Certification in Search Markets

Teddy Mekonnen[*] Zeky Murra-Anton[†] Bobak Pakzad-Hurson[‡]

April 17, 2022

Abstract

We consider a firm seeking to fill a single vacancy by searching over a sequence of workers who are ex-post differentiated in their productivity but are ex-ante identical. Prior to its hiring decision, the firm may acquire information about a worker's productivity by paying an intermediary to certify the worker. The intermediary, through the certification tests and fees it offers, affects how much surplus is generated in each period as well as how long the firm searches. We characterize the intermediary's profit-maximizing spot-contract and show that the contract (*i*) induces efficient hiring standards, (*ii*) extracts the full-surplus, and (*iii*) strings along the firm, i.e., keeps the firm searching for longer than the firm would like. We also consider the case in which the worker pays the certification fees and show that (*iv*) the intermediary may be able to create a demand for certification, even when the certificate conveys little to no information, and (*v*) the worker benefits when disclosure of test results is mandatory.

JEL Classifications: D42, D83, L12, L15

Keywords: search market, labor market, information design, certification

[*]Department of Economics, Brown University. Contact: mekonnen@brown.edu

[†]Department of Economics, Brown University. Contact: zeky_ahmed_murra_anton@brown.edu

[‡]Department of Economics, Brown University. Contact: bph@brown.edu

1. Introduction

Search and matching markets often contend with frictions stemming from incomplete information. For example, a firm seeking to fill a vacancy may be uncertain about a candidate's productivity, even after conducting an interview. In many cases, the candidates themselves may not know their own productivity at any one firm. Since the lack of information could lead to inefficient matches and search outcomes, it is unsurprising that information intermediaries are frequently utilized to reduce or resolve the inherent uncertainty in search markets.

One service that intermediaries provide is *certification*, which offers firms a signal about workers' productivity or ability. Yet, the goals of a firm or a worker—filling a vacancy by hiring a productive worker or finding employment—likely differ from those of an intermediary—maximizing profit from selling certification tests. Therefore, to understand the economic value intermediaries generate in a market, we must first study how the incentives of the information suppliers (intermediaries) interact with the incentives of the information consumers (firms and workers).

As a motivating example, consider the Society of Actuaries (SOA) which certifies actuaries during the early stages of their careers. Certification involves passing seven examinations sequentially, and typically requires as many years to complete. SOA publishes the names of those who pass each exam, and (continued) passage of exams is typically required for employment. Early in the sequence, examination fees are typically paid by workers, and later examination fees are typically paid by employers; fees increase by roughly four fold from the early examinations to the later ones, despite nearly identical numbers of workers taking each examination and the overall pass rate.[1] Total examination revenue in 2019 exceed $33 million, and account for nearly 60% of revenues for the SOA.[2] For comparison, this figure accounts for over 20 times the annual dues collected by the American Economic Association, a similarly sized entity.[3]

This market raises several questions about how an intermediary optimally certifies workers: How much information should certification convey? How much should the intermediary charge for its services? Does the information conveyed by a certificate

[1]See https://www.soa.org/Education/Exam-Req/Syllabus-Study-Materials/Exam-and-Module-Fees.aspx and https://www.soa.org/education/general-info/exam-results/edu-exam-results-detail/.
[2]See https://theactuarymagazine.org/2019-soa-annual-report/2018-2019-financial-results/.
[3]See https://www.aeaweb.org/content/file?id=15511.

change if a worker pays for it instead of the firm? How should certification results be communicated to an interested employer? We study these questions in this paper within an information design framework.

We present a model with a long-lived intermediary, a long-lived firm, and a series of short-lived workers who are ex-post differentiated in their productivity but are ex-ante identical. The firm considers one worker per period, and exits the market upon hiring a worker. Prior to its hiring decision, the firm can solicit information about a worker's productivity by paying an information intermediary to certify the worker. We assume that the intermediary lacks commitment power and can only offer a spot contract in each period.

The intermediary has offsetting desires. On the one hand, it can charge higher prices by offering more informative signals of worker's productivity. On the other hand, providing more informative signals leads to faster matching, which implies shorter duration of rent extraction.

Despite these offsetting desires, our main result shows that the intermediary can indeed induce the efficient hiring standards while also extracting the surplus generated in the search market. The equilibrium we characterize is stationary in that the intermediary offers the same certification test and price in each period and after every history. Thus, our main result suggests that in our setting, the intermediary neither gains from considering more complex history-dependent contracts nor has any value for commitment.

Interestingly, the optimal per-period price the intermediary charges could be a small fraction of the maximal surplus. While doing so lowers the amount of rent the intermediary can extract in any one period, the prospect of future low prices inflates the firm's continuation value, thereby incentivizing the firm to keep searching, and importantly, keep paying the intermediary. Additionally, the intermediary forces the firm to adopt a higher standard of hiring than the firm would prefer by pooling low productivity workers with some of the intermediate productivity workers that the firm would have hired if it could distinguish them. Consequently, the firm remains in the search market for a long time, allowing the intermediary to extract the maximum surplus through drips and drabs.

We next consider the case in which the worker pays the certification fee. Once again, the intermediary has competing interests. On the one hand, it can charge higher prices by offering tests that increase the probability of employment. On the other hand, such tests lead to faster matching, which implies shorter duration of rent extraction. Nonetheless, we show that the first interest dominates; in a stationary equilibrium, the intermediary designs a test that maximizes the probability the firm hires the worker. Consequently, the hiring standard is inefficiently low when the worker pays the certification fee. We show that in any stationary equilibrium, the intermediary is either shut out of the market, or provides just enough information to make a certificate a requirement for employment.

We further distinguish between a mandatory disclosure setting—the worker's test results are directly disclosed to the firm—and a voluntary disclosure setting—the worker can disclose her results at her discretion. We show that if the firm prefers to never hire a worker absent information, then the stationary equilibrium outcomes under mandatory and voluntary disclosure are the same. In contrast, if the firm prefers to hire a worker absent information, then workers prefer mandatory disclosure because it is credible for a worker to approach a firm without a certificate and claim to have never been tested.

While we approach the intermediary's profit-maximization problem within the framework of information design, the search market setting we consider introduces inter-temporal considerations. In particular, the firm's optimal choice in the current period depends not only on its current posterior belief but also on its continuation value which is tied to the distribution of posterior beliefs it will entertain in the future. Thus, the intermediary must consider how future certification tests on the equilibrium path affect both the firm's or worker's decision to solicit information (the extensive margin) and the firm's decision to hire a worker (the intensive margin) in the current period. Due to these inter-temporal considerations, the intermediary's profit maximization problem in a stationary equilibrium is formulated not as the maximization of a linear functional on a convex set of distributions, as is the case in most of the information design literature, but as the maximization of a non-linear functional.

Related Literature

Coming soon

2. Model

We consider a search market comprised of a long-lived firm seeking to fill a single vacancy, a sequence of short-lived workers seeking employment, and a long-lived information intermediary. In each period $t = 0, 1, 2\ldots$ (as long as the position remains vacant), a new worker is "born" and interviews with the firm. The worker's productivity, denoted by $\theta_t \in \mathbb{R}$, is distributed according to a continuous cumulative distribution function (CDF) F with compact support $[\underline{\theta}, \bar{\theta}] \triangleq \Theta$, where $\bar{\theta} > 0$, and mean $m_{\varnothing}$.[4] No one in the market directly observes a worker's productivity. Thus, workers are ex-post differentiated but are ex-ante identical.

Before the firm makes its hiring decision, it may solicit additional information by having the information intermediary test the worker. The intermediary has limited commitment and can only offer the firm a spot contract which consists of a price $p_t \geq 0$ and a certification test $\pi_t : \Theta \to \Delta(S)$, where S is a set of possible test scores. If the firm accepts the contract ($a_t = 1$), the intermediary administers the test on the worker. The firm pays p_t to the intermediary and observes the worker's test score $s_t \in S$. If the firm instead rejects the contract ($a_t = 0$), it observes no additional information.

Finally, the firm decides whether to hire the worker. If the firm hires the worker ($d_t = 1$), the game ends and the firm's payoff from the match is the worker's productivity. If the firm instead rejects the worker ($d_t = 0$), the period-t worker perishes and the game continues on to $t+1$ with the firm earning a period-t payoff of zero. We treat workers as non-strategic for the moment and abstract away from their payoffs until Section 5. The firm and the intermediary are both risk-neutral and share a common discount factor $\delta \in (0, 1)$.

Our goal is to characterize the intermediary's profit maximizing contract. We focus on stationary equilibrium outcomes, which are characterized by a contract proposal strategy for the intermediary, a contract acceptance and a hiring strategy for the firm, and a belief-updating process such that (*i*) strategies are sequentially rational, (*ii*) beliefs are derived by Bayes rule, and (*iii*) strategies are history-independent—a stronger requirement than stationarity on the equilibrium path. We provide a more formal definition after we define the contract space.

[4]We assume that F is continuous for ease of exposition. Our results extend to the case where F is a discrete distribution, but some care is required in extending our proofs.

	$a_t = 1$	$a_t = 0$
$d_t = 1$	$\theta_t - p_t, p_t$	$\theta_t, 0$
$d_t = 0$	$-p_t, p_t$	$0, 0$

Table 1: Ex-post period-t payoffs for the firm and the intermediary respectively.

3. Preliminary Analysis

Before we characterize stationary equilibrium outcomes, we first formalize the space of contracts, and characterize the firm's search problem.

3.1. Simplifying certification tests

Suppose the intermediary proposes a period-t contract (p_t, π_t). If the firm rejects the contract and observes nothing, it gets an expected payoff of $\mathbb{E}[\theta_t] = m_\varnothing$ by hiring the worker. In contrast, if the firm accepts the contract and observes a test score $s_t \in S$, it updates its expectation of the worker's productivity by Bayes rule to a posterior mean $m_t = \mathbb{E}[\theta_t | s_t]$, and gets an expected payoff of $m_t - p_t$ by hiring the worker.

In any given period, the intermediary is unrestricted in the certification tests that it can design. However, because a certification test π_t impacts the period-t payoffs only through the posterior mean, it suffices for our purposes to only consider the distribution over posterior means that is induced by a test.[5] Formally, each certification test π induces a CDF $G : \Theta \to [0, 1]$ over posterior means, with G being a mean-preserving contraction of F, i.e., $\int_x^{\bar{\theta}} G(m)dm \geq \int_x^{\bar{\theta}} F(m)dm$ for all $x \in \Theta$ with equality at $x = \underline{\theta}$. For example, a test that fully reveals a worker's productivity induces the distribution F while an uninformative test induces the degenerate distribution $G_\varnothing$ given by

$$G_\varnothing(m) = \begin{cases} 0 & \text{if} \quad m < m_\varnothing \\ 1 & \text{if} \quad m \geq m_\varnothing \end{cases}.$$

Let $\mathcal{G}$ be the set of distributions that are mean-preserving contractions of F. **?**

[5]This approach to information design was developed by Gentzkow and Kamenica (2016) and further extended by Dworczak and Martini (2019) and Kolotilin (2018).

and Rothschild and Stiglitz (1970) show that any certification test induces a distribution over posterior means in $\mathcal{G}$, and any distribution in $\mathcal{G}$ can be induced by some certification test. Thus, without loss of generality, we assume that for each period t, the intermediary can propose any contract of the form (p_t, G_t) from the set $\mathbb{R}_+ \times \mathcal{G}$.

Before we discuss the firm's search problem, we introduce a function $c_G : \Theta \to \mathbb{R}$ given by

$$\begin{aligned} c_G(x) &= \int_x^{\bar{\theta}} (m - x) dG(m) \\ &= \int_x^{\bar{\theta}} (1 - G(m)) dm, \end{aligned}$$

where the last equality follows from integration by parts. The value $c_G(x)$ represents the firm's added value from continuing to search when it gets a payoff of x from an "outside option".[6] Because the function c_G is a useful analytical tool, we provide details on its properties. For any $G \in \mathcal{G}$,

(*a*) c_G is continuous, weakly decreasing, and convex,

(*b*) $c_G(\underline{\theta}) = m_\varnothing - \underline{\theta} > 0 = c_G(\bar{\theta}) = 0$, and

(*c*) c_G is right differentiable with $\partial_+ c_G(\theta) = G(\theta) - 1$.

Furthermore, G is a mean-preserving spread of G' if and only if $c_G \geq c_{G'}$. Hence, $c_{G_\varnothing} \leq c_G \leq c_F$ for all $G \in \mathcal{G}$.

Let $\mathcal{C}$ be the set of weakly decreasing, continuous, and convex functions from Θ to $\mathbb{R}$ such that any $c \in \mathcal{C}$ is bounded below by $c_{G_\varnothing}$ and bounded above by c_F. Thus, $c_G \in \mathcal{C}$ for all $G \in \mathcal{G}$. Additionally, for each $c \in \mathcal{C}$, there exists a distribution $G \in \mathcal{G}$ such that $c = c_G$ (Proposition 1, Gentzkow and Kamenica (2016)). Therefore, we can represent a certification test not only as a distribution over posterior means but also as a continuous, convex, and decreasing function from Θ to $\mathbb{R}$.

[6]The function c_G is often used in characterizing the optimal stopping rule in search models. For example, it is used in **?** to characterize the reservation wage in job search markets, in **?** to characterize reservation values when searching over alternatives, and in Wolinsky (1986) to characterize optimal consumer search behavior. **?** refer to c_G as the *incremental-benefit function*.

3.2. Firm's search problem

We first consider the firm's search problem when it faces a null-contract $(0, G_\varnothing)$ in every period. This is akin to rejecting any contract the intermediary proposes in every period. When the firm has no information and $m_\varnothing > 0$, the firm hires the worker and gets a payoff of $m_\varnothing$. In contrast, when the firm has no information and $m_\varnothing < 0$, the firm never hires the worker and gets a payoff of 0. Therefore, the firm can guarantee itself a payoff of $\underline{u} = \max\{0, m_\varnothing\}$, which we refer to as the firm's *autarky payoff.*

We next consider the firm's search problem when it faces a contract $(0, G)$ in every period for some $G \in \mathcal{G}$. Let $u(G)$ be the firm's expected payoff in this setting and let $r(G) = \delta u(G)$ denote the firm's reservation value. Given some posterior mean $m \in \Theta$, the firm hires the worker if $m > r(G)$ and continues searching if $m < r(G)$. Hence, the firm's payoff $u(G)$ is given by

$$u(G) = r(G) + \int_\Theta \max\{m - r(G), 0\} dG(m). \tag{1}$$

If $r(G) < \underline{\theta}$, the firm would hire even the least productive worker; in this case, the firm's expected payoff and reservation value would both independent of G with $u(G) = m_\varnothing$ and $r(G) = \delta m_\varnothing$. Thus, if $\delta m_\varnothing < \underline{\theta}$, the firm never engages in active search for any $G \in \mathcal{G}$, making the contracting problem trivial. We therefore make the following assumption:

Assumption 1 $\underline{\theta} < \delta m_\varnothing$.

Given Assumption 1, $r(G)$ is in the open interval $(\underline{\theta}, \bar{\theta})$ for any $G \in \mathcal{G}$. Additionally, from (1), $r(G)$ can be characterized as the unique solution to the fixed point (as a function of r)

$$r\left(\frac{1-\delta}{\delta}\right) = c_G(r). \tag{2}$$

The fixed point problem in (2) is particularly amenable to a geometric representation as shown in Figure 1. Equivalently, $u(G)$ can be characterized as the unique solution to the fixed point problem (as a function of u)

$$u = \frac{c_G(\delta u)}{1-\delta}. \tag{3}$$

Since $c_{G'} \leq c_G$ whenever G' is a mean-preserving contraction of G, the firm's highest payoff is the full-information payoff $u(F) \triangleq \bar{u}$, which coincides with the maximum surplus that can be generated from search, and its lowest payoff is the no-information payoff $u(G_\varnothing)$, which coincides with the autarky payoff $\underline{u}$. Consequently, $r(G) \in [\delta\underline{u}, \delta\bar{u}]$ for all $G \in \mathcal{G}$, with the socially efficient outcome acheieved by the firm hiring any worker whose true productivity exceeds $r(F) = \delta\bar{u}$.

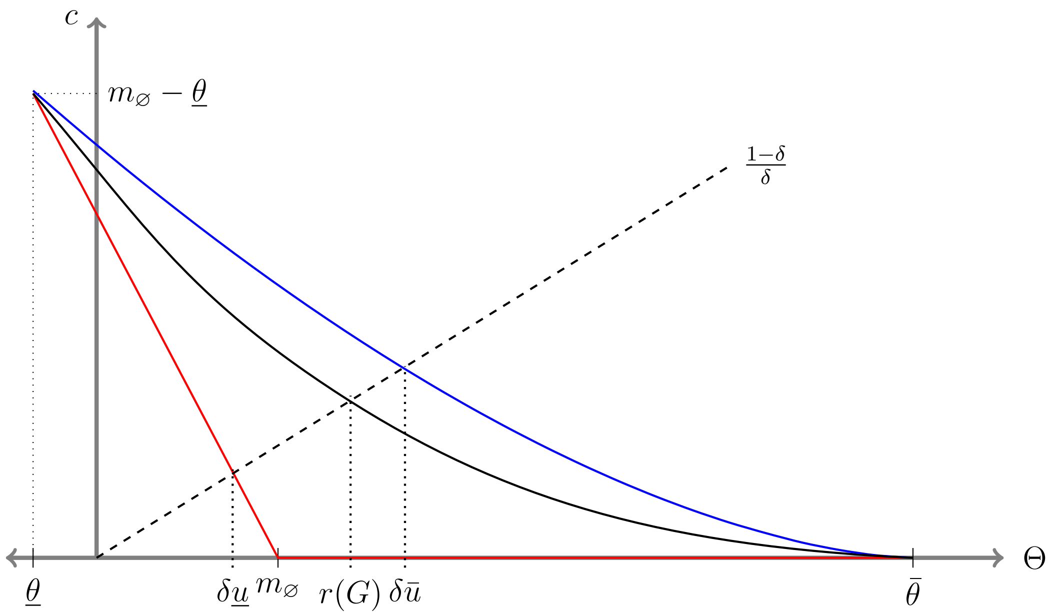

Figure 1: Example with $\underline{\theta} < 0 < m_\varnothing < \bar{\theta}$ so that $\underline{u} = m_\varnothing$. The dashed line has a slope of $(1-\delta)/\delta$. The red curve is $c_{G_\varnothing}$, the black curve is c_G for some arbitrary $G \in \mathcal{G}$, and the blue curve is c_F. The intersection of each respective curve with the dashed line represents the solution to the fixed point problem in (2) with $r(G_\varnothing) = \delta\underline{u}$ and $r(F) = \delta\bar{u}$.

Finally, we consider a problem in which the firm faces a contract $(p, G) \in \mathbb{R}_+ \times \mathcal{G}$ in the current period and some exogenously given continuation value of $U \in [\underline{u}, \bar{u}]$. The firm finds it optimal to hire the worker whenever $m > \delta U$ and keeps searching whenever $m \leq \delta U$.[7] If the firm accepts the contract, it gets a payoff of

$$\delta U + \int_\Theta \max\{m - \delta U, 0\} dG(m) - p$$

[7] Technically, the firm is indifferent when $m = \delta U$ but the intermediary prefers for the firm to continue searching because hiring a worker ends the game.

$$=\delta U + c_G(\delta U) - p.$$

If it rejects the contract, the firm gets a payoff of

$$\begin{aligned}&\delta U + \max\{m_\varnothing - \delta U, 0\}\\=&\delta U + c_{G_\varnothing}(\delta U).\end{aligned}$$

Thus, it is optimal for the firm to accept the contract whenever $c_G(\delta U) - c_{G_\varnothing}(\delta U) \geq p$.

4. Stationary Contracts

In any equilibrium, the firm's payoff cannot be lower than its autarky payoff $\underline{u}$ and cannot exceed the maximum surplus $\bar{u}$. Similarly, the intermediary's payoff cannot be lower than zero and cannot exceed $\bar{u} - \underline{u}$, which is equivalent to extracting the firm's maximal value of search.

Since all players are risk neutral and the space of contracts $\mathbb{R}_+ \times \mathcal{G}$ is convex, the intermediary does not gain from randomizing over spot contracts. To guarantee equilibrium existence, we also assume that whenever the firm is indifferent between any two actions, it chooses the one that maximizes the intermediary's payoff. Therefore, it suffices for our purposes to only consider pure strategies. Additionally, since the firm rejecting a contract is payoff-equivalent to the intermediary instead proposing the null-contract $(0, G_\varnothing)$ and the firm accepting, we can, without loss of generality, restrict attention to contracts that induce the firm to accept in equilibrium.

Given the simplifications above, a stationary equilibrium is defined by a contract $(p, G) \in \mathbb{R}_+ \times \mathcal{G}$, a continuation value for the firm $U \in [\underline{u}, \bar{u}]$, and a continuation value for the intermediary $V \in [0, \bar{u} - \underline{u}]$ such that

- (i) Given any worker with expected productivity $m \in \Theta$, the firm hires if and only if $m > \delta U$.

- (ii) Given any contract $(\widehat{p}, \widehat{G}) \in \mathbb{R}_+ \times \mathcal{G}$ the intermediary has proposed, the firm accepts if and only if

$$c_{\widehat{G}}(\delta U) - c_{G_\varnothing}(\delta U) \geq \widehat{p}. \tag{PC}$$

(*iii*) It is optimal for the intermediary to propose contract (p, G), i.e.,

$$(p, G) \in \underset{(\widehat{p}, \widehat{G}) \in \mathbb{R}_+ \times \mathcal{G}}{\arg\max} \; \widehat{p} + \widehat{G}(\delta U)\delta V \quad \text{subject to (PC)}. \tag{4}$$

(*iv*) Payoffs are self-generating, i.e.,

$$U = \delta U + c_G(\delta U) - p \tag{5}$$

and

$$V = p + G(\delta U)\delta V. \tag{6}$$

In words, in each period of a stationary equilibrium, the firm and the intermediary anticipate the continuation values U and V, respectively. Sequential rationality then implies that the firm hires any worker whose expected productivity exceeds the firm's discounted continuation payoff. Additionally, it is optimal for the firm to accept any contract that gives a higher payoff than $\max\{m_{\varnothing}, \delta U\}$, which is the payoff the firm gets either by hiring the worker in the current period without soliciting information or by waiting for the next period. Given the firm's strategies, the intermediary then proposes a contract that maximizes the sum of its per-period revenue and, *conditional on the firm continuing to search*, its discounted continuation value. Finally, the anticipated continuation values must themselves be generated from the firm's and intermediary's stationary equilibrium strategies.

To gain some intutition of the economic forces at play, it is instructive to separate the *per-period trade-offs*—trade-offs the intermediary faces when proposing a contract in the current period while holding contracts in subsequent periods fixed—from the *continuation trade-offs*—the trade-offs the intermediary faces when considering different future contracts.

Let us first consider the per-period trade-offs. Suppose the intermediary proposes the contract (p, G) in periods $t + 1$ and onward. In period t, each contract $(\widehat{p}, \widehat{G})$ that the intermediary proposes has an impact on the firm's search problem along an extensive and an intensive margin. First, the contract affects whether the firm even wants to solicit information from the intermediary. This impact on the extensive margin is captured by the participation constraint (PC). Second, conditional on the firm soliciting information, the contract affects the probability with which the firm hires in period t. This impact on the intensive margin of the firm's search problem is captured by $\widehat{G}(\delta U)$. Thus, holding U and V fixed, the intermediary proposes a period-t contract that maximizes its profits by appropriately tailoring the contract's impact on the firm along the extensive and intensive margins.

Next, let us consider the continuation trade-offs. Suppose the intermediary instead considers proposing a contract (p', G) in period $t+1$ and onward with $p' > p$. From (5) and (6), assuming the firm accepts the new contract, we can compute the associated continuation values U' and V' for the firm and intermediary, respectively. A higher price p' has a direct positive effect on the intermediary's continuation payoff V'. However, a higher price also has a direct negative effect on the firm's continuation value U', which lowers the threshold a worker must clear to be hired. Thus, an increase in p' has an indirect negative effect on the intermediary's continuation value by decreasing the probability that the firm continues searching. Thus, the intermediary must trade-off having high prices and a short contractual relationship with the firm against having low prices and a long contractual relationship. A similar, albeit less transparent, trade-off exists when changing the certification test. For example, proposing a contract (p, G') instead of (p, G) in period $t+1$ and onward could change the firm's continuation value from U to U' when G and G' differ in their informativeness. Thus, the intermediary's continuation value is both directly and indirectly affected as $G(\delta U)$ changes to $G'(\delta U')$.

When the contract that resolves the per-period trade-offs coincides with the contract that resolves the continuation trade-offs, then we have in fact found a stationary equilibrium.

In order to state our main result, we first define a distribution G^* over posterior means given by

$$G^*(m) = \begin{cases} 0 & \text{if} \quad m < \mathbb{E}_F[\theta | \theta \leq \delta \bar{u}] \\ F(\delta \bar{u}) & \text{if} \quad \mathbb{E}_F[\theta | \theta \leq \delta \bar{u}] \leq m < \mathbb{E}_F[\theta | \theta > \delta \bar{u}] \\ 1 & \text{if} \quad m \geq \mathbb{E}_F[\theta | \theta > \delta \bar{u}] \end{cases},$$

which is induced by a "pass-fail" test that passes any worker whose true productivity is $\theta > \delta \bar{u}$ and fails any worker whose true productivity is $\theta \leq \delta \bar{u}$; see Figure 3.

Proposition 1 *A contract $(p, G) \in \mathbb{R}_+ \times \mathcal{G}$ along with a pair of continuation values $(U, V) \in [\underline{u}, \bar{u}] \times [0, \bar{u} - \underline{u}]$ constitute a stationary equilibrium if and only if*

(i) $U = \underline{u}$,

(ii) $V = \bar{u} - \underline{u}$,

(iii) $p = (\bar{u} - \underline{u})\big(1 - \delta F(\delta \bar{u})\big)$,

(iv) $G(\delta \underline{u}) = G^*(\delta \underline{u})$, *and*

(v) $c_G(x) \geq c_{G^*}(x)$ *for all* $x \in \Theta$ *with equality at* $x = \delta \underline{u}$.

The first implication of Proposition 1 is that there exists a stationary equilibrium and it is essentially unique in the sense that all stationary equilibria are payoff equivalent. The equilibrium can be sustained by the intermediary charging a price $p = (\bar{u} - \underline{u})(1 - F(\delta \bar{u}))$ for a simple "pass-fail" test that certifies workers if and only if their true productivity exceeds $\delta \bar{u}$, which is the socially efficient hiring threshold. Consequently, the equilibrium contracts yields the maximal surplus in the search market.

In equilibrium, the intermediary extracts the maximal value of information $\bar{u} - \underline{u}$ from the firm. However, it does not do so by charging an extractive price in each period—the per-period price p is strictly less than $\bar{u} - \underline{u}$. Instead, the intermediary imposes a higher standard of hiring on the firm by withholding information. In particular, the firm would be willing to hire any worker whose true productivity is above $\delta \underline{u}$ but the intermediary pools the desirable workers whose productivity lies in the interval $[\delta \underline{u}, \delta \bar{u}]$ with undesirable workers whose productivity is lower than $\delta \underline{u}$ by failing all of them. This forces the firm to search for longer than it otherwise would like, allowing the firm to extract the maximal value of information through drips and drabs.

Proof. Suppose a tuple (p, G, U, V) constitutes a stationary equilibrium. Since the objective function in (4) is increasing in the price, the constraint (PC) must bind at the maximum, i.e., $p = c_G(\delta U) - c_{G_\varnothing}(\delta U)$. The firm's continuation value in (5) then becomes

$$U = \frac{c_{G_\varnothing}(\delta U)}{1 - \delta},$$

which is the same fixed point problem in (3) with a unique solution $U = \underline{u}$.

The intermediary's continuation value can now be written as

$$V = \max_{\widehat{G} \in \mathcal{G}} \frac{c_{\widehat{G}}(\delta \underline{u}) - c_{G_\varnothing}(\delta \underline{u})}{1 - \delta \widehat{G}(\delta \underline{u})}. \tag{6'}$$

Notice that $V > 0$ because the objective function in (6′) is strictly positive when evaluated at $\widehat{G} = F$, and G is a solution to (6′) because (p, G) is a stationary equilibrium contract. Furthermore, any $G' \in \mathcal{G}$ with $G'(\delta \underline{u}) = G(\delta \underline{u})$ and $c_{G'}(\delta \underline{u}) = c_G(\delta \underline{u})$ is also a solution to (6′).

There is no stationary equilibrium in which the firm either always or never hires a worker in each period. To see why, first assume to the contrary that the firm always hires a worker, i.e., $G(\delta \underline{u}) = 0$, which implies that $\partial_+ c_G(\delta \underline{u}) = -1$. Since c_G is a decreasing and convex function, it must be the case that $\partial_+ c_G(\theta) = -1$ for all $\theta \leq \delta \underline{u}$. Equivalently, $G(\theta) = 0$ for all $\theta \leq \delta \underline{u}$. Hence,

$$\begin{aligned} c_G(\delta \underline{u}) &= c_G(\underline{\theta}) - \int_{\underline{\theta}}^{\delta \underline{u}} 1 - G(m) dm \\ &\leq c_{G_\varnothing}(\underline{\theta}) - \int_{\underline{\theta}}^{\delta \underline{u}} 1 - G_\varnothing(m) dm \\ &= c_{G_\varnothing}(\delta \underline{u}), \end{aligned}$$

where the inequality follows from the fact that $c_G(\underline{\theta}) = c_{G_\varnothing}(\underline{\theta})$ and $0 = G(\theta) \leq G_\varnothing(\theta)$ for all $\theta \leq \delta \underline{u}$. However, $c_G \geq c_{G_\varnothing}$, which implies that $c_G(\delta \underline{u}) = c_{G_\varnothing}(\delta \underline{u})$. Consequently, the intermediary's payoff would be $V = 0$, which establishes a contradiction.

Next, assume that the firm never hires a worker, i.e., $G(\delta \underline{u}) = 1$, which implies that $\partial_+ c_G(\delta \underline{u}) = 0$. Since c_G is a decreasing and convex function, it must be the case that $\partial_+ c_G(\theta) = 0$ for all $\theta \geq \delta \underline{u}$. Equivalently, $G(\theta) = 1$ for all $\theta \geq \delta \underline{u}$. Therefore $c_G(\delta \underline{u}) = c_{G_\varnothing}(\delta \underline{u}) = 0$, and once again, the intermediary's payoff would be $V = 0$, establishing a contradiction. Thus, $G(\delta \underline{u}) \in (0, 1)$ and the expectations $\mathbb{E}_G[\theta | \theta \leq \delta \underline{u}]$ and $\mathbb{E}_G[\theta | \theta > \delta \underline{u}]$ are well-defined.

Let us define a new distribution $G^* \in \mathcal{G}$ with

$$G^*(m) = \begin{cases} 0 & \text{if} \quad m < \mathbb{E}_G[\theta|\theta \leq \delta\underline{u}] \\ G(\delta\underline{u}) & \text{if} \quad \mathbb{E}_G[\theta|\theta \leq \delta\underline{u}] \leq m < \mathbb{E}_G[\theta|\theta > \delta\underline{u}] \\ 1 & \text{if} \quad m \geq \mathbb{E}_G[\theta|\theta > \delta\underline{u}] \end{cases}.$$

The distribution G^* is *binary* as its support only contains two posterior means. By definition, $G^*(\delta\underline{u}) = G(\delta\underline{u})$ and $c_G \geq c_{G^*}$ with

$$\begin{aligned} c_{G^*}(\delta\underline{u}) &= \Big(\mathbb{E}_G[\theta|\theta > \delta\underline{u}] - \delta\underline{u}\Big)(1 - G(\delta u)) \\ &= \int_{\delta\underline{u}}^{\bar{\theta}} (m - \delta\underline{u})dG(m) \\ &= c_G(\delta\underline{u}). \end{aligned}$$

Thus, G^* is also a solution to (6′), and

$$\begin{aligned} V &= \frac{c_{G^*}(\delta\underline{u}) - c_{G_\varnothing}(\delta\underline{u})}{1 - \delta G^*(\delta\underline{u})} \\ &= \frac{c_{G^*}(r(G^*)) + \displaystyle\int_{\delta\underline{u}}^{r(G^*)} (1 - G^*(m))dm - c_{G_\varnothing}(\delta\underline{u})}{1 - \delta G^*(\delta\underline{u})} \\ &= \frac{r(G^*)\left(\frac{1-\delta}{\delta}\right) + \Big(r(G^*) - \delta\underline{u}\Big)\big(1 - G(\delta\underline{u})\big)dm - \underline{u}(1-\delta)}{1 - \delta G^*(\delta\underline{u})} \\ &= \frac{r(G^*) - \delta\underline{u}}{\delta}, \end{aligned}$$

where the second equality follows from breaking up the limits of the integral in $c_{G^*}(\delta\underline{u})$, and the third equality follows from computing $c_{G_\varnothing}(\delta\underline{u})$ and the fact that $r(G^*)$ is the unique reservation value that solves (2) when the distribution is G^*. However, $r(G) \leq r(F) = \delta\bar{u}$ for all $G \in \mathcal{G}$. Since G^* is a solution to (6′), it must be that $r(G^*) = r(F) = \delta\bar{u}$, and $V = \bar{u} - \underline{u}$.

Note that $r(G^*) = r(F)$ if and only if $c_{G^*}(\delta\bar{u}) = c_F(\delta\bar{u})$. Moreover, c_{G^*} must be tangent to c_F at $\delta\bar{u}$ because $c_F \geq c_{G^*}$. We conclude the proof by characterizing the unique $G^* \in \mathcal{G}$ such that (i) G^* is binary, and (ii) c_{G^*} is tangent to c_F at $\delta\bar{u}$, as shown in Figure 2.

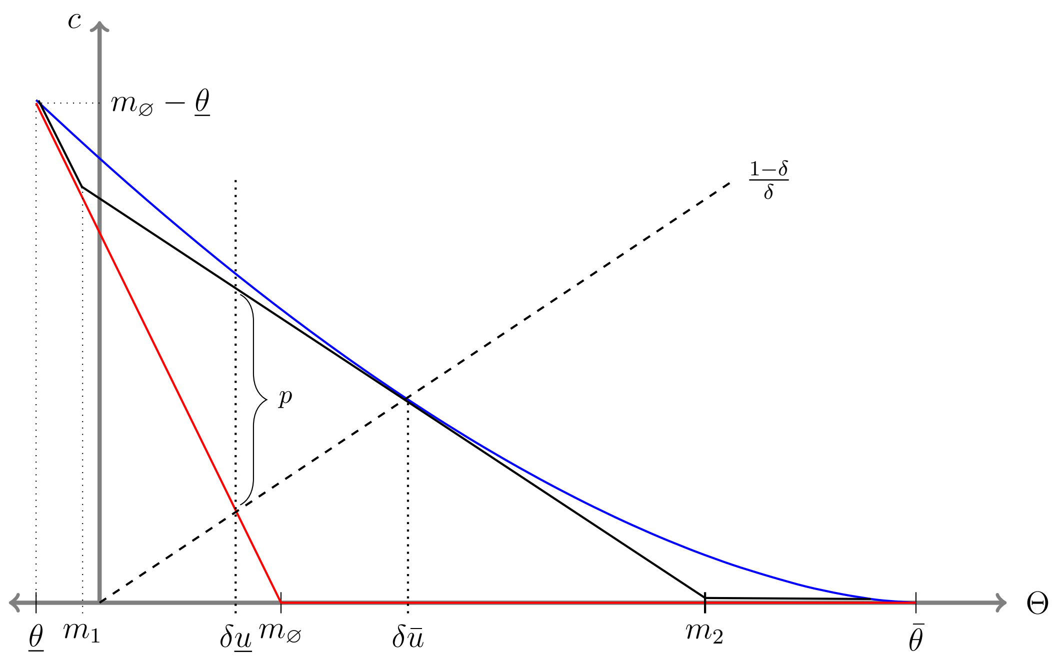

Figure 2: The dashed line has a slope of $(1-\delta)/\delta$. The red curve is $c_{G_\varnothing}$, the black curve is c_{G^*} for the optimal pass-fail $G^* \in \mathcal{G}$, and the blue curve is c_F. The optimality condition is that c_{G^*} is tangent to c_F at $\delta\bar{u}$. The optimal price the intermediary charges for G^* is given by p, the difference between c_{G^*} and $c_{G_\varnothing}$ at $\delta\underline{u}$.

By definition, the support of G^* is $\{m_1, m_2\} \subset \Theta$ with $m_1 < \delta\underline{u} < \delta\bar{u} < m_2$. Thus, c_{G^*} is smooth on the open interval (m_1, m_2), and c_F is smooth on Θ given our assumption that F is absolutely continuous. Hence, tangency of c_{G^*} and c_F at $\delta\bar{u}$ implies that $\partial_+ c_{G^*}(\delta\bar{u}) = \partial_+ c_F(\delta\bar{u}) = F(\delta\bar{u}) - 1$, and we can conclude that $G^*(m_1) = G^*(\delta\underline{u}) = G^*(\delta\bar{u}) = F(\delta\bar{u})$. Additionally,

$$
\begin{aligned}
& c_{G^*}(\delta\bar{u}) = c_F(\delta\bar{u}) \\
\Leftrightarrow\ & (m_2 - \delta\bar{u})\big(1 - F(\delta\bar{u})\big) = \int_{\delta\bar{u}}^{\bar{\theta}} (m - \delta\bar{u}) dF(m) \\
\Leftrightarrow\ & m_2 = \mathbb{E}_F[\theta | \theta > \delta\bar{u}],
\end{aligned}
$$

and

$$m_1 G^*(m_1) + m_2(1 - G^*(m_1)) = m_\varnothing,$$

implies that $m_1 = \mathbb{E}_F[\theta|\theta \leq \delta\bar{u}]$. Finally, $p = c_{G^*}(\delta\underline{u}) - c_{G_\varnothing}(\delta\underline{u}) = (\bar{u} - \underline{u})(1 - \delta F(\delta\bar{u}))$. ■

While the price and the continuation values in a stationary equilibrium are uniquely pinned down, the pass-fail test that induces G^* is merely the least informative equilibrium test. For example, a lower-censorship test, which fully reveals a worker's productivity if it exceeds $\delta\bar{u}$ and pools workers in the interval $[\underline{\theta}, \delta\bar{u}]$ by giving them a "fail" grade, is also an equilibrium test because it induces a distribution

$$G^{LC}(m) = \begin{cases} 0 & \text{if} \quad m < \mathbb{E}_F[\theta|\theta \leq \delta\bar{u}] \\ F(\delta\bar{u}) & \text{if} \quad \mathbb{E}_F[\theta|\theta \leq \delta\bar{u}] \leq m < \delta\bar{u} \\ F(m) & \text{if} \quad m \geq \delta\bar{u} \end{cases}$$

which satisfies the last two conditions of Proposition 1. Figure 3 compares G^* to G^{LC}. There are in fact an uncountable number of tests that can sustain a stationary equilibrium. Nevertheless, a stationary equilibrium test must yield the same distribution of scores for workers whose productivity falls in the interval $[\delta\underline{u}, \delta\bar{u}]$. In other words, an equilibrium test that is more informative than G^* either reveals information about workers whose productivity is high enough that the firm would have hired them even with a pass-fail test, e.g., lower-censorship test, or reveals information about workers whose productivity is low enough that the firm would not have hired them even with a pass-fail test.

Corollary 1 *If $G \in \mathcal{G}$ is an equilibrium test, then $G(m) = F(\delta\bar{u})$ for all $m \in [\delta\underline{u}, \delta\bar{u}]$.*

Proof. By Proposition 1, if the tuple (p, G, U, V) constitutes a stationary equilibrium, then $G(\delta\underline{u}) = G^*(\delta\underline{u}) = F(\delta\bar{u})$. Hence, $\partial_+ c_G(\delta\underline{u}) = F(\delta\bar{u}) - 1$. Furthermore, $c_F \geq c_G \geq c_{G^*}$ with equality at $\delta\bar{u}$, and $\partial_+ c_F(\delta\bar{u}) = \partial_+ c_{G^*}(\delta\bar{u}) = F(\delta\bar{u}) - 1$. Thus, $\partial_+ c_G(\delta\bar{u}) = F(\delta\bar{u}) - 1$. Since c_G is convex, this implies that $\partial_+ c_G(x) = F(\delta\bar{u}) - 1$ for all $x \in [\delta\underline{u}, \delta\bar{u}]$. In other words, $G(x) = F(\delta\bar{u})$ for all $x \in [\delta\underline{u}, \delta\bar{u}]$. ■

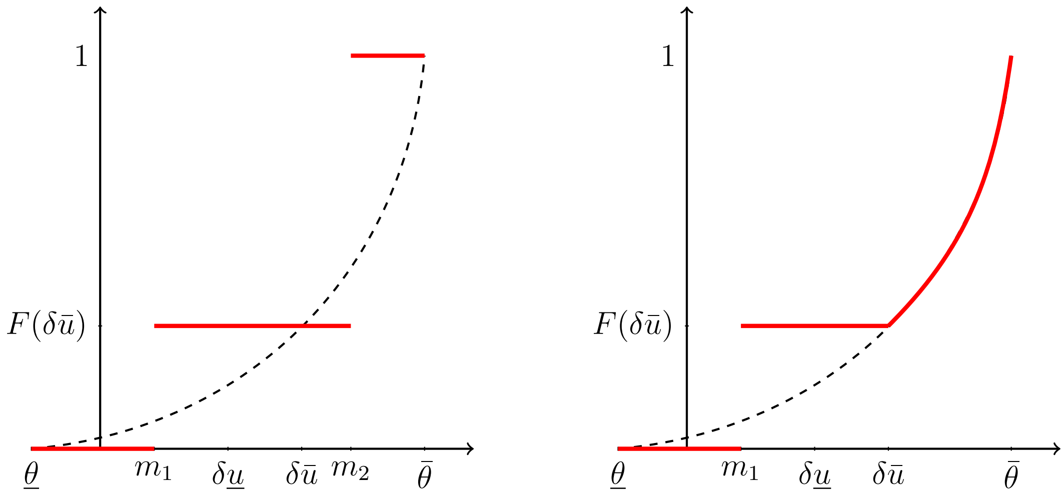

(a) G^* induced by optimal pass-fail.

(b) G induced by optimal lower-censorship.

Figure 3: The dashed curve corresponds to an arbitrary prior distribution F, the red curves correspond to two different tests that can sustain a stationary equilibrium, and $m_1 = \mathbb{E}_F[\theta|\theta \leq \delta\bar{u}]$ and $m_2 = \mathbb{E}_F[\theta|\theta > \delta\bar{u}]$. Note that both tests are constant on $[\delta\underline{u}, \delta\bar{u}]$.

We conclude this section by commenting on stationarity and our assumption of limited commitment. Plainly, restricting attention to stationary strategies is limiting as the intermediary could have opted for a non-stationary strategy that conditions the current contract proposal on past contracts, on the firm's past choices, and on the scores of previous workers. Yet, we know from Proposition 1 that restricting attention to stationary strategies is without loss of optimality, as the intermediary is able to sustain the highest possible equilibrium payoffs through stationary strategies alone.

Similarly, an intermediary with long-term commitment power can sustain a larger set of contracts in equilibrium. For example, the intermediary could propose a sequence of contracts $(p_t, G_t)_{t\geq 0}$ such that $G_t = F$ for all $t \geq 0$, $p_t = 0$ for all $t \geq 1$, and $p_0 = \bar{u} - \underline{u}$, which is akin to an intermediary that "sells" its testing technology to the firm in period $t = 0$, thereby extracting the maximal value of information at once. Such a sequence of contracts cannot be implemented without commitment; the intermediary would want to deviate away from charging $p_t = 0$ in periods $t \geq 1$, and thus, the firm would never accept the contract (p_0, F) in the initial period. Yet, by Proposition 1, the intermediary has no value for commitment as it is able to sustain the highest possible equilibrium payoffs via spot contracts alone.

5. Contracting with Workers

Coming Soon

6. Conclusion

Coming Soon

References

Dworczak, P. and Martini, G. . The simple economics of optimal persuasion. *Journal of Political Economy*, 127(5):1993–2048, 2019.

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

Kolotilin, A. . Optimal information disclosure: A linear programming approach. *Theoretical Economics*, 13(2):607–635, 2018.

Rothschild, M. and Stiglitz, J. E. . Increasing risk: I. a definition. *Journal of Economic theory*, 2(3):225–243, 1970.

Wolinsky, A. . True monopolistic competition as a result of imperfect information. *The Quarterly Journal of Economics*, 101(3):493–511, 1986.