Hard information design*

Sulagna Dasgupta Ilia Krasikov Rohit Lamba

July 2022

Abstract

Many transactions in the marketplace rely on hard (or verifiable) information about the underlying value of the intended exchange, typically through certification—housing, diamonds, bonds being cases in point. What is the class of Pareto efficient certifications for such scenarios? This paper studies the canonical monopolistic screening problem, and models certification as hard information produced through a test to be flexibly chosen pre-trade. It argues that Pareto efficient tests take a simple form—they produce certification with a partitional structure, often with one or two thresholds. This claim is shown to be true for both the linear trading model and the non-linear pricing model.

1 Introduction

Certification of the value or quality of products pre-trade is ubiquitous. We see this in housing markets, market for jewellery, bond markets, and labor markets. Certification is typically verifiable, its outcome is hard information that cannot be fudged. And it often employs simple tests that result in a coarse set of attainable grades and corresponding prices. The objective of this paper is to model *hard information design* in the context of a standard trading problem, i.e. monopolistic screening, and provide a foundation for such simplicity.

To that end, we study the following stylistic model: A buyer and a seller start with a common prior on the buyer's value from trade, which will often be referred to as the

*Dasgupta: University of Chicago, sulagna@uchicago.edu; Krasikov: Higher School of Economics University Moscow, krasikovis.main@gmail.com; Lamba: Pennsylvania State University, rlamba@psu.edu. We are grateful to Nageeb Ali, Yeon-Koo Che, Henrique de Oliveira, Nima Hagpanah, Navin Kartik, Vijay Krishna, Stephen Morris, Daniel Quigley, Ron Siegel, Denis Shishkin, Philipp Strack, Roland Strausz, Balázs Szentes, Rakesh Vohra and seminar participants at Columbia University and Pennsylvania State University for comments.

buyer's type. The seller's cost of production is common knowledge. The buyer chooses a statistical test, the outcome of which is privately observed.[1] The test produces hard information in the sense of Dye [1985]— with some probability it fails to produce a certificate, and with complementary probability it generates a certificate with a verifiable score. The probability of failure and the choice of scores can be arbitrarily related with the underlying value of the buyer. An outcome has two components—whether a certificate is generated, and the score printed on it, in case the certificate is generated. The seller observes the chosen test but not the outcome, and commits to a pricing mechanism as a function of the information presented by the buyer. If test does indeed generate a certificate, the buyer decides whether or not to present it to the seller. Finally, prices and allocations are implemented as a function of what the buyer presents, and payoffs are realized.

The problem has two connected components— the information production part and the price discrimination part. The main result of the paper is that the information production part is partitional, typically with one or two thresholds, which divide the valuation space into two or three regions, respectively. In the first case, the optimal test produces a "high" score or no score. All buyer types above the single threshold get a high score and all buyer types below the threshold do not get a certificate. In the second case, the optimal test produces a "high" score, "low" score or no score. All buyer types in the top most region get a high score, in the bottom most region get a low score and those in the middle get no certificate.

In the second part of the problem, the seller responds to the chosen test through a set of score contingent prices. Compute the expected value of the buyer types who do not get a certificate and call this the base price. When a test with a single threshold is chosen, the seller optimally sets the score contingent price for the top region to be the same as the price of trade for no certificate. This culminates essentially in a posted price which is set equal to the base price. By construction, this extracts full surplus from the buyer types in the bottom region with no certificate, since the posted price exactly equals their expected valuation. It also provides (information) rent to those in the top region with the high score, since their expected valuation is strictly higher than the posted price.

When a test with two thresholds is chosen, the seller still offers the base price, and in addition, may offer another price with a discount. Buyer types in the top and middle region trade at the base price, and those in the bottom region are targeted with the discounted price, which is set exactly equal to the expected valuation of the bottom region. The seller, of course, trades with the bottom region only if the discounted price is greater

[1] Later we will vary the bargaining power of the two parties on the choice of test, thereby mapping the Pareto frontier.

than the cost of production. Full surplus is extracted from the types in the middle and bottom regions, and rent is paid to those in the top region.

The intuition for the result is as follows. Since information is verifiable, if the buyer shows the certificate, the both parties have a common posterior of the buyer's valuation, and the seller sets a price exactly equal to the posterior mean extracting full surplus. Thus, the measure of values for which the test proves inconclusive (i.e. produces no certificate) must be positive; otherwise the seller extracts the entire surplus.

In addition, note that the seller does not have to trade with buyer types who do not show a certificate. So, the buyer must also choose a test in response to which the seller finds it profitable to trade with buyers who do not have a certificate at what we called the base price. Otherwise, again, only buyers with a certificate trade and full surplus is extracted by the seller. Call the enticing of the seller to trade with buyers who do not produce a certificate the seller's constraint.

Now, assume the cost of the seller is zero, so to trade is always efficient. At the optimum, trade will indeed be efficient, i.e. happen with certainty. For starters, ask: for any arbitrary test, what is the pattern of trade? All buyer types whose (expected) value exceeds the base price will not show the certificate, instead trade at the base price, and gain some rent. Those with scores that imply a posterior mean lower than the base price will be offered a discounted price that makes them exactly indifferent, i.e. binds their participation constraint, so they trade too. At the optimum, the value of the base price is then pinned down by binding the seller's constraint: the seller's profit from trading with types that do not have a certificate equals the rent she gives up to the buyers whose value exceeds the base price. Thus, the seller and all buyer types want to trade, hence trade is efficient.

The best way for the buyer to generate efficient trade while maximizing his payoff and respecting the seller's constraint is to select a one threshold test. Since all buyer types below the base price get full surplus extracted by the seller, the buyer can instead combine these types in the region with no-certificate. This helps the buyer by reducing the base price, so that the high types get the highest possible information rent. The seller's constraint then determines the single threshold that partitions the type space.

What score should be printed on the certificate for the types in the top region? The exact value is irrelevant, because the buyer always chooses the base price, since it can be claimed that the test was inconclusive. So coarsely reporting that the values lie in this top region is sufficient. We call this the "high score". Thus effectively, the buyer endogenously turns the problem into a two-types screening problem where the individual rationality of the "low type" and the incentive constraint of "high type" bind.

Finally, suppose that the seller's cost is positive. Then trade with buyers whose value is lower than the seller's cost is inefficient. If we continue with the single threshold test, then at some high enough cost, the seller's constraint cannot be satisfied because the measure of inefficient trades is now too large and the seller will choose not to trade with buyers in the bottom region, that is those who do not show a certificate. This leads to the creation of a third bottom most region where a certificate is issued, which is meant preclude (at least some) inefficient trades.

Thus, for a low enough cost, the optimal information structure partitions the valuation space into two regions and the seller simply posts one price. For a high enough cost, the optimal information partitions the valuation space into three regions, where the top two regions trade at the posted price and the bottom region may be offered discounts if trade with those types is profitable for the seller.[2]

Our model departs from both the traditional monopolistic screening problem (Mas-Colell, Whinston, and Green [1995] Chapters 13-14) and the more recently explored information design problem (Bergemann and Morris [2019]). In the former, the initial allocation of information is fixed and also *soft*, in that the buyer can report any information to the seller. In the latter, the extent of information asymmetry before physical exchange is chosen as part of the design problem. The information though is still soft, released to the agent in the form of a signal, which can, again, be misreported. In a closely related paper, Roesler and Szentes [2017] study the monopolistic screening problem with soft information design. They find that the buyer optimally designs an information structure that generates a truncated Pareto distribution of posterior means, in response to which a seller chooses an optimal posted price.

To varying degrees, the question of hard information design has been explored in the context of disclosure games (eg. recently in Ben-Porath, Dekel, and Lipman [2018], DeMarzo, Kremer, and Skrzypacz [2019], Shishkin [2022]). However, we know of no study that situates the endogenous generation of such hard information in a standard price discrimination or mechanism design framework. Or any study that allows the choice of hard information to be fully flexible. These forms our starting point. We bring together two canonical models—monopolistic screening and Dye-hard evidence, using tools from

[2]To be sure, the optimal information structure is not unique. The paritional nature of it divides the type space into three regions. The middle region does not produce any certificate. Then the Blackwell most informative optimal information structure prints on the certificate the exact value of the buyer in the top and bottom regions. Correspondingly the seller offers the base price in the middle region and a price exactly equal to the score printed on the certificates in the top and bottom regions. Any independent garbling of the top and bottom regions is also optimal. The particular optimal information structure we emphasize is the Blackwell least informative–it prints on the certificate only the information whether the value is the top or bottom region.

information design.

Two features of our model lead to the simplicity of the optimal information structure— the hardness of information production and the seller's ability to commit to a pricing mechanism as a function of the test. Relaxing either takes the model back to the soft information benchmark, wherein the optimal information structure constitutes a distinct signal distribution that targets a truncated Pareto distribution of posterior means (as in Roesler and Szentes [2017]). Intuitively, the hardness of information simplifies the buyer's underlying incentive problem of what can be misreported to the seller, and the commitment power makes the seller's prices - which are tagged to verifiable information - credible form the buyer's perspective.

Broadly, in comparison to the soft information benchmark, the hardness of information hurts the buyer and benefits the seller in an *ex-ante* sense, even though the buyer is allowed to choose any feasible test. In the mechanism design lexicon, whatever gain accrues to the buyer from choosing the optimal test is offset by the fact that instead of a continuum of incentive constraints, the seller now has to pay him the shadow price of only one, i.e. whether the buyer produces the certificate with the test score or not.

So far we have talked about the buyer-optimal information and associated pricing mechanism. What if due to protocols outside the purview of the model (e.g. for regulatory reasons), the seller can influence the chosen test? To accommodate such a line of inquiry, we characterize the Pareto frontier of the test design problem. The finding here too is that, any pair of buyer and seller payoffs on the Pareto frontier can be generated by a partitional informational structure with few thresholds, typically two, where scores are produced for types in the top and bottom regions, and the middle region does not produce any certificate. The structure of the thresholds moves with the extent of bargaining power of the seller. So does the total value of the surplus and its split between the buyer and the seller.

The Pareto frontier has two parts— a linear part and a potentially concave part. In the linear part, the total value of the surplus is at the efficient level. If the cost of the seller is zero, trade is efficient and the Pareto frontier is simply a straight line. If the seller has most of the bargaining power in the choice of the test, the efficient surplus is again attained irrespective of the cost, and most of it is extracted by the seller through prices. In fact, the bottom region now produces two scores, one below and one above the marginal cost. This ensures the seller knows whether trades are above or below the cost, so the ones below can be precluded.

When the cost is high enough and buyer has most of the bargaining power in the choice of the test, a trade-off emerges between revealing enough information about valuations below the marginal cost so that trade is efficient, and hiding enough information about the buyer's value from the seller. This trade-off gives curvature to the Pareto frontier even when preferences are linear, and culminates in a total surplus below the efficient level.

Finally, we consider a general model that subsumes within it the linear model discussed so far, and the non-linear pricing problem in the spirit of Mussa and Rosen [1978] which features the intensive margin of interior supply by the seller. It also accommodates arbitrary prior distributions - allowing for discrete, continuous or mixed supports - and costs of certification for the buyer. We completely characterize the class of information structures that achieve payoffs on the Pareto frontier for this general model. These maximally partition the buyer's type space with four thresholds or into five regions. The extra regions can potentially be used when the seller has enough bargaining power to destroy some information rent for the buyer while increasing or in the least not shrinking the total surplus.

With some more structure on the problem, the optimal information structure can be more precisely characterized. To that end, we study the uniform linear-quadratic model, where the prior is uniform, preferences are quasi-linear in transfers, and the seller's cost is quadratic. The class of tests that characterizes the Pareto frontier is two-threshold. However, the added curvature ensures that the test implementing a specific point on the Pareto frontier is now unique. It necessarily reveals the exact value in the bottom and top regions and does not produce any certificate in the middle region. So, the seller supplies the efficient quantity in the bottom and top regions, and supplies the associated "virtual value", in the Myersonian sense, in the middle region. As before, total expected surplus is extracted by the seller in the middle and bottom regions, but information rent is provided to the buyer in the top region.

In summary, we introduce a model of screening with endogenous hard information and show that the universe of tests on the Pareto frontier is simple. The reader can view this as a conceptual counterpart to the soft information design benchmark. Both models are perhaps an extreme version of the underlying reality of information production in various transactions which involves some combination of soft and hard information. We view our analysis as an invitation to further explore the space in between, and we believe the monopolistic screening problem is a reasonable framework for this exercise.

Related Literature. Asymmetric information is central to our understanding of markets or lack thereof. Hayek [1945]'s foundational framing and Stigler [1961]'s call to greater exploration were arguably turning points in the explosion of a field of inquiry oft referred to as information economics. One of the workhorse tools in information economics is the classical monopolistic screening model (Mas-Colell, Whinston, and Green [1995] Chapters 13-14). The goal of this paper is to add to this model the feature of pre-trade production of hard information using statistical tests.

As in the simplest models of competitive equilibrium where the initial endowments of goods drives prices (Arrow and Debreu [1954], and Varian [2020] Chapter 16), in information economics the initial endowments of information drives the extent of inefficiency in markets (Akerlof [1970], Spence [1973], and Rothschild and Stiglitz [1976]). In the mechanism design literature, the initial endowment of information is commonly held fixed (Myerson [1981]). In the rapidly growing field of information design the extent of information asymmetry before physical exchange is intermediate and chosen as part of the design problem (Kamenica [2019], Bergemann and Morris [2019], and Mathevet, Perego, and Taneva [2020]). The information though is still soft, released to the agent(s) as a signal, which is to be revealed truthfully under incentive and/or obedience constraints.

Several recent papers have explored distinct aspects of soft information design in price or contract theoretic settings.[3] As discussed, in this family of studies, Roesler and Szentes [2017] is closest to our work. Also related is a concurrent paper, Bergemann, Heumann, and Morris [2022], that studies a non-linear pricing problem á la Mussa and Rosen [1978] (as in Section 8 here), where, in addition to the pricing mechanism, the seller can also design soft information for the buyer without observing it directly. It finds sufficient conditions under which a coarse partition of the buyer's type space is optimal; though the finer structure of information revelation and subsequent pricing contracts are quite different from ours.

In modeling hard information, we follow the classical approach à la Dye [1985] in that verifiable information is produced with some probability bounded away form one. This is a commonly invoked methodology that allows us to sidestep the unraveling benchmark of Grossman [1981] and Milgrom [1981]. In addition, this is also practically relevant in that it provides the agent some property rights over their own information, as often discussed in the literature on the economic of privacy (Acquisti, Taylor, and Wagman [2016]).

The role of hard information in the context of monopolistic screening has been explored most notably by Sher and Vohra [2015]. It assumes the buyer's valuation to be privately known but allows him access to a rich evidence structure to present to the seller. It then characterizes sufficient conditions on evidence structures under which simple mechanisms are optimal. The key difference between that model and ours is the endogenous

[3] See, for example, Bergemann, Brooks, and Morris [2015], Li and Shi [2017], Smolin [2020], Segura-Rodriguez [2021], Halac, Lipnowski, and Rappoport [2021], Wei and Green [2022], Ravid, Roesler, and Szentes [2022], Yang [2022], and Haghpanah and Siegel [2022] amongst others.

generation of information and its subsequent connection with pricing, which is the focus of this paper.

In recent work, verifiable information has been modelled through Dye hard evidence in a disclosure game by Ben-Porath et al. [2018]. In their model, an agent (and also a challenger) chooses amongst projects with varying risk and the payoff depends on strategically revealing information to an outside observer. Unlike in our framework, the structure of evidence there is exogenous. DeMarzo et al. [2019] also study a disclosure game with an exogenous set of tests in which sender follows a minimum principle: she chooses a test which minimizes the asset's value conditional on non-disclosure. This also culminates is a simple pass-fail test being optimal in many situations. Shishkin [2022] partially endogenizes the evidence structure in the context of cheap talk style sender-receiver games. In the language of our paper, the probability of no certificate (i.e an inconclusive test) is exogenous and the main result there is that if this probability is below some threshold the sender chooses a pass-fail test. In all of these papers, a fully flexible choice of test would make the problem uninteresting, yielding the unraveling result in some form or another.[4]

All of these studies are complementary to ours, in that they model hard information in the spirit of Dye and establish an underlying simplicity to information disclosure. Though, to the best of our understanding, the full flexibility of the information structure, and its interaction with a price discrimination problem is new here. Also, as mentioned earlier, it is both the hardness of information and the seller's commitment power that lends simplicity to the tests in our model. If we took away the seller's commitment, the model reduces to a disclosure game with hard information. However, the optimal information structure is no longer simple in the sense we describe, and goes back to the soft information benchmark of Roesler and Szentes [2017]. Moreover, our proof techniques are quite different, which may be useful in analyzing other models of screening with hard information in contexts such as regulation, insurance, and markets for lemons.

Design of certification by a profit maximizing intermediary has been studied in Lizzeri [1999] and Ali, Haghpanah, Lin, and Siegel [2022], where the former looks at mandatory disclosure and latter at voluntary disclosure. Lizzeri [1999] finds full information disclosure with a high testing fee to be optimal. Ali et al. [2022] study the worst-case equilibrium and find a truncated exponential distribution of posterior means to be optimal.[5]

Modeling hard information or evidence is a classical and still growing area of research

[4] See also Jovanovic [1982], Verrecchia [1983], Matthews and Postlewaite [1985], Jung and Kwon [1988], and Shavell [1994] for early contributions on disclosure games with (exogenous) hard information, which propose ways to break the unraveling result of Grossman [1981] and Milgrom [1981].

[5] See also Harbaugh and Rasmusen [2018] and Kartik, Lee, and Suen [2021] for related models of information disclosure with certification and through intermediation.

in economic theory. See Dekel [2016] for broad overview, Ben-Porath, Dekel, and Lipman [2021] for a general treatment on mechanism design with evidence acquisition and Schweighofer-Kodritsch and Strausz [2022] for a statement of a general revelation principle for evidence games.

2 Model

2.1 Primitives

Consider a risk-neutral seller (she) who can produce an object for which there is a single risk-neutral buyer (he). The buyer's value from trade $\theta \in [0,1]$ is distributed according to a continuous cumulative distribution function F supported on the unit interval. The seller's marginal cost $c \in \left[0, \mathbb{E}[\theta]\right)$ is commonly known.

At the outset the buyer is uninformed about his valuation. He can learn θ by designing a test which generates verifiable or *hard* information. Formally, a *test* is an experiment which draws a signal $s \in [0,1] \cup \{\emptyset\}$ as a function of the underlying true state $\theta \in [0,1]$, according to a right-continuous, non-decreasing function $H(\cdot|\theta) : [0,1] \to [0,1]$ as follows:[6]

$$\mathbb{P}(s \in [0,t]|\theta) = H(t|\theta) \;\; \forall t \in [0,1], \quad \mathbb{P}(s = \emptyset|\theta) = 1 - H(1|\theta).$$

Each signal $s \in [0,1]$ comes in the form of a certificate with the value of s written on it. As a result, it can be verified ex-post by the seller. We shall refer to such signals as *scores* to emphasize their ex-post verifiability. In contrast, the *null signal* $s = \emptyset$ is *soft* in the sense that the buyer can always claim to have it.[7]

Given the test, the seller commits to a *base price* p and a set of *score-contingent prices* $\mathbf{p} := \left(p(s)\right)_{s\in[0,1]}$. For each $s \in [0,1]$, the buyer who wishes to purchase the good for the price $p(s)$ must show a certificate with the respective score. He can always opt for the price p or choose not to trade. The restriction of the space of mechanisms to this form is without loss of generality — we show in the appendix that the seller cannot gain by using more complicated selling procedures.

To sum up, the timing of the game is as follows: 1) the buyer chooses a test and obtains a signal, 2) the seller observes the test chosen by the buyer, but not his signal, and picks a base price and a score-contingent price schedule, 3) the buyer with a score $s \in [0,1]$ can trade for $p(s)$, p or walk away, whereas the buyer with the null signal can either trade for p or not trade at all. As is standard, we assume tie-breaking in favor of the buyer at the

[6] For each $s \in [0,1]$, the function $H(s|\cdot)$ is assumed to be measurable.

[7] Note that the null signal is informative about the buyer's valuation whenever $H(1|\cdot)$ is non-constant.

second stage, i.e., when indifferent among multiple prices, the seller charges the lowest price.

2.2 Interpretation

The model here is admittedly stylized, chosen to provide a conceptual counterpart to both standard monopolistic screening in mechanism design and optimal learning through soft signals pre-trade in information design. In both types of models, there are two levels of negotiations happening in the background. The first is who gets to choose the contract and second, what learning technology is agreed upon. To keep things simple, the bargaining power of setting the trading terms, i.e., prices, typically rests with the seller and the private information on the surplus from trade is held by the buyer. The main departure here is that this private information is neither fixed nor soft—it is learnt by the buyer through a commonly known test whose result can be hidden but not falsified.

Since we eventually map the Pareto frontier of all tests, a simple interpretation is as follows. A rating or certification agency is chosen to generate information on the surplus. In the benchmark case we discuss first, the buyer is given the authority to choose the agency, in other cases the agency is picked as a function of who amongst the seller and buyer has a greater say in the negotiation. After the agency has been chosen, the seller offers a menu of prices, which must always include a base price, at which the buyer can trade by discarding the outcome of the test.

An added interpretation of arming the buyer with the ability to claim the null signal or not produce a certificate is on grounds of privacy. After observing the outcome of the test, the buyer may not want to share the information with the seller or in the least may not want to allow it to be used in the transaction. This is in line with several recent analyses of privacy as being the ability of agents to control, at least a part of, their information.[8]

3 Reduced tests

First, we reformulate the buyer's learning problem. Note that both players' utilities only depend on the posterior mean of θ. Therefore, it is without loss of generality to focus on distributions of posteriors means that can be generated by tests.[9]

[8] For example Acquisti, Taylor, and Wagman [2016] write "Privacy is not the opposite of sharing—rather, it is control over sharing. For the individual, therefore, the potential benefits of strategically sharing certain data while protecting other data are quite apparent."

[9] Gentzkow and Kamenica [2016] term this the Rothschild-Stiglitz approach to information design.

A *reduced test* is a non-decreasing, right-continuous function $L : [0,1] \to [0,1]$ that records cumulative probabilities of posterior means when some certificate was obtained, that is

$$\mathbb{P}\big(\mathbb{E}[\theta|s] \in [0,t], s \neq \emptyset\big) = L(t) \;\; \forall t \in [0,1].$$

Reduced tests contain all payoff-relevant information. For example, under L, the null signal occurs with probability $1-L(1)$; and, if this probability is positive, then the buyer's posterior mean conditional on $s = \emptyset$, say $\theta(L)$, can be derived using the Bayes rule as:

$$\theta(L) := \frac{\mathbb{E}[\theta] - \int_0^1 s dL(s)}{1 - L(1)}. \tag{1}$$

We shall refer to $\theta(L)$ as the *base score*.

Of course, not every reduced test can be obtained from the prior. A reduced test L is said to be *feasible* if there exists a test $\big(H(\cdot|\theta)\big)_{\theta\in[0,1]}$ such that

$$L(s) = \int_0^1 H(s|\theta)dF(\theta) \;\; \forall s \in [0,1], \tag{2}$$

$$\int_0^s t dL(t) = \int_0^1 \theta H(s|\theta)dF(\theta) \;\; \forall s \in [0,1]. \tag{3}$$

Eq. (2) ensures that the distribution of signals under the test $\big(H(\cdot|\theta)\big)_{\theta\in[0,1]}$ is exactly the one implied by L, and Eq. (3) ensures that every score is an unbiased signal, i.e., $s = \mathbb{E}[\theta|s]$ for all s in the support of L.

Simplifying further, a specific family of reduced tests is sufficient for characterizing the whole set of feasible reduced tests. Given two thresholds $0 \leqslant \alpha \leqslant \beta \leqslant 1$, define a *dual threshold reduced test* $L_{\alpha,\beta}$ as follows:

$$L_{\alpha,\beta}(s) := \begin{cases} F(s) & \text{if} \;\; F(s) < \alpha, \\ \alpha & \text{if} \;\; \alpha \leqslant F(s) < \beta, \\ F(s) + \alpha - \beta & \text{if} \;\; F(s) \geqslant \beta. \end{cases} \tag{4}$$

Remark 1. *Every dual threshold reduced test is feasible. In fact, $L_{\alpha,\beta}$ corresponds to the test in which the buyer receives a perfectly informative score for $F(\theta) < \alpha$ and $F(\theta) \geqslant \beta$, and the null signal otherwise, see Figure 1.*

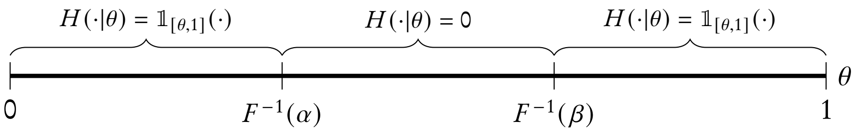

Figure 1: Generating the reduced test $L_{\alpha,\beta}$.

The proposition below pins down the family of feasible reduced tests in terms of dual threshold tests.

Proposition 1. *A reduced test L is feasible iff there exist a dual threshold reduced test $L_{\alpha,\beta}$ such that*

$$L(1) = L_{\alpha,\beta}(1), \tag{5}$$

$$\int_0^1 s dL(s) = \int_0^1 s dL_{\alpha,\beta}(s), \tag{6}$$

$$\int_0^t L(s)ds \leqslant \int_0^t L_{\alpha,\beta}(s)ds \quad \forall t \in [0,1]. \tag{7}$$

To visualize the conditions of Proposition 1, suppose that the buyer's value is uniformly distributed on the unit interval. In Figure 2, the black line is the prior cdf and the blue one represents a feasible reduced test $L(s) = \lambda s$ in which the buyer perfectly learns his valuation with a certain constant probability $\lambda \in [0,1]$.

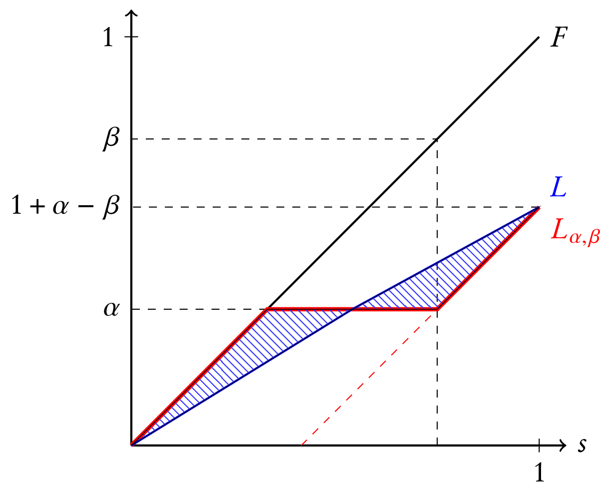

Figure 2: Constructing feasible reduced tests.

According to Proposition 1, this reduced test can alternatively by obtained by garbling some dual threshold reduced test $L_{\alpha,\beta}$. Specifically, the thresholds $\alpha \leqslant \beta$ are uniquely determined by Eq. (5) and Eq. (6), which pin down the right-end point and the average score, that is

$$\lambda = 1 + \alpha - \beta, \quad \frac{\lambda}{2} = \frac{\alpha^2}{2} + \frac{1-\beta^2}{2} \quad \Longrightarrow \quad \alpha = \frac{\lambda}{2}, \quad \beta = 1 - \frac{\lambda}{2}.$$

More generally, the total certification probability is the same under L and $L_{\alpha,\beta}$. It follows that Eq. (6) is equivalent to the fact that their integrals coincide, i.e., the two blue triangular regions in Figure 2 have equal areas. Eq. (7) is equivalent to requiring that L - normalized by the total certification probability - second order stochastically dominates $L_{\alpha,\beta}$, similarly normalized. Finally, we note that Eq. (1) implies that both L and $L_{\alpha,\beta}$ have the same base score, i.e., $\theta(L) = \theta(L_{\alpha,\beta})$.

Proposition 1 can be seen as a generalization of the classical Blackwell result in which $L(1) = 1$ and $\int_0^1 s dL(s) = \mathbb{E}[\theta]$, thus $\alpha = \beta = 1$. As in Blackwell [1953] and Gentzkow and Kamenica [2016], the statement of Proposition 1 is true even if F has atoms and/or its support is not connected. The complete proof can be found in the appendix.

4 Seller's optimal pricing

We now solve for the seller's optimal prices given a feasible reduced test L. It is easy to see that the buyer with a score s will not take the score-contingent price if $p(s) > \min\{s, p\}$. If $p(s) > s$, this gives him negative payoff and if $p(s) > p$, he can always choose the price p associated with the null score. For a fixed base price, the seller optimally sets her score-contingent price at the highest level at which the buyer is incentivized to trade, provided she makes a profit on this transaction, that is

$$p(s) = \begin{cases} \min\{s, p\} & \text{if} \quad \min\{s, p\} \geqslant c, \\ c & \text{if} \quad \min\{s, p\} < c. \end{cases} \tag{8}$$

The seller's profit can be written as a function of the base price and base score as follows:

$$\int_0^1 \max\{0, \min\{s, p\} - c\} dL(s) + (p - c)(1 - L(1)) \cdot \mathbb{1}_{[0,\theta(L)]}(p). \tag{9}$$

The first term in Eq. (9) is the expected profit that the seller makes by the selling the good for the score-contingent prices as defined above. The second term corresponds to the profit that she obtains from selling the good to the buyer who received the null signal

— the probability of this event is $1 - L(1)$ and the associated posterior mean is $\theta(L)$. So the buyer trades as long as $\theta(L) \geqslant p$, which forms the argument of the indicator function.

While the score contingent price $p(s)$ is given by Eq. (8), which is a function of the base price p; the base price itself depends on two criteria. First, suppose the probability of the null signal is zero, i.e., $1 - L(1) = 0$. Then, the buyer shouldn't claim the null signal. So the base price can be set at some $p \geqslant 1$ to deter him from making the "off-path" claim.

Consider the case $1 - L(1) > 0$. It is easy to see that when the seller offers a tradable base price ($p < 1$), it should be equal to the buyer's reservation value conditional on the null signal, so $p = \theta(L)$. But, does the seller indeed find it profitable to offer a tradable base price? It depends on the information rents she has to pay for allowing the buyer to "misreport" having the null signal. Assuming $c = 0$ for simplicity, the seller's profits in the absence and presence of a tradable base price are given respectively by

$$\int_0^1 s dL(s) \quad \text{and} \quad p(1 - L(1)) + \int_0^p s dL(s) + \int_p^1 p dL(s).$$

Substituting $p = \theta(L)$ and writing the above as an inequality, we get the following: $\theta(L)(1 - L(1)) \geqslant \int_{\theta(L)}^1 (s - \theta(L)) dL(s)$. So, the seller should offer $p = \theta(L)$, whenever this inequality holds and offer $p \geqslant 1$ otherwise. Incorporating positive costs in the above calculation and assuming ties are broken in favor of the buyer gives us the following result.

Proposition 2. *Fix a reduced test L.*

1. *If $1 - L(1) > 0$ and $(\theta(L) - c)(1 - L(1)) \geqslant \int_{\theta(L)}^1 (s - \theta(L)) dL(s)$, then the seller's optimal pricing scheme is given by*

$$p = \theta(L) \text{ and } p(s) = \max\{\min\{s, \theta(L)\}, c\} \quad \forall s \in [0, 1].$$

2. *If either $1 - L(1) = 0$ or $(\theta(L) - c)(1 - L(1)) < \int_{\theta(L)}^1 (s - \theta(L)) dL(s)$, then the seller's optimal pricing scheme is given by*

$$p \geqslant 1 \text{ and } p(s) = \max\{s, c\} \quad \forall s \in [0, 1].$$

Proposition 2 completely characterizes the seller's optimal pricing mechanism as a best response to the buyer's choice of a reduced test. In the first case, the seller posts the base price that equals to the base score, that is the posterior mean conditional on the null signal, and offers discounts for showing a certificate with a score below the base price. In the second case, the seller serves only the buyer with a certificate thereby extracting the whole surplus conditional on some certificate being printed.

The inequality in Proposition 2, for $1 - L(1) > 0$, can be re-written

$$\underbrace{\theta(L) - \frac{\int_{\theta(L)}^{1}(s - \theta(L))dL(s)}{1 - L(1)}}_{\text{virtual value}} \geqslant c. \tag{10}$$

In the language of Myersonian mechanism design, $\int_{\theta(L)}^{1}(s - \theta(L))dL(s)$ is essentially the information rent to be paid to all types $s \geqslant \theta(L)$, and total term on the left hand side is simply the virtual value that captures through the inequality whether it is indeed worthwhile to pay these type the information rent associated with the null signal.

In the case when the virtual value is higher than the cost of production, trade happens more often because the buyer is served even if he receives the null signal. But such trades are ex-post inefficient when the buyer's valuation is below the marginal cost and he has the null signal. Yet, having a tradable base price is better for the seller as long as the probability of the null signal, $1 - L(1)$, and base score, $\theta(L)$, are large enough.

5 Buyer's optimal reduced test

We now have all the ingredients to describe the problem of buyer's optimal learning under hard information. For a given reduced test L, the buyer's expected payoff is zero when the seller finds it optimal to set $p \geqslant 1$. This is so because the case corresponds to the second part of Proposition 2, where $p(s) = s$ whenever there is trade. That the buyer avoids such reduced tests becomes a constraint, and his problem can be succinctly expressed as

$$\sup_{0 \leqslant \alpha < \beta \leqslant 1,\, L} \int_{\theta(L)}^{1}(s - \theta(L))dL(s) \quad \text{subject to} \tag{11}$$
$$\text{Eq. (5) - Eq. (7)}, \quad (\theta(L) - c)(1 - L(1)) \geqslant \int_{\theta(L)}^{1}(s - \theta(L))dL(s).$$

The buyer's objective in Problem (11) integrates over the outcomes of the reduced test that generate a score higher than the base score. This is because in all other cases the buyer's payoff is exactly zero. Owing to Proposition 1, the choice is restricted to reduced tests with two thresholds α and β which satisfy Eq. (5) - Eq. (7). And lastly the inequality constraint ensure we are in the first case of Proposition 2 so that the seller finds it optimal to offer a base price of $\theta(L)$. A reduced form test that solves this problem is said to be *buyer's optimal*. In what follows we completely characterize the family of such reduced tests.

5.1 Two threshold optimum

Consider the class of dual threshold reduced tests introduced in Proposition 1. The following result states the existence of a unique buyer-optimal reduced test in this class and that every buyer's optimal reduced test is a garbled version of this dual threshold reduced test.

Theorem 1. *There are two unique thresholds $0 \leqslant \alpha^\star < F(c) \leqslant \beta^\star \leqslant 1$ such that $L_{\alpha^\star,\beta^\star}$ is a buyer's optimal reduced test. Moreover, a reduced test L is buyer's optimal if and only if it can be obtained from $L_{\alpha^\star,\beta^\star}$ by garbling scores independently in $\left[0, F^{-1}(\alpha^\star)\right)$ and $\left[F^{-1}(\beta^\star), 1\right]$.*

Remark 2. *Amongst all buyer's optimal reduced tests, the most (Blackwell) informative reduced test is $L_{\alpha^\star,\beta^\star}$; and the least (Blackwell) informative one corresponds to the test under which the buyer receives $s_L := \mathbb{E}[\theta|F(\theta) < \alpha^\star]$ for $F(\theta) < \alpha^\star$, $s_H := \mathbb{E}[\theta|F(\theta) \geqslant \beta^\star]$ for $F(\theta) \geqslant \beta^\star$, and the null signal otherwise, see Figure 3.*

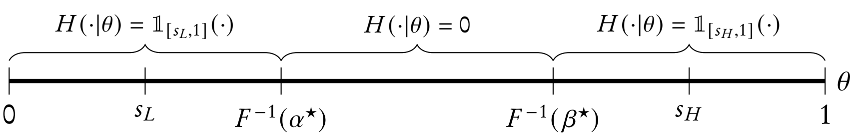

Figure 3: Generating the least informative buyer's optimal reduced test.

Theorem 1 underscores the simplicity of the endogenously generated optimal tests in the large universe of all possible learning procedures. They are partitional with at most two thresholds. Two specific reduced tests described in Remark 2 are easy to implement. Partition the type space into three intervals. Either perfectly reveal the buyer's valuation at the bottom and the top or reveal only the interval from which the value is drawn. The pricing mechanism used by the seller follows immediately from Proposition 2.

It is obvious that the buyer wants a positive measure of types to not produce a certificate, i.e. $1 - L(1) > 0$, else the seller extracts the entire surplus. Intuitively, when the cost of the seller is zero, the buyer pools all these types with no certificate at the bottom, i.e. $\alpha^* = 0$. It ensures the price at which the seller trades with the types with no-certificate, i.e. the base price given by $\theta(L_{\alpha^*,\beta^*}) = \mathbb{E}[\theta|\theta \leqslant F^{-1}(\beta^*)]$, is the smallest possible. This, in turn, generates the maximal (information) rent for the types that do have a certificate in the top region, $[F^{-1}(\beta^*), 1]$ because they too want to trade at the base price. Linearity of the model ensures that any garbling of types in the top region continues to be optimal and the Blackwell least informative one simply prints a unique score in the top region. Finally, trade is efficient and always takes place at the base price.

If the cost of the seller is positive and $\alpha^* = 0$, then at some high enough value of c, the seller is no longer willing to trade with the buyers in the region $[0, F^{-1}(\beta^*)]$, the ones without a certificate. The virtual value constraint described in Eq. (10) is violated and we fall into part 2 of Proposition 2. So the seller will exclude these buyers by setting a base price greater than 1, and trade only with those that produce a certificate in the region $[F^{-1}(\beta^*), 1]$, extracting the whole surplus using score contingent prices. In order to avoid this, the buyer produces some information at the bottom by setting $\alpha^* > 0$, restores some efficiency in trade and increases the base price to $\theta(L_{\alpha^*,\beta^*}) = \mathbb{E}[\theta | F^{-1}(\alpha^*) \leqslant \theta \leqslant F^{-1}(\beta^*)]$. Again this bottom interval can be independently garbled and the Blackwell least informative one simply prints a unique score in the region. Trade takes place in the intervals $[F^{-1}(\alpha^*), F^{-1}(\beta^*))$ and $[F^{-1}(\beta^*), 1]$ at the base price, and a discount is offered in the interval $[0, F^{-1}(\alpha^*))$ if it is profitable for the seller to trade with these types.

The formal proof of Theorem 1 is sketched in the next subsection. Before that we document the comparative static of the information structure with the respect to the seller's cost.

Corollary 1. *There exists $\underline{c}, \overline{c} \in \big(0, \mathbb{E}[\theta]\big)$ such that the value of $\alpha^\star$ as described in Proposition 1 is 0 for all $c \in \big[0, \underline{c}\big]$, and it is positive for all $c \in \big[\overline{c}, \mathbb{E}[\theta]\big)$.*

5.2 Proof sketch of Theorem 1 and Corollary 1

The buyer's problem (11) is high-dimensional as it involves the choice of a function L. In what follows we will show that this infinite-dimensional program can be reduced to the choice of a single variable from which the optimal dual threshold reduced test claimed in Theorem 1 can be constructed. Our argument is split in three separate steps. To start with, note that the buyer's objective can be re-written as follows:

$$\int_{\theta(L)}^{1} (s - \theta(L))dL(s) = \int_0^1 s dL(s) - \theta(L)L(1) + \delta(L), \quad \text{where } \delta(L) := \int_0^{\theta(L)} L(s)ds. \tag{12}$$

Since L must be feasible, its total certification probability, average score and base score are pinned down by the associated dual threshold reduced test $L_{\alpha,\beta}$. Thus, the only free variable in Problem (11), which is independent of α and β, is $\delta(L)$. Our first step is to reduce the choice $\delta(L)$ to a constant δ. Then, we will show that both β and δ can be imputed as functions of α, thus further reducing the buyer's optimization problem to the choice a single variable.

Step I: Reduce the choice of $\delta(L)$ to a single variable δ. Feasibility immediately imposes restrictions on the choice of $\delta(L)$; in fact, it can be bounded from above and below. First, we note that Eq. (7) implies $\delta(L) \leqslant \delta(L_{\alpha,\beta})$. As for the lower bound, we consider the least (Blackwell) informative feasible reduced test, i.e., one that assigns all mass to the (conditional) average score.[10] By construction, $\delta(L)$ must be larger than the corresponding value for this reduced test, that is

$$\delta(L) \geqslant \max\left\{0, \theta(L_{\alpha,\beta})(1+\alpha-\beta) - \int_0^1 s dL_{\alpha,\beta}(s)\right\}. \tag{13}$$

Now, fix a pair of thresholds $0 \leqslant \alpha < \beta \leqslant 1$, and some δ between the upper and lower bounds mentioned above. We will construct a feasible reduced test L that attains this value of δ. For $\underline{\alpha} \in [0, \alpha]$ and $\overline{\beta} \in [\beta, 1]$, define a reduced test $L_{\alpha,\beta}^{\underline{\alpha},\overline{\beta}}$ as follows:

$$L_{\alpha,\beta}^{\underline{\alpha},\overline{\beta}}(s) := \begin{cases} F(s) & \text{if} \quad F(s) < \underline{\alpha}, \\ \underline{\alpha} & \text{if} \quad F(s) \geqslant \underline{\alpha} \ \wedge \ s < \theta(L_{\alpha,\beta}), \\ \overline{\beta} + \alpha - \beta & \text{if} \quad F(s) < \overline{\beta} \ \wedge \ s \geqslant \theta(L_{\alpha,\beta}), \\ F(s) + \alpha - \beta & \text{if} \quad F(s) \geqslant \overline{\beta}. \end{cases} \tag{14}$$

According to Proposition 1, a reduced test $L_{\alpha,\beta}^{\underline{\alpha},\overline{\beta}}$ is feasible if and only if it has the same mean as $L_{\alpha,\beta}$.[11] The reader can verify that the average score under $L_{\alpha,\beta}^{\underline{\alpha},\overline{\beta}}$ is increasing in $\underline{\alpha}$ and decreasing in $\overline{\beta}$. Thus, by varying $\underline{\alpha}$ and adjusting $\overline{\beta}$ accordingly, one can attain any value of δ between the lower and upper bounds.

[10]The lowest value of $\delta(L)$ amongst all reduced test satisfying Eq. (5) and Eq. (6) is attained at $L(s) = (1+\alpha-\beta)\mathbb{1}_{\left[\int_0^1 t dL(t)/(1+\alpha-\beta), 1\right]}(s)$.

[11]Such reduced test can be obtained from the test in which the buyer's receives a perfectly informative score for $\theta \in \left[0, F^{-1}(\underline{\alpha})\right) \cup \left[F^{-1}(\overline{\beta}), 1\right]$, the score of $\theta(L_{\alpha,\beta})$ for $\theta \in \left[F^{-1}(\underline{\alpha}), F^{-1}(\alpha)\right) \cup \left[F^{-1}(\beta), F^{-1}(\overline{\beta})\right)$ and the null signal, otherwise.

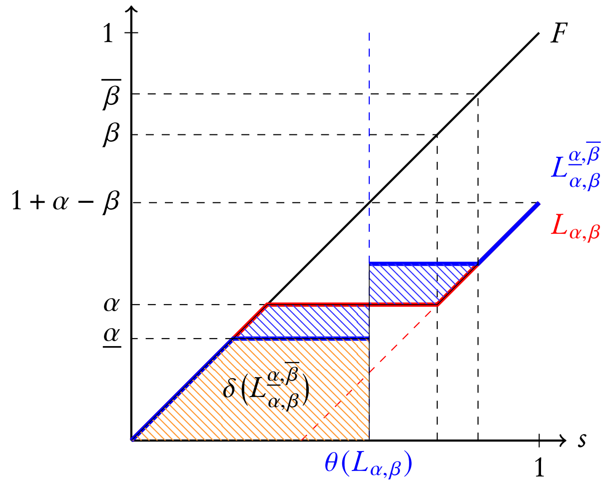

Figure 4: Attaining various values of δ.

To see the construction in action, consider the example of the uniform distribution depicted in black in Figure 4. Given two thresholds α and β, the dual threshold reduced test $L_{\alpha,\beta}$, which is shown in red, is obtained by cutting out all valuations between these thresholds in the quantile space. By construction, the base score $\theta(L_{\alpha,\beta})$ lies between $F^{-1}(\alpha)$ and $F^{-1}(\beta)$. Thus, $\delta(L_{\alpha,\beta})$ equals to the area of the trapezoid with the following vertices: $(0,0)$, $(\theta(L_{\alpha,\beta}),0)$, $(\theta(L_{\alpha,\beta}),\alpha)$ and (α,α). In order to attain lower values of δ, we decrease informativeness of $L_{\alpha,\beta}$ by pooling together scores nearby $\theta(L_{\alpha,\beta})$ in a way that the average score of these pooled types is exactly $\theta(L_{\alpha,\beta})$. Graphically, we construct $L_{\alpha,\beta}^{\underline{\alpha},\overline{\beta}}$ by lowering $L_{\alpha,\beta}$ to the left of $\theta(L_{\alpha,\beta})$ to $\underline{\alpha}$ and raising it to the right of $\theta(L_{\alpha,\beta})$ to $\overline{\beta}+\alpha-\beta$ such that the areas of two blue trapezoids are the same, which ensures that the new reduced test has the same mean as the original dual threshold reduced test. Then, the value of $\delta(L_{\alpha,\beta}^{\underline{\alpha},\overline{\beta}})$ is the area of the orange trapezoid.

Building on these observations, the following lemma simplifies Problem (11). It reduces the dimensionality of the buyer's learning decision to just three variables, namely α, β and δ.

Lemma 1. *Fix $0 \leqslant \alpha < \beta \leqslant 1$ and $u \in \mathbb{R}$. There exists a reduced test L that satisfies the constraints of Problem (11) and yields the buyer's expected payoff of u if and only if there exists*

$\delta \in \mathbb{R}$ *such that*

$$u = \int_0^1 s dL_{\alpha,\beta}(s) - \theta(L_{\alpha,\beta})(1 + \alpha - \beta) + \delta, \tag{15}$$

$$\delta(L_{\alpha,\beta}) \geqslant \delta \geqslant \max\left\{0, \theta(L_{\alpha,\beta})(1 + \alpha - \beta) - \int_0^1 s dL_{\alpha,\beta}(s)\right\}, \tag{16}$$

$$\theta(L_{\alpha,\beta}) - \int_0^1 s dL_{\alpha,\beta}(s) - c(\beta - \alpha) \geqslant \delta. \tag{17}$$

Eq. (15) states the buyer's expected payoff as a function of these three variables. Eq. (16) ensures that there exists a feasible reduced test that can attain the given value of δ. Finally, Eq. (17) ensures that the seller wants to trade with the buyer who has the base score.

To sum up, the problem of buyer's optimal learning can be split into two parts. First, find two thresholds α, β and the value of δ that gives the highest possible payoff u to the buyer. This is a finite-dimensional optimization problem with the objective given by Eq. (15) and constrained by Eq. (16) and Eq. (17). Once the solution $(\alpha^\star, \beta^\star, \delta^\star)$ is known, the whole family of buyer's optimal reduced tests can be reconstructed from it: Eq. (15) - (17) ensure that any reduced test L that is feasible for $(\alpha^\star, \beta^\star)$ and has the same value of $\delta(L) = \delta^\star$ is buyer's optimal.

Step II: Reduce δ further to be a function of α and β. We now show that it is optimal for the buyer to choose δ at its maximal feasible level, i.e., $\delta = \delta(L_{\alpha,\beta})$. At the conceptual level, this means that the family of dual threshold reduced tests is sufficient to solve for the buyer's highest payoff.

The mechanics of this result is as follows. Take a tuple $(\alpha, \beta, \delta, u)$ satisfying the conditions of Lemma 1 and assume that δ is below its upper bound. As discussed in the previous step, the payoff u can be attained by some feasible test $L_{\alpha,\beta}^{\underline{\alpha},\overline{\beta}}$ with $\underline{\alpha} < \alpha \leqslant \beta < \overline{\beta}$, where two additional thresholds $\underline{\alpha}$ and $\overline{\beta}$ are chosen to match the given value of δ. Using the parametric form of $L_{\alpha,\beta}^{\underline{\alpha},\overline{\beta}}$, the reader can verify that the buyer's payoff u under this reduced test can be re-written as

$$u = \int_{F^{-1}(\overline{\beta})}^1 (\theta - \theta(L_{\alpha,\beta}))dF(\theta). \tag{18}$$

In the appendix we show that we can increase both α, β by a small $\varepsilon > 0$, push $\underline{\alpha}$ and $\overline{\beta}$ closer to α and β so that Eq. (16) and Eq. (17) still hold. Although, the perturbation marginally decreases each "existing" type's information rent as $\theta(L_{\alpha,\beta})$ goes up, it allows many more types to have rents, because of lower $\overline{\beta}$. The second effect is always dominant and thus the expected payoff, as defined in Eq. (18), goes up.

To sum up, the perturbation of $L_{\alpha,\beta}^{\underline{\alpha},\overline{\beta}}$, which corresponds to the tuple $(\alpha, \beta, \delta, u)$, yields a new tuple $(\alpha', \beta', \delta', u')$ that still satisfies constraints of the buyer's optimization problem, and the expected payoff u' is strictly higher than u. Since we can keep raising the buyer's expected payoff by further perturbing $L_{\alpha,\beta}^{\underline{\alpha},\overline{\beta}}$ as long as the value of δ is below its maximal upper bound, the claim follows.

Lemma 2. *Let $(\alpha, \beta, \delta, u) \in \mathbb{R}^4$ satisfy Eq. (15) - Eq. (17) with $0 \leqslant \alpha < \beta \leqslant 1$ and $\delta < \delta(L_{\alpha,\beta})$. Then, there exists $(\alpha', \beta', \delta', u')$ satisfying the same constraints but $\delta' = \delta(L_{\alpha',\beta'})$, and $u' > u$.*

Step III: Rewrite β as a function of α Using Steps I and II, we can finally reduce the problem of the buyer's optimal learning to just two thresholds and one inequality constraint that guarantees that the seller is willing to trade with the buyer who has the base score. It is convenient to work with the inverse of the prior distribution $R := F^{-1}$. Then, using Lemma 1 and the fact that $\delta = \delta(L_{\alpha,\beta})$ (Lemma 2), Problem (11) can now be written as

$$\sup_{0 \leqslant \alpha < \beta \leqslant 1} \int_{\beta}^{1} (R(s) - \theta(L_{\alpha,\beta}))ds \text{ subject to } (\theta(L_{\alpha,\beta}) - c)(1 - \alpha) \geqslant \int_{\beta}^{1} (R(s) - c)ds, \tag{19}$$

where $\theta(L_{\alpha,\beta}) = \frac{\int_{\alpha}^{\beta} R(s)ds}{\beta - \alpha}$, which follows directly from Eq. (1).

We now argue that the constraint in Problem (19) must hold as an equality, and there is a unique value of β for which it is true. To see the first point, note that the objective is decreasing in β. Since the inequality constraint is necessarily violated as $\beta \searrow \alpha$, any solution to Problem (19) satisfies the constraint as an equality — the seller's is made exactly indifferent between trading with the buyer who has the base score and ignoring this buyer.

As for the second point, the constraint requires $\theta(L_{\alpha,\beta}) \geqslant c$, thus $R(\beta) \geqslant c$ (resp. $\beta \geqslant F(c)$). Since for each $\alpha \in [0, 1)$, the left-hand side is increasing and the right-hand side is decreasing in β for $\beta \in [F(c), 1]$, there exists a unique number $\beta(\alpha) \in (\alpha, 1)$ such that the constraint indeed holds as an equality, that is

$$(\theta(L_{\alpha,\beta(\alpha)}) - c)(1 - \alpha) = \int_{\beta(\alpha)}^{1} (R(s) - c)ds. \tag{20}$$

This discussion is summarized in the following lemma.

Lemma 3. *For each $\alpha \in [0, 1)$, there exists a unique number $\beta(\alpha) \in [F(c), 1]$ satisfying Eq. (20). Moreover, if (α, β) solves Problem (19), then $\beta = \beta(\alpha)$.*

Thus, Lemmas (1) - (3) show that the buyer's maximal expected payoff can be found by solving a "simple" maximization program — Problem (19) in which the inequality constraint binds. Let $V(\alpha)$ denote the objective in this problem as a function of α holding $\beta = \beta(\alpha)$, that is

$$V(\alpha) := \int_{\alpha}^{\beta(\alpha)} (R(s) - c) ds. \tag{21}$$

So, the buyer's problem finally reduces to the choice of α to maximize $V(\alpha)$. The rest of the proof precisely characterizes the optimal value of α and backward inducts the family of buyer's optimal reduced tests. Specifically, we establish that $V(\alpha)$ is single-peaked, or more precisely that the derivative of $V(\alpha)$ crosses zero at most once and from above. This gives us the necessary and sufficient condition for the optimal threshold in terms of the primitives, see Eq. (25) in the appendix. In the proof of Corollary 1, we show that this condition is necessarily satisfied at $\alpha^\star = 0$ ($\alpha^\star > 0$) for small (resp. large) c.

5.3 Example

We now look at a specific numerical example to see the workings of Theorem 1.

Example 1. *Suppose that the buyer's value is uniform on the unit interval, that is $F(\theta) = R(\theta) = \theta$. By Eq. (20), given $\alpha \in [0, 1)$, the value of $\beta(\alpha) \in (\alpha, 1]$ is a root of the following quadratic equation: $\beta^2 + \beta(1 - \alpha - 2c) + (\alpha(1 - \alpha + 2c) - 1) = 0$. Only the largest root lies in $(\alpha, 1]$, thus*

$$\beta(\alpha) = \frac{1}{2}\left(\alpha + 2c - 1 + \sqrt{5\alpha^2 - 4\alpha c - 6\alpha + 4c^2 - 4c + 5}\right).$$

The objective in Problem (19) evaluated at $\beta = \beta(\alpha)$ is given by Equation (21):

$$V(\alpha) = \frac{1}{2}\left(\beta(\alpha) - \alpha\right)\left(\beta(\alpha) + \alpha - 2c\right).$$

As can be see from Figure 5, V is not concave but it is single-peaked. For $c = 0$, the peak is at $\alpha^\star = 0$; however, for large c the peak is positive. Since V is single-peaked, the positivity of $\alpha^\star$ is completely determined by the sign of the derivative of V at $\alpha = 0$, that is

$$V'(0) = \big(R(\beta(0)) - c\big)\beta'(0) + c.$$

The reader can verify that both $R(\beta(0)) - c$ and $\beta'(0)$ are increasing in c; and, $V'(0)$ equals

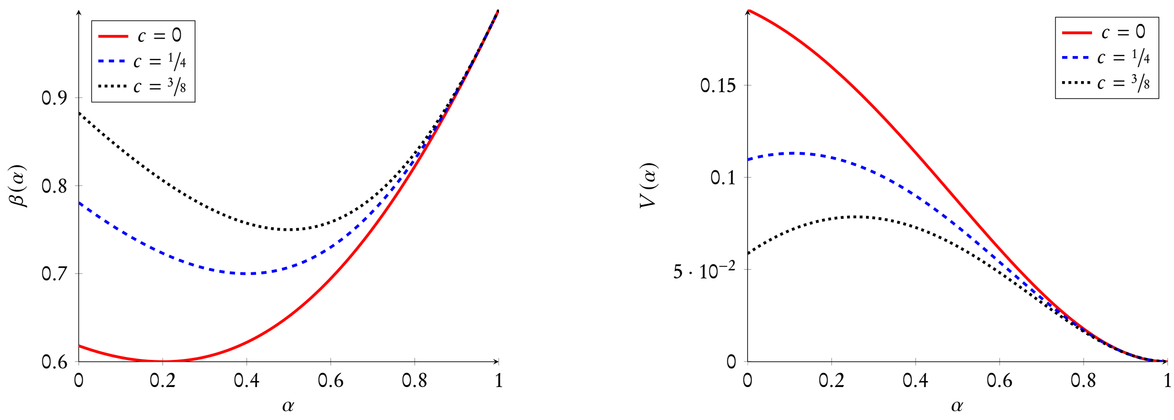

Figure 5: $\beta(\cdot)$ and V as a function of α.

zero for $c = 0.1585$. *It follows that for* c *below this value,* $\alpha^\star = 0$ *and the single threshold test is optimal. Otherwise,* $\alpha^\star > 0$, *and we need two regions, in addition to the no-certificate region. In fact, the optimal value of* $\alpha^\star$ *as well as* $\beta(\alpha^\star)$ *are non-decreasing in the seller's marginal cost as shown in Figure 6.*

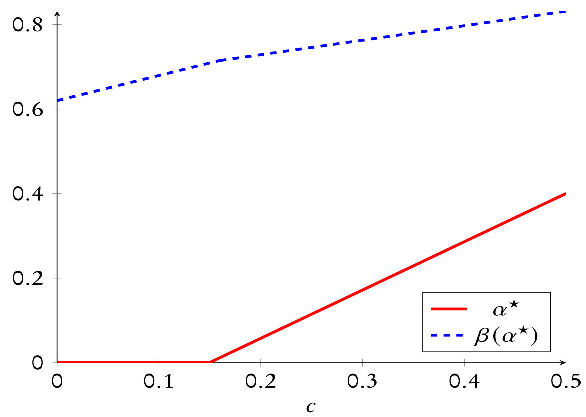

Figure 6: Thresholds in the buyer's optimal dual threshold reduced test as a function of c.

Finally, to give a fuller picture for the simplest case, note that when seller's marginal cost is zero, i.e. $c = 0$, *then* $\alpha^\star = 0$, *and* $\beta(0) = \frac{\sqrt{5}-1}{2} \approx 0.62$. *So the buyer receives a certificate in the region* $[\beta, 1]$ *with 'high' printed on it and receives no certificate in the low region* $[0, \beta)$. *The seller in turn picks a price mechanism that binds the participation constraint of the low region and the incentive constraint of the high region. This culminates in a posted price of* $p = \mathbb{E}[\theta | \theta \in [0, \beta)] = 0.31$. *It extracts all the surplus from the low region buyers and keeps the high region buyers indifferent between showing or not showing the certificate. Trade always takes place, and hence is efficient. The profit of the seller is fixed at 0.31 and the expected profit of the buyer is* $\mathbb{E}[\theta] - p = 0.19$.

6 Discussion

In this section we explore the implications of the hardness of information and seller's commitment power for the design of the optimal information structure. We show that it is the combination of these two which makes the production of information simpler than in the classical soft information design benchmark. Intuitively, hard information allows the seller to tag prices to a verifiable object (score), which cannot be falsified, thereby drastically shrinking the set of (buyer's) incentive constraints; and, her commitment power is necessary for such contingent prices to be credible from the buyer's point of view.

6.1 Role of hard information

To understand the role of hard information in the choice of optimal test, we now consider the opposite case in which information is completely soft — the buyer can fake a certificate with any score at no costs. To make things simple, we assume that the marginal cost of production is zero, i.e., $c = 0$.

Given that the buyer can always claim to have any score, the seller cannot gain by offering score-contingent prices strictly below her base price. So, it is without loss of generality to assume that $p(s) = p$ for all $s \in [0, 1]$, where the base price is chosen to maximize the seller's profit for the given distribution of posterior means, which is specified by a reduced test. Furthermore, since the seller's incentives depend only on the distribution of posterior means, it is without loss to consider reduced tests in which the null signal is drawn with probability zero.[12] The problem can be stated as follows:

$$\sup_{0 \leqslant p \leqslant 1,\, L} \int_p^1 (s-p)\,dL(s) \quad \text{subject to}$$

$$\text{Eq. (5) - Eq. (7) for } \alpha = \beta = 1, \quad p(1 - L(p-)) \geqslant p'(1 - L(p'-)) \quad \forall p' \in [0, 1].$$

The solution to this is well-known, and we document it mainly for completeness.

Proposition 3 (Roesler and Szentes [2017]). *Suppose that $c = 0$ and information is soft. Then, there exist two numbers π and B such that the following reduced test L is buyer's optimal:*

$$L(s) := \begin{cases} \max\left\{0, 1 - \frac{\pi}{s}\right\} & \textit{if} \quad s < B, \\ 1 & \textit{if} \quad s \geqslant B. \end{cases}$$

[12] Formally, let L be a feasible reduced test with $1 - L(1) > 0$. Define a reduced test L' as $L'(s) := L(s) + (1 - L(1))\mathbb{1}_{[\theta(L),1]}(s)$. Then, L' is feasible, and the seller's profit as a function of the base price is the same under L' and under L.

The seller best responds to this test by offering the base price of π at which trade always takes place. Furthermore, any reduced test with the same distribution of posterior means as L is also buyer's optimal.

As can be seen from the contrast in Theorem 1 and Proposition 3, the buyer's optimal test takes a markedly different forms under hard and soft information. To fix ideas, recollect that in the fixed information benchmark studied in the classical screening problem, the seller would choose p to maximize $p(1 - F(p))$, which is achieved at, say, p^{classic}. Then, no trade takes place for $\theta < p^{\text{classic}}$ and all buyer types $\theta \geqslant p^{\text{classic}}$ trade at the posted price. In contrast, in both soft and hard information design models, trade always occurs, creating a total surplus of $\mathbb{E}[\theta]$. In both cases, it is optimal for the buyer to achieve the efficient surplus at the cheapest possible price. But the optimal information structure and the induced trading rule are different.

In the soft information benchmark, as in the fixed information case, the underlying mechanism design problem has a continuum of incentive constraints, since the buyer can misreport any signal. This leads to a posted price mechanism being the best a seller can do to provide the right incentives. So, the buyer's learning technology aims to generate the lowest possible posted price, while ensuring efficient trade. In fact the buyer can ensure this price is strictly less than p^{classic}.

Proposition 3 essentially states that this information structure must guarantee that the seller is indifferent between any two prices in $[\pi, B]$, i.e., $p(1 - L(p-)) = \pi$ for all $p \in [\pi, B]$. The buyer chooses the smallest π and B so that the truncated-Pareto reduced test L can be obtained from the prior. The buyer's posterior lies in $[\pi, B]$ and he never learns his valuation precisely.

In the hard information design problem we study, the buyer who has a certificate can deviate in only one way - by hiding it. This changes the seller's problem in two fundamental ways. First, she can tag prices to verifiable scores when the certificate is presented, and second she has to pay the shadow price of a dramatically lower number of underlying incentive constraints for the buyer. In fact buyer types whose value is below the base price are offered a discount and incentivized to trade. Thus, due to the shrunken set of deviations available to the buyer and the increased set of instruments available to the seller, the buyer responds by choosing a very simple information structure.

It must be noted that while both soft and hard information design problems lead to an implementation with a unique posted price (when $c = 0$), the economic forces driving this result in the two cases are quite different. In soft information design, the best the seller can do is to post a single price. In the hard information design problem, the seller has a richer set of instruments, in that prices can be tagged to the scores shown on the certificate. But, in response to this richness the buyer chooses an information structure in which a posted price, in fact, achieves the optimum.

As the reader can perhaps guess, the buyer (resp. seller) always prefers to learn through soft (resp. hard) information. To see it, consider the one threshold test discussed above; the implied reduced test is as follows:

$$L(s) = \begin{cases} 1-\beta(0) & \text{if} \quad s \geqslant \int_0^1 s d\left(\frac{L_{0,\beta(0)}(s)}{1-\beta(0)}\right), \\ 0 & \text{if} \quad s < \int_0^1 s d\left(\frac{L_{0,\beta(0)}(s)}{1-\beta(0)}\right). \end{cases}$$

By Theorem 1, L is buyer's optimal reduced test under hard information. The reader can verify that, given this reduced test, the seller prefers to offer the base of price of $\theta(L)$ even if information is soft. It follows that the buyer's highest expected payoff under hard information is lower than under soft information. Since the trade is efficient under both hard and soft information whenever $c = 0$ (see Theorem 1 and Proposition 3) the sum of buyer's expected payoff and seller's expected profit equals to $\mathbb{E}[\theta]$. Thus, we can conclude that the seller prefers the buyer to learn through hard information.

Corollary 2. *Suppose that $c = 0$. The buyer's expected payoff is smaller and the seller's profit is larger under hard information than under soft information.*

To understand the workings of this result consider the setting of Example 1 where prior is uniform, $F(\theta) = \theta$, and marginal cost of the seller is zero, $c = 0$. We know from the fixed information benchmark that the seller maximizes $p(1-p)$, picking a price of $\frac{1}{2}$. This generates an expected profit of 0.25 for her and an expected rent of 0.125 for the buyer. In the soft information design benchmark considered here, the buyer chooses a truncated Pareto distribution (of posterior means) of the kind stated in Proposition 3 with $\pi \approx 0.2$ and $B \approx 0.87$. Given this distribution, the seller posts a price of $p = \pi \approx 0.2$, and trade always takes place. So the seller's profit is fixed at 0.2 and buyer's expected rent is 0.3. As shown in Section 5.3, in the hard information model we study, the buyer partitions the type space into two regions at the threshold $\beta = 0.62$. So the seller posts a price of $p = 0.31$, which generates a profit of 0.31 and an expected payoff of 0.19 for the buyer.

6.2 Role of commitment

We now take a different perspective to the problem of the buyer's optimal learning, and explore the role of seller's commitment. Consider the extreme case in which the seller cannot commit to prices. Instead the players play the following disclosure game: 1) the buyer decides to show his certificate (if he has one) or not, 2) the seller posts a price for the good, 3) the buyer decides to take the trade or not. To keep things simple, the marginal cost is taken to be zero.

As usual, we solve the game backwards for the given reduced test L. If the buyer shows her certificate with a score $s \in [0, 1]$, then, since information is hard, the seller is certain that the buyer's posterior mean is s. It follows that the seller offers the highest score-contingent price that the buyer can accept, that is $p(s) = s$. As for the base price p - offered when no certificate is shown - since there is no cost, it is always profitable for the seller to trade with the buyer with the null signal. So, p is chosen to maximize the seller's expected profit given the distribution of posterior means conditional on no disclosure.

Further, since disclosing a certificate yields zero surplus to the buyer, a buyer with a score strictly above the base price will never reveal it. Moreover, concealing scores below the base price "jams" the seller's ability to infer information on lower valuations, giving the incentive to lower the base price. We conclude therefore that there is no loss of generality in assuming that the buyer never shows his certificate. Specifically, for any equilibrium of the disclosure game with base price p, there exists another equilibrium in which the buyer never shows his certificate, and the seller's optimal base price is weakly lower than p.

Given that there is no disclosure, the problem becomes identical to the soft information design counterpart, since the best the seller can do at the outset is to use a posted price mechanism. We document this observation in the following corollary. The formal description of the disclosure game and the proof of the result are presented in the appendix.

Corollary 3. *Suppose that $c = 0$ and the seller cannot commit. Then, the set of buyer's optimal reduced tests is exactly the same as in Proposition 3.*

7 Pareto frontier

In this section we characterize the Pareto frontier with learning through hard information. Recall that for the given the reduced test L, the seller best-responds by either extracting the whole surplus conditional on $s \neq \emptyset$ or offers the base price that equals the posterior mean conditional on the null signal. Using Proposition 2, the buyer's expected payoff, say $U(L)$, can be expressed as follows:

$$U(L) := \begin{cases} \int_0^1 s dL(s) - \theta(L)L(1) + \delta(L) & \text{if } 1 - L(1) > 0,\ \theta(L) - \int_0^1 s dL(s) - \delta(L) \geqslant c(1 - L(1)), \\ 0 & \text{otherwise,} \end{cases}$$

where $\theta(L)$ is the posterior mean conditional on the null signal (Eq. (1)) and $\delta(L)$ is given by Eq. (12). Combining the expression for the seller's profit function (Eq. (9)) and Proposition 2, the seller's expected profit, say $\Pi(L)$, can be succinctly written as

$$\Pi(L) := \mathbb{E}[\theta] - c - U(L) + \int_0^c L(s)ds.$$

Our goal is to characterize the set of buyer's payoff-seller's profit pairs on Pareto frontier and the set of reduced tests which implement such pairs. We will call such tests *efficient*. Formally, a feasible reduced test L is said to be *efficient* if there is no feasible reduced test L' such that both players prefer L' to L, and at least one of them strictly prefers L', that is

$$\Big(U(L'), \Pi(L')\Big) \gneqq \Big(U(L), \Pi(L)\Big).$$

Note that both the buyer's expected payoff and the seller's expected profit depend on L through (i) its certification probability, (ii) the average score, (iii) $\delta(L) = \int_0^{\theta(L)} L(s)ds$, and (iv) $\int_0^c L(s)ds$. According to Proposition 1, feasibility pins down (i) and (ii) as a function of certain thresholds $1 \leqslant \alpha \leqslant \beta \leqslant 1$ associated with the dual threshold reduced test $L_{\alpha,\beta}$. In addition, (iii) and (iv) must be lower than their analogues in the dual threshold reduced test $L_{\alpha,\beta}$, i.e., $\delta(L_{\alpha,\beta})$ and $\int_0^c L_{\alpha,\beta}(s)ds$. Thus, starting from the high dimensional family of all feasible tests, we show that, in fact, a much smaller the family of dual threshold tests is sufficient to trace the Pareto frontier. In particular, in these tests, the seller is indifferent between trading and not trading with the buyer who has the null signal.

Proposition 4.

1. *A feasible reduced test L with $U(L) = 0$ is efficient if and only if $\Pi(L) = \Pi(F)$.*
2. *A feasible reduced test L with $U(L) > 0$ is efficient if and only if there exists $\alpha \in [\alpha^\star, 1]$, where $\alpha^\star$ is as defined in Proposition 1, such that $U(L) = U(L_{\alpha,\beta(\alpha)})$ and $\Pi(L) = \Pi(L_{\alpha,\beta(\alpha)})$.*

Proposition 4 completely characterizes the payoff-profit pairs achievable through efficient tests. It also provides a simple and constructive way to trace the whole frontier through dual thresholds reduced tests, as follows. Starting with the buyer's optimal threshold $\alpha^{\star}$, increase this threshold and adjust the second threshold β according to Eq. (20) so that the seller is still indifferent between trading with the buyer who has no certificate and ignoring this buyer. The contents of the result are best illustrated through an example.

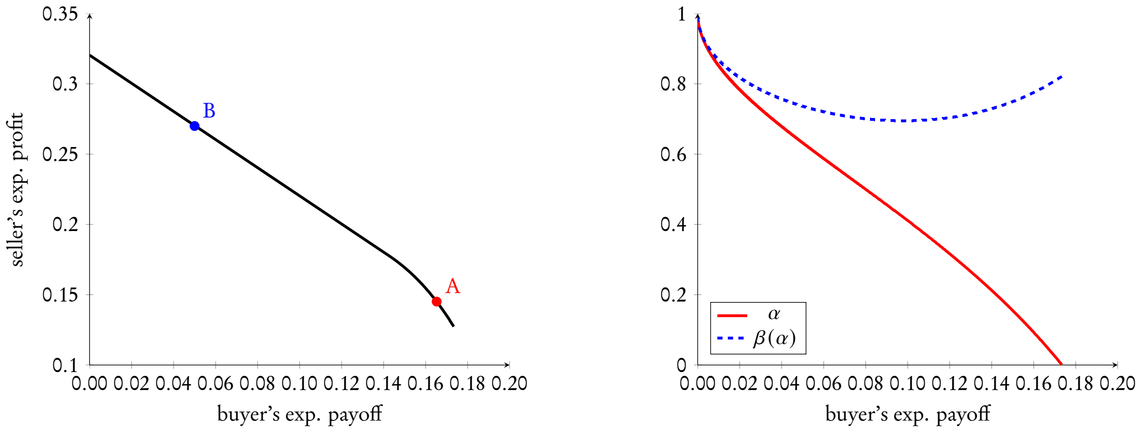

Figure 7: Pareto frontier (left panel) and $(\alpha, \beta(\alpha))$ (right panel) as functions of the buyer's expected payoff on the frontier for uniformly distributed valuation, $F(\theta) = \theta$ and marginal cost $c = 0.2$.

Example 2. *Assume that the prior is uniformly distributed, $F(\theta) = \theta$, and in addition marginal cost is $c = 0.2$. Then the total efficient surplus is given by $\mathbb{E}\left[(\theta - c)\mathbb{1}_{[c,1]}(\theta)\right] = 0.32$. Figure 7 plots the Pareto frontier, and the corresponding dual thresholds that implement each point on it.*

The right most point on the Pareto frontier corresponds to the buyer's optimal information structure and the left most point corresponds to the seller's optimal information structure. One can immediately note that when the seller has all the bargaining power on what test to choose, she chooses $\alpha = \beta = 1$, that is full information revelation, and thereby attains the entire efficient surplus as profit. As we gradually move down the Pareto frontier, we get $0 < \alpha < \beta < 1$, so the buyer is given information rent for obtaining the null signal in the region $\theta \in \left[F^{-1}(\alpha), F^{-1}(\beta)\right)$, which in the uniform case is simply $\theta \in [\alpha, \beta)$.

The Pareto frontier is at first linear and continues to achieve the efficient surplus, that is, $U + \Pi = 0.32$. For a large enough value of the buyer's payoff, however, the Pareto frontier is concave and does not achieve the efficient surplus. Comparing the two graphs in Figure 7, it can be noted that the curvature flips exactly at $\alpha = 0.2$. Since $\alpha < c$ in the right most part of the Pareto frontier, it is clear that we may have trade even for types $\alpha \leqslant \theta < c$, which is inefficient.

To grasp the reasoning better, consider the two points on the Pareto frontier in Figure 7, marked by the red dot "A" on the curved region and blue dot "B" on the flat part. The coarsest *information structures that achieve these points on the Pareto frontier are depicted in Figure 8. The one on the right corresponds to the blue dot, where $\alpha > c$. As before, the buyer does not receive a certificate in the region $[\alpha, \beta)$; receives a certificate in the region $[\beta, 1]$ with s_H printed on it; and - in a departure from previously discussed cases - the region $[0, \alpha)$ is split into two regions $[0, c)$ and $[c, \alpha)$, where $\underline{s}_L$ is printed on the certificate in the region $[0, c)$ and $\overline{s}_L$ is printed on the certificate in the region $[c, \alpha)$. This allows for efficient trade, since both parties know (on path) when the value lies in interval $[0, c)$ and no trade takes place then. A discount is offered in the region $[c, \alpha)$, and a posted price equal to $\frac{\beta+\alpha}{2}$ incentivizes trade in the regions $[\alpha, \beta)$ and $[\beta, 1]$.*

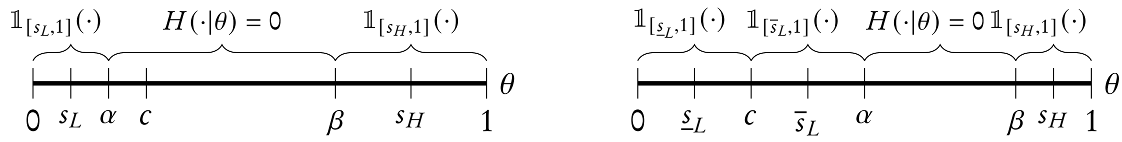

Figure 8: Least informative buyer's optimal reduced tests corresponding to points "A" (left panel) and "B" (right panel) in Figure 7, where $F(\theta) = \theta$ and marginal cost $c = 0.2$.

The information structure on the left in Figure 8 corresponds to the red point on the Pareto frontier in Figure 7. Here $\alpha < c$. Trade again takes place at the posted price $\frac{\beta+\alpha}{2}$ in the regions $[\alpha, \beta)$ and $[\beta, 1]$. However, in the bottom region, $[0, \alpha)$, there is no trade. As the figure shows, the region $[\alpha, c)$ now features trade, even though it is inefficient. What should be size of this inefficient region? The answer to this, in calculating the optimal α and β gives curvature to the Pareto frontier.

The curvature in the Pareto frontier in Figure 7 represents a trade-off that can be formalized more generally. Recollect that in Example 1 we had $c = 0$. The Pareto frontier in that case is a straight line with $U + \Pi = 0.5$, the efficient value, at all points. Following Example 2, more generally, we note that full efficiency typically cannot be attained with $c > 0$ whenever the buyer's expected payoff is large enough. The reason is the learning-information rent trade-off discussed in Section 5.1. In order to provide a large expected payoff to the buyer, we need to decrease the quality of learning for small valuations, which might lead to ex-post inefficiencies. To see it more formally, write down the total surplus as a function of α, i.e.,

$$\Pi(L_{\alpha,\beta(\alpha)}) + U(L_{\alpha,\beta(\alpha)}) = \underbrace{\mathbb{E}[\max\{0, \theta - c\}]}_{\text{maximal surplus, } \Pi(F)} - \underbrace{\int_0^c (F(s) - L_{\alpha,\beta(\alpha)}(s))ds}_{\text{loss from imperfect learning}} .$$

The first term is the maximal efficient surplus under perfect information. The second term captures the loss from imperfect learning. It is positive if and only if $\alpha < F(c)$, which, in the uniform case, is simply $\alpha < c$. If learning is imperfect then the Pareto frontier is non-linear, as seen in Figure 7.

8 A general result

We now extend our characterization of efficient tests to a more general setting. First, we allow the distribution of the buyer's valuation F to be supported on any subset of the unit interval (finite or continuum or some mixture) that includes $\{0, 1\}$. Second, generating hard evidence is potentially costly for the buyer. We assume there exists some function $\iota : [0, 1] \to \mathbb{R}_+$ with $\iota(0) = 0$, which maps evidence probabilities to costs for the buyer. Specifically, a test H costs $\iota\left(1 - \int_0^1 H(1|\theta)dF(\theta)\right)$. Third, the seller can supply any non-negative quantity q up to $\overline{q} \geqslant 0$ at a cost of $C(q) \geqslant 0$, which is assumed to be continuous and convex with $C(0) = 0$.[13]

This third feature allows us to extend the model to non-linear pricing á la Mussa and Rosen [1978]. Specifically, the seller commits to score-contingent quantity-price pairs $(\mathbf{q}, \mathbf{p}) := \left(q(s), p(s)\right)_{s\in[0,1]}$ and a base quantity-price pair (q, p).[14] Given the seller's offer, the buyer with a score $s \in [0, 1]$ can take $(q(s), p(s))$, (q, p) or walk away, and the buyer with the null signal can either take (q, p) or refuse to trade.

The notion of a reduced test is the same as before; furthermore, Proposition 1 still pins down the set of feasible reduced tests in this more general environment. Given the reduced test L, the seller selects her score-contingent quantity-price pairs and the base quantity-price pair. To formalize the solution to the seller's problem, we need two auxiliary objects: for each $s \in [0, 1]$, let $Q(s)$ be the surplus-maximizing quantity and $W(s)$ be the associated maximal surplus, that is

$$Q(s) := \max \arg \max_{q\in[0,\overline{q}]} sq - C(q), \quad W(s) := sQ(s) - C(Q(s)).$$

Proposition 5. *Fix a reduced test L.*

1. *If $1 - L(1) > 0$, then the seller's optimal quantity-price pairs are given by*

$$p = \theta(L)q, \; q = Q\left(v(L)\right) \; \textit{and} \; p(s) = sq(s) - \max\{0, s - \theta(L)\}q, \; q(s) = Q(s) \quad \forall s \in [0, 1],$$

[13] Continuity and convexity can be substantially relaxed to only requiring that the cost function is lower-semi continuous. Then, slightly abusing notations, we identify C with its convex envelope.

[14] As shown in the appendix, this space of mechanisms is without loss of generality.

where $v(L)$ is the virtual value associated with the null signal, that is

$$v(L) := \theta(L) - \frac{\int_{\theta(L)}^{1}(s-\theta(L))dL(s)}{1-L(1)}. \tag{22}$$

2. *If $1-L(1)=0$, then the seller's optimal quantity-price pairs are given by*

$$p \geqslant q \text{ and } p(s) = sq(s),\ q(s) = Q(s) \quad \forall s \in [0,1].$$

The result can be absorbed as follows. Analogous to the linear case, the information rent given to types $s \geqslant \theta(L)$ who can misreport $s = \emptyset$ is captured by $\int_{\theta(L)}^{1}(s-\theta(L))dL(s)$, which in turn constitutes the virtual value for the buyer normalized by the weight of null types, $1-L(1)$. The allocation for the null signal is then simply the surplus maximizing quantity at the virtual value, viz. $q = Q(v(L))$, and since the seller will optimally choose to make the IR constraint of the buyer bind, $p = \theta(L)q$.

The buyer who produces the certificate is allocated the pointwise surplus maximizing quantity, $q(s) = Q(s)$. As long as the score is below $\theta(L)$, the seller can make the buyer's IR constraint bind, charging $p(s) = sq(s)$. For types $s \geqslant \theta(L)$, because of the intensive margin of allocation, the information rent must reflect a sliding discount in prices, which is captured by $\max\{0, s-\theta(L)\}q$. Finally, if the probability of the null certificate is zero, $L(1) = 1$, the seller produces the efficient surplus and extracts all of it.

Analogous to Section 7, we now characterize the Pareto frontier. For the reduced test L and the seller's optimal response to it described in Proposition 5, denote the buyer's expected payoff net his learning cost by $U(L)$ and the seller's expected profit by $\Pi(L)$, resp. Using Proposition 5, these can be expressed as follows:

$$U(L) := \begin{cases} Q(v(L))\left(\int_0^1 s dL(s) - \theta(L)L(1) + \delta(L)\right) - \iota(1-L(1)) & \text{if} \quad 1-L(1)>0, \\ 0 & \text{otherwise;} \end{cases}$$

$$\Pi(L) := \begin{cases} \int_0^1 W(s)dL(s) + (1-L(1))W(v(L)) & \text{if} \quad 1-L(1)>0, \\ \int_0^1 W(s)dL(s) & \text{otherwise,} \end{cases}$$

where $\theta(L)$ is the posterior mean conditional on the null signal (Eq. (1)), $\delta(L)$ is given by Eq. (12) and $v(L)$ is the virtual value defined by Eq. (22). A feasible reduced test L is *buyer's optimal* if there no feasible reduced test L' that the buyer's strictly prefers to L, that is $U(L') > U(L)$. The notion of efficiency is the same as before: a feasible reduced test L is *efficient* if there is no feasible reduced test L' such that both players prefer L' to

L, and at least one of them strictly prefers L', that is $(U(L'), \Pi(L')) \gneqq (U(L), \Pi(L))$.

Theorem 2.

1. *Let L be efficient with $1 - L(1) = 0$. Then, the completely informative reduced test F is efficient, and $U(L) = U(F)$, $\Pi(L) = \Pi(F)$.*

2. *Let L be buyer's optimal with $1 - L(1) > 0$. Then, there exists two thresholds $0 \leqslant \alpha < \beta \leqslant 1$ such that the dual threshold reduced test $L_{\alpha,\beta}$ is buyer's optimal, and $U(L) = U(L_{\alpha,\beta})$.*

3. *Let L be efficient with $U(L) > 0$. Then, there exists four thresholds $0 \leqslant \underline{\alpha} \leqslant \alpha < \beta \leqslant \overline{\beta} \leqslant 1$ such that the reduced test $L_{\alpha,\beta}^{\underline{\alpha},\overline{\beta}}$ described in Eq. (14) is efficient, and $U(L) = U(L_{\alpha,\beta}^{\underline{\alpha},\overline{\beta}})$, $\Pi(L) = \Pi(L_{\alpha,\beta}^{\underline{\alpha},\overline{\beta}})$.*

Theorem 2 says that under very mild assumptions on the players' payoffs, any efficient test can be replicated by a reduced test in a certain parametric family $L_{\alpha,\beta}^{\underline{\alpha},\overline{\beta}}$, which is specified in Eq. (14) and is a function of four thresholds. For any continuous cdf F, the efficient test can be obtained as follows: the buyer receives a perfectly informative score for $\theta \in [0, F^{-1}(\underline{\alpha})) \cup [F^{-1}(\overline{\beta}), 1]$, the score of $\theta(L_{\alpha,\beta})$ for $\theta \in [F^{-1}(\underline{\alpha}), F^{-1}(\alpha)) \cup [F^{-1}(\beta), F^{-1}(\overline{\beta}))$ and the null signal in the middle region $[F^{-1}(\alpha), F^{-1}(\beta))$; represented in Figure 9.

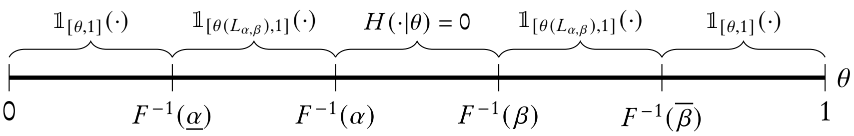

Figure 9: Generating the reduced test $L_{\alpha,\beta}^{\underline{\alpha},\overline{\beta}}$.

In other words, the bottom and top intervals have certification with the buyer's exact value printed on it, and the middle region does not get a certificate (i.e., the test is deemed inconclusive). The two new intermediate regions, to left and right of the middle region, get a certificate with a score $\theta(L_{\alpha,\beta})$ printed on them. To fix ideas, if F is uniform, we have $\frac{\beta+\alpha}{2} = \frac{\overline{\beta}+\underline{\alpha}}{2}$. So the four thresholds preserve the average score with full information revelation, and split the middle region into three regions, two of which do produce a certificate but with no new information at the intensive margin. If the buyer receives a certificate with $\theta(L_{\alpha,\beta})$ printed on it, he knows that his values lies in $[\underline{\alpha}, \alpha) \cup [\beta, \overline{\beta})$.

Note in Theorem 2 that the buyer-optimal test still has at most two thresholds. So, it is when the seller has some bargaining power on the choice of the test, she can use this extra instrument of creating two new thresholds to potentially increase the profit. With linear preferences, splitting the middle region into three regions does not help the seller because in equilibrium the prices that are accepted by the buyer do not change in response to this richer information structure, and hence the (bang-bang) allocation too remains the same. Thus, the total surplus and the buyer's payoff also are the same. However, with convex production costs, this extra instrument can potentially allow the seller to increase total surplus and her share in the split. In terms of the allocation rule in Proposition 5, $q(s) \geqslant Q(v(L))$ for s corresponding to the two new regions. Whenever this inequality is strict, the seller will find it optimal to have the extra thresholds.

Our result applies to situations in which the learning cost function is not monotone or even continuous, and C is just lower-semi continuous (see Footnote 13). If more structure is imposed, then the space of dual threshold reduced tests can be sufficient to trace the entire Pareto frontier, as happens in the uniform-quadratic setting.

Example 3. *Suppose that the buyer's value is uniform on the unit interval, certification is costless and the production cost is quadratic. That is, $F(\theta) = R(\theta) = \theta$, $\iota \equiv 0$, and $C(q) = \frac{q^2}{2}$ with $\overline{q} \geqslant 1$. The solution to this example is depicted in Figure 10.*

First note that the Pareto frontier, left panel in Figure 10, is strictly concave. The extra curvature, driven by the intensive margin of allocation (as opposed to the bang-bang nature in the linear model), implies that efficiency can in general not be attained, and the distortion away from efficiency, captured by Eq. (22)*, changes continuously in the bargaining power of the choice of test. At the bottom end of the Pareto frontier, when the buyer has most of the bargaining power, the curvature is extra sharp for reasons similar to the linear model—the trade-off between learning and information rents is non-trivial and leads to imperfect learning.*

The optimal information structures that achieve the points on the Pareto frontier are captured by the right panel of Figure 10. For the linear-quadratic model, at most two thresholds are sufficient to trace entire frontier. The figure plots the two thresholds, α and β as we move down the Pareto frontier. In comparison to the linear model, the measure for the null signal or no certificate, i.e., $\beta - \alpha$, is smaller (compare right panel in Figure 7). Thus, more information is produced with convex costs then with linear. Since the information is hard, its hurts the buyer and hence his share of the total surplus too is lower in the concave model than in the linear one.

9 Final remarks

Production of information in the marketplace is invariably some combination of "soft" and "hard". While we use the terms in a specific theoretical sense, their colloquial un-

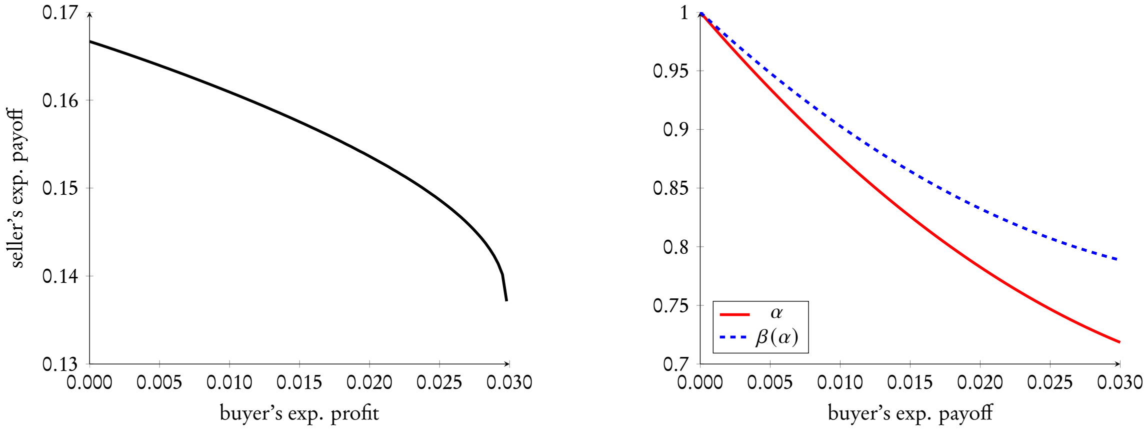

Figure 10: Pareto frontier (left panel) and $(\alpha, \beta(\alpha))$ (right panel) as functions of the buyer's expected payoff on the frontier for uniformly distributed valuation, $F(\theta) = \theta$, and quadratic cost, $C(q) = \frac{q^2}{2}$.

derstanding has been steadily evolving. For example, Liberti and Petersen [2018] meticulously document the kinds of information in use in the finance industry, and classify them as soft or hard. They note the traditional role of soft information in finance, and then define hard information in contrast as satisfying certain criteria. The rise in information technology, they contend, has enabled the processing of hard information, and hence its greater use.

While the theoretical literature in economics has responded with gusto to the rise of information technology through the study of soft information design, hard information design is arguably lagging behind its softer cousin. Our aim in this paper has been to take inspiration from (i) the rise in the use of hard information in various markets, (ii) plethora of classical studies in (exogenous) hard information and evidence, (iii) the extensive development of information design tools; and use these to study the impact of hard information design in the context of the canonical monopolistic screening problem.

We find that despite the large universe of available tests, those that map the Pareto frontier take a simple partitional form, often with at most two thresholds. This points towards a foundation for the coarseness of certifications prevalent in real market transactions. The simplicity is generated by the interaction of hardness of information and the seller's ability to commit to a pricing mechanism. Relaxing either takes the model back to the soft information design benchmark, where a more detailed distribution of signals is required at the optimum.

We start with the buyer optimal information structure and then map the Pareto frontier by varying the bargaining power on the choice of test, while keeping fixed the party who offers the contracts, i.e. the seller. It is also plausible to ask how the optimal information structure and corresponding pricing mechanism change as we vary both the bargaining power of who gets to choose the test and who gets to offer contracts. This is an interesting question for future work.[15]

Finally, as Laffont and Martimort [2009] note, the screening model can be adapted to study a wide array of economic interactions—price discrimination, regulation, insurance and financial markets. It would be interesting to extend the ideas developed here to these contexts. It would also be interesting to write down models where both soft and hard information co-exist and have to be designed as part of the larger transaction. In addition, as briefly mentioned in Section 2.2, but not fully pursued, thinking about privacy by varying the method of disclosure, i.e. what mixture of hard and soft information, seems useful. We leave the development of these ideas for future work.

10 Appendix

The proofs from Sections 5, 6.1, and 7 are provided here. In the Online Appendix, we provide the missing proofs for results on the general model described in Section 8, and also a revelation principle for the space of mechanisms the seller can employ and a formalization of the disclosure game presented in Section 6.2.

Proof of Lemma 2. First of all, we note that $\beta < 1$, because the upper and lower bounds in Eq. (16) coincide whenever $\beta = 1$, that is

$$\delta(L_{\alpha,1}) = \theta(L_{\alpha,1})\alpha - \int_0^1 s dL_{\alpha,1} \quad \forall \alpha \in [0,1).$$

The argument is constructive, for each $\varepsilon \in [0, 1-\beta]$, define $\alpha^\varepsilon := \alpha + \varepsilon$, $\beta^\varepsilon := \beta + \varepsilon$ and δ^ε as

$$\delta^\varepsilon := \delta + \left(\theta(L_{\alpha^\varepsilon,\beta^\varepsilon}) - \int_0^1 s dL_{\alpha^\varepsilon,\beta^\varepsilon}(s)\right) - \left(\theta(L_{\alpha,\beta}) - \int_0^1 s dL_{\alpha,\beta}(s)\right).$$

Let u^ε be such that Eq. (15) holds for $(u^\varepsilon, \alpha^\varepsilon, \beta^\varepsilon, \delta^\varepsilon)$.

By construction, $1-L_{\alpha^\varepsilon,\beta^\varepsilon}(1) = \beta-\alpha = 1-L_{\alpha,\beta}$ and Eq. (17) holds for all $\varepsilon \in [0, 1-\beta]$; moreover, $\int_0^1 s dL_{\alpha^\varepsilon,\beta^\varepsilon}(s)$ is decreasing in ε. It follows from Eq. (1) that $\theta(L_{\alpha^\varepsilon,\beta^\varepsilon})$ is

[15] For standard fixed information screening models to have bite, typically one party is given the bargaining power to offer contracts (seller) and the other access to some private information (buyer).

increasing in ε. As a result, u^{ε} is increasing in ε, and for each $\varepsilon \in [0, 1-\beta]$,

$$\delta^{\varepsilon} \geqslant \max\left\{0, \theta(L_{\alpha^{\varepsilon},\beta^{\varepsilon}})(1+\alpha-\beta) - \int_0^1 s dL_{\alpha^{\varepsilon},\beta^{\varepsilon}}(s)\right\}.$$

By assumption, $\delta(L_{\alpha,\beta}) > \delta$; therefore, starting from $\varepsilon = 0$ and increasing this number, we must necessarily reach the point at which $\delta^{\varepsilon} = \delta(L_{\alpha^{\varepsilon},\beta^{\varepsilon}})$. This is guaranteed to happen for some $\varepsilon \in [0, 1-\beta]$, because, as shown above, the upper and lower bounds in Eq. (16) coincide whenever $\beta = 1$. □

Proof of Theorem 1. Part 1. We shall show that there exists a unique number $0 \leqslant \alpha^{\star} < F(c)$ which maximizes $V(\alpha)$ defined by Eq. (21). It will follow then from Eq. (20) that $L_{\alpha^{\star},\beta^{\star}}$ for $\beta^{\star} := \beta(\alpha^{\star})$ is the unique buyer's optimal dual threshold reduced test. To establish the claim, we will show that the derivative of $V(\alpha)$ crosses zero at most once and from above. Although the single-peakedness of $V(\alpha)$ is not necessary to prove Theorem 1, it will allows us to precisely pin down $\alpha^{\star}$ and prove the claim of Corollary 1.

Recall that under the dual threshold reduced test $L_{\alpha,\beta}$, the base score is given by $\theta(L_{\alpha,\beta}) = \frac{\int_{\alpha}^{\beta} R(s)ds}{\beta-\alpha}$. Then, Eq. (20) can be re-stated as follows:

$$(1-\alpha)\int_{\alpha}^{\beta(\alpha)}(R(s)-c)ds = (\beta(\alpha)-\alpha)\int_{\beta(\alpha)}^{1}(R(s)-c)ds.$$

The implicit function theorem implies that $\beta(\alpha)$ is differentiable; moreover, its derivative is given by the following expression:

$$\beta'(\alpha) = \frac{(1-\alpha)\big(R(\alpha)-c\big) - \frac{1-\beta(\alpha)}{1-\alpha}\int_{\beta(\alpha)}^{1}(R(s)-c)ds}{(1-\alpha)\big(R(\beta(\alpha))-c\big) + (\beta(\alpha)-\alpha)\big(R(\beta(\alpha))-c\big) - \int_{\beta(\alpha)}^{1}(R(s)-c)ds}, \tag{23}$$

where we used the binding constraint defining $\beta(\alpha)$ to substitute for $\int_{\alpha}^{\beta(\alpha)}(R(s)-c)$. By Eq (20), $\int_{\beta(\alpha)}^{1}(R(s)-c)ds = (1-\alpha)\big(\theta(L_{\alpha,\beta(\alpha)})-c\big)$; thus, since $R(\beta(\alpha)) > \theta(L_{\alpha,\beta(\alpha)} > c$, the denominator is always positive. The function $V(\alpha)$ is clearly differentiable with the derivative $V'(\alpha) = \big(R(\beta(\alpha))-c\big)\beta'(\alpha) - \big(R(\alpha)-c\big)$. Since the denominator in the equation defining $\beta'(\alpha)$ is positive and $R(\beta(\alpha)) - c > 0$, the sign of $V'(\alpha)$ is given by

$$\big(R(\alpha)-c\big)\Delta(\alpha) - \big(R(\beta(\alpha))-c\big)\frac{1-\beta(\alpha)}{1-\alpha}\int_{\beta(\alpha)}^{1}(R(s)-c)ds, \tag{24}$$

$$\text{where } \Delta(\alpha) := \int_{\beta(\alpha)}^{1}(R(s)-c)ds - (\beta(\alpha)-\alpha)\big(R(\beta(\alpha))-c\big).$$

We now unpack the right-hand side of the above equation. First of all, note that $R(\alpha) - c < R(\beta(\alpha)) - c$, and

$$\Delta(\alpha) - \frac{1-\beta(\alpha)}{1-\alpha}\int_{\beta(\alpha)}^{1}(R(s)-c)ds = \int_{\alpha}^{\beta(\alpha)}(R(s)-c)ds - (\beta(\alpha)-\alpha)\big(R(\beta(\alpha))-c\big) < 0,$$

where the last inequality follows from the definition of $\beta(\alpha)$. As a result, $V'(\alpha) < 0$ whenever $R(\alpha) \geqslant c$ and/or $\Delta(\alpha) \geqslant 0$. It follows that the maximizer of $V(\alpha)$, which exists due to continuity and compactness, must be below $F(c)$. So, consider α such that $R(\alpha) < c$, i.e., $\alpha < F(c)$ and $V'(\alpha) \leqslant 0$ with $\Delta(\alpha) < 0$. Let $\alpha' \in (\alpha, F(c))$ with $\Delta(\alpha') < 0$, we want to show that $V'(\alpha') < 0$. Indeed, note that $\beta'(\alpha'') < 0$ for every $\alpha'' < F(c)$, thus $\beta(\alpha') < \beta(\alpha)$ that implies

$$\begin{aligned}&(R(\alpha')-c)\frac{\Delta(\alpha')}{R(\beta(\alpha'))-c} - (R(\alpha)-c)\frac{\Delta(\alpha)}{R(\beta(\alpha))-c} = \\ &= \underbrace{(R(\alpha)-c)}_{<0}\underbrace{\left(\frac{\Delta(\alpha')}{R(\beta(\alpha'))-c} - \frac{\Delta(\alpha)}{R(\beta(\alpha))-c}\right)}_{>0} + \underbrace{\frac{\Delta(\alpha')}{R(\beta(\alpha'))-c}}_{<0}\underbrace{(R(\alpha')-R(\alpha))}_{>0} < 0, \text{ and}\end{aligned}$$

$$\begin{aligned}&\frac{1-\beta(\alpha')}{1-\alpha'}\int_{\beta(\alpha')}^{1}(R(s)-c)ds - \frac{1-\beta(\alpha)}{1-\alpha}\int_{\beta(\alpha)}^{1}(R(s)-c)ds = \\ &= \underbrace{\frac{1-\beta(\alpha)}{1-\alpha}}_{>0}\underbrace{\int_{\beta(\alpha')}^{\beta(\alpha)}(R(s)-c)ds}_{>0} + \underbrace{\int_{\beta(\alpha)}^{1}(R(s)-c)ds}_{>0}\underbrace{\left(\frac{1-\beta(\alpha')}{1-\alpha'} - \frac{1-\beta(\alpha)}{1-\alpha}\right)}_{>0} > 0.\end{aligned}$$

Taking all pieces together, we conclude that $V'(\alpha') < 0$.

Since $V'(\alpha)$ is a continuous function of α, $V'(\alpha)$ can cross zero at most once and from above. As a result, there is a unique threshold $\alpha^\star$ that satisfies the first-order necessary condition, that is

$$V'(\alpha^\star)\begin{cases} = 0 & \text{if} \quad \alpha^\star > 0, \\ \leqslant 0 & \text{if} \quad \alpha^\star = 0. \end{cases}$$

Combining Lemmas 1 - 3, we conclude that $L_{\alpha^\star,\beta^\star}$ for $\beta^\star := \beta(\alpha^\star)$ is buyer's optimal.

It is instructive to express the first-order condition characterizing $\alpha^\star$ in terms of primitives. Substituting for $\beta'(\alpha)$ from Eq. (23) in $V'(\alpha) = \big(R(\beta(\alpha))-c\big)\beta'(\alpha) - \big(R(\alpha)-c\big)$, the reader can verify that first-order condition can be re-written as follows:

$$\overbrace{\frac{\int_{\alpha^\star}^{\beta(\alpha^\star)}(R(s)-c)ds}{\beta(\alpha^\star)-\alpha^\star}}^{\text{average net type in } (R(\alpha^\star),R(\beta(\alpha^\star)))} \cdot \overbrace{\frac{(1-\alpha^\star)(R(\alpha^\star)-c)-(1-\beta(\alpha^\star))(R(\beta(\alpha^\star))-c)}{\beta(\alpha^\star)-\alpha^\star}}^{\text{average net virtual type in } (R(\alpha^\star),R(\beta(\alpha^\star)))}$$

$$\begin{cases} = (R(\alpha^\star)-c)(R(\beta(\alpha^\star))-c) & \text{if} \quad \alpha^\star > 0, \\ \leqslant (R(\alpha^\star)-c)(R(\beta(\alpha^\star))-c) & \text{if} \quad \alpha^\star = 0. \end{cases} \tag{25}$$

Part 2. Recall that the set of buyer's optimal reduced test is defined as solutions to Problem (11). Suppose that a reduced test L is buyer's optimal. Proposition 1 implies that there are unique numbers $0 \leqslant \alpha < \beta \leqslant 1$ such that Eq. (5) - Eq. (7) are satisfied. Then, Lemmas 1 - 3 jointly imply that the buyer's expected payoff under L is at most $V(\alpha)$. By the first part of this proposition, $V(\alpha)$ is uniquely maximized at $\alpha^\star$. It follows that $\alpha = \alpha^\star$, $\beta = \beta^\star = \beta(\alpha^\star)$, which implies that $L(1) = L_{\alpha^\star,\beta(\alpha^\star)}(1)$ and $\int_0^1 sdL(s) = \int_0^1 sdL_{\alpha^\star,\beta(\alpha^\star)}(s)$; in addition, $\delta(L) = \delta(L_{\alpha^\star,\beta(\alpha^\star)})$.

By Eq. (7), we must have

$$\int_s^{\theta(L)} L(t)dt \geqslant \int_s^{\theta(L)} L_{\alpha^\star,\beta^\star}(t)dt = \alpha^\star(\theta(L)-s) \qquad \forall s \in [F^{-1}(\alpha^\star),\theta(L)],$$
$$\int_{\theta(L)}^{s} L(t)dt \leqslant \int_{\theta(L)}^{s} L_{\alpha^\star,\beta^\star}(t)dt = \alpha^\star(s-\theta(L)) \qquad \forall s \in [\theta(L),F^{-1}(\beta^\star)],$$

where the equalities follow from the fact that $L_{\alpha^\star,\beta^\star}$ is constant in the middle region. By taking s to $\theta(L)$ on each side, we conclude that $L(\theta(L)) = \alpha^\star$. Then, since L is non-decreasing, it must be constant throughout the whole middle region, that is $L(s) = \alpha^\star$ for all $s \in [F^{-1}(\alpha^\star),F^{-1}(\beta^\star)]$. As a result, $L(F^{-1}(\alpha^\star)) = L_{\alpha^\star,\beta^\star}(F^{-1}(\alpha^\star)) = \alpha^\star$ and $\int_0^{F^{-1}(\alpha^\star)} L(s)ds = \int_0^{F^{-1}(\alpha^\star)} L_{\alpha^\star,\beta^\star}(s)ds$ which combined with Eq. (7) and Blackwell [1953] precisely means that restricted to the lower interval, i.e., $[0,F^{-1}(\alpha^\star)]$, L is a garbled version of $L_{\alpha^\star,\beta^\star}$. The same argument applies to the upper interval, i.e., $[F^{-1}(\beta^\star),1]$.

Conversely, let L be feasible and satisfy the assumptions of the second part of the proposition. As shown in the first two parts of the proposition, the dual threshold reduced test $L_{\alpha^\star,\beta(\alpha^\star)}$ is buyer's optimal, and it yields him $V(\alpha^\star)$. The fact that L is a garbled version of $L_{\alpha^\star,\beta(\alpha^\star)}$ on each of the two extreme intervals, i.e., $[0,F^{-1}(\alpha^\star)]$ and $[F^{-1}(\beta^\star),1]$, taken together with the fact that $L_{\alpha^\star,\beta(\alpha^\star)}$ is constant in the middle region implies that L is constant in the middle region as well. As a result, $\delta(L) = \delta(L_{\alpha^\star,\beta(\alpha^\star)})$. Since the buyer's expected payoff and the constraint in Problem (11) only depend on L through its average score, certification probability, posterior mean condition on the null signal and $\delta(L)$, L gives the same expected payoff to the buyer and provides the same incentives to trade to the seller as $L_{\alpha^\star,\beta(\alpha^\star)}$. To sum up, L is a solution to Problem (11) as well. $\square$

Proof of Corollary 1. We now establish that the optimal threshold $\alpha^\star$ is zero/positive whenever the production cost is small/large. As shown above, $\alpha^\star = 0$ if and only if the derivative of the objective at zero, i.e., $V'(0)$, is non-positive.

Let $c = 0$. Then, the value of $\beta(0) \in (0,1)$ is a unique root of the following equation: $\frac{\int_0^{\beta(0)} R(s)ds}{\beta(0)} = \int_{\beta(0)}^1 R(s)ds$. It is easy to see that $\beta'(0)$ defined in Eq. (23) is negative; thus, $V'(0) = R(\beta(0))\beta'(0) < 0$. Continuity of $V'(0)$ with respect to c implies $V'(0)$ is still negative for $c > 0$ that is sufficiently close 0.

Next, note that $\beta(0) \to 1$ and $\beta'(0) \to -\frac{-\mathbb{E}[\theta]}{2(1-\mathbb{E}[\theta])}$ as $c \to \mathbb{E}[\theta]$. It follows that $V'(0) \to \frac{1}{2}\mathbb{E}[\theta] > 0$. Continuity of $V'(0)$ with respect to c implies $V'(0)$ is still positive for $c < \mathbb{E}[\theta]$ that is sufficiently close $\mathbb{E}[\theta]$. $\square$

Proof of Corollary 2. By Theorem 1 and Corollary 1, the dual threshold reduced test $L_{0,\beta(0)}$ is buyer's optimal under hard information. Moreover, the second part of this theorem implies that the reduced test L in which the buyer receives his (conditional) average score with certainty attains the same expected payoff, that is

$$L(s) = \begin{cases} 1-\beta(0) & \text{if} \quad s \geqslant \int_0^1 sd\left(\frac{L_{0,\beta(0)}(s)}{1-\beta(0)}\right), \\ 0 & \text{if} \quad s < \int_0^1 sd\left(\frac{L_{0,\beta(0)}(s)}{1-\beta(0)}\right). \end{cases}$$

Given this test, the seller will be incentivized to offer the base of price of $\theta(L)$ even if information is soft. Thus, the buyer's best expected payoff under hard information is weakly lower than under soft information. Since under both hard and soft information with $c = 0$ trade happens with probability 1 (see Theorem 1, Corollary 1 and Proposition 3), the sum of buyer's expected payoff and seller's expected profit equals to $\mathbb{E}[\theta]$. Conclude that the seller weakly prefers the buyer's to learn through hard information. $\square$

Proof of Proposition 4. Part 1. First of all, we note that the buyer's payoff is zero under the completely informative reduced test , whereas the seller obtains the whole expected surplus, that is

$$U(F) = 0, \quad \Pi(F) = \mathbb{E}[\max\{0, \theta - c\}] = \mathbb{E}[\theta] - c + \int_0^c F(s)ds.$$

It follows that the seller's profit under a feasible reduced test can be re-written as

$$\Pi(L) = \Pi(F) - U(L) - \int_0^c (F(s) - L(s))ds,$$

where the last term is loss from imperfect learning — it is non-negative by Condition (7). Now, Let L be a feasible reduced test with $U(L) = 0$. Since $\Pi(F) \geqslant \Pi(L')$ for all feasible reduced tests L', the reduced test L is efficient if and only if $\Pi(F) = \Pi(L)$.

Part 2. To characterize the set of efficient reduced tests that yield a positive expected payoff to the buyer, take $\underline{u} \in \big(0, U(L_{\alpha^\star,\beta(\alpha^\star)})\big]$ and consider the following program, which is similar to Problem (11):

$$\sup_{0\leqslant\alpha<\beta\leqslant 1,\, L} \quad \Pi(L) \quad \text{subject to} \quad U(L) \geqslant \underline{u}, \tag{26}$$
$$\text{Eq. (5) - Eq. (7)}, \quad (\theta(L) - c)(1 - L(1)) \geqslant \int_{\theta(L)}^{1} (s - \theta(L))dL(s).$$

In words, we aim at maximizing the seller's expected profit amongst feasible reduced tests subject to providing the buyer with at least $\underline{u} > 0$. Similarly to the buyer's optimal reduced tests, we will characterize the set of solutions to the above problem in four steps.

Step 1. Take any feasible reduced test L that satisfies the constraints of Problem (26) for certain $0 \leqslant \alpha < \beta \leqslant 1$. In the main text, see Eq. (12), we established that $U(L) = \int_0^1 sdL(s) - \theta(L)L(1) + \delta(L)$, where $\delta(L) = \int_0^{\theta(L)} L(s)ds$. Furthermore, there exists a unique pair of thresholds $\underline{\alpha} \in [0, \alpha]$ and $\overline{\beta} \in [\beta, 1]$ such that the reduced test $L_{\alpha,\beta}^{\underline{\alpha},\overline{\beta}}$, which is defined below, is feasible with $\theta(L) = \theta(L_{\alpha,\beta}^{\underline{\alpha},\overline{\beta}})$ and $\delta(L) = \delta(L_{\alpha,\beta}^{\underline{\alpha},\overline{\beta}})$,

$$L_{\alpha,\beta}^{\underline{\alpha},\overline{\beta}}(s) = \begin{cases} F(s) & \text{if} \quad F(s) < \underline{\alpha}, \\ \underline{\alpha} & \text{if} \quad F(s) \geqslant \underline{\alpha} \ \wedge \ s < \theta(L_{\alpha,\beta}), \\ \overline{\beta} + \alpha - \beta & \text{if} \quad F(s) < \overline{\beta} \ \wedge \ s \geqslant \theta(L_{\alpha,\beta}), \\ F(s) + \alpha - \beta & \text{if} \quad F(s) \geqslant \overline{\beta}. \end{cases}$$

It follows that the buyer's expected payoff under $L_{\alpha,\beta}^{\underline{\alpha},\overline{\beta}}$ is the same as under the original reduced test L. By construction, since $c < \theta(L) = \theta(L_{\alpha,\beta})$, we have that $\int_0^c L_{\alpha,\beta}^{\underline{\alpha},\overline{\beta}}(s) \geqslant \int_0^c L(s)ds$. In other words, the seller weakly prefers $L_{\alpha,\beta}^{\underline{\alpha},\overline{\beta}}$ to L. By construction, $L_{\alpha,\beta}^{\underline{\alpha},\overline{\beta}}$ and $L_{\underline{\alpha},\overline{\beta}}$ have the same certification probability, average score, posterior mean conditional on the null signal, $\delta(L_{\alpha,\beta}^{\underline{\alpha},\overline{\beta}}) = \delta(L_{\underline{\alpha},\overline{\beta}})$ and $\int_0^c L_{\alpha,\beta}^{\underline{\alpha},\overline{\beta}}(s)ds = \int_0^c L_{\underline{\alpha},\overline{\beta}}(s)ds$. To sum up, if there is a solution to Problem (26), then there exists a dual threshold reduced test that yields the same expected payoff-profit pair.

Step 2. Consider a dual threshold reduced test $L_{\alpha,\beta}$. Using the representation from the first part of the proof, the seller's expected profit can be expressed as follows:

$$\Pi(L_{\alpha,\beta}) = \Pi(F) - U(L_{\alpha,\beta}) - \int_{\alpha}^{F(c)\vee\alpha} (c - R(s))ds,$$

where $U(L_{\alpha,\beta}) = \int_\beta^1 (R(s) - \theta(L_{\alpha,\beta}))ds$, $\theta(L_{\alpha,\beta}) = \frac{\int_\alpha^\beta R(s)ds}{\beta-\alpha}$ and R is the inverse of the prior. Recall that $\beta(\alpha)$, see Eq. (20), is the unique threshold for which the seller is indifferent between trading with the buyer who has the null signal and ignoring this buyer. It is easy to see that Problem (26) can be re-stated as

$$\sup_{0\leqslant\alpha<\beta\leqslant 1} \Pi(L_{\alpha,\beta}) \quad \text{subject to} \quad U(L_{\alpha,\beta}) \geqslant \underline{u}, \quad \beta \geqslant \beta(\alpha),$$

We now show that there is always a solution in which both constraint bind. First, suppose that the thresholds (α^0, β^0) are such that $U(L_{\alpha^0,\beta^0}) > \underline{u}$ and $\beta^0 \geqslant \beta(\alpha^0)$. Clearly, we must have $\beta^0 < 1$, because $\underline{u} > 0$. Let $\beta^\varepsilon = \beta + \varepsilon$ for $\varepsilon \geqslant 0$. It is easy to see that the buyer's expected payoff is decreasing, thus the seller's expected profit is increasing, in ε. Starting from $\varepsilon = 0$, we can keep increasing ε and improving the objective in Problem (26) till $U(L_{\alpha^0,\beta^\varepsilon}) = \underline{u}$. Second, suppose that the thresholds (α^0, β^0) are such that $U(L_{\alpha^0,\beta^0}) = \underline{u}$ and $\beta^0 > \beta(\alpha^0)$. Let $\alpha^\varepsilon = \alpha+\varepsilon$ for $\varepsilon \geqslant 0$. Since $U(L_{\alpha,\beta})$ is decreasing in both thresholds, for sufficiently small ε there exists a unique $\beta(\alpha^\varepsilon) \leqslant \beta^\varepsilon < \beta^0$ such that $U(L_{\alpha^\varepsilon,\beta^\varepsilon}) = \underline{u}$. Clearly, the seller's expected profit is non-decreasing in ε; thus, we can (weakly) improve the objective in Problem (26) by increasing ε from 0 to the point at which $\beta^\varepsilon = \beta(\alpha^\varepsilon)$.

Step 3. Consider the family of reduced tests $L_{\alpha,\beta(\alpha)}$, where $\beta(\alpha)$ is chosen according to Eq. (20). In the proof of Proposition 1, we showed that the buyer's expected payoff $V(\alpha)$ as a function of α is maximized at a unique value $\alpha^\star \in [0, F(c))$; moreover, $V(\alpha)$ is decreasing on $[\alpha^\star, 1)$. It follows that for each $\underline{u} \in \big(0, U(L_{\alpha^\star,\beta(\alpha^\star)})\big]$ there exists a unique $\alpha \in [\alpha^\star, 1)$ such that $U(L_{\alpha,\beta(\alpha)}) = \underline{u}$. Since $\alpha^\star < F(c)$, even if there exists another threshold $\alpha' < \alpha^\star$ with $U(L_{\alpha,\beta(\alpha)}) = \underline{u} = U(L_{\alpha',\beta(\alpha')}) = \underline{u}$, the seller would strictly prefer $\alpha \geqslant \alpha^\star$. To sum up, the family of dual threshold reduced tests $L_{\alpha,\beta(\alpha)}$ with $\alpha \in [F(c), 1)$ traces the Pareto frontier (except the seller's optimum).

Step 4. We are finally ready to prove the second part of the proposition. Take an efficient reduced test L with $U(L) > 0$. By definition, this must be a solution to Problem (26) for $\underline{u} = U(L) \in \big(0, U(L_{\alpha^\star,\beta(\alpha^\star)})\big]$. The three previous steps jointly imply that there exists the dual threshold reduced test $L_{\alpha,\beta(\alpha)}$ with $\alpha \geqslant \alpha^\star$ such that $U(L_{\alpha,\beta(\alpha)}) = U(L)$ and $\Pi(L_{\alpha,\beta(\alpha)}) = \Pi(L)$. Conversely, take a feasible reduced test L such that for some $\alpha \geqslant \alpha^\star$, we have $U(L_{\alpha,\beta(\alpha)}) = U(L)$ and $\Pi(L_{\alpha,\beta(\alpha)}) = \Pi(L)$. Since $L_{\alpha,\beta(\alpha)}$ is efficient, L must be efficient as well. $\square$

References

A. Acquisti, C. Taylor, and L. Wagman. The economics of privacy. *Journal of Economic Literature*, 54(2):442–492, 2016.

G. Akerlof. The market for "lemons": Quality uncertainty and the market mechanism. *Quarterly Journal of Economics*, 84(3):488–500, 1970.

S. N. Ali, N. Haghpanah, X. Lin, and R. Siegel. How to sell hard information. *The Quarterly Journal of Economics*, 137(1):619–678, 2022.

K. J. Arrow and G. Debreu. Existence of an equilibrium for a competitive economy. *Econometrica*, 22(3):265–290, 1954.

E. Ben-Porath, E. Dekel, and B. L. Lipman. Disclosure and choice. *The Review of Economic Studies*, 85(3):1471–1501, 2018.

E. Ben-Porath, E. Dekel, and B. L. Lipman. Mechanism design for acquisition of stochastic evidence. Hebrew University, Northwestern University and Boston University, 2021.

D. Bergemann and S. Morris. Information design: a unified perspective. *Journal of Economic Literature*, 57(1):44–95, 2019.

D. Bergemann, B. Brooks, and S. Morris. The limits of price discrimination. *American Economic Review*, 105(3):921–957, 2015.

D. Bergemann, T. Heumann, and S. Morris. Screening with persuasion. *Yale University, Pontificia Universidad Católica de Chile and MIT*, 2022.

D. Blackwell. Equivalent comparisons of experiments. *The Annals of Mathematical Statistics*, 24(2):265–272, 1953.

E. Dekel. On evidence in games and mechanism design. *Econometric Presidential Address based on joint work with Elchanan Ben-Porath and Barton L. Lipman*, 2016.

P. M. DeMarzo, I. Kremer, and A. Skrzypacz. Test design and minimum standards. *American Economic Review*, 109(6):2173–2207, 2019.

R. A. Dye. Strategic accounting choice and the effects of alternative financial reporting requirements. *Journal of Accounting Research*, 23(2):544–574, 1985.

M. Gentzkow and E. Kamenica. A Rothschild-Stiglitz approach to bayesian persuasion. *American Economic Review P&P*, 106(5):597–601, 2016.

S. J. Grossman. The informational role of warranties and private disclosure about product quality. *The Journal of Law & Economics*, 24(3):461–483, 1981.

N. Haghpanah and R. Siegel. The limits of multi-product discrimination. *American Economic Review: Insights, forthcoming*, 2022.

M. Halac, E. Lipnowski, and D. Rappoport. Rank uncertainty in organizations. *American Economic Review*, 111(3):757–786, 2021.

R. Harbaugh and E. Rasmusen. Coarse grades: Informing the public by withholding information. *American Economic Journal: Microeconomics*, 10(1):210–235, 2018.

F. A. Hayek. The use of knowledge in society. *American Economic Review*, 35(4):519–530, 1945.

B. Jovanovic. Truthful disclosure of information. *The Bell Journal of Economics*, 13(1): 36–44, 1982.

W.-O. Jung and Y. K. Kwon. Disclosure when the market is unsure of information endowment of managers. *Journal of Accounting Research*, 26(1):146–153, 1988.

E. Kamenica. Bayesian persuasion and information design. *Annual Review of Economics*, 11:249–272, 2019.

N. Kartik, F. X. Lee, and W. Suen. Information validates the prior: A theorem on bayesian updating and applications. *American Economic Journal: Microeconomics*, 3(2):165–182, 2021.

J.-J. Laffont and D. Martimort. *The theory of incentives: the principal-agent model*. Princeton university press, 2009.

H. Li and X. Shi. Discriminatory information disclosure. *American Economic Review*, 107(11):3363–3385, 2017.

J. M. Liberti and M. A. Petersen. Information: Hard and soft. *The Review of Corporate Finance Studies*, 8(1):1–41, 2018.

A. Lizzeri. Information revelation and certification intermediaries. *The RAND Journal of Economics*, 30(2):214–231, 1999.

A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic Theory*. Oxford University Press, 1995.

L. Mathevet, J. Perego, and I. Taneva. On information design in games. *Journal of Political Economy*, 128(4):1370—1404, 2020.

S. Matthews and A. Postlewaite. Quality testing and disclosure. *The RAND Journal of Economics*, 16(3):328–340, 1985.

P. R. Milgrom. Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics*, 12(2):380–391, 1981.

M. Mussa and S. Rosen. Monopoly and product quality. *Journal of Economic Theory*, 18 (2):301–317, 1978.

R. B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.

D. Ravid, A.-K. Roesler, and B. Szentes. Learning before trading: on the inefficiency of ignoring free information. *Journal of Political Economy*, 130(2):000–000, 2022.

A. K. Roesler and B. Szentes. Buyer-optimal learning and monopoly pricing. *American Economic Review*, 107(7):2072–2080, 2017.

M. Rothschild and J. Stiglitz. Equilibrium in competitive insurance markets: An essay on the economics of imperfect information. *Quarterly Journal of Economics*, 90(4): 629–649, 1976.

S. Schweighofer-Kodritsch and R. Strausz. Principled mechanism design with evidence. *Humboldt-Universität zu Berlin, Institute for Economic Theory*, 2022.

C. Segura-Rodriguez. Selling data. *Available at SSRN 3385500*, 2021.

S. Shavell. Acquisition and disclosure of information prior to sale. *The RAND Journal of Economics*, 25(1):20–36, 1994.

I. Sher and R. Vohra. Price discrimination through communication. *Theoretical Economics*, 10(2):597–648, 2015.

D. Shishkin. Evidence acquisition and voluntary disclosure. University of California, San Diego, 2022.

A. Smolin. Disclosure and pricing of attributes. *Available at SSRN 3318957*, 2020.

M. Spence. Job market signaling. *The Quarterly Journal of Economics*, 87(3):355–374, 1973.

G. J. Stigler. The economics of information. *Journal of Political Economy*, 69(3):213–225, 1961.

H. R. Varian. *Intermediate Microeconomics: A Modern Approach, Ninth Edition*. W.W. Norton and Company, 2020.

R. E. Verrecchia. Discretionary disclosure. *Journal of Accounting and Economics*, 3:179–194, 1983.

D. Wei and B. Green. (reverse) price discrimination with information design. *Available at SSRN 3263898*, 2022.

K. H. Yang. Selling consumer data for profit: Optimal market-segmentation design and its consequences. *American Economic Review*, 112(4):1364–1393, 2022.

11 Online Appendix

11.1 Feasibility

Proposition 1 is still true even if the distribution of buyer's valuation has atoms and/or its support is not connected but includes $\{0, 1\}$. We shall prove the proposition for such general priors, and we will later use this more general characterization in Section 8.

Proof of Proposition 1. We first establish that every dual threshold reduced test is feasible by constructing a specific test that generates a given dual threshold reduced test $L_{\alpha,\beta}$. Then, we will prove the proposition.

Generating dual threshold reduced tests. Take $0 \leqslant \alpha \leqslant \beta \leqslant 1$, and consider the test $(H(\cdot|\theta))_{\theta\in[0,1]}$ in which the buyer receives a perfectly informative score for low/high valuations till the certification probability at the bottom/top reaches α/β, that is $H(s|\theta) = \mathbb{1}_{\{s\geqslant\theta\}} \cdot \lambda_{\alpha,\beta}(\theta)$, where

$$\lambda_{\alpha,\beta}(\theta) := \begin{cases} 1 & \text{if} \quad F(\theta) < \alpha \ \vee \ F(\theta-) > \beta, \\ \frac{\alpha - F(\theta-)}{F(\theta)-F(\theta-)} & \text{if} \quad F(\theta-) \leqslant \alpha \leqslant F(\theta) \ \wedge \ F(\theta) \neq F(\theta-), \\ \frac{F(\theta)-\beta}{F(\theta)-F(\theta-)} & \text{if} \quad F(\theta-) \leqslant \beta \leqslant F(\theta) \ \wedge \ F(\theta) \neq F(\theta-), \\ 0 & \text{if} \quad F(\theta-) > \alpha \ \wedge \ F(\theta) < \beta. \end{cases}$$

By construction, every score generated by this test is an unbiased signal. To guarantee that the buyer has no certificate for intermediate valuations with probability $\beta - \alpha$, randomization might be needed when the prior distribution has mass points.

We claim that the distribution of signals under $(H(\cdot|\theta))_{\theta\in[0,1]}$ is given by the dual threshold reduced test $L_{\alpha,\beta}$ as defined in Eq. (4). Indeed, direct computations yield that

- if $F(s) < \alpha$, then

$$\int_0^1 H(s|\theta)dF(\theta) = \int_0^1 \mathbb{1}_{[\theta,1]}(s)dF(\theta) = F(s) = L_{\alpha,\beta}(s);$$

- if $\alpha \leqslant F(s) < \beta$, then

$$\int_0^1 H(s|\theta)dF(\theta) = \int_{\{\theta:F(\theta)<\alpha\}} dF(\theta) + \left(\alpha - F(F^{-1}(\alpha)-)\right) = \alpha = L_{\alpha,\beta}(s),$$

  where F^{-1} is the generalized inverse defined as $F^{-1}(\tau) := \inf\{\theta \in [0,1] | F(\theta) \geqslant \tau\}$;

- if $F(s) \geqslant \beta$, then

$$\int_0^1 H(s|\theta)dF(\theta) = \alpha + \int_{\{\theta : F(\theta-) > \beta\}} \mathbb{1}_{[\theta,1]}(s)dF(\theta) + \big(F(F^{-1}(\beta+)) - \beta\big) = F(s) + \alpha - \beta = L_{\alpha,\beta}(s).$$

"If" direction. We now are in position to prove the "if" part of the proposition. Consider a reduced test L, and suppose that there are $0 \leqslant \alpha \leqslant \beta \leqslant 1$ such that Eq. (5) - Eq. (7) hold. We need to show that the reduced test L is feasible.

Suppose that $L(1) = 0$, thus $L \equiv 0$. Then, by construction, $L = L_{0,1}$, which is feasible. Next, suppose that $L(1) > 0$. By Eq. (5), $L(1) = L_{\alpha,\beta}(1) = 1 + \alpha - \beta$, therefore Eq. (6) and Eq. (7) imply that $\frac{L(s)}{1+\alpha-\beta}$ is a mean-preserving contraction of $\frac{L_{\alpha,\beta}(s)}{1+\alpha-\beta}$. It is well-known, i.e., see Gentzkow and Kamenica [2016], that there exists a family of distributions $(H(\cdot|\theta))_{\theta\in[0,1]}$ such that

$$L(s) = \int_0^1 H(s|\theta)dL_{\alpha,\beta}(\theta) \ \ \forall s \in [0,1],$$
$$\int_0^s t dL(t) = \int_0^1 \theta H(s|\theta)dL_{\alpha,\beta}(\theta) \ \ \forall s \in [0,1].$$

As seen from the definition of $L_{\alpha\,\beta}$ and $\lambda_{\alpha,\beta}$, we have $L_{\alpha,\beta}(\theta) = \int_0^\theta \lambda_{\alpha,\beta}(s)dF(s)$ for all $\theta \in [0,1]$. Substituting for $L_{\alpha,\beta}(\theta)$ in the two equations above, we see that that the test $(H(\cdot|\theta) \cdot \lambda_{\alpha,\beta}(\theta))_{\theta\in[0,1]}$ generates the reduced test L, i.e., Eq. (2) and Eq. (3) hold.

"Only if" direction. Let L be a feasible reduced test. We shall show that there exist two thresholds $0 \leqslant \alpha \leqslant \beta \leqslant 1$ such that Eq. (5) - Eq. (7) hold.

Suppose that $L(1) = 0$, thus $L \equiv 0$. Recall that $L_{0,1} \equiv 0$, therefore the required conditions hold for $\alpha = 0$ and $\beta = 1$. Finally, suppose that $L(1) > 0$. Since L is feasible, there exists a test H such that Eq. (2) and Eq. (3) hold. By definition, $L(1) = \int_0^1 H(1|\theta)dF(\theta)$ and for each $\theta \in [0,1]$, the function $s \mapsto H(s|\theta) \in [0,1]$ is non-decreasing. We claim that this gives us the following bounds:

$$L_{0,1-L(1)}(s) \leqslant \int_0^s H(1|\theta)dF(\theta) \leqslant L_{L(1),1}(s) \ \ \forall s \in [0,1]. \tag{27}$$

Indeed, for each $s \in [0,1]$, $\int_0^s (1 - H(1|\theta))dF(\theta) \leqslant \int_0^1 (1 - H(1|\theta))dF(\theta)$ which can be re-arranged as follows:

$$F(s) - 1 + L(1) \leq \int_0^s H(1|\theta)dF(\theta).$$

By definition, $L_{0,1-L(1)}(s) = \max\{0, F(s) - 1 + L(1)\}$ which combined with the above inequality implies the lower bound in Eq. (27). The upper bound in Eq. (27) can be shown analogously.

Since Eq. (27) holds for all $s \in [0,1]$, integrating throughout by parts, we obtain

$$\int_0^1 s dL_{0,1-L(1)}(s) ds \geqslant L(1) - \int_0^1 \left(\int_0^s H(1|\theta) dF(\theta) \right) ds \geqslant \int_0^1 s dL_{L(1),1}(s) \tag{28}$$

Using Eq. (4), the reader can verify that for each $0 \leqslant \alpha < \alpha' \leqslant L(1)$, we have $L_{\alpha',1+\alpha'-L(1)}(s) \geqslant L_{\alpha,1+\alpha-L(1)}(s)$ for all $s \in [0,1]$, and the inequality is strict on some subset of $[0,1]$ of positive (Lebesgue) measure. It follows that $\alpha \in [0, L(1)] \mapsto \int_0^1 s dL_{\alpha,1+\alpha-L(1)}(s)$ is increasing, thus there exists a unique number $\alpha \in [0, L(1)]$ such that for $\beta := 1 + \alpha - L(1)$, we have

$$\int_0^1 s dL_{\alpha,\beta}(s) = L(1) - \int_0^1 \left(\int_0^s H(1|\theta) dF(\theta) \right) ds$$

We now show that Eq. (5) - Eq. (7) hold for $\int_0^s H(1|\theta) dF(\theta)$. The first of this equations, namely Eq. (5), is satisfied by construction as $L_{\alpha,\beta}(1) = 1 + \alpha - \beta = L(1) = \int_0^1 H(1|\theta) dF(\theta)$, where the last equality is due to Eq. (2). Using this fact, Eq. (3) and integration by parts, we obtain that

$$L(1) - \int_0^1 \left(\int_0^s H(1|\theta) dF(\theta) \right) ds = \int_0^1 t d \left(\int_0^t H(1|\theta) dF(\theta) \right),$$

which establishes Eq. (6). As for Eq. (7), if $F(s) < \alpha$, then $L_{\alpha,\beta}(s) = L_{L(1),1}(s)$ that implies

$$\int_0^s \left(\int_0^t H(1|\theta) dF(\theta) \right) dt \leqslant \int_0^s L_{\alpha,\beta}(t) dt.$$

Since the average score under $L_{\alpha,\beta}$ is the same as under $\int_0^s H(1|\theta) dF(\theta)$; if $F(s) \geqslant 1 + \alpha - L(1)$, then $L_{\alpha,\beta}(s) = L_{0,L(1)}(s)$ that implies

$$\int_s^1 \left(\int_0^t H(1|\theta) dF(\theta) \right) dt \geqslant \int_s^1 L_{\alpha,\beta}(t) dt.$$

Combining the above inequalities with the fact that $L_{\alpha,\beta}(s)$ is constant for s such that $\alpha \leqslant F(s) < 1 + \alpha - L(1)$, we conclude that Eq. (7) is satisfied for $\int_0^s H(1|\theta) dF(\theta)$.

It remains to show that Eq. (5) - Eq. (7) hold for the original reduced test L. Using the fact that $L(1) = \int_0^1 H(1|\theta)dF(\theta)$, Eq. (2) and Eq. (3) can be re-written as follows:

$$\frac{L(s)}{L(1)} = \int_0^1 \frac{H(s|t)}{H(1|t)} d\left(\frac{\int_0^t H(1|\theta)dF(\theta)}{\int_0^1 H(1|\theta)dF(\theta)}\right) \quad \forall s \in [0,1],$$

$$\int_0^s td\left(\frac{L(t)}{L(1)}\right) = \int_0^1 t\frac{H(s|t)}{H(1|t)} d\left(\frac{\int_0^t H(1|\theta)dF(\theta)}{\int_0^1 H(1|\theta)dF(\theta)}\right) \quad \forall s \in [0,1],$$

where we use the convention that $\frac{0}{0} = 1$. Blackwell [1953], Gentzkow and Kamenica [2016] imply that $\frac{L(s)}{L(1)}$ is a mean-preserving contraction of $\frac{\int_0^s H(1|\theta)dF(\theta)}{\int_0^1 H(1|\theta)dF(\theta)}$; equivalently,

$$\int_0^1 tdL(t) = \int_0^1 td\left(\int_0^t H(1|\theta)dF(\theta)\right) = \int_0^1 tdL_{\alpha,\beta}(t),$$
$$\int_0^s L(t)dt \leqslant \int_0^s \left(\int_0^t H(1|\theta)dF(\theta)\right)dt \leqslant \int_0^s L_{\alpha,\beta}(t)dt \quad \forall s \in [0,1].$$

The first line follows from the definition of these thresholds, and the second line follows from the fact that Eq. (7) is satisfied by $\int_0^s H(1|\theta)dF(\theta)$. □

11.2 Proofs for the general model

Proof of Proposition 5. Part 1. Fix a feasible reduce test L with $1 - L(1) > 0$. Then, the posterior mean conditional on the null signal is well-defined, and it is determined by Eq. (1). As shown in Section 9.4, there is no loss of generality to focus on quantity-price pairs in which the buyer with a score prefers to take the contingent quantity-price pair $(p(s), q(s))$, and the buyer with the null signal takes the base quantity-price pair, that is

$$\theta(L)q - p \geqslant 0, \quad sq(s) - p(s) \geqslant \max\{0, sq - p\} \ \ \forall s \in [0,1].$$

For the given quantities, the seller sets her prices at the highest level so that the above constraints are still satisfied, i.e.,

$$p = \theta(L)q, \quad p(s) = sq(s) - \max\{0, s - \theta(L)\}q \ \ \forall s \in [0,1].$$

Then, the seller's expected profit can be expressed as follows:

$$\int_0^1 (p(s) - C(q(s)))dL(s) + (1 - L(1))(p - C(q)) =$$

$$= \int_0^1 (sq(s) - C(q(s)) - \max\{0, s - \theta(L)\}q)dL(s) + (1 - L(1))\theta(L)q - C(q)) =$$

$$= \int_0^1 (sq(s) - C(q(s)))dL(s) + (1 - L(1))(\nu(L)q - C(q)),$$

where $\nu(L) = \theta(L) - \frac{\int_{\theta(L)}^1 (s-\theta(L))dL(s)}{1-L(1)}$ is the virtual value associated with the null signal. The claim of the proposition follows from pointwise maximization of the above expression.

Part 2. Now, let L be a feasible reduced test in which $1-L(1) = 0$. Since the likelihood of the null signal is zero, there is no value from offering a base price-quantity pair — the seller can extract the whole surplus conditional on $s \neq \emptyset$ by setting $q(s) = Q(s)$, $p(s) = sq(s)$ and choosing (p, q) in a way that the buyer with a certificate will never prefer the this base price-quantity pair, i.e., any (p, q) such that $q - p \leqslant 1$ will do.

□

Proof of Proposition 2. Part 1. Since $\iota(0) = 0$, the buyer's expected payoff under the completely informative reduced test is zero; on the other hand, the seller obtains the whole expected surplus. Then, we note that any feasible reduced test L with $1 - L(1) = 0$ yields the same payoff-profit pair as the completely informative reduced test. As a result, if L is efficient, then F must be efficient as well.

Part 2. Take a feasible reduced test L with $1 - L(1) > 0$. Since L is feasible, it must satisfy Conditions (5) - (7) for some $0 \leqslant \alpha < \beta \leqslant 1$. We now construct a dual threshold test that is weakly better for the buyer than L. If $\delta(L) = \delta(L_{\alpha,\beta})$, then $L_{\alpha,\beta}$ is a dual threshold test which yields the same payoff as L. If $\delta(L) < \delta(L_{\alpha,\beta})$, then we must have $\beta < 1$. Repeating the construction used in Lemma 2, define $(\alpha^0, \beta^0) := (\alpha, \beta)$, and $\alpha^\varepsilon := \alpha + \varepsilon$, $\beta^\varepsilon := \beta + \varepsilon$ for $\varepsilon \geqslant 0$. Moreover, let δ^ε be given by

$$\delta^\varepsilon := \delta + \left(\theta(L_{\alpha^\varepsilon,\beta^\varepsilon}) - \int_0^1 s dL_{\alpha^\varepsilon,\beta^\varepsilon}(s)\right) - \left(\theta(L_{\alpha,\beta}) - \int_0^1 s dL_{\alpha,\beta}(s)\right).$$

Then, for every $\varepsilon \in [0, 1 - \beta]$ with $\delta^\varepsilon \leqslant \delta(L_{\alpha^\varepsilon,\beta^\varepsilon})$, there exists a feasible reduced test L^ε such that the virtual value under L^ε is the same as under L, i.e., $\nu(L^\varepsilon) = \nu(L)$. Note that the buyer's information rents $\int_{\theta(L^{\varepsilon})}^{1}(s-\theta(L^{\varepsilon}))dL^{\varepsilon}(s)$ is an increasing function of ε which contradicts optimality of L. We conclude that if L is buyer's optimal then $\delta(L)=\delta(L_{\alpha,\beta})$ that implies that $L_{\alpha,\beta}$ is buyer's optimal as well.

Part 3. Take an efficient reduced test L with $1-L(1)>0$. Since L is feasible, it must satisfy Conditions (5) - (7) for some $0\leqslant\alpha<\beta\leqslant 1$. Moreover, as shown in the main text, there exists a unique pair of thresholds $\underline{\alpha}\in[0,\alpha]$ and $\overline{\beta}\in[\beta,1]$ such that the reduced test $L_{\alpha,\beta}^{\underline{\alpha},\overline{\beta}}$, which is defined in Eq. (14), is feasible with $\theta(L)=\theta(L_{\alpha,\beta}^{\underline{\alpha},\overline{\beta}})$ and $\delta(L)=\delta(L_{\alpha,\beta}^{\underline{\alpha},\overline{\beta}})$.

Below we show that the reduced test $L_{\alpha,\beta}^{\underline{\alpha},\overline{\beta}}$ is more (Blackwell) informative than L, i.e. $\int\limits_0^s L_{\alpha,\beta}^{\underline{\alpha},\overline{\beta}}(t)\geqslant\int\limits_0^s L(t)dt$ for all $s\in[0,1]$. By construction, $L_{\alpha,\beta}^{\underline{\alpha},\overline{\beta}}(s)=L_{\alpha,\beta}(s)$ for $F(s)<\underline{\alpha}$ and $L_{\alpha,\beta}^{\underline{\alpha},\overline{\beta}}(s)$ is constant for $F(s)\geqslant\underline{\alpha}$ and $s<\theta(L_{\alpha,\beta})$. Then, $\int_0^{\theta(L_{\alpha,\beta})}L_{\alpha,\beta}^{\underline{\alpha},\overline{\beta}}(s)ds=\int_0^{\theta(L_{\alpha,\beta})}L_{\alpha,\beta}ds$ and the fact $L_{\alpha,\beta}$ is more informative than L jointly imply that $\int\limits_0^s L_{\alpha,\beta}^{\underline{\alpha},\overline{\beta}}(t)\geqslant\int\limits_0^s L(t)dt$ holds for $s<\theta(L_{\alpha,\beta})$. The argument is identical for $s\geqslant\theta(L_{\alpha,\beta})$, because the average scores under $L_{\alpha,\beta}^{\underline{\alpha},\overline{\beta}}$ and L coincide.

Therefore, by convexity of the surplus function $W(s)$ we have that

$$\int_0^1 W(s)dL_{\alpha,\beta}^{\underline{\alpha},\overline{\beta}}(s)\geqslant\int_0^1 W(s)dL(s).$$

Since $L_{\alpha,\beta}^{\underline{\alpha},\overline{\beta}}$ and L yield the same payoff to the buyer, $L_{\alpha,\beta}^{\underline{\alpha},\overline{\beta}}$ is (weakly) preferred by the seller to L. It follows that $L_{\alpha,\beta}^{\underline{\alpha},\overline{\beta}}$ is efficient as well. □

Proof of Example 3. We shall show Problem (26) always has a solution that belongs to the family of dual threshold reduced test. The claim is trivial for $\underline{u}=0$, because the completely informative reduced test necessarily maximizes the seller's expected profit. And, since certification is costless such test yields zero to the buyer. Moreover, the result has already been established for the buyer's maximal payoff, see Theorem 1. So, in what follows we consider $\underline{u}\in\left(0,U(L_{\alpha^\star,\beta(\alpha^\star)}\right)$.

By Proposition 2, the family of tests $L_{\alpha,\beta}^{\underline{\alpha},\overline{\beta}}$ can trace the Pareto frontier. Clearly, $\theta(L_{\alpha,\beta}^{\underline{\alpha},\overline{\beta}})=\frac{\alpha+\beta}{2}$; thus, $L_{\alpha,\beta}^{\underline{\alpha},\overline{\beta}}$ is feasible if and only if $\underline{\alpha}+\overline{\beta}=\alpha+\beta$. Direct computation gives that the buyer's information rent is $\frac{1}{2}(1-\underline{\alpha})(1-\overline{\beta})$. Importantly, both the score conditional on the null signal and buyer's information rent only depend on $\underline{\alpha}$ and $\overline{\beta}$. The other two threshold, namely α and β, enter the seller's profit and buyer's payoff only through the virtual value $\nu(L_{\alpha,\beta}^{\underline{\alpha},\overline{\beta}})=\frac{\underline{\alpha}+\overline{\beta}}{2}-\frac{\frac{1}{2}(1-\underline{\alpha})(1-\overline{\beta})}{\beta-\alpha}$ and probability of the null signal $\beta-\alpha$. Letting $\kappa:=\beta-\alpha\in[0,\overline{\beta}-\underline{\alpha}]$, Problem (26) in our setting can be re-written as

$$\sup_{0\leqslant\underline{\alpha}\leqslant\overline{\beta}\leqslant 1,\,\kappa} \int_0^{\underline{\alpha}} \left(\frac{s^2}{2}\right) ds + \int_{\overline{\beta}}^{1} \left(\frac{s^2}{2}\right) ds + (\overline{\beta}-\underline{\alpha})\frac{(\underline{\alpha}+\overline{\beta})^2}{8} - \frac{\kappa}{8}\left((\underline{\alpha}+\overline{\beta})^2 - \left((\underline{\alpha}+\overline{\beta}) - \frac{(1-\underline{\alpha})(1-\overline{\beta})}{\kappa}\right)^2\right)$$

$$\text{subject to } \left(\underline{\alpha}+\overline{\beta} - \frac{(1-\underline{\alpha})(1-\overline{\beta})}{\kappa}\right)\cdot(1-\underline{\alpha})(1-\overline{\beta}) \geqslant 4\underline{u},$$

$$0 \leqslant \kappa \leqslant \overline{\beta}-\underline{\alpha} \text{ and } \underline{\alpha}+\overline{\beta} - \frac{(1-\underline{\alpha})(1-\overline{\beta})}{\kappa} \geqslant 0.$$

It is easy to see that the above problem admits a maximizer such that $\overline{\beta} < 1$ and $\kappa > 0$; moreover, the first constraint binds at the optimum, because the objective is decreasing in κ, whereas the buyer's payoff is increasing. Solve for κ from the binding constraint and substitute it into the objective to obtain the following equivalent formulation:

$$\sup_{0\leqslant\underline{\alpha}<\overline{\beta}\leqslant 1} \int_0^{\underline{\alpha}} \left(\frac{s^2}{2}\right) ds + \int_{\overline{\beta}}^{1} \left(\frac{s^2}{2}\right) ds + (\overline{\beta}-\underline{\alpha})\frac{(\underline{\alpha}+\overline{\beta})^2}{8} - (\underline{\alpha}+\overline{\beta})\frac{(1-\underline{\alpha})(1-\overline{\beta})}{8} - \frac{\underline{u}}{2} \tag{29}$$

$$\text{subject to } \left(\underline{\alpha}+\overline{\beta} - \frac{(1-\underline{\alpha})(1-\overline{\beta})}{\overline{\beta}-\underline{\alpha}}\right)\cdot(1-\underline{\alpha})(1-\overline{\beta}) \geqslant 4\underline{u}.$$

Clearly, at the optimum the constraint in (29) binds if and only if a dual reduced threshold test is efficient for this given value of $\underline{u}$. By the way of contradiction, suppose that $\underline{\alpha}$ and $\overline{\beta}$ solves Problem (29) and the constraint is slack. As discussed above, we necessarily have $\underline{\alpha} < \overline{\beta} < 1$. The reader can verify that the second derivative of the objective with respect to $\overline{\beta}$ equals to $\frac{1-\overline{\beta}}{4} > 0$, which contradicts optimality of $\underline{\alpha}$, $\overline{\beta}$ with the constraint being slack. □

11.3 The game without commitment

In section 6.2 we argued that when the seller does not have commitment power, for any equilibrium of the ensuing disclosure game there exists another equilibrium in which the buyer discloses nothing and is weakly better off. We provide a formal proof of this claim below.

Proof of Corollary 3. Consider an equilibrium of the disclosure game discussed in the text, that is i) the buyer's disclosure strategy $d^* : [0,1] \to [0,1]$ that maps his score to a disclosure probability and ii) the seller's selling strategy that consists of the base price p^* and score-contingent prices $(p^*(s))_{s\in[0,1]}$. As in the setting with commitment, we shall break ties in favor of the buyer whenever the seller is indifferent, First of all, since the seller cannot commit, upon observing a score $s \in [0, 1]$, she optimally sets $p^*(s) = s$ thereby extracting the whole surplus. As a response, the buyer with a score $s > p^*$ will never find it optimal to disclose, i.e., $d^*(s) = 0$. The buyer's with a score $s < p^*$ is indifferent between disclosing his certificate and concealing it but buying the good.

Taking into an account the buyer's response, the seller's ex-ante profit $\pi^*(p)$ as a function of the base price p alone is given by

$$\pi^*(p) := p\Big(L(1) - L(p-) + (1 - L(1)) \cdot \mathbb{1}_{[0,\theta(L)]}(p) - \int_{[p,p^*]} d(s)dL(s)\Big).$$

By definition, $\pi^*(p^*) \geqslant \pi^*(p)$ for all $p \in [0, 1]$.

We now show that there always exists a different equilibrium in which the buyer never discloses his certificate and the seller's sets a weakly lower base price. Let $\pi^{**}(p)$ be the seller's profit from offering the base price p whenever the buyer's always conceals his certificate:

$$\pi^{**}(p) := p\Big(L(1) - L(p-) + (1 - L(1)) \cdot \mathbb{1}_{[0,\theta(L)]}(p)\Big).$$

By construction, the function π^{**} is upper-semi continuous; thus, the set of its maximizer is non-empty and compact. Let p^{**} be the smallest such maximizer. Examination of π^{**} shows that it coincides with π^* for all prices above p^*, thus $\pi^{**}(p^*) \geqslant \pi^{**}(p)$ for all $p \in [p^*, 1]$. It follows that p^{**} is lower p^*.

To sum up, in our search for buyer optimal equilibria and tests, it is without loss to consider equilibria in which a certificate is never shown, i.e., information is always soft. By Roesler and Szentes [2017]'s arguments we get the same buyer's optimal reduced tests as in Proposition 3. □

11.4 Space of mechanisms

Consider the general setting of Section 8, we shall show that offering price-quantity pairs is without loss of generality for the seller. Formally, the seller commits to a *mechanism* that specifies a message space M and, for each message $m \in M$, the following objects: a probability of the seller requesting a certificate $\chi(m)$, a quantity-price given that a certificate is not requested $(q(0|m), p(0|m))$, a quantity-price given that a certificate is requested $(q(1, \emptyset|m), p(1, \emptyset|m))$ but not shown and a quantity-price pair given that a certificate is requested and shown $(q(1, s|m), p(1, s|m))_{s\in[0,1]}$.

For the given informational endowment, which is specified by L, and mechanism, the buyer is facing a simple decision problem in which for every signal $s \in [0, 1] \cup \{\emptyset\}$ he first selects a message $m \in M$ to report. Then, if a certificate is not requested, the buyer can take $(q(0|m), p(0|m))$ or walk away. Similarly, if it is requested, then the buyer with the null signal can take $(q(1, \emptyset|m), p(1, \emptyset|m))$ or walk away, the buyer with a score s can take $(q(1, s|m), p(1, s|m))$ in addition to the previous two options.

Take the buyer's strategy that maximizes his expected payoff. Then, we define a new mechanism in which the message space is $[0, 1] \cup \{\emptyset\}$, $\chi(s) = 1$ for all $s \in [0, 1]$ and $\chi(\emptyset) = 0$. If $1 - L(1) > 0$, then let $(q(0|\emptyset), p(0|\emptyset))$ be the expected quantity-price pair received by the buyer with the null signal under the original mechanism and his optimal strategy; otherwise, let this quantity-price pair be as in the original mechanism. Similarly, for each $s \in [0, 1]$, let $(q(0, s|s), p(0, s|s))$ be the expected quantity-price pair received by the buyer with the signal s under the previous mechanism and his optimal strategy. The remaining quantity-price pairs are such that the quantity is zero and price is large enough so that the buyer will never take it. By construction, under the new mechanism the buyer's optimal strategy is to report her signal and take the trade that is designed for him, i.e., the buyer with a score $s \in [0, 1]$ receives $(q(0, s|s), p(0, s|s))$, that is

$$sq(0, s|s) - p(0, s|s) \geqslant \max\{0, sq(0|\emptyset) - p(0|\emptyset)\}.$$

Similarly, if $1 - L(1) > 0$, then $\theta(L)$ is well-defined, and the buyer with the null signal receives a non-negative expected payoff from taking $(q(0|\emptyset), p(0|\emptyset))$,

$$\theta(L)q(0|\emptyset) - p(0|\emptyset) \geqslant 0.$$

We note that the new mechanism delivers each buyer's type the expected quantity-price pair which he obtains under the first mechanism. The seller's expected profit is then weakly higher under the new mechanism, because her production cost is convex C.