

Test Design and Minimum Standards

Peter M. DeMarzo, Ilan Kremer and Andrzej Skrzypacz[†]

September 2018

ABSTRACT. We analyze test design and certification standards when an uninformed seller has the option to generate and disclose costly information regarding asset quality. We characterize equilibria by a *minimum principle*: the test and disclosure policy are chosen to minimize the asset's value conditional on nondisclosure. Thus, when sellers choose the information provided, simple pass/fail certification tests are likely to dominate the market. A social planner could raise informational and allocative efficiency, and lower deadweight testing costs, by raising the certification standard. Monopolist certifiers also satisfy the minimum principle but set a higher standard and reduce testing rates to maximize revenue.

[†] We thank seminar participants at Columbia, Hebrew University, Michigan State University, NYU, Stanford, Technion, U.C. Berkeley, University of Bonn, University of Michigan, University of Toronto, Warwick, the 2017 Summer FTG Conference, the 2018 Texas A&M Young Scholars Finance Consortium, and the 2018 European Summer Symposium in Economic Theory in Gerzensee, as well as Piotr Dworzak, Giorgio Martini and Richard Saouma for comments and suggestions. DeMarzo: Graduate School of Business, Stanford University; Kremer: Department of Economics and School of Business Administration, Hebrew University, and Department of Economics, University of Warwick; Skrzypacz: Graduate School of Business, Stanford University.

1. Introduction

Transacting parties often generate and disclose information to improve their terms of trade. Decisions about what to collect and what to disclose impact the quality of information in the market. Consider, for example, a Ph.D. candidate who submits her paper to a journal before entering the academic job market. If the student receives a positive outcome, such as a ‘revise and resubmit,’ then she will disclose it in an attempt to enhance her prospects. But if the paper is rejected, she may choose to conceal the outcome so as not to harm her chances. Because there are other reasons a candidate might not disclose a response from a journal – the paper may not have been submitted, or the journal has not yet responded – employers will not know whether the paper was submitted and rejected if the student reveals nothing.

Situations where the seller of an asset chooses both what information to generate and whether or not to disclose it are ubiquitous. A technology firm that seeks financing can beta-test a prototype to learn and demonstrate product quality. Very simple tests or minimal prototypes yield results with high probability, but the results may not be very informative. More complicated tests or prototypes provide more information if completed, but require more time and so firms may be less likely to obtain clear results before they need financing. Again, the decisions about what prototype to build and whether to disclose the results are both strategic.

This paper focuses on the equilibrium outcomes of such testing and disclosure games. We study the nature of information that is generated and disclosed in a rational market. We also consider the extent to which the amount of information that becomes publicly available is efficient.

With this goal in mind, we consider a single agent holding an asset whose market value is a function of its unknown quality.¹ Initially, the agent is uninformed and shares the same prior belief as the market. Before selling the asset, the agent can choose from a set of available tests. Each test may produce a verifiable result regarding the asset’s quality, or it may yield a null or ‘no-result’ outcome, in which case the agent obtains no verifiable information. Agents may also have no result if the test was too costly (or infeasible) to undertake. Tests differ in the probability of obtaining a

¹ While we cast much of the paper in terms of a single asset being traded, other applications – introduced in Section 2.1 – include disclosing information to increase the quantity of a good sold, as well as managers disclosing information to influence their firm’s investment decisions.

verifiable result and in the information contained in the results. For example, a binary pass/fail test yields with some probability a result that implies that the asset quality is above a certain threshold.

Because a null outcome cannot be verified, the agent can conceal a negative test result by claiming instead that he received the null result and has nothing to disclose. When the agent does not disclose information, there are two sources of information asymmetry: the market does not know whether the agent indeed received a null result, nor which test he took, and both play an important role in our analysis. The market learns what test the agent selected only upon disclosure.

We assume the market is competitive, so it sets a price equal to the expected value of the asset conditional on the information disclosed (or not disclosed). Upon nondisclosure, the market accounts for the possibility that the agent hid an unfavorable outcome and forms expectations (consistent with the equilibrium strategy of the agent) over which test the agent chose and which results he could be hiding. The agent chooses the test and disclosure policy to maximize expected revenue.

In an environment with a single test, the equilibrium expected price of the asset is equal to its ex-ante expected value. Even though the agent strategically conceals negative outcomes, the market makes correct inferences in the event of nondisclosure, and so this result follows immediately from the law of iterated expectations. Therefore, if the choice of test were publicly observable (even in the event of nondisclosure), the agent would be indifferent between tests, and any test choice would be equilibrium. Moreover, if test-taking were costly, the agent would prefer not to take any test.

In contrast, we show that when the agent's choice of test is not known unless he discloses it, this indifference result disappears. Even though in equilibrium the agent does not benefit on average from generating and disclosing information, he will seek to acquire information and disclose it, even at a cost. Moreover, there is a generically unique equilibrium test choice.²

Which test will the agent choose in equilibrium? Because we put no constraints on the form, complexity, or reliability of available tests, a priori there is no obvious ranking of tests in terms of their informativeness. Nonetheless, we show that there exists a one-dimensional index that determines which test is chosen in equilibrium. We call that index the *nondisclosure price* of the test. It is the expected value of the asset that the market would infer if the test were known but the

² By “generically unique” we mean that the equilibrium choice of test is unique for any finite set of tests with generic parameters.

agent does not disclose any information and claims a null result (either honestly or strategically). This nondisclosure price is naturally described as a fixed-point in the single-test game: it is the expected value of the asset conditional on the agent not receiving a test result that would imply a higher value.

Extending results from Acharya, DeMarzo and Kremer (2011), we show that the nondisclosure price for any test can be characterized as the minimum nondisclosure price that is attained over all possible disclosure strategies. We show that in the equilibrium of the original multi-test game, the agent chooses the test with the lowest nondisclosure price. Thus, the equilibrium outcome of the joint test choice and disclosure game satisfies what we refer to as the *minimum principle*: it minimizes the expected price upon nondisclosure over all possible tests *and* disclosure strategies.

The intuition for this result is as follows. Suppose that in equilibrium the agent picks test A, but there exists another test, B, with a lower nondisclosure price. As we noticed before, on the equilibrium path, the expected price equals the ex-ante expected value of the asset. Consider a deviation in which the agent picks test B and discloses every outcome that leads to a higher reward than B's nondisclosure reward. In the event of nondisclosure, if the market were able to observe this deviation to test B the agent would receive test B's nondisclosure price. In that case, the law of iterated expectations again implies that the expected price received would be the ex-ante expected value of the asset, and the agent would not gain from the deviation. But because the test taken is not revealed when the agent chooses not to disclose, then despite the deviation the market would offer the higher nondisclosure price for test A, and the agent would strictly gain by deviating to test B. To avoid such a profitable deviation, the equilibrium test must, therefore, be the one with the lowest nondisclosure price.

Based on the above characterization, we examine the quality of the information revealed in equilibrium, which depends on both dimensions of the strategy: the choice of the test and the disclosure policy. We first show that in equilibrium the agent chooses more accurate and more reliable tests; that is, tests with less noisy outcomes or with a lower probability of a null result. The selection of such tests improves both the quality of the disclosed information and the likelihood of disclosure.

Despite this positive result, we show that too little information may be provided to the market. In particular, the minimum principle favors tests offering greater precision about the left tail of the

distribution, near the disclosure threshold. As a result, sellers may be driven towards simple pass/fail tests that certify only that the expected quality of the asset is above a given threshold. Because the nondisclosure price depends only on the set of signals that are hidden, pooling disclosed signals into one “passing” result does not change the nondisclosure price. Thus, if simple tests are cheaper to administer or more likely to produce verifiable results, they would be chosen in equilibrium.

We next consider the welfare consequences of such an equilibrium. We introduce natural settings in which allocative efficiency and hence welfare increases with the informativeness of prices. We then show that within the class of simple pass/fail certification tests, the agent chooses tests that are too easy to pass. Although the agent would prefer to set the test threshold equal to the nondisclosure price (an implication of the minimum principle), it would be more informative to pool those types who “just pass” such a test (and so have a value close to the nondisclosure price) with those who fail. Prices would thus reflect true quality better if a regulator imposed minimum certification standards.

We then extend the model to consider the case of heterogeneous private costs of test-taking. In that case, an agent may choose not to take a test and have a null result simply because the available tests are too costly. In equilibrium, agents with low testing costs earn rents at the expense of those with high costs. In that setting, imposing minimum certification standards improves the equilibrium outcome in two dimensions. In addition to improving informativeness, minimum standards also reduce the number of agents who incur certification costs.

Thus far, we have allowed the agent (seller) to determine the test design. Our analysis can also be applied when certification is provided by an external organization that designs a test that agents can take to certify that their products or assets meet certain standards. Sellers who pass the test can then use their certificates to advertise quality. In practice, it is not uncommon for certified and non-certified sellers to co-exist, and while non-certified firms receive a lower price they are not presumed to have necessarily failed the certification (e.g., certification may have been too costly, or still be pending).³

³ Examples of such certification organizations are the National Committee of Quality Assurance (NCQA) that since 1991 offers HMOs voluntary certification program (see Jin (2005) for a detailed description of the NCQA program); and the National Association for the Education of Young Children that provides accreditation of child care centers. Another classic example is the Underwriters Laboratories (UL), which is the largest independent, not-for-profit testing

To analyze such markets, we model a monopolist certifier who chooses a certification test and charges a fee for certification to maximize revenue. The monopolist revenue is equal to the fee times the fraction of agents who take the test. We show that *conditional* on the number of agents the monopolist entices to testing, its choice of the test also minimizes the nondisclosure price. However, since the monopolist raises the cost of obtaining certification, the monopolist induces lower participation, a higher nondisclosure price and a higher certification threshold (compared to the case where the agent chooses the test design and/or certification is competitive).

We finish the paper discussing two extensions. First, we consider an agent who has some private information about the quality of the asset before he chooses the test. We show that in a separating equilibrium, agents with negative information choose tests that satisfy the minimum principle, while agents with positive information separate by choosing tests that are harder to pass. Second, we consider agents whose payoffs are non-linear functions of the perceived posterior quality of the asset. If payoffs are concave in ex-post beliefs, for example, due to risk aversion, agents choose equilibrium tests that are easier to pass than in the linear case. If payoffs are convex, for example, because the agent has an alternative use for the asset besides selling it, then in equilibrium higher standards will be imposed.

The rest of the paper is organized as follows. We finish the introduction by discussing related literature. In Section 2 we provide the model and equilibrium characterization that features the minimum principle. We then apply this principle to describe further which tests are used in equilibrium and demonstrate the bias toward simple certification tests. In Section 3 we show that the equilibrium certification standard is informationally inefficient. In Section 4 we endogenize the possibility that the agent obtains no test results by including heterogeneous private costs of test-taking and show that the insights from Sections 2 and 3 apply to this richer framework. In Section 5 we analyze third-party certification. Finally, Section 6 considers extensions to the model by relaxing some of our assumptions.

1.1. Literature Review

Our paper is related to a few strands of the literature. First, it belongs to a growing literature that examines communication via strategic/voluntary disclosure of verifiable information. If it is

laboratory in the world that conducts safety and quality tests on a broad range of products, from fire doors to CCTV cameras.

known that the agent has information, there is unraveling in equilibrium (as in Grossman (1981) or Milgrom (1981)), but as shown by Dye (1985) and Jung and Kwon (1988), if there is a possibility that the agent has no information, then the equilibrium is only partially revealing with low types choosing to pool with the uninformed agents and not disclose information.⁴ We can interpret the minimum principle as a broad generalization of the “unraveling” result of Grossman (1981) and Milgrom (1981): in equilibrium, nondisclosure is interpreted in the worst possible way.

The question of the informational efficiency of disclosure is considered in Glazer and Rubinstein (2006), Sher (2011), and Hart, Kremer and Perry (2016), among others. In contrast to the conclusion of our paper, these papers argue that there exists an informationally efficient equilibrium, which may be unique subject to a certain refinements. The key difference is that these papers assume that the agent is endowed with information (pieces of evidence) that he can disclose. In our paper, the agent chooses both the distribution of signals he obtains and the disclosure policy. Our results on the informational inefficiency of the equilibrium choice of tests show that endogenizing test design has first-order consequences for the results.⁵

In a recent paper, Ben-Porath, Dekel and Lipman (2017) study a related problem of project choice with voluntary disclosure. In their model, the agent (a manager, politician, etc.) chooses a project from a set of projects with different payoff distributions. If the project’s payoff is realized, the agent has the option to disclose it. When the two projects have the same mean, have the same probability of early payoff realization, and are ranked by second-order-stochastic dominance, they show that the option to disclose makes the agent prefer the riskier project. This result can be viewed as a special case of our result (Proposition II) that the agent will pick more informative tests, as more informative tests imply a riskier conditional valuation for the asset.

The bulk of their analysis is on projects with different means, and the focus, interpretation and other results of the two papers are very different (for example, they consider disclosure by the agent’s adversary). Still, their result that the agent in equilibrium may choose an inefficient project has an analog in our setting: If tests have different costs, then in our setting the agent might choose a costlier test if it has a lower nondisclosure payoff.

⁴ This model has been extended to a dynamic setting with information arriving gradually and exogenously in Acharya, DeMarzo and Kremer (2011), and to multidimensional problems (for example, Pae (2005) or Guttman, Kremer and Skrzypacz (2014)).

⁵ A separate literature asks about how the regulation of accounting disclosures affect competition. See for example Friedman, Hughes, and Saouma (2016) that focuses on the effects of mandating conservatism in financial reports.

Our paper is also related to papers where information disclosure is costly, as in Jovanovic (1982), Verecchia (1983), Einhorn and Ziv (2008) or Marinovic, Skrzypacz and Varas (2015). In Section 4, testing is costly and hence the market beliefs about agent's type when he discloses no information depend on this cost. The main contribution of our paper to this literature is again to endogenize the type of evidence the agent acquires.

Finally, our paper is related to Lizzeri (1999) who focuses on certification by a monopolistic third party that commits ex-ante to both a certification rule and a fee structure.⁶ The main result in that paper is that in equilibrium the monopolist charges for a test that provides no information to the market. We focus instead on the choice of the test (from some set of tests) and allow the possibility that the test yields non-verifiable information (or are too costly for some agents to take). We obtain a less paradoxical result – some information is provided in equilibrium, and some equilibrium forces favor more rather than less information. At the same time, our result that binary tests tend to have too-low passing thresholds is related to that in Lizzeri (1999). Indeed, if tests *always* yielded a verifiable result (no possibility of a null outcome), the minimum principle would call for the test with the smallest possible threshold, and so passing could become uninformative in our model as well.

2. General Tests

2.1. The Model

Consider a risk-neutral agent who holds an asset and plans to sell it in a competitive market of risk-neutral investors. The asset's market value $v(q)$ depends on its quality q , which is initially unknown but is distributed according to some common prior distribution. We assume $v(q)$ has finite mean and variance.

Though initially uninformed, the agent has the opportunity to certify and verify the asset's quality by taking a test. There is a set of K tests available, and each test $k \in K$ is described by a random variable S_k , which is the outcome of the test. The distribution of S_k depends on q , so the signal is informative about the asset's quality. Conditional on taking test k and realizing outcome S_k ,

⁶ See also Albano and Lizzeri (2001) for analysis of how intermediaries affect incentives to invest in quality.

the expected value of the asset is equal to $E[v(q) | S_k]$, which we assume is finite for all outcomes of the test.

For each test, there is the possibility that the result will provide no verifiable information. We call this outcome a “null result,” denoted by \emptyset , and assume that for each test k , there is positive probability of such an outcome: $\Pr(S_k = \emptyset) > 0$ for all k . Because a null result cannot be verified, the agent can always *claim* to have received a null result rather than disclose his information. This possibility prevents equilibrium unraveling to full disclosure, as in Grossman (1981) and Milgrom (1981), so that in our context both the agent’s test choice and disclosure rule are non-trivial.

We do not require that the null outcome be independent of q , so $E[v(q) | S_k = \emptyset]$ need not equal $\bar{v} \equiv E[v(q)]$. To avoid trivial outcomes, we do impose the minimal requirement that there exists at least one test k that might produce verifiable “good news” so that

$$\Pr(S_k \neq \emptyset \cap E[v(q) | S_k] > \bar{v}) > 0. \quad (1)$$

The game proceeds as follows. First, the agent privately chooses a test k . Second, he learns whether the test yields a verifiable result and if so, what the result is. Then he decides whether to disclose the test’s result. Finally, he sells the asset at a price equal to the market’s expectation of its value conditional on the information released to the market. That is, we assume that investors are competitive, risk-neutral, and symmetrically informed so that the price of the asset -- which is the agent’s payoff -- is the asset’s conditional expected value based on its perceived quality.⁷

The agent’s strategy thus has two components: he chooses which test, k , to take and whether to reveal the result, S_k . The agent’s strategy can thus be described by a test k and a disclosure policy $\theta(S_k) \in \{0, 1\}$. If $\theta(S_k) = 0$, the agent discloses and investors learn both the test and its outcome. If $\theta(S_k) = 1$, the agent does not disclose and claims a null result (which is impossible for investors to verify). We allow for mixed strategies, so that given S_k , $\theta(S_k)$ is a random variable independent of q ; if the agent mixes, then $E[\theta(S_k) | S_k]$ is the probability the agent does not disclose given

⁷ One way to interpret the agent is that a non-profit self-regulatory organization (SRO) designs voluntary certification program with the goal of maximizing the participants’ expected payoffs.

outcome S_k .⁸ Because the agent has nothing to disclose in the event of a null result, a feasible disclosure policy must have $\theta(\emptyset) = 1$, a requirement we impose throughout the paper.

For now, we assume the agent always takes a test (but the set of tests can include a test that always yields the null result). We will also consider settings in which testing might not be available, or be too costly. If the agent does not take a test, he also has nothing verifiable to report; we can think of the null result as subsuming this possibility as well.

The agent chooses a test and disclosure policy to maximize the expected price he will receive for the asset. Let $\pi^D(k, s)$ be the price the agent expects upon disclosing $S_k = s$ and π^N be the price the agent expects in the event of nondisclosure. Then an equilibrium can be described as follows:

DEFINITION. An equilibrium with general tests is a collection $\{k_E, \theta_E, \pi_E^N, \pi_E^D(k, s)\}$ that satisfies

- (i) $\pi_E^D(k, S_k) = E[v(q) | S_k]$, the expected value given the test and outcome,
- (ii) $\pi_E^N = E[v(q) | \theta_E(S_{k_E}) = 1]$, the expected value given nondisclosure, and
- (iii) (k_E, θ_E) solves $\max_{(k, \theta)} E[(1 - \theta(S_k))\pi_E^D(k, S_k) + \theta(S_k)\pi_E^N]$.

The first two conditions state that prices offered by the market are equal to asset's expected value given the equilibrium beliefs and disclosed information. Condition (iii) is that the agent maximizes jointly over the choice of tests and disclosure policies given market prices.

Alternative Applications

Although we have formulated our model in terms of an agent selling an asset, there are many other applications that our model captures. Two important alternative settings include quantity competition and agency problems within organizations.

Quantity Maximization: The agent is selling a product at a fixed price r (for example, a regulated price, or a fixed subsidy per student in case of charter schools) and thus earns a profit proportional to the number of units sold. There is a continuum of risk-neutral potential buyers indexed by l ,

⁸ We could also allow for a mixed strategy in the choice of tests, however, as we discuss later in the proof of PROPOSITION I, there is always an equilibrium in pure strategies, and any mixed strategy equilibria are non-generic and payoff equivalent. To simplify notation we assume that the agent does not randomize in his choice of tests.

distributed uniformly over $[0, L]$. We may interpret l as a taste parameter, or as the value of the buyer's outside option, so that type l has valuation $v(q) - l$ for the asset. We assume the support of $v(q)$ is within $[r, r + L]$, in which case demand is linear in the posterior mean of v . In summary, the agent in this model maximizes expected quantity sold, and that in turn is a linear function of the posterior mean of v , so that preferences are the same as in the benchmark model.

Internal Organization: The agent is an employee who tries to produce information in order to influence the decisions of the principal. In particular, suppose the agent gathers information regarding the return on investment $v(q)$ of a project. Given the information provided by the agent, the principal makes an investment decision I at a cost $I + cI^2$ in order to maximize the expected profit $E[1 + v(q)]I - I - cI^2$. If we assume the agent enjoys private benefits that are proportional to the scale of investment (a common assumption in internal budgeting literature), then the agent's payoff is linear in the posterior expectation of v .

2.2. Equilibrium Tests: The Minimum Principle

Which test will the agent take in equilibrium? Which results will he reveal and which will he hide? Note that given any *fixed* test k and disclosure policy θ , the martingale property of beliefs implies that the expected price the agent will receive is equal to \bar{v} . Thus it might seem that the agent will be indifferent between the choice of tests.

Instead, we show that generically (for a finite set of available tests) the equilibrium is unique and that it can be characterized by a “minimum principle”: the equilibrium test and disclosure rule are chosen to minimize the price upon nondisclosure, π_E^N . Specifically, define for any test k and disclosure policy θ the corresponding nondisclosure price⁹ if the test and disclosure policy were known:

$$\pi^N(k, \theta) \equiv E[v \mid \theta(S_k) = 1].$$

Then we have the following result.

⁹ Throughout, we restrict (k, θ) to be feasible and interpret v as the random variable $v(q)$.

PROPOSITION I (MINIMUM PRINCIPLE). In an equilibrium with general tests the agent earns \bar{v} , and chooses a test and disclosure policy that minimize the nondisclosure price:

$$(k_E, \theta_E) \in \arg \min_{k, \theta} \pi^N(k, \theta).$$

PROOF: For an arbitrary strategy (k, θ) , given equilibrium condition (i) and the definition of $\pi^N(k, \theta)$, the law of iterated expectations implies:

$$E\left[(1 - \theta(S_k))\pi^D(k, S_k) + \theta(S_k)\pi^N(k, \theta)\right] = \bar{v}.$$

In words, if the market knows the test and disclosure policy, the expected price must equal the asset's expected value.

Because investors correctly anticipate the agent's choice of (k_E, θ_E) , the agent's equilibrium payoff must be \bar{v} . Moreover, for an arbitrary deviation (k', θ') , if the deviation were observed by the market, then nondisclosure price would change to $\pi^N(k', \theta')$, and the expected agent's payoff would remain \bar{v} . However, because the market does not observe the deviation, the nondisclosing agent would still receive π_E^N . Hence, if $\pi_E^N > \pi^N(k', \theta')$, the deviation is profitable. On the other hand, if $\pi_E^N \leq \pi^N(k', \theta')$ for all (k', θ') then no profitable deviation exists.

Though we have (for simplicity) restricted the agent to pure strategies in the test choice, the result is unchanged if we allow randomization. If the agent were expected to randomize over a set of tests, the price upon nondisclosure would equal the expected value conditional on a mixture of events where each event corresponds to the choice of a specific test. The resulting non-disclosure price would be a weighted average of the nondisclosure prices for each of the tests in the support of the randomization, and so could not be lower than the minimum over all tests, π_E^N . The agent would therefore randomize only over tests with non-disclosure payoff equal to π_E^N , and the proposition identifies the nondisclosure price and the support of possible randomization. ♦

We can interpret the minimum principle as a broad generalization of the “unraveling” result of Grossman (1981) and Milgrom (1981). In the classic unraveling result, nondisclosure is interpreted to mean the agent has the worst possible information. Here too the market interprets nondisclosure as the worst possible outcome given the strategies available to the agent.

This characterization of the equilibrium has several useful implications. First, we can decompose the minimization problem in PROPOSITION I into two steps: finding the minimizing disclosure policy for any test and then minimizing over the tests. Doing so allows us to define, for any test k , the *nondisclosure price* of that test as

$$\pi_k^N \equiv \min_{\theta} \pi^N(k, \theta). \quad (2)$$

In other words, π_k^N is the equilibrium nondisclosure price in the simpler, pure disclosure game with only a single test. Given this definition, we can restate PROPOSITION I as follows:

COROLLARY: In equilibrium the agent chooses test k with the smallest nondisclosure price π_k^N .

Next, consider the optimal disclosure policy for a given test. Given any nondisclosure price, π^N , the agent optimizes by disclosing results that would lead to a better valuation than the nondisclosure price and concealing results otherwise. That is, the following threshold disclosure policy is optimal:¹⁰

$$\theta(s) = 0 \text{ if } \pi^D(k, s) > \pi^N \text{ and } \theta(s) = 1 \text{ if } \pi^D(k, s) \leq \pi^N. \quad (3)$$

Thus, given the single test k , if the agent does not disclose then investors can infer that either $S_k = \emptyset$ or $E[v | S_k] \leq \pi_k^N$. This observation allows us to alternatively characterize the nondisclosure price of test k as the solution to a fixed-point problem:

COROLLARY: For any test k the nondisclosure price, π_k^N , is the unique solution to the fixed-point problem:

$$\pi_k^N = E\left[v | S_k = \emptyset \text{ or } E[v | S_k] \leq \pi_k^N\right]. \quad (4)$$

PROOF: Any solution to (4) clearly defines an equilibrium in the single test game, and thus by PROPOSITION I must equal π_k^N , which is unique. To see directly that (4) must correspond to the

¹⁰ When $\pi^D(k, s) = \pi^N$, the agent is indifferent and can randomize. But because in equilibrium this condition implies $E[v | S_k = s] = \pi^N$, such randomization will have no impact on payoffs or the equilibrium nondisclosure price. We thus adopt, without loss of generality, the convention that the agent does not disclose in case he is indifferent.

minimum in (2), note that any other disclosure policy would potentially remove from the set of non-disclosers types with expected value below π_k^N , or add to the set types with values above π_k^N , either of which would raise the nondisclosure price. ♦

Combining these observations, we have the following useful characterization of the equilibrium nondisclosure price for a given test (which generalizes the characterization of the disclosure threshold for a single test given by Acharya, DeMarzo, and Kremer (2011)).

COROLLARY: For any test k the equilibrium nondisclosure price satisfies

$$\pi_k^N = \min_{\pi} E[v | S_k = \emptyset \text{ or } E[v | S_k] \leq \pi]. \quad (5)$$

Comparing (5) with the definition (2), we see that it is sufficient to consider only threshold disclosure policies in the minimization to determine π_k^N .

In words, the fix-point characterization is that the equilibrium nondisclosure price for test k equals both the average expected value of the asset conditional on non-disclosure and the expected value of the asset conditional on receiving a marginal test result. Since the equilibrium disclosure policy is a threshold policy, when the marginal concealed expected value equals the average concealed value, the average is minimized (this logic is analogous to the case of production theory with increasing marginal costs and positive fixed costs that the average cost and marginal cost cross exactly once at the minimum of the average cost curve). We illustrate this result in Figure 1.

2.3. Applying the Minimum Principle

We now present a concrete example illustrating the minimum principle and its implications (which we generalize in the next subsection). Suppose the asset is either good or bad, with $v = q \in \{0, 1\}$ and $\bar{v} = \Pr(q = 1) \in (0, 1)$. Consider a test k that reveals with probability p_k the verifiable signal $S_k = q + \epsilon_k$, where the error ϵ_k is independent of q with distribution F_k and log-concave density f_k . With probability $1 - p_k$, also independent of q , the test fails and no verifiable information is revealed ($S_k = \emptyset$). Then the expected value of the asset given disclosure can be calculated using Bayes' rule as:

$$\pi_E^D(k, s) = E[v | S_k = s] = \Pr(q = 1 | S_k = s) = \frac{f_k(s-1) \Pr(q = 1)}{f_k(s-1) \Pr(q = 1) + f_k(s) \Pr(q = 0)}. \quad (6)$$

Note that the log-concavity of the error density implies a monotone likelihood ratio, so that the disclosure payoff is increasing in the signal s . Therefore, the optimal disclosure rule will be a threshold rule based on the signal. The minimum principle then implies that to determine the equilibrium disclosure threshold and corresponding nondisclosure payoff given test k , we can directly calculate:

$$\begin{aligned} \pi_k^N &= \min_s E[v | S_k = \emptyset \text{ or } S_k \leq s] \\ &= \min_s \frac{[1 - p_k + p_k F_k(s-1)] \Pr(q=1)}{1 - p_k + p_k [F_k(s-1) \Pr(q=1) + F_k(s) \Pr(q=0)]}. \end{aligned} \quad (7)$$

Figure 1 below shows the equilibrium threshold and nondisclosure payoff given a test with a normally distributed standard error of $\sigma_\epsilon = 100\%$ and a failure rate of $1 - p_k = 2\%$. The blue curve shows the price if the agent reveals signal s , and the red curve shows the price if the agent reveals nothing and the market believes his disclosure rule is to reveal signals above s . In this case, the equilibrium disclosure threshold is a signal value of $s^* = -0.9265$, with a nondisclosure payoff of 19.4%.¹¹ As the figure illustrates, consistent with equations (4) and (5) the minimum nondisclosure payoff is equal to the disclosure payoff at the threshold.¹²

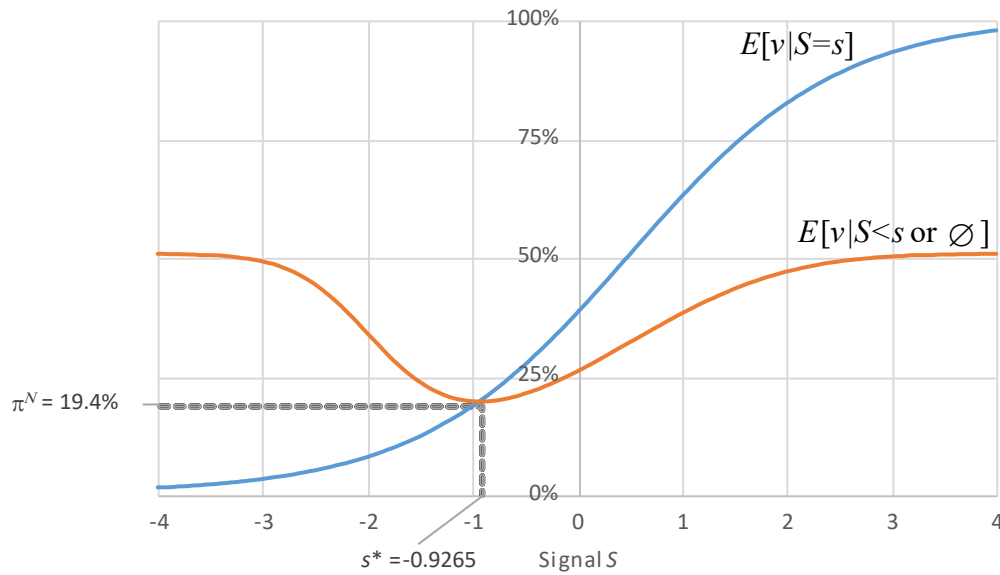


Figure 1: Disclosure Payoff and Nondisclosure Payoffs.

¹¹ We express prices in terms of the probability the value is high.

¹² The intuition for why the blue curve crosses the orange curve at its minimum is analogous to the intuition why in production theory the marginal cost curve crosses the average cost curve at its minimum.

Shown are the expected value conditional on signal S and the nondisclosure payoff if that signal were the disclosure cutoff. These payoffs coincide at the equilibrium disclosure threshold, which minimizes the nondisclosure payoff. (Test has noise volatility equal to the quality differential and a failure rate of 2%.)

2.4. Reliability and Informativeness

Suppose there are multiple tests with different failure rates and that generate signals with differing precision. How would the agent select between these tests? The minimum principle implies that the tests can be ranked according to single statistic, their equilibrium nondisclosure payoff π_k^N . From (7), in our example the nondisclosure payoff increases with the failure rate $1 - p_k$ of the test and the standard error of the signal (σ_ϵ). This observation that the agent would choose less noisy tests with lower failure rates is quite general, as we show in the following result:

PROPOSITION II (RELIABILITY AND INFORMATIVENESS).

Consider tests S_1 and S_2 such that

- (i) $S_1 = \emptyset \Rightarrow S_2 = \emptyset$, i.e., S_1 is *more reliable* than S_2 ; and,
- (ii) $E[v | S_1, S_2] = E[v | S_1]$, i.e., S_1 is *more informative* than S_2 .

Then $\pi_1^N \leq \pi_2^N$, and so in equilibrium the agent will choose more reliable and more informative tests.

PROOF: For any test k , and any nondisclosure price π^N , the agent chooses a disclosure policy to maximize the expected payoff,

$$\begin{aligned} E\left[(1 - \theta(S_k))E[v | S_k] + \theta(S_k)\pi^N\right] &= \pi^N + E\left[(1 - \theta(S_k))(E[v | S_k] - \pi^N)\right] \\ &= \pi^N + E\left[\left(E[v | S_k] - \pi^N\right)^+ 1[S_k \neq \emptyset]\right], \end{aligned} \quad (8)$$

where the final expression follows from the optimality of a threshold disclosure policy. Note that by the law of iterated expectations and equation (4) if we use π_k^N as the nondisclosure price in (8) the agent's expected payoff is \bar{v} .

Because the term $\left(E[v | S_k] - \pi^N\right)^+$ is non-negative as well as convex in $E[v | S_k]$, (8) is weakly increasing in the set of non-null outcomes, as well as in the variability (in the sense of mean-

preserving spread) of the posterior $E[v | S_k]$. Therefore, for the same π^N , the payoff in (8) is weakly higher for test 1 than for test 2. However, because (8) is increasing in π^N , to obtain the same expected payoff \bar{v} , we must have $\pi_1^N \leq \pi_2^N$. ♦

In reality, we might expect there to be a tradeoff between the accuracy of a test and the likelihood of failure (more accurate tests might take longer and thus not be completed by the time of sale, or, as we will model later, might be too costly for some agents to undertake). Returning to our parametric example, we can calculate the agent’s “indifference curve” for these test attributes by computing, for any increase in signal noise, the corresponding decrease in the failure rate that would make the test equally attractive. Figure 2 below shows the set of tests that lead to the same nondisclosure payoff as the test illustrated in Figure 1. By PROPOSITION I, the agent would prefer any test below the curve in Figure 2 to the test in Figure 1 (in the sense that if any test below the curve is available, in equilibrium the agent does not choose tests on the curve).




Figure 2: Equivalent Tests

Shown are the combinations of test noise and failure rate that have the same minimum nondisclosure payoff as the test in Figure 1. Tests below this curve would be strictly preferred by an agent (given equilibrium prices).

2.5. Simple Tests and Certification Standards

Suppose the set of tests is rich in the sense that a new test can be constructed by coarsening an existing signal. For example, given a test S with multiple outcomes, the agent can select a simple pass/fail version of this test that would only certify whether the test result is in a given subset of the possible verifiable outcomes. A special case of a simple test is one which certifies whether the asset is of sufficiently high quality. We can define such a test according to a certification standard which specifies a threshold, x , for the asset's valuation:

$$S_x \equiv \begin{cases} 1 & \text{if } S \neq \emptyset \text{ and } E[v(q) | S] > x, \\ \emptyset & \text{if } S = \emptyset \text{ or } E[v(q) | S] \leq x. \end{cases}$$

We denote by $\pi^N(x)$ the nondisclosure price for the certification test S_x with standard x .

Because any simple test is a coarsening of the original test, we know from PROPOSITION II that it must have a weakly higher nondisclosure payoff. Somewhat more surprising is the fact that there exists a certification standard with the same nondisclosure payoff as the original test.

PROPOSITION III (SUFFICIENCY OF CERTIFICATION STANDARDS).

Consider any test S with nondisclosure payoff π_S^N . Then the certification test with standard $x^* = \pi_S^N$ has the same nondisclosure payoff; that is

$$\pi^N(x^*) = \min_x \pi^N(x) = \pi_S^N.$$

Therefore, given a choice of all possible certification standards for a given signal S , $x_E = \pi_S^N$ is an equilibrium.

PROOF: From PROPOSITION II, for any x , $\pi^N(x) \geq \pi_S^N$. Now let $x^* = \pi_S^N$, and suppose the agent discloses iff $S_{x^*} = 1$. Then $\theta(S_{x^*}) = 1$ iff $S_k = \emptyset$ or $E[v | S] \leq x^*$. Because the nondisclosure price is the minimum over all disclosure policies, we get:

$$\pi^N(x^*) \leq E[v | S = \emptyset \text{ or } E[v | S] \leq \pi_S^N] = \pi_S^N.$$

The last equality follows from (4). The fact that x^* is an equilibrium certification standard then follows from the minimum principle. ♦

As an example, consider again Figure 1. We can interpret the blue curve as the certification standard x , and the orange curve as the nondisclosure price given that standard. The equilibrium standard is again the minimum nondisclosure price over all standards x , which coincides with the nondisclosure price of the original test. Figure 1 also illustrates a more general property:

LEMMA. $\pi^N(x)$ is quasiconvex.

PROOF: Given an optimal disclosure policy, for any standard x , $\pi^N(x) \leq \bar{v}$. For $\pi^N(x) < \bar{v}$, because $\pi^N(x) = E[v | S = \emptyset \text{ or } E[v | S] \leq x]$, then $\pi^N(x)$ is strictly decreasing (increasing) only if $x < (>) \pi^N(x)$. Quasiconvexity then follows from the uniqueness of the fixed point $x^* = \pi^N(x^*) = \pi_S^N$ from (4). ♦

Thus, combining PROPOSITION II and PROPOSITION III, we see that market forces only drive sellers to produce information that is useful in distinguishing themselves from non-disclosers. There is no benefit for refining further results that would clearly be disclosed (or not disclosed). As a result, if refined information is more costly or less reliably produced, we should expect that when sellers choose the information to provide, simple certification standards will dominate the market. We explore next the optimal choice of such standards from a welfare perspective.

3. Minimum Certification Standards

As we have seen in the previous section, market forces may lead sellers to choose coarse pass/fail certification tests, and in that sense tests need not be informationally efficient. In this section we restrict attention to certification tests and ask whether, within that class, the equilibrium certification standard is socially optimal. To do so, we first define a standard measure of informational efficiency which also corresponds to social welfare in the alternative applications introduced in Section 2.1. We then show that a planner could improve welfare by raising the minimum certification standard.

3.1. Information and Welfare

Consider a social planner whose objective is to align market prices with actual quality, measured by the square distance between the true value of the asset and its market price. Given a signal S

(that has some informative outcomes in the sense of (1)) and a certification standard x such that $E[v | S_x = \emptyset] \leq \bar{v}$, market prices in the event of disclosure or nondisclosure are given by

$\pi^D(x) \equiv E[v | S_x = 1]$, the expected value conditional on passing S_x , and

$\pi^N(x) \equiv E[v | S_x = \emptyset]$, the expected value conditional on failure/nondisclosure.

If $E[v | S_x = \emptyset] > \bar{v}$, then disclosure is not optimal and we define $\pi^N(x) \equiv \pi^D(x) \equiv \bar{v}$.¹³

Then the planner would like to choose the threshold x to minimize the residual variance,¹⁴

$$L(x) \equiv E\left[(v - E[v | S_x])^2\right] = E\left[S_x(v - \pi^D(x))^2 + (1 - S_x)(v - \pi^N(x))^2\right]. \quad (9)$$

This objective is a simple way to capture the value of information. In our benchmark model, where agents are risk-neutral and buyers compete away all of their surplus, information has no direct value to the players and $L(x)$ must be interpreted as an external social benefit. We can easily extend our model to settings in which the price also plays an allocative role so that minimizing $L(x)$ coincides with maximizing total surplus.

For example, consider the alternative application (described in Section 2.1) of a producer selling at a regulated price to consumers with heterogeneous tastes (or outside options). With a linear demand curve, the seller's profit is linear in the posterior value, $E[v | S_x]$. However, total consumer surplus is quadratic in the posterior value. As a result, expected welfare is maximized by maximizing the variance of the posterior. From the law of total variance, the variance of the posterior is equal to the total variance less the residual variance:

$$E\left[E[v | S_x]^2\right] = E[v^2] - E\left[(v - E[v | S_x])^2\right] = E[v^2] - L(x). \quad (10)$$

Hence, the certification standard that minimizes $L(x)$ also maximizes welfare. Here, better information improves the allocation of the good to consumers.

The same is true for principal-agent setting (also described in Section 2.1) in which the agent tries to influence the manager's investment decision. In that case, the project's quality determines the

¹³ This case is only relevant if a null outcome of S is good news ($E[v | S = \emptyset] > \bar{v}$) and x is sufficiently small. Because disclosure would not be optimal, the market price will equal \bar{v} .

¹⁴ Here we treat \emptyset mathematically as zero. We prove existence of a social optimum in the appendix.

return on investment. With quadratic adjustment costs, investment is linear in the expected return $E[v | S_x]$. But the total NPV of the project is again quadratic in the posterior value, and the same conclusion holds: Better information improves the allocation of investment, and so minimizing $L(x)$ maximizes welfare.

Finally, even within our standard model in which the seller sells an asset to one of several competing buyers, we can generate a similar result if one of the buyers is a strategic buyer with the opportunity to invest and increase the value of the asset. While the price of the asset will be determined by $E[v | S_x]$, the strategic buyer's expected surplus will increase with better information.¹⁵

3.2. Socially Optimal Standards

We have previously characterized the agent's threshold choice in PROPOSITION III. Suppose now that a social planner cannot set prices (which must satisfy investor optimality), but can choose a minimum testing standard x . What standard would the social planner choose?¹⁶

Given an arbitrary standard x , define $\bar{\pi}$ to be the average of the disclosure and nondisclosure prices:

$$\bar{\pi}(x) \equiv \frac{\pi^N(x) + \pi^D(x)}{2}.$$

Now consider a standard x and outcome S such that $E[v | S] \in (x, \bar{\pi}(x))$. Because the asset's conditional value is closer to π^N than π^D , at the given prices informational efficiency would be improved if an agent with this outcome did not pass the test. Alternatively, if the standard and outcome of the test are such that $E[v | S] \in (\bar{\pi}(x), x)$, informational efficiency would be improved

¹⁵ Specifically, assume the strategic buyer can invest i to earn $(1+i)v - ci^2$. In a second price auction, the strategic buyer would pay $E[v | S]$ when competing against financial buyers, but would earn a surplus proportional to $E[v | S]^2$.

¹⁶ Because the market price is equal to the conditional expectation of v given the information provided to the market, under this objective the planner has no incentive to distort prices ex post (i.e. if they do not affect the choice of tests and disclosure policy). One may wonder whether for a given test, the planner may want to commit to distort prices ex ante to affect the disclosure policy and reduce L this way. Hart, Kremer and Perry (2016) prove that for a fixed test it is not possible. Our result in this section can be interpreted as showing that price distortions could be useful if they can affect the kind of evidence the agent collects since prices could be a way to implement the planner-optimal test we describe.

if the agent were to pass the test. Therefore, in either case, changing the standard from x to $\bar{\pi}(x)$ improves price efficiency:

$$\begin{aligned}
L(x) &= E \left[S_x (v - \pi^D(x))^2 + (1 - S_x)(v - \pi^N(x))^2 \right] \\
&\geq E \left[S_{\bar{\pi}(x)} (v - \pi^D(x))^2 + (1 - S_{\bar{\pi}(x)})(v - \pi^N(x))^2 \right] \\
&\geq E \left[S_{\bar{\pi}(x)} (v - \pi^D(\bar{\pi}(x)))^2 + (1 - S_{\bar{\pi}(x)})(v - \pi^N(\bar{\pi}(x)))^2 \right] = L(\bar{\pi}(x)).
\end{aligned} \tag{11}$$

The first inequality follows because agents pass only if their value is closer to π^D than π^N , while the second inequality follows since setting prices equal to conditional expectations given the test $\bar{\pi}(x)$ minimizes the mean-squared error. Both inequalities are strict as long as there are outcomes for which $E[v | S]$ lies between x and $\bar{\pi}(x)$, and so prices change. This observation leads to the following characterization:

PROPOSITION IV (SOCIALY OPTIMAL STANDARDS). There exists a socially optimal test threshold $x_{SO} = \bar{\pi}(x_{SO}) > x_E$. That is, the socially optimal threshold x_{SO} exceeds the unconstrained market equilibrium x_E . If the planner sets x_{SO} as the minimal allowable certification, agents will choose x_{SO} in equilibrium.

PROOF: In the appendix we establish the existence of a social optimum $\hat{x} \in \arg \min_x L(x)$. Then from (11), $x_{SO} = \bar{\pi}(\hat{x})$ is also a social optimum with the same prices as \hat{x} , and therefore $x_{SO} = \bar{\pi}(x_{SO})$. Because the test is informative, at the optimum $\pi^D(x_{SO}) > \pi^N(x_{SO})$. Therefore, from (11),

$$x_{SO} = \bar{\pi}(x_{SO}) > \pi^N(x_{SO}) \geq \pi^N(x_E) = x_E,$$

where the last inequality follows from the minimum principle. If the planner sets x_{SO} as the minimal available test, then since $\pi^N(x)$ is increasing above x_{SO} (by quasiconvexity), the minimum principle implies the agent will choose x_{SO} in equilibrium. ♦

Figure 3 illustrates the comparison between x_E and x_{SO} .




Figure 3: Informationally-Efficient Threshold x_S versus Equilibrium Threshold x_E . (Parameters as in Figure 1.)

Our result that minimum passing standards can improve informativeness of binary tests can be generalized to any coarse test with discrete outcomes. Suppose that a simple test is characterized by thresholds $x_1 < x_2 < \dots < x_n$ so that if $E[v|S] \leq x_1$ the agent fails the test, and otherwise the interval in which the expected value lies is disclosed. In that case, the nondisclosure payoff is determined solely by the passing threshold x_1 , and the equilibrium is characterized by $x_1 = \pi^N(x_1)$ as before. The same reasoning as above implies that raising x_1 by a positive amount would reduce the mean square error.¹⁷

4. Costly Tests and Endogenous Participation

In our model thus far, testing has been costless. Suppose instead there is a fixed cost c to take a test. If π_E^N is the minimum non-disclosure price, then the corresponding test remains an

¹⁷ Specifically, we can raise it to the average of x_1 and $E[v|E[v|S] \in (x_1, x_2)]$. How regulation of the higher thresholds would affect welfare is ambiguous. Having more grades is better, as we argued above, but the specific optimal location of the thresholds other than the lowest passing one depends on the details of the problem.

equilibrium as long as $c \leq \bar{v} - \pi_E^N$. In that case, the agent prefers to pay the cost to take the test (and thus receive $\bar{v} - c$) rather than not take a test and have nothing to report (and receive π_E^N).¹⁸

On the other hand, if the cost $c > \bar{v} - \pi_E^N$, then the agent would choose not to take any test. Indeed, in practice, a common reason an agent may not provide certification is that he found the test too costly to undertake. In this section, we consider the possibility that test-taking is costly and that the cost faced by the agent is drawn from a distribution. Because the agent will only take a test if the expected benefit outweighs the cost, the probability that a test is taken arises endogenously as part of the equilibrium. Moreover, a low cost agent will strictly gain from testing, earning an expected payoff, net of testing costs, in excess of $\bar{v} = E[v(q)]$.

4.1. Costly Testing Equilibrium

Let the random variable c be the private cost to the agent of taking a test, which we assume is continuously distributed with full support on $[0, \infty)$ and is independent of the test taken, its result, or the asset's ultimate quality. We assume that if the agent does not take the test, there is nothing he can disclose. Therefore, the agent is willing to take a test only if the expected gain from the option to disclose the test outcome exceeds the agent's cost. Given a reward for no disclosure, π_E^N , the gain from taking test k and having disclosure policy θ is given by:

$$G(k, \theta, \pi_E^N) \equiv E\left[(1 - \theta(S_k))\pi_E^D(k, S_k) + \theta(S_k)\pi_E^N\right] - \pi_E^N.$$

An agent facing cost c for taking test k would only do so if the benefit from selecting the optimal disclosure policy is at least as high:

$$c \leq \max_{\theta} G(k, \theta, \pi_E^N). \quad (12)$$

Given the agent's incentive constraint (12), we have the following definition of a costly testing equilibrium:

DEFINITION. A costly testing equilibrium is a collection $\{k_E, \theta_E, \pi_E^N, \pi_E^D(k, s), \bar{c}_E\}$ that satisfies

¹⁸ If $c < \min_i \Pr(E[v | S_i] \leq \pi_i^N)(\bar{v} - \pi_i^N)$ then the test-taking equilibrium is unique. Otherwise, there is also an equilibrium in which the agent does not take any test.

- (i) $\pi_E^D(k, s) = E[v(q) | S_k = s]$,
- (ii) $\pi_E^N = E[v(q) | \theta_E(S_{k_E}) = 1 \text{ or } c > \bar{c}_E]$,
- (iii) $(k_E, \theta_E) \in \arg \max G(k, \theta, \pi_E^N)$, and
- (iv) $\bar{c}_E = G(k_E, \theta_E, \pi_E^N)$

Condition (i) is straightforward and identical to our earlier definition. In condition (ii), we recognize that an agent with nothing to report may have found testing too costly to undertake.¹⁹ Condition (iii) is based on the agent selecting the optimal test and disclosure policy. Condition (iv) states that the highest testing cost the agent would pay is equal to the maximum expected gain from the option to disclose test results.

We then have the following characterization of an equilibrium, which also features the minimum principle. First, define

$$\pi^N(k, \theta, \bar{c}) \equiv E[v | \theta(S_k) = 1 \text{ or } c > \bar{c}]$$

to be the implied nondisclosure price if only agents with cost below \bar{c} take the test k and use disclosure rule θ . As we show next, the equilibrium test minimizes the nondisclosure payoff given the endogenously determined set of test-takers.

PROPOSITION V (COSTLY TESTING). Costly testing equilibria exist and are characterized by

- (i) The test and disclosure policy (k_E, θ_E) satisfy the minimum principle so that

$$\pi_E^N = \pi^N(\bar{c}_E) \equiv \min_{(k, \theta)} \pi^N(k, \theta, \bar{c}_E). \quad (13)$$

- (ii) The equilibrium maximal testing cost \bar{c}_E is the fixed point of:

$$\bar{c}_E = \frac{\bar{v} - \pi^N(\bar{c}_E)}{\Pr(c \leq \bar{c}_E)}. \quad (14)$$

PROOF: See Appendix.

¹⁹ When some agents have nothing to report because they find testing too costly, we no longer need to require that tests have a null outcome in order to avoid full revelation (although we continue to allow it for generality).

The intuition for condition (13) is analogous to our previous results, the only difference is that now we keep the set of test-takers fixed based on the equilibrium threshold \bar{c}_E . The intuition for condition (14) is that all agents get paid on average \bar{v} while non-test takers get paid π_E^N . That means that the test-takers earn an expected premium of

$$\frac{\bar{v} - \pi^N(\bar{c}_E)}{\Pr(c \leq \bar{c}_E)} = \max_{(k, \theta)} G(k, \theta, \pi^N(\bar{c}_E)) \equiv G(\bar{c}_E) \quad (15)$$

from taking the test. In equilibrium that premium must equal the cost of the marginal test-taker; that is, $G(\bar{c}_E) = \bar{c}_E$.

In the model with costly tests, the expected payoff of the agent is no longer \bar{v} , but is reduced by expected testing costs to $\bar{v} - E[c1_{c \leq \bar{c}_E}]$ on average. Yet despite the deadweight loss associated with testing, and in contrast to Section 2, now some agents strictly gain in equilibrium, earning more than \bar{v} . To see why, note that the expected payoff of an agent who takes the test is $\pi_E^N + G(\bar{c}_E) - c$, which exceeds \bar{v} as long as

$$c < \pi_E^N + G(\bar{c}_E) - \bar{v} = \frac{\Pr(c > \bar{c}_E)}{\Pr(c \leq \bar{c}_E)} (\bar{v} - \pi_E^N), \quad (16)$$

where the last equality follows from (14) and (15). Thus, testing allows low cost sellers to profit at the expense of high cost sellers.

Figure 4 illustrates an equilibrium with costly tests with an exponential distribution of costs (to take a test). Except for agents who choose not to take it, the test itself has no null outcomes; the remaining parameters match Figure 1. The horizontal axis represents the fraction of types that choose to take the test in equilibrium. The blue curve is the nondisclosure price $\pi^N(\bar{c})$, which decreases in the fraction of agents the market expects to take the test. The orange curve is the value $G(\bar{c})$ of the optimal test given \bar{c} from (14). Equilibrium condition (iv) implies that the fixed point is where the orange curve crosses the gray curve (that is the cost of the marginal test-taker given the fraction of test-takers).

While in this example the equilibrium is unique, multiple equilibria are possible unless we impose further conditions on the distribution of costs. Specifically, if $F(\bar{c}) = \Pr(c \leq \bar{c})$, then the

equilibrium will be unique as long as the function F^{-1} (the gray line in Figure 4) is sufficiently steep. In the extreme case in which costs are zero for a fraction p of the population, and prohibitively high for the remaining fraction, F^{-1} is vertical at p and the equilibrium will be unique and will coincide with the analysis of Section 2 with an additional proportion p of agents with null results.

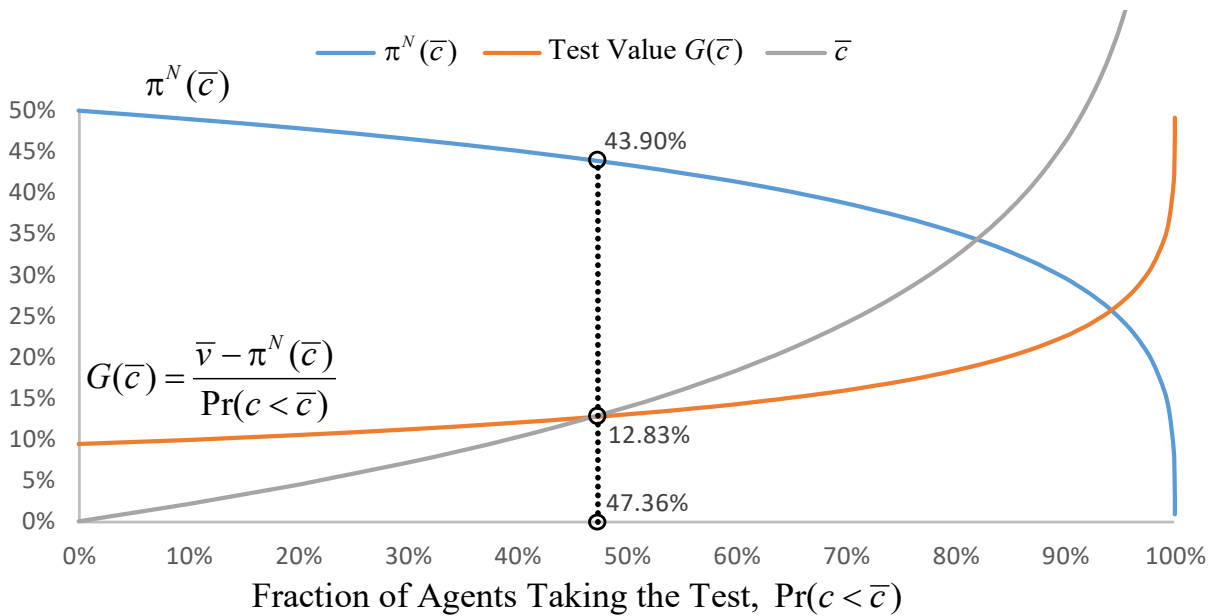


Figure 4: Example of equilibrium construction when c is distributed according to an exponential distribution with mean 20%. The remaining parameters are the same as in Figure 1. The test value is the RHS of (14); an equilibrium is found when this value is equal to \bar{c} . In this example, the equilibrium $\pi^N = 43.90\%$ and $\bar{c} = 12.83\%$. The 47.36% of agents who take the test receive an average price of $\pi^N + G(\bar{c}) = 56.73\%$.

We can also consider comparative statics with respect the cost of testing. Figure 5 shows the change in the nondisclosure price, the participation threshold (\bar{c}), and the aggregate testing costs of those who take the test, as we raise the cost of the test (we keep the distribution of costs to be exponential and vary the mean of that distribution).

First, consider the case in which testing costs approach zero. Then nearly all agents would take the test. If the test itself has no null outcomes, then the non-disclosure price converges to the minimal possible posterior for the asset value, and all types would disclose except for those with the worst possible test outcome. (This result is similar to Lizzeri (1999), though in a competitive environment.)

On the other hand, when the testing costs are very high, very few agents are willing to take the test, and so having nothing to report carries little information. The nondisclosure price therefore approaches the expected value of the asset.

More generally, an increase in testing costs raises the nondisclosure price and lowers the value of taking the test.²⁰ The reduced participation rate can cause the aggregate testing costs to fall, as shown in Figure 5 when the mean test cost exceeds 8%.




Figure 5: Change in nondisclosure price, participation threshold, and aggregate testing costs as a function of the average cost of the test. All parameters, except the mean test cost, match Figure 4.

4.2. Testing Costs and Informational Efficiency

Following our previous discussion, we now focus on certification tests based on the signal S . Again, we consider whether a minimum standard may improve informational efficiency. Because testing is both costly and endogenous, however, any change to the standard will also impact the amount of testing that occurs in equilibrium. We will show that setting a minimal standard has the additional benefit of reducing aggregate testing costs.

²⁰ When the equilibrium is not unique these comparative statics apply to the equilibria with the highest and lowest levels of participation.

Assume that the certification tests are based on a signal such that $E[v|S]$ is continuously distributed with full support. This assumption guarantees that for any test $x > x_E$, there is a positive probability the agent would pass the easier test x_E but fail the more difficult test x . Let x_E denote the threshold of the equilibrium certification test.

PROPOSITION VI (MINIMUM STANDARDS WITH COSTLY TESTING). Consider a family of certification tests $\{S_x\}$ and let x_E be the costly testing equilibrium threshold. (If the equilibrium is not unique, consider the largest equilibrium x_E .²¹) Generically, there exists a minimum standard $x_{MS} > x_E$ that would lead to a more informative equilibrium with lower overall costs of certification (because fewer agents would be tested).

PROOF: See Appendix.

The intuition behind this result is as follows. We know from **PROPOSITION IV** that if we increase x_{MS} slightly above x_E it leads to a direct improvement of informativeness for a fixed fraction of types that take tests, \bar{c} . However, as x_S increases above x_E the nondisclosure price goes up for the given \bar{c} and that reduces incentives to take the test. As fewer agents take the test, the informativeness of the test decreases. However, we claim that at least for small increases of x_{MS} increases above x_E the first effect dominates. The reason is that the nondisclosure price is minimized for the given \bar{c} at x_E so the gains in informativeness are of first order while the change in the nondisclosure price (and hence participation) is only of second order.

To summarize, compared to the test that is chosen in equilibrium, by raising the minimum certification standard a social planner could both increase informativeness *and* reduce testing costs.²²

²¹ In the case of multiplicity, the equilibrium with the smallest \bar{c} corresponds to the highest π_E^N (by (14)) and thus the highest x_E .

²² As we can see from Figure 5, if average testing costs are above 8%, the planner could also lower aggregate testing costs by imposing a tax on testing, reducing participation and raising π^N . But because π^N would continue to be determined according to the minimum principle, the planner would still gain by imposing a minimum standard.

5. Third-Party Certification

Thus far, we have assumed that the seller of the asset chooses the test design. In many markets, such as for credit ratings, audits, licensing exams, LEED certification, organic produce certification, or ISO 9000, the information that is acquired and disclosed is not determined by the agent himself, but by a trusted third party that acquires the information and charges for certification. The agent decides whether to pay for the certification or not, as well as whether to disclose the results, but in many situations cannot determine the design of the certification test.

In this section, we consider environments in which the design of the test is determined by a third-party. When designing a test, the incentives of a third-party certifier differ from that of an individual agent: the certifier seeks to maximize the total fees collected from the agents, rather than the price paid to the agent by investors. We now ask what certification test the certifier would choose to maximize its expected profits. We focus on the case of a monopolist certifier (though the results in the oligopoly case are similar).

We model monopoly certification as follows. There is a monopolist certifier who chooses a test k from some set K of available tests and posts a price m that the agent has to pay for taking the test.²³ The agent learns his private cost realization c (which, as in the previous section, is independent of the quality of the asset or the probability of obtaining the null outcome) and then decides whether to take the test (and pay m). If the agent takes the test, he learns the outcome of the test, S_k and decides whether to disclose it or not. Finally, the market can observe the test and price offered by the certifier and whether the agent disclosed any results. Conditional on this information, the market sets prices $\pi^N(k, m)$ and $\pi^D(k, s)$ equal to the conditional expectations of the value of the asset. An equilibrium of this game as follows.

DEFINITION. A monopoly certification equilibrium is a collection $\{k_M, m_M, \theta_M(k, m), \pi_M^N(k, m), \pi_M^D(k, s), \bar{c}(k, m)\}$ that satisfies

²³ We could allow the monopolist to set a fee conditional on receiving a non-null outcome, or even conditional on the agent choosing to disclose the test result, and the result that is disclosed. Because the agent's decision to take the test will depend only on the expected fee charged, setting a single upfront fee for taking the test is without loss of generality. Also, because the asset's value and the agent's cost of taking a test are independent, there is no value to the monopolist of using a menu of tests to screen agents based on their cost.

- (i) $\pi_M^D(k, s) = E[v(q) | S_k = s]$,
- (ii) $\pi_M^N(k, m) = E[v(q) | \theta_M(S_k) = 1 \text{ or } c > \bar{c}(k, m)]$,
- (iii) $\theta_M(k, m) \in \arg \max G(k, \theta, \pi_M^N(k, m))$,
- (iv) $\bar{c}(k, m) = G(k, \theta_M(k, m), \pi_M^N(k, m)) - m$, and
- (v) $(k_M, m_M) \in \arg \max m \Pr(c \leq \bar{c}(k, m))$

Conditions (i) - (iii) match those in the definition of a costly testing equilibrium from Section 4, with the only change that the agent no longer chooses the test k . Note in particular that unlike our previous settings, if the monopolist deviates and offers an out-of-equilibrium test or price, the market and agent observe it and adjust their nondisclosure price and disclosure policy accordingly. Condition (iv) pins down the marginal type taking the test, who has cost \bar{c} equal to the expected gain from taking the test net of the monopolist's fee. (In comparison to the previous model, the cost of taking the test now has two components: the physical cost of taking the test, c and the price m charged by the monopolist.)

Finally, condition (v) states that the monopolist chooses a test and its price to maximize expected revenue, which is given by the price per test-taker times the probability the agent takes the test. Indeed, the natural interpretation of this model, which we will maintain throughout, is that the monopolist faces a continuum of agents with private costs distributed across the population. Thus, $\Pr(c \leq \bar{c})$ represents the fraction of agents that take the test if the marginal type is \bar{c} .

We have defined the monopoly equilibrium by allowing the monopolist to set a price m per test, with the quantity of test-takers (agents with $c \leq \bar{c}$) determined in the market. We could alternatively allow the monopolist to commit to a quantity of test-takers (for example, by committing to a capacity of inspectors or testing facilities) and then let the price m be determined in the market. These two settings will coincide when the equilibrium is unique, for which a sufficient condition is that indifference condition for the marginal test-taker,

$$m = G(\bar{c}) - \bar{c} = \frac{\bar{v} - \pi^N(\bar{c})}{\Pr(c \leq \bar{c})} - \bar{c}, \quad (17)$$

has a unique solution \bar{c} for the equilibrium choice of the fee m .²⁴ If this condition is not satisfied, there may be multiple participation rates consistent with a given test price, where a low turnout is justified because with a low turnout, π^N will be high and the value of the test will be low. When the equilibrium with a price-setting monopolist is not unique, the equilibrium with the highest payoff for the monopolist (and the highest participation for a given price m) will be uniquely selected if the monopolist can instead set a quantity.

Recall that $\pi^N(k, \theta, \bar{c}) \equiv E[v(q) | \theta(S_k) = 1 \text{ or } c > \bar{c}]$ is the rational nondisclosure price for an arbitrary test, disclosure policy, and testing threshold. Then the following proposition describes the equilibrium of the monopoly certification problem and shows that the minimum principle is still useful in understanding what drives the test design in markets, though the monopolist also considers the surplus provided to the agents.

PROPOSITION VII (MONOPOLY CERTIFICATION). There exists a monopoly certification equilibrium. In the equilibrium with the highest monopolist's payoff, the test, disclosure policy, and testing threshold jointly minimize

$$\{k_M, \theta_M, \bar{c}_M\} \in \arg \min_{(k, \theta, \bar{c})} \pi^N(k, \theta, \bar{c}) + \bar{c} \Pr(c \leq \bar{c}), \quad (18)$$

and the monopolist sets $m_M = G(k_M, \theta_M, \pi_M^N) - \bar{c}_M$ with $\pi_M^N = \pi^N(k_M, \theta_M, \bar{c}_M)$.

PROOF: We know from the previous section that a different way to express the equilibrium condition (iv) is that the equilibrium testing threshold is given by the fixed point

$$\bar{c}_M(k, m) + m = \frac{\bar{v} - \pi_M^N(k, m)}{\Pr(c \leq \bar{c}_M(k, m))}, \quad (19)$$

where $\pi_M^N(k, m)$ is defined in condition (ii) and depends on the testing threshold. The expected revenue of the monopolist is $m \Pr(c \leq \bar{c}(k, m))$ which, given the equilibrium condition (19), can be re-written as

$$\bar{v} - \pi_M^N(k, m) - \bar{c}_M(k, m) \Pr(c \leq \bar{c}_M(k, m)). \quad (20)$$

²⁴ In Figure 4, the RHS of (17) corresponds to the gap between the orange and gray curves. A sufficient condition is that $G'(\bar{c}) < 1$ if $G(\bar{c}) > \bar{c}$.

So, given the equilibrium prices and participation decisions, the monopolist chooses (k, m) to minimize:

$$\pi_M^N(k, m) + \bar{c}_M(k, m) \Pr(c \leq \bar{c}_M(k, m)). \quad (21)$$

From the previous analysis, we know that given the market's beliefs about participation, $\bar{c} = \bar{c}_M(k, m)$, the equilibrium disclosure policy minimizes the nondisclosure price,

$\pi_M^N(k, m) = \min_{\theta} \pi^N(k, \theta, \bar{c})$. If we are considering the best equilibrium for the monopolist, we can choose the largest fixed point in (19) and that corresponds to the monopolist being able to choose the threshold \bar{c} directly and then adjust prices, so that condition (iv) is satisfied. Note that since the nondisclosure price in condition (ii) is continuous in the testing threshold, for every \bar{c} there exists a (unique) m such that condition (iv) is satisfied. That allows us to re-write minimization (21) for the best equilibrium as the monopolist choosing jointly:

$$\min_{(k, \theta, \bar{c})} \pi^N(k, \theta, \bar{c}) + \bar{c} \Pr(c \leq \bar{c}). \quad \blacklozenge$$

PROPOSITION VII highlights the two forces that determine the monopolist's test choice. First, given the level of participation (the monopolist's "market share"), the monopolist chooses the test with the lowest non-disclosure price, as that determines the value of the agent's outside option from not taking the test. Second, to ensure a given level of participation $p = \Pr(c \leq \bar{c})$ in the test, the monopolist must set a low enough fee so that the gain from taking the test is $\bar{c} = F^{-1}(p)$. Thus, the monopolist seeks to minimize the total rents given away: $\pi^N(k, \theta, \bar{c})$ to all types and $\bar{c} \Pr(c \leq \bar{c})$ to test takers. This result has several consequences.

First, suppose the distribution of costs is binary, so that some fraction $F(0) > 0$ of agents have a low cost c_0 to take the test, and all others have a prohibitively high cost of taking the test.²⁵ Then it is optimal for the monopolist to choose $\bar{c} = c_0$, and so the fraction of agents taking the test is fixed at $F(0)$ independent of the test chosen. Hence, in this case, the test designed by the

²⁵ We have assumed so far that the distribution of costs has full support. However, this case of binary costs maps into the model in the first part of the paper if we interpret the agents with high costs as those that have nothing to disclose.

monopolist would be the same as if the agent were choosing: the monopolist would choose the test that minimizes the nondisclosure price π^N .²⁶

Moreover, conditional on a given participation threshold \bar{c} , all of our earlier results continue to apply. For example, the monopolist also prefers more reliable and more informative tests, as they would lead to a lower π^N , increasing the gain from the test and allowing him to charge a higher price. The monopolist is also satisfied to restrict attention to simple certification tests, as his revenues depend only on π^N . Finally, given any certification test chosen by the monopolist, a social planner would prefer a higher standard, *if* the planner could also prevent the monopolist from increasing his price. If the monopolist can adjust his price, then in response to the higher standard, he will raise the price and restrict participation, which will have a potentially offsetting effect on informativeness of equilibrium disclosure.

When the distribution of costs is continuous, so that $Pr(c \leq \bar{c})$ increases strictly and continuously with \bar{c} , then the monopolist would induce less testing than the equilibrium with the highest testing threshold, \bar{c}_E , when agents choose the test on their own. The intuition is that in the equilibrium with test value \bar{c}_E the monopolist's fee would be $m = 0$. As the monopolist in equilibrium charges positive prices, $\bar{c}_M \leq \bar{c}_E$ and fewer agents will find it optimal to take the test.

Third, in case of simple certification tests, the monopolist will choose a higher threshold for passing than the agents choose in equilibrium. That is because $\pi^N(\bar{c})$ is decreasing in \bar{c} , and therefore $\pi^N(\bar{c}_M) \geq \pi^N(\bar{c}_E)$. By the minimum principle, the passing threshold is equal to the nondisclosure price, and thus the monopolist sets a higher passing threshold than the agents would choose.

Figure 6 illustrates the monopolist equilibrium in the same setting we considered in Figure 4 of Section 4. The monopolist charges a fee equal to the difference between the gain from the test $G(\bar{c})$ and the cost \bar{c} of the marginal test-taker. The monopolist maximizes expected revenue by lowering participation (\bar{c}) relative to when the agent chooses, which raises the no-disclosure payoff π^N and lowers the value of the test. Because π^N coincides with the testing threshold for

²⁶ A special case is Lizzeri (1999). In his model there is no cost to certification, and no possibility of a null result. He shows that the monopolist would choose the pass-fail test with the lowest possible threshold (which, with no null result, minimizes the non-disclosure price).

the optimal certification test (see PROPOSITION III), the monopolist raises the certification standard above the level that would be chosen by the agent.




Figure 6: Monopoly Certification Equilibrium. Compared with a competitive environment in which the agent selects the test, the monopolist sets a higher standard ($\pi^N = 46.6\%$ versus 43.9%) and restricts participation compared (30% vs. 47.4%) in order to charge a positive fee m (4.2%). The monopolist earns expected revenue equal to $1.26\% = 30\% \times 4.2\% = 50\% - 48.73\%$.

6. Discussion and Conclusions

We conclude the paper by considering some implications of dropping two assumptions in our main model that could be important in certain applications, namely that the agent may (i) already be privately informed, or (ii) have non-linear payoffs.

6.1. Privately Informed Agent

We have assumed that the agent becomes privately informed only *after* taking the test. However, in some situations, it would be natural to assume that the agent has some private information before choosing the test. For example, we can envision a model with gradual information gathering, in which the agent would perform a preliminary exploration that generated private unverifiable information about the likely results he could get from different kinds of tests, and only then he would choose a test.

To provide some intuition about such situations, we discuss the following example. There are two possible states of the world, $\omega \in \{L, H\}$ (with some prior distribution). The value of the asset is distributed over $[0,1]$ with conditional distribution $F_\omega(v)$, such that $F_H(v)$ is a stronger distribution than $F_L(v)$ in the sense of monotone likelihood ratio property (MLRP) and both distributions have full support. The agent knows privately the state before he chooses a test, so ω is his private type. The agent has access to binary tests with thresholds $x \in [0,1]$ that reveal whether the value of v is above or below x . For simplicity suppose all tests have the same probability $1-p$ of obtaining the null outcome, independent of v or ω . The rest of the game is as before.

The difficulty in analyzing such a model is that upon disclosure of a result (that v is above some x), the market price depends on the market's belief about the agent's type. This game hence has elements of a signaling game. Specifically, we need to specify off-equilibrium beliefs about the agent's type when the agent chooses an out-of-equilibrium test. Instead of providing a comprehensive analysis of equilibria of this game, we discuss only properties of separating equilibria.

An equilibrium is described by two thresholds for the two types, x_E^L, x_E^H , beliefs $\mu(x)$ that assign the probability that the agent is the high type given disclosure of passing a test with difficulty x (the beliefs have to be consistent with equilibrium strategy on the equilibrium path), and prices $\pi^N, \pi^D(x)$ that are consistent with $\mu(x)$ and the two thresholds. In a separating equilibrium $x_E^L < x_E^H$ and $\mu(x_E^L) = 0, \mu(x_E^H) = 1$, so that the tests used on the equilibrium path reveal the type of the agent. Given any cutoffs (x^L, x^H) , the nondisclosure price is given by:

$$\pi^N(x^L, x^H) \equiv E[v \mid S = \emptyset, \text{ or } \omega = L \text{ and } v \leq x^L, \text{ or } \omega = H \text{ and } v \leq x^H].$$

So the equilibrium nondisclosure price is: $\pi_E^N = \pi^N(x_E^L, x_E^H)$. In an equilibrium, the expected payoff of a type $i \in \{L, H\}$ when he picks test x is given by:

$$F_i(x)\pi^N + (1 - F_i(x))E[v \mid v > x, \mu(x)].$$

In a separating equilibrium, when type i picks x_E^j this equals:

$$F_i(x_E^j)\pi^N + (1 - F_i(x_E^j))E[v | v > x_E^j, j].$$

We first argue that in a separating equilibrium, the equilibrium choice of the low type minimizes the nondisclosure price given the equilibrium strategy of the high type: $x_E^L = \pi_E^N = \min_{x^L} \pi^N(x^L, x_E^H)$. Suppose by way of contradiction that, $x_E^L \neq \pi_E^N$ and consider the deviation to $x = \pi_E^N$. Such a deviation would be profitable because:

$$\begin{aligned} F_L(\pi_E^N)\pi_E^N + (1 - F_L(\pi_E^N))E[v | v > \pi_E^N, \mu] &\geq \\ F_L(\pi_E^N)\pi_E^N + (1 - F_L(\pi_E^N))E[v | v > \pi_E^N, L] &\geq \\ F_L(x)\pi_E^N + (1 - F_L(x))E[v | v > x, L] &\text{ for any } x \end{aligned}$$

The first inequality follows from the MLRP ranking of the high and low conditional distributions. The second inequality follows from our previous reasoning that for a fixed nondisclosure price the best response of the agent is to disclose results that yield prices above that nondisclosure price and hide all other results. Formally, we can re-write the last equation as:

$$\pi_E^N + (1 - F_L(x))E[v - \pi_E^N | v > x, L] = \pi_E^N + E\left[(v - \pi_E^N)1_{v>x} | L\right],$$

which is maximized by setting $x = \pi_E^N$. By the same reasoning as in the proof of the minimum principle, $x_E^L = \pi_E^N$ implies $\pi_E^N = \min_{x^L} \pi^N(x^L, x_E^H)$.

How about the high type? In any separating equilibrium that satisfies standard refinements of off-path beliefs, the high type chooses a threshold that minimizes $\pi^N(x_E^L, x^H)$ subject to satisfying the incentive compatibility of the low type (i.e., the smallest threshold above the minimum of $\pi^N(x_E^L, x^H)$ such that the low type prefers to separate choosing x_E^L instead of mimicking the high type).

6.2. Non-Linear Agent Payoffs

Another assumption we maintained in the paper is that the agent's payoffs are linear in the market price. In some situations, when the agent owns the asset he is trying to sell, linearity may be a good approximation of his incentives. However, in other situations the agent's payoffs may not be linear in the price, for example, because he is selling the asset on behalf of somebody else and has a

contract that rewards him as a function of the price he receives. To illustrate what happens in that case, we focus attention on certification tests.

PROPOSITION VIII (NON-LINEAR PAYOFFS). Suppose that the agent has access to a set of certification tests S_x with any threshold x . Let x_E be the equilibrium threshold when the agent's payoff is linear in the market prices, and assume x_E is in the interior of the support of $E[v|S]$. Then,

- i) If the agent's payoff is **concave** in market prices, the equilibrium threshold x^* is smaller than in the linear case: $x^* < x_E$
- ii) If the agent's payoff is **convex** in market prices, the equilibrium threshold x^* is larger than in the linear case: $x^* > x_E$

In either case, the equilibrium nondisclosure price is higher than in the linear case.

Proof: We provide the proof for the case of concave payoffs; the proof for the convex case is analogous and hence omitted.

Suppose the market expects the agent to pick x^* . A deviation to some threshold x induces a lottery with two outcomes: $\pi^N(x^*)$ and $\pi^D(x)$.

We first argue by way of contradiction that it cannot be that $x^* > x_E$. Consider the deviation to x_E . This would induce a lottery with outcomes $\pi^N(x^*)$ and $\pi^D(x_E)$. Since $\pi^N(x^*) > \pi^N(x_E) = x_E$, this lottery would lead to a higher expected price (as the agent earns a higher price than he deserves for nondisclosure). Moreover, because the low reward is the same, $\pi^N(x^*)$, and the high reward is lower upon deviation, $\pi^D(x_E) < \pi^D(x^*)$, the agent's risk is also reduced. Hence, if the agent's payoff is concave in prices, the agent would strictly gain from this deviation.²⁷

Next, assume by way of contradiction that $x^* = x_E$. Consider a deviation to a slightly lower threshold, $x_E - \varepsilon$. By the same reasoning as in the previous case, this deviation also reduces agent's

²⁷ Consider two lotteries with outcomes $\{A,B\}$ and $\{A,C\}$, such that $A < B < C$ and the first lottery has weakly higher mean. Then the first lottery second-order stochastically dominates the first.

risk but the mean effect is zero (second order in ε) because $\pi^N(x_E) = 0$. Thus there exists $\varepsilon > 0$ so that the agent would find this deviation to be profitable, a contradiction. ♦

A special case of payoff convex in prices is if the agent has an outside option to sell the asset at a fixed price. In that case, it is straightforward to show that if the outside option is lower than π_E^N from the linear case, it would have no impact on the equilibrium. However, if it were higher, the equilibrium threshold would be increased to that outside option (increasing equilibrium nondisclosure price but not as much).

6.3. Conclusions

In this paper, we have analyzed an agent who makes strategic decisions about what information to collect and what results to disclose, assuming that negative results can be hidden from the market while successful results can be credibly revealed. Our model allows us to understand better decisions made by individual agents as well as by independent organizations that provide quality certification within an industry.

We showed that the equilibrium choice can be characterized by the minimum principle: equilibrium tests minimize the nondisclosure price over all possible joint choices of test and disclosure policy. As a result, market forces push sellers to choose tests that are more likely to provide verifiable information regarding the left tail of the value distribution. Because information regarding the right tail is not valuable in equilibrium, in the extreme, agents will be driven toward simple pass-fail certification tests, especially if they are cheaper or easier to perform than more refined tests.

When information plays an additional allocative role, we have shown that the certification threshold chosen in equilibrium is informationally inefficient. The passing threshold is too low in the sense that the planner would benefit from imposing a minimum certification standard. In a model with endogenous costly participation in the acquisition of information, such minimum requirements would both reduce the costs of test-taking and improve informativeness of equilibrium disclosure. We finally discussed how the analysis can be extended to monopoly certification and how the methods can be used even when we start relaxing some of our assumptions, allowing the agent to be privately informed before choosing the test or have payoffs which are non-linear in prices.

Many questions remain open. In particular, a natural question is what would happen if the agent could take multiple tests. If the agent takes two tests and decides to disclose the results of only one of them, investors would form beliefs regarding the results of the other test. Therefore both nondisclosure and partial-disclosure prices would be equilibrium objects that satisfy a fixed point. Another natural direction would be to study tests with heterogeneous costs, so that agents face a tradeoff between the test's cost and its nondisclosure price.

References


- Acharya, Viral V, Peter DeMarzo and Ilan Kremer. (2011). "Endogenous Information Flows and the Clustering of Announcements." *American Economic Review*, 101(7): 2955-79.
- Ben-Porath, Elchanan, Eddie Dekel and Barton L. Lipman (2017) "Disclosure and Choice." Forthcoming in the *Review of Economic Studies*.
- Dye, Ronald A. (1985). "Disclosure of Nonproprietary Information." *Journal of Accounting Research*, 23(1): 123–45.
- Einhorn, Eti, and Amir Ziv. (2008). "Intertemporal Dynamics of Corporate Voluntary Disclosures." *Journal of Accounting Research* 46(3): 567–589.
- Friedman, Henry, John Hughes, and Richard Saouma (2016). "Implications of biased reporting: conservative and liberal accounting policies in oligopolies." *Review of Accounting Studies* 21, 251–279.
- Glazer, Jacob and Rubinstein, Ariel (2006). "A Study in the Pragmatics of Persuasion: A Game Theoretical Approach," *Theoretical Economics* 1, 395–410.
- Grossman, S. (1981). "The Informational Role of Warranties and Private Disclosure about Product Quality," *Journal of Law and Economics*, Vol. 24, pp. 461-489.
- Guttman, Ilan, Ilan Kremer and Andrzej Skrzypacz. (2014). "Not Only What but also When - A Theory of Dynamic Voluntary Disclosure." *American Economic Review* 104(8): 2400-2420.
- Hart, Sergiu, Ilan Kremer and Motty Perry. (2017). "Evidence Games: Truth and Commitment." *American Economic Review*, 107(3): 690-713.

- Jin, Ginger (2005). "Competition and Disclosure Incentives: An Empirical Study of HMOs," *RAND Journal of Economics*, 26, 93-113.
- Jovanovic, Boyan. (1982). "Truthful Disclosure of Information." *Bell Journal of Economics* 13(1): 36-44.
- Jung, Woon-Oh, and Young K. Kwon. (1988). "Disclosure When the Market Is Unsure of Information Endowment of Managers." *Journal of Accounting Research*, 26(1): 146–53.
- Lizzeri, Alessandro. (1999). "Information Revelation and Certification Intermediaries." *The RAND Journal of Economics* 30, no. 2: 214-31.
- Marinovic, Ivan, Andrzej Skrzypacz and Felipe Varas. (2017) "Dynamic Certification and Reputation for Quality" forthcoming in *American Economic Journal: Microeconomics*.
- Milgrom, P.R. (1981). "Good News and Bad News: Representation Theorems and Applications," *The Bell Journal of Economics*, Vol. 12, pp. 380-391.
- Pae, Suil. (2005) "Selective Disclosures in the Presence of Uncertainty About Information Endowment." *Journal of Accounting and Economics* 39(3): 383–409.
- Sher, I. and Vohra, R. (2015), "Price discrimination through communication." *Theoretical Economics*, 10: 597–648.
- Verrecchia, Robert E. (1983). "Discretionary Disclosure." *Journal of Accounting and Economics* 5 (3): 179–94.

Appendix

PROOF OF EXISTENCE OF SOCIAL OPTIMUM (SECTION 3.1): For $x < x_E$, $\pi^N(x)$ is decreasing in x (by quasiconvexity). Because $\pi^D(x)$ is increasing in x , the variance of $E[v|S_x]$ is increasing, and hence for $x < x_E$, $L(x)$ is decreasing, in x .²⁸ Because S is informative, $\pi^N(x_E) < \bar{v} < \pi^D(x_E)$ and so $L(x_E) < E[(v - \bar{v})^2]$. As $x \rightarrow \infty$, S_x becomes uninformative, and $L(x) \rightarrow E[(v - \bar{v})^2] > L(x_E)$. Hence there exists a finite upper bound B such that the infimum for L occurs in the compact range $[x_E, B]$. If L were continuous, the existence of a social optimum would thus be automatic.

Given the definition of S_x , the functions π^N , π^D and L are right-continuous with left limits. Discontinuities can only occur at atoms of the distribution of $E[v|S]$. Therefore, to establish that the social optimum $\inf_x L(x)$ is attained, we must rule out the possibility that $L(x)$ strictly decreases as we approach x^* from the left, but then jumps up at x^* :



For any function f define $f(x^{*-}) \equiv \lim_{x \uparrow x^*} f(x)$, the left limit at x^* . Next define

$$\hat{L}(x) \equiv E\left[S_x(v - \pi^D(x^{*-}))^2 + (1 - S_x)(v - \pi^N(x^{*-}))^2\right].$$

Note that $\hat{L}(x) \geq L(x)$ and $\hat{L}(x^{*-}) = L(x^{*-})$ (because in $L(x)$ prices adjust to conditional expectations while in $\hat{L}(x)$ they are kept constant at their left limits at x^*). Therefore $\hat{L}(x)$ is

²⁸ To verify this claim, recall that since the distribution of $E[v|S_x]$ is binary, its variance given by $(\pi^D(x) - \bar{v})(\bar{v} - \pi^N(x))$, and recall that minimizing $L(x)$ is equivalent to maximizing the variance of $E[v|S_x]$.

strictly decreasing as we approach x^* from the left. Because an increase in x assigns types with expected value x to the nondisclosure price, $\hat{L}(x)$ strictly decreasing implies x is closer to $\pi^N(x^{*-})$ than $\pi^D(x^{*-})$; that is, $x < \bar{\pi}(x^{*-})$. But then $x^* \leq \bar{\pi}(x^{*-})$, and so by the same logic, $\hat{L}(x^*) \leq \hat{L}(x^{*-})$. Thus, $L(x^*) \leq \hat{L}(x^*) \leq \hat{L}(x^{*-}) = L(x^{*-})$. In other words, if L is strictly decreasing to the left of x^* , then any jump at x^* cannot be upward. ♦

PROOF OF PROPOSITION V: We begin by arguing that these conditions are necessary. Using condition (i) of an equilibrium, the agent's gain from taking test k with disclosure policy θ is

$$\begin{aligned} E\left[(1-\theta(S_k))\pi_E^D(k, S_k) + \theta(S_k)\pi_E^N\right] - \pi_E^N &= E\left[(1-\theta(S_k))(\pi_E^D(k, S_k) - \pi_E^N)\right] \\ &= E\left[(1-\theta(S_k))(v - \pi_E^N)\right]. \end{aligned} \quad (22)$$

Where the second equality follows from the law of iterated expectations.

For any (k, θ, \bar{c}) :

$$\bar{v} - \pi_E^N = E\left[\left((1-\theta(S_k))1_{c \leq \bar{c}} + \theta(S_k)1_{c \leq \bar{c}} + 1_{c > \bar{c}}\right)(v - \pi_E^N)\right],$$

and therefore the agent's expected gain in (22) can be rewritten as

$$\frac{\bar{v} - \pi_E^N + E\left[\pi_E^N - v \mid \theta(S_k) = 1 \text{ or } c > \bar{c}\right] \Pr(\theta(S_k) = 1 \text{ or } c > \bar{c})}{\Pr(c \leq \bar{c})}. \quad (23)$$

Condition (iii) of the equilibrium definition is thus equivalent to maximizing

$$\max_{(k, \theta)} \left(\pi_E^N - E\left[v \mid \theta(S_k) = 1 \text{ or } c > \bar{c}\right] \right) \Pr(\theta(S_k) = 1 \text{ or } c > \bar{c}). \quad (24)$$

When $\bar{c} = \bar{c}_E$, equilibrium condition (ii) implies (24) must be zero, and therefore (13) must hold (otherwise, there would be a profitable deviation from the supposed equilibrium). Similarly, setting the final term in the numerator in (23) to zero implies that (14) must hold.

This construction also implies that these conditions are sufficient for an equilibrium. Given the equilibrium prices, once an agent decides to take a test, choosing the test and disclosure policy that satisfy (13) is a best response. Moreover, (14) implies that types with costs lower than \bar{c}_E find it optimal to take a test and higher types do not. To show existence, note first that

$$E[v | \theta(S_k) = 1 \text{ or } c > \bar{c}] = \frac{\Pr(c > \bar{c})\bar{v} + \Pr(c \leq \bar{c})\Pr(\theta(S_k) = 1)E[v | \theta(S_k) = 1]}{\Pr(c > \bar{c}) + \Pr(c \leq \bar{c})\Pr(\theta(S_k) = 1)} \quad (25)$$

is continuous in \bar{c} and that

$$\frac{\partial}{\partial \bar{c}} E[v | \theta(S_k) = 1 \text{ or } c > \bar{c}] \propto g(\bar{c})(E[v | \theta(S_k) = 1] - \bar{v}) \quad (26)$$

where g is the density of c . Then, because $\pi^N(0) = E[v | \theta(S_k) = 1 \text{ or } c > 0] = \bar{v}$, we apply l'Hôpital's rule to show that

$$\lim_{\bar{c} \rightarrow 0} \frac{\bar{v} - \pi^N(\bar{c})}{\Pr(c \leq \bar{c})} = \frac{-\pi^{N'}(0)}{g(0)} > 0.$$

In addition,

$$\lim_{\bar{c} \rightarrow \infty} \frac{\bar{v} - \pi^N(\bar{c})}{\Pr(c \leq \bar{c})} = \bar{v} - \min_{(k, \theta)} E[v | \theta(S_k) = 1] < \bar{v}.$$

Together these imply that the right-hand side (14) must cross the 45-degree line from above, establishing the existence of a fixed point \bar{c}_E . ♦

PROOF OF PROPOSITION VI: We start with showing that there exists a test with a threshold $x_{MS} > x_E$ such that if all agents use that test, the resulting equilibrium would be both more informative and have less testing than under the original equilibrium threshold x_E .

Let $\pi^N(x, \bar{c}) \equiv E[v | S_x = 0 \text{ or } c > \bar{c}]$ be the nondisclosure price when the test is x and the testing threshold is \bar{c} . It is continuous, decreasing in \bar{c} , and at $\bar{c} = \bar{c}_E$ it is minimized and at $x = x_E$ with $\frac{\partial}{\partial x} \pi^N(x_E, \bar{c}_E) = 0$.

Consider imposing a restriction of tests to a test with threshold $x_{MS} > x_E$. The new testing threshold, \bar{c}_{MS} is a solution to the fixed-point problem:

$$\bar{c}_{MS} = \frac{\bar{v} - \pi^N(x_{MS}, \bar{c}_{MS})}{\Pr(c \leq \bar{c}_{MS})}.$$

By our earlier existence argument (in the proof of **PROPOSITION V**), there exists a fixed point with $\bar{c}_{MS} < \bar{c}_E$ because at x_{MS} and $\bar{c} = \bar{c}_E$ the right hand side of the above expression is lower than \bar{c}_E (we are using here the assumption that x_E corresponds to the equilibrium with the lowest testing threshold \bar{c}_E , which is also the equilibrium with the largest equilibrium pass-fail threshold). Therefore, if the testing standard is raised to x_{MS} there exists an equilibrium with a lower cost of certification. Moreover, by the implicit function theorem, because $\frac{\partial}{\partial x} \pi^N(x_E, \bar{c}_E) = 0$, generically we have that $\frac{\partial \bar{c}_{MS}}{\partial x_{MS}} = 0$ at $x_{MS} = x_E$.²⁹

By the reasoning in **PROPOSITION IV**, increasing x_{MS} above x_E leads to a direct improvement of informativeness for a fixed \bar{c} . But by the argument above, there is also an indirect effect that a higher x_{MS} implies less certification and hence less information. To see which effect dominates, write the planner's objective as

$$L(x, \bar{c}) = E \left[1_{c \leq \bar{c}} S_x (v - \pi^D(x))^2 + (1 - 1_{c \leq \bar{c}}) S_x (v - \pi^N(x, \bar{c}))^2 \right]$$

It is continuous in both variables and $\frac{\partial}{\partial x} L(x_E, \bar{c}_E) > 0$ (as we argued in **PROPOSITION IV**). A small increase of the test threshold implies a first-order direct improvement of the planner's objective, but because $\frac{\partial \bar{c}_{MS}}{\partial x_{MS}} = 0$ at $x_{MS} = x_E$, it has only a second order (negative) effect of lesser participation. Therefore, it is possible to find a test such that in equilibrium there is both less testing and more information.

Finally, we claim that even if the agents could use any test with $x \geq x_{MS}$, the unique equilibrium is the one in which they choose x_{MS} and the threshold is \bar{c}_{MS} . If there were another equilibrium with threshold $\hat{x}_E > x_{MS} > x_E$, then by the quasi-convexity of $\pi^N(x, \bar{c}_E)$, it would also be an equilibrium of the original game with no restrictions on the set of certification tests the agents can choose, contradicting the fact that x_E was the largest unconstrained equilibrium. ♦

²⁹ We are using here our assumption that $E[v | S]$ is continuously distributed with full support, so that small changes in the threshold change the probabilities and conditional expectations continuously.