}{p(1-q_i)} = 1$ as $q_i = p$, so the agent is indifferent and flips a coin. A H signal is more positive, so the agent endorses. By symmetry, where $q_i = 1 - p$, following a H signal the agent flips a coin, while following a L signal the agent rejects. \square

Proof of Lemma 3. With a neutral test, so setting $\phi = \frac{1}{2}$, we have:

$$\begin{aligned} q_P(\phi = \frac{1}{2}) &= \frac{p^2 + p(1-p)}{p^2 + (1-p)^2 + 2p(1-p)} = p \\ q_F(\phi = \frac{1}{2}) &= \frac{(1-p)^2 + p(1-p)}{(1-p)^2 + p^2 + 2p(1-p)} = 1 - p \end{aligned}$$

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

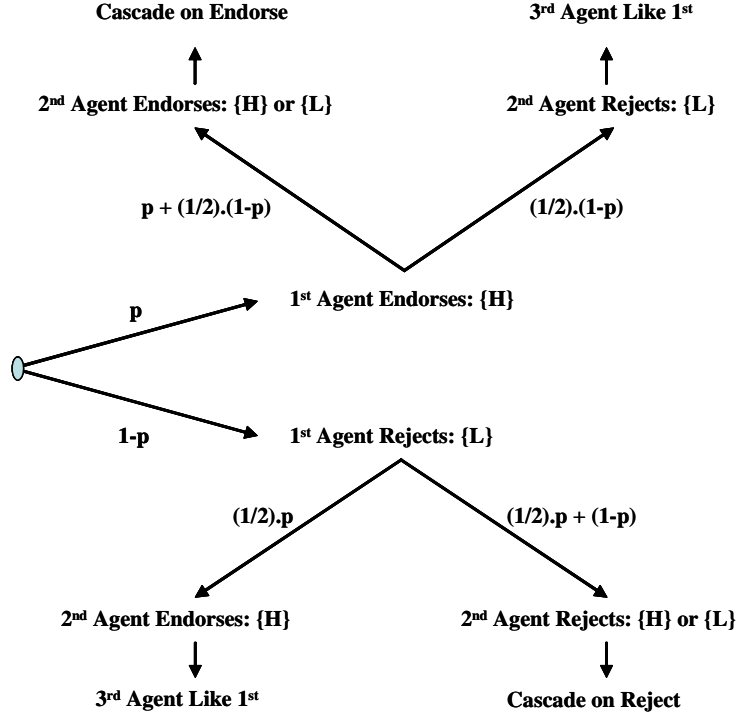
$$\begin{aligned} q_P(\phi = 1) &= \frac{p^2 + 2p(1-p)}{p^2 + (1-p)^2 + 4p(1-p)} > \frac{1}{2} \Leftrightarrow p^2 > (1-p)^2 \\ q_F(\phi = 0) &= \frac{(1-p)^2 + 2p(1-p)}{(1-p)^2 + p^2 + 4p(1-p)} < \frac{1}{2} \Leftrightarrow (1-p)^2 < p^2 \end{aligned}$$

Using the fact that q_P and q_F are both strictly decreasing in ϕ , it follows that $q_P > \frac{1}{2}$ and $q_F < \frac{1}{2}$ $\forall \phi$. \square

Proof of Lemma 4. (i) $q_1 = \frac{1}{2}$. The first agent will follow his signal, i.e., will endorse iff $X_1 = H$. Suppose the first agent endorses. This reveals his signal to be H to the second agent. If the second agent also gets a H signal, he therefore also endorses, but if he gets a L signal the H and L signals cancel, so he is indifferent and flips a coin. If the third agent observes two endorse decisions, a cascade on endorse starts. If he gets $X_3 = L$, his signal and that of the first agent cancel, but because the second agent endorsed, he is more likely to have observed $X_2 = H$ than $X_2 = L$. Formally, for the third agent $\frac{\Pr[V=1|A_2=Y]}{\Pr[V=-1|A_2=Y]} = \frac{p + \frac{1}{2}(1-p)}{(1-p) + \frac{1}{2}p} > 1$. Thus the third agent endorses if he receives a bad signal, and *a fortiori* endorses with a good signal, so a cascade on endorse has started.

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

The case where the first agent rejects is the symmetric opposite. If the second agent also rejects, a cascade on reject starts. If the second agent endorses, then the third agent is back to exactly the same situation as the first agent. All this can be illustrated in the following decision tree:⁹



We can use this tree to calculate the expected number of endorsements for the principal in this case by finding the expected endorsements down various branches of the tree and multiplying by the probability of the relevant branch. Note that we have a recursive structure, whereby the expected number of endorsements from various points further down the tree are equivalent to those from points higher up in the tree. Letting $\pi_{q_1=\frac{1}{2}}$ be the expected number of endorsements at the beginning of the tree, we get:

$$\pi_{q_1=\frac{1}{2}} = p \left[p + \frac{1}{2}(1-p) \right] \left(\frac{1}{1-\theta} \right) + p \left[\frac{1}{2}(1-p) \right] \left(1 + \theta^2 \pi_{q_1=\frac{1}{2}} \right) + (1-p) \left(\frac{1}{2} p \right) \left(\theta + \theta^2 \pi_{q_1=\frac{1}{2}} \right)$$

which solves to give the value for $\pi_{q_1=\frac{1}{2}}$ in the lemma.

(ii) $q_1 \in (\frac{1}{2}, p)$. From Lemma 1 the first agent will endorse iff $X_1 = H$. Following endorsement by the first agent, a cascade on endorse starts. The second agent can infer that the first one got a H signal. Thus, $q_2 = \frac{pq_1}{pq_1 + (1-p)(1-q_1)} > p$ as $q_1 > pq_1 + 1 - p - q_1 + pq_1$ iff $2q_1(1-p) > 1-p$ or $q_1 > \frac{1}{2}$, which of course we have assumed. Hence by Lemma 2 a cascade on endorse starts.


⁹Remember that the principal knows it is a good type. The branch probabilities are predicated upon this assumption.

If the first agent rejects, then the second agent endorses iff $X_2 = H$. The second agent can infer $X_1 = L$, so if $X_2 = H$, the two signals cancel, and hence the second agent endorses as $q_1 > \frac{1}{2}$. If $X_2 = L$, then the agent has effectively seen two negative signals, and so rejects given the first agent with a single negative signal does so.

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

Following rejection by the first two agents, a cascade on reject starts. If the third agent receives a H signal, this cancels one of the two inferred L signals, so the agent is left with just one L signal which, just as for the first agent with L , leads to rejection. *A fortiori*, he also rejects if he receives a L signal.

All this information can be summarized in the following decision tree, with branch probabilities conditional on the principal being the good type.



Thus, for any specific $q_1 \in (\frac{1}{2}, p)$, we can find expected number of endorsements for the good type of principal,

$$\pi_{q_1 \in (\frac{1}{2}, p)} = p \left(\frac{1}{1-\theta} \right) + (1-p)p \left[\theta + \theta^2 \pi_{q_1 \in (\frac{1}{2}, p)} \right]$$

which solves to give the value for $\pi_{q_1 \in (\frac{1}{2}, p)}$ in the lemma.

(iii) $q_1 \in (1-p, \frac{1}{2})$. The shape of the decision tree, which is determined by agents who do not know the type of principal, will be the symmetric opposite of the one in case (ii). A L signal starts a cascade on reject, just like before a H signal started a cascade on endorse, while two H signals start a cascade on endorse, just like before two L signals started a cascade on reject. If the first agent endorses but the second does not, the inferred signals cancel. Thus, for any specific $q_1 \in (1-p, \frac{1}{2})$,

$$\pi_{q_1 \in (1-p, \frac{1}{2})} = p^2 \left(\frac{1}{1-\theta} \right) + p(1-p) \left[1 + \theta^2 \pi_{q_1 \in (1-p, \frac{1}{2})} \right]$$

which solves to give the value for $\pi_{q_1 \in (\frac{1}{2}, p)}$ in the lemma. \square

Proof of Proposition 1. To show that any $\phi_T \in [0, \frac{1}{2})$ gives a larger number of expected endorsements than $\phi = \frac{1}{2}$ we simply need to show that $\Pi[\phi = 0] > \Pi[\phi = \frac{1}{2}]$, given $\frac{d(\Pi[\phi \in [0, \frac{1}{2}]])}{d\phi} > 0$. Using (4.7), (4.3) and Lemma 2, $\Pi[\phi = 0] = p^2 \left(\frac{1}{1-\theta} \right) + (1-p)^2 \frac{p[p+(1-p)(1-\theta)]}{[1-p(1-p)\theta^2](1-\theta)}$, which simplifies to:

$$(A.1) \quad \Pi[\phi = 0] = \frac{p[1+p(1-p)+p^2(1-p)\theta(1-\theta)-(1-p)\theta]}{[1-p(1-p)\theta^2](1-\theta)}$$

From (4.6), $\Pi[\phi = \frac{1}{2}] = p\pi_{q_1=p} + (1-p)\pi_{q_1=1-p}$. Thus, using (4.4), (4.5) and (4.1):

$$\Pi[\phi = \frac{1}{2}] = p \left[p + \frac{1}{2}(1-p) \right] \left(\frac{1}{1-\theta} \right) + p \frac{1}{2}(1-p)\theta \frac{p[2-(1-p)\theta^2]}{2[1-p(1-p)\theta^2](1-\theta)} + (1-p) \frac{1}{2}p \left(1 + \theta \frac{p[2-(1-p)\theta^2]}{2[1-p(1-p)\theta^2](1-\theta)} \right)$$

which simplifies to

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

From (A.1) and (A.2),

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

The denominator is strictly positive, as is $p(1-p)$, so to show $\Pi[\phi = 0] > \Pi[\phi = \frac{1}{2}]$ we just need to show that $2p+2p^2\theta(1-\theta)-\theta-2p\theta(1-\theta) > 0$. Thus, a sufficient condition is that $2p-\theta-2p\theta(1-\theta) > 0$. This hold iff $2p[1-\theta(1-\theta)] > \theta$. But $2p > 1$, so a further sufficient condition is that $1-\theta(1-\theta) > \theta$, or $(\theta-1)^2 > 0$, which is clearly true.

Our final task is to show that $\phi_T \in [0, \frac{1}{2})$ gives a larger number of expected endorsements than $\phi_E \in (\frac{1}{2}, 1]$. From (4.8), $\Pi[\phi \in (\frac{1}{2}, 1]] = \{p^2 + 2p(1-p)\phi\} \pi_{q_1 \in (\frac{1}{2}, p)}$. Now, $\frac{d(\Pi[\phi \in (\frac{1}{2}, 1]])}{d\phi} > 0$ as $\Pr[P] = \{p^2 + 2p(1-p)\phi\}$ is strictly increasing in ϕ and from (4.2), $\pi_{q_1 \in (\frac{1}{2}, p)}$ is constant in ϕ . Thus, we simply need to show that $\Pi[\phi = 0] > \Pi[\phi = 1]$. Using (4.2):

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

Using (A.1) and (A.3), we can derive

$$\Pi[\phi = 0] - \Pi[\phi = 1] = \frac{p(1-p)(1-\theta)[1-2p(1-p)\theta]}{[1-p(1-p)\theta^2](1-\theta)}$$

The denominator is strictly positive, as is $p(1-p)(1-\theta)$, so the sign of $\Pi[\phi = 0] - \Pi[\phi = 1]$ and $1-2p(1-p)\theta$ must be the same. Now, since $p > \frac{1}{2}$, it must be that $p(1-p) < \frac{1}{4}$, so $2p(1-p)\theta$ must always remain smaller than a half. Thus, $1-2p(1-p)\theta > 0$. \square

Proof of Proposition 3. With no test, the first agent's q_1 equals the prior belief $\frac{1}{2}$. Therefore, Π [No Test] = $\pi_{q_1=\frac{1}{2}}$. Thus, using (4.1) and (A.2), $\Pi [\phi = \frac{1}{2}] > \Pi$ [No Test] iff:

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

which holds. □

REFERENCES

- Albano, G.L., Lizzeri, A., 2001. Strategic Certification and Provision of Quality. *Int. Econ. Rev.* 42, 267-283.
- Banerjee, A.V., 1992. A Simple Model of Herd Behavior. *Quart. J. Econ.* 107, 797-817.
- Bikhchandani, S., Hirshleifer, D., Welch, I., 1992. A Theory of Fads, Fashion, Custom and Cultural Change as Informational Cascades. *J. Polit. Economy* 100, 992-1026.
- Bose, S., Orosel, G., Ottaviani, M., Vesterlund, L., 2005. Dynamic Monopoly Pricing and Herding. Mimeo. Forthcoming in *RAND J. Econ.*
- Calvert, R. L., 1985. The Value of Biased Information: A Rational Choice Model of Political Advice. *Journal of Politics* 47, 530-555.
- Chamley, C.P., 2004. *Rational Herds: Economic Models of Social Learning*. Cambridge University Press.
- Chiao, B., Lerner, J., Tirole, J., 2005. The Rules of Standard Setting Organizations: An Empirical Analysis. NBER Working Paper 11156.
- Farhi, E., Lerner, J., Tirole, J., 2005. Certifying New Technologies. *Journal of the European Economic Association* 3, 734-744.
- Gill, D., SgROI, D., 2004. The Superiority of Biased Reviewers in a Model of Simultaneous Sales. Department of Economics Discussion Paper 206, University of Oxford.
- Lerner, J., Tirole, J., 2004. A Model of Forum Shopping, with Special Reference to Standard Setting Organizations. Harvard NOM Research Paper No. 04-31. Forthcoming in *Amer. Econ. Rev.*
- Ottaviani, M., Prat, A., 2001. The Value of Public Information in Monopoly. *Econometrica* 69, 1673-1683.
- Sah, R.K., Stiglitz, J.E., 1986. The Architecture of Economic Systems: Hierarchies and Polyarchies. *Amer. Econ. Rev.* 76, 716-727.
- SgROI, D., 2002. Optimizing Information in the Herd: Guinea Pigs, Profits and Welfare. *Games Econ. Behav.* 39, 137-166.
- Smith, L., Sorensen, P., 2000. Pathological Outcomes of Observational Learning. *Econometrica* 68, 371-398.
- Taylor, C.R., 1999. Time-on-the-Market as a Sign of Quality. *Rev. Econ. Stud.* 66, 555-578.