

Market Signaling with Grades*

Brendan Daley

Duke University
The Fuqua School of Business
bd28@duke.edu
<http://faculty.fuqua.duke.edu/~bd28>

Brett Green

Northwestern University
Kellogg School of Management
b-green@kellogg.northwestern.edu
www.kellogg.northwestern.edu/faculty/green

October 27, 2011

Abstract

We consider a market signaling model in which receivers observe both the sender's costly signal as well as a stochastic grade that is correlated with the sender's type. In equilibrium, the sender resolves the trade-off between using the costly signal versus relying on the noisy grade to distinguish himself. When grades are sufficiently informative, separating equilibria do not survive stability-based refinements. The prediction depends on the prior—e.g., with two types the equilibrium involves full-pooling when the prior puts sufficient weight on the high type and partial pooling otherwise. Further, the equilibrium converges to the complete-information outcome as the prior tends to a degenerate one—resolving a long-standing paradox within the signaling literature.

JEL Classification: D82, D83, D41, C72

Keywords: Signaling, Asymmetric Information, Information Economics, Refinements

*The authors are indebted to Andrzej Skrzypacz, Paul Milgrom, and Jeremy Bulow for their continual guidance. We also thank James Anton, Terrence August, Giuseppe Lopomo, Leslie Marx, David McAdams, Marc Meredith, Michael Ostrovsky, Tomasz Sadzik, Yuval Salant, Larry Samuelson, Ron Siegel, Bruno Strulovici, Robert Wilson, and seminar participants at Stanford, Duke and Washington Universities and Yahoo! Research for useful conversations and comments. Brendan gratefully acknowledges the support of the Koret Graduate Dissertation Fellowship through a grant to the Stanford Institute for Policy Research. Brett gratefully acknowledges the support of the SIEPR Dissertation Fellowship.

1 Introduction

Signaling models typically assume observable, costly actions are the only channels that can convey information about the sender. These models implicitly overlook the existence of *grades*.¹ A *grade* refers to any imperfect public message about the sender's type. Grades are prevalent in many signaling environments such as education, financial markets, advertising, and warranties. For an example that amends the canonical signaling story (Spence, 1973), potential employers observe not only how long a student attends school at his chosen course of study (his costly signal), but also his G.P.A. (his grade).

In this paper, we study how strategic agents behave when both channels are available for information transmission. Specifically, we consider a signaling environment in which receivers (also referred to as *the market*) observe a stochastic grade in addition to the signal chosen by the sender. The likelihood of each grade depends on the sender's privately-known type and (potentially) on his chosen signal. After observing the chosen signal and the realized grade, receivers undertake actions in response.

Separating equilibria have been the primary focus of signaling models since their inception. Indeed, in most (gradeless) signaling models, the least cost separating equilibrium (LCSE) is the unique stable outcome (Cho and Kreps, 1987; Cho and Sobel, 1990; Ramey, 1996). That grades will alter equilibrium behavior is not an immediate implication of their presence. One might be tempted to reason that if types are perfectly separating based on observable actions, then any additional noisy information will be ignored and the prediction should remain unchanged. We show that this reasoning is incorrect; the very presence of informative grades renders pooling equilibria more plausible. More specifically, we demonstrate that separating equilibria do not survive stability-based refinements when the grade-generating technology, known as the *test*, is sufficiently informative. The intuition is that if the sender were to choose a very costly signal, attempting to convince the market he is a high type, receivers would infer that he is trying to minimize the importance of the test. When the test is informative, lower types have more incentive to do this. A costly signal backfires by showing that the sender is too eager to de-emphasize the results of the test. No sender wants to spend resources to convince the market he is a low type, so separating equilibria do not survive. Thus, the presence of a potentially informative test leads to (at least) some degree of pooling at signals less costly than in the LCSE.

When types pool in equilibrium, the grade conveys relevant information about the sender's type to the market. Better grades indicate a greater likelihood of high value. Therefore,

¹See Riley (2001) for an extensive survey. Exceptions include Weiss (1983), Fang (2001), Feltovich et al. (2002), Kremer and Skrzypacz (2007), Angeletos et al. (2006); Angeletos and Pavan (2011), Alos-Ferrer and Prat (2011), and Daley and Green (2011), see Section 6.1 for further discussion.

higher types expect more favorable responses than the lower types they pool with, in contrast to pooling equilibria in gradeless models. In short, grades convey meaningful information to the market, affect payoffs and improve welfare in equilibrium.

To formally demonstrate these results, we first analyze a canonical application of our model: job-market signaling (Spence, 1973). In conveying his private information, the high type faces a trade-off. He must decide how much to exploit his cost advantage and how much to rely on his expected grade advantage (i.e., the test). The more informative the test, the stronger is the high type's grade advantage. We derive a condition (*RC-Informativeness*), which states precisely when the test is informative enough relative to the *cost* advantage to induce the high type's reliance on the test. For RC-Informative tests, the relevant indifference curves do not satisfy the single-crossing property, and hence our analysis differs from that of a typical signaling model.

We characterize the set of Perfect Bayesian Equilibria (PBE) for this application and demonstrate that all Pareto efficient equilibria involve some degree of pooling when the test is RC-Informative. To gain further insights, we refine the set of equilibria using D1 (Cho and Kreps, 1987; Banks and Sobel, 1987). We show that (generically) a unique equilibrium outcome emerges. The equilibrium depends both on the market's prior belief as well as the informativeness of the test. If the test is RC-Informative and the probability that the prior assigns to the sender being the high type is below a given threshold, the equilibrium involves partial pooling. The high type chooses a less costly signal than in the LCSE, while the low type mixes between revealing himself (at no expense) and imitating the high type. Both the signal and the grade convey information. As the prior puts more weight on the high type, the equilibrium shifts from partial pooling to full pooling at a less costly signaling level. When the prior is more favorable, the sender needn't work as hard to convince the market of his type. As a result, the expected utility of both types is increasing in the prior. As the prior goes to one, the equilibrium converges to the complete-information outcome, where the sender devotes no resources to costly signaling.²

We then extend our analysis to a general two-type model. We consider a more general class of preferences and also allow the accuracy of the test to vary with the signal—encapsulating a larger class of applications. In this setting, the sender decides how informative a test to subject himself to, knowing the market will observe how accurate a test he chose as well as its result. This serves not only as a robustness check, but also to separate the key forces behind the results from mere artifacts of a particular application. We identify a key property (analogous to the single-crossing property), which facilitates tractable

²The complete-information outcome refers to the equilibrium of the game in which the sender has no private information—his type is commonly known.

equilibrium analysis and generates sharp predictions. We then consider a model with N types and show that informative tests again lead to pooling as well as convergence to the complete-information outcome.

1.1 Evidence and Implications: Do Grades Matter?

If (as is commonly predicted by gradeless models) types are in fact perfectly distinguishing themselves based on observable actions, then any additional noisy information should be ignored. That is, if the prediction of separation is correct, grades should be irrelevant. Empirical evidence suggests otherwise. For example, Jones and Jackson (1990) find that higher G.P.A.'s translate to higher salaries among college graduates. Conversely, in some situations, such as in assessing the quality of bonds through rating agencies, it may seem that *only* grades carry meaningful information, with costly signaling playing no role. However, Hsueh and Kidwell (1988) find that the decision to hire more bond raters is correctly interpreted by the market as a signal of strength—independent of scores. Our model examines the interaction between these two potential sources of information, explaining the importance of both in market outcomes.³

Warranties for new cars provide an example where signaling behavior varies with the market's prior in the manner predicted by the model with grades and in contrast to the prediction of the gradeless model. Warranties have been cited as an example of signaling (Grossman, 1981; Lutz, 1989; Gal-Or, 1989)—makers of more reliable cars can offer warranties more cheaply (in expectation) than makers of less reliable cars. Prior to purchase, consumers observe the maker's choice of warranty as well as ratings and awards bestowed by consumer groups (i.e., grades). If the gradeless model is believed, more reliable cars should come with better warranties. Why then did Honda, long regarded as the gold standard in reliability, offer a much less comprehensive warranty for its 2006 Civic than Hyundai did for its 2006 Elantra?

According to data collected by *Consumer Reports*, the 2006 models of these cars have very similar reliability.⁴ The heuristic explanation for the differing warranties is that Honda

³An alternative explanation for the relevance of grades is that costly signaling and grades are meant to convey different aspects of the sender's type, both of which affect his market value. For example, a sender's market value is $\theta + \tau$, where θ determines the sender's cost of signaling, and τ determines his grade distribution. In the simplest case, θ and τ are independent, and the sender only knows θ . As in our model, grades will affect receivers' actions. However, unlike in our model, the equilibrium will still be least cost separating and have no reliance on the prior. Further, if the sender is risk-neutral, the informativeness of the grade will have no bearing on the equilibrium. If τ is privately-known or correlated with θ , then the model qualitatively returns toward ours—the sender is privately informed about his market value and attempts to rely on both a costly signal and a grade to convey that type to the market.

⁴Summaries are available online to subscribers at www.consumerreports.org. This data was collected in subsequent years, and therefore was not part of a consumer's information set at the time of purchase, but indicates that the Civic and the Elantra were of roughly the same "type" with regard to reliability.

needn't signal as vigorously because of its superior "reputation." When the new Civic is introduced, consumers are already reasonably sure of its high quality. Therefore, Honda does not need to expend much on signaling. It can rely on this confidence and reviews/ratings/awards to sell its products at a high price. Consumers are less sure about the quality of an upstart company like Hyundai. In the words of Hyundai CEO, John Krafcik, "...we were seeing indications that [Hyundai's] quality was really much better than the perception."⁵ Correspondingly, the market has a more favorable prior for a new-issue Civic than a new-issue Elantra. While the gradeless model predicts that this should be irrelevant, the model with grades matches the observed facts. Firms with a better reputation (as measured by the prior) incur smaller signaling costs.

Applied to academic settings, our results have policy implications for several issues currently facing educators. Recently, elite business schools have grappled with grade disclosure policy. In 1998, Harvard Business School adopted a policy that prohibited students from revealing grades to potential employers. Seven years later, administrators at Harvard reversed the policy citing the need for more transparency and accountability in the classroom.⁶ At several other top business schools, the student body has adopted a grade non-disclosure policy (Chicago Booth, Stanford GSB and Wharton). Advocates of non-disclosure argue that it leads to a more collegial learning environment. But at what cost? Our results suggest that non-disclosure may lead students to try to distinguish themselves through other, perhaps very costly, channels such as additional or joint degrees, certificate programs, and extra-curricular involvement.

Grade inflation at secondary schools, colleges and universities is another relevant issue. Since the 1980's, the average GPA at American colleges and universities has risen at a rate between 0.1 and 0.15 points per decade on a 4.0 point scale (Rojstaczer, 2009). Qualitatively, grade inflation has an effect similar to non-disclosure because it reduces the informativeness of grades observed by admissions offices and potential employers. This makes it more difficult to distinguish between candidates based on grades and provides an explanation for the concurrent increase in extra-curricular involvement that has become an essential part of competitive applications for higher education and employment.

Finally, in gradeless models the existence of separating equilibria relies crucially on the single-crossing property. Consequently, the theory has been limited to environments in which single-crossing is justifiable. Many applications of interest do not fit this criterion, especially when the signal corresponds to a monetary expenditure. Because our analysis does not

⁵From interview in *Newsweek* March 24, 2010.

⁶www.thecrimson.com/article/2005/12/15/in-reversal-hbs-to-allow-grade/ (accessed October 24, 2011)

rely on (strict) single-crossing, our results apply equally well to such settings and provides insight as to how and when agents can undertake costly dissipative actions to convey private information even in the absence of type-dependent signaling costs. One example of such an environment is advertising. The use of advertising campaigns containing little or no obvious informational content has been well documented (Nelson, 1974). Prior explanations for this behavior have been based on repeat purchase considerations (Nelson, 1974; Kihlstrom and Riordan, 1984; Milgrom and Roberts, 1986). Our results suggest that the presence of grades can provide alternative explanation for dissipative advertising. Our theory is also consistent with the empirical finding that quality and advertising expenditures are positively correlated (Thomas et al., 1998).

This paper is organized as follows. In Section 2, we present the two-type model and introduce our solution concept. Section 3 analyzes a canonical application of the model, which provides intuition for the more general results presented in Section 4. Section 5 extends the results to the model with a larger type space. Section 6 discusses our results within the context of the literature and concludes. Proofs are located in Appendix A.

2 The Two-Type Model

We model the signaling environment as a game of incomplete information. There is one sender and a set of receivers. The sender is privately informed of his type $t \in \{L, H\}$. Receivers share a commonly-known prior $\mu_0 \equiv \Pr(t = H) \in (0, 1)$. The sender chooses a signal $x \in \mathbb{R}_+$. A strategy for each type of sender is a probability distribution, denoted Υ_t , with support $S_t \subseteq \mathbb{R}_+$. If the strategy of the type- t sender contains mass points, we use $\sigma_t(x)$ to denote the probability assigned to $x \in S_t$. In addition, a grade $g \in \mathbb{R}$ is realized according to the probability density function $f_t(g|x)$. The signal, x , and the grade, g , are publicly observed.

Based on the observed signal and grade, the receivers then update their (common) belief about the sender's type to a final belief $\mu_f(x, g) = \Pr(t = H|x, g)$. The expected payoff of a type- t sender is then $U_t(x, \mu_f)$.⁷ We assume that U_t is differentiable in both arguments, with $U_{t,2} > 0$, where $U_{t,i}$ denotes the partial derivative of U_t with respect to its i^{th} argument. In addition, we put the following structure on the sender's preferences.

A.1 $-U_{H,1}/U_{H,2} \leq -U_{L,1}/U_{L,2}$ for all (x, μ_f) .

⁷Similar to Mailath (1987), this is a reduced-form representation of an environment in which following the observation of (x, g) , the sender and the receivers participate in a continuation game with the key feature that, given (x, μ_f) , all equilibria yield the same expected payoff for a type- t sender. For example, in Spence (1973) receivers make simultaneous wage offers to the sender, who decides which offer (if any) to accept (see Section 2.2 for more examples).

A.2 $U_t(x, \mu_{t'})$ is strictly quasiconcave in x for all t, t' , where μ_t denotes the degenerate belief that places probability one on type t .

A.3 There exists an $\hat{x} \geq 0$ and $d > 0$ such that $U_{t,1}(x, \mu_f) < -d$ for all t, μ_f and $x \geq \hat{x}$.

A.1 is a weak version of the Spence-Mirrlees condition. A.3 states that a higher signal is *eventually* costly for the sender. Combined with A.2, it ensures that the complete-information outcome is both well-defined and unique. Let $x_t^* \equiv \arg \max_x U_t(x, \mu_t)$ denote the signal chosen by the sender in the complete information setting. A.2 and A.3 also imply that there exists a unique $\bar{x} > x_L^*$, such that $U_L(x_L^*, 0) = U_L(\bar{x}, 1)$. Finally, we will assume that

A.4 $x_H^* < \bar{x}$.

In combination with A.1, A.4 implies that (i) the least-cost-separating strategy profile $\sigma_L(x_L^*) = 1, \sigma_H(\bar{x}) = 1$ is part of a Perfect Bayesian Equilibrium (PBE), and (ii) the fully efficient outcome $\sigma_L(x_L^*) = 1, \sigma_H(x_H^*) = 1$ cannot be sustained as part of a PBE.⁸

2.1 Grades and Testing Technologies

Given his choice of x , the sender is subjected to a *test*. A test is a pair of probability density functions $\{f_L, f_H\}$ associated with a continuous random variable G . Given $t \in \{L, H\}$, G has density f_t , where f_t is continuous almost everywhere and $\Pr(a \leq G \leq b | t) = \int_a^b f_t(g) dg$.⁹ A realization of a test is a *grade*, $g \in \mathbb{R}$. Let $R(g) \equiv f_L(g)/f_H(g)$.¹⁰ $R(g)$ measures the informativeness of the grade g . If $R(g) = 1$, then the grade offers no information about the sender's type. Conversely, if $R(g) > 1$ (< 1), then the grade causes a Bayesian to decrease (increase) the probability assigned to the high type. We restrict attention to tests that do not contain perfectly informative grades, $R(g) \in (0, \infty)$ for all g , unless otherwise specified. Without loss of generality, order the grades such that the Monotone Likelihood Ratio Property (MLRP) holds with R weakly decreasing over the common support of f_L and f_H . A test is *statistically informative* if there exists a set of grades $\mathcal{G} \subseteq \mathbb{R}$, such that $R(g) \neq 1$ for all $g \in \mathcal{G}$ and $\int_{\mathcal{G}} f_t(g) dg > 0$ for $t \in \{L, H\}$.¹¹ Because our primary interest is to study an environment in which grades have the potential to reveal information, we focus on statistically informative tests and henceforth use *gradeless model* in reference to the model with a test that is not statistically informative (or equivalently, the model in which g is not observed).

⁸The profile in (i) can be sustained by assigning probability one to $t = L$ for all off-path signals.

⁹In many environments the set of grades is finite: *Pass/Fail*, number of *stars*, letter grades *A* to *F*. To encompass a situation with a countable set of grades $\{y_1, y_2, \dots\}$, with probabilities $p_t(y_n)$, let $f_t(g) = p_t(y_n)$ for $g \in [n, n+1)$ and $f_t(g) = 0$ for all other g .

¹⁰If $f_H(g) = f_L(g) = 0$, we adopt the convention that $R(g) = 1$.

¹¹Statistical informativeness is a property of the test. It does not imply that the grade will convey meaningful information to the receivers in equilibrium.

While $R(g)$ measures the informativeness of a *grade*, we are also interested in comparing the informativeness of different *tests*. Blackwell (1951) and Lehmann (1988) provide the predominant notions for what it means for one test to be more informative than another. With two types, the notions are equivalent (Jewitt, 2007) and can be stated as a condition on the expectations of convex functions of R .¹² However, these notions endow only a partial ordering over tests—many pairs of tests will not be ranked. The low type’s expected likelihood ratio, $\mathbb{E}[R(g)|t = L]$, is a measure of test informativeness that endows a complete ordering and will play a key role in the analysis that follows. The lowest possible value of $\mathbb{E}[R(g)|L]$ is 1, which obtains when the test is not statistically informative. The higher is $\mathbb{E}[R(g)|L]$, the more informative the test. This measure is consistent with the notions of Blackwell and Lehmann in the following sense.

Fact 1. *If the test $\{f_L, f_H\}$ is more informative than the test $\{\hat{f}_L, \hat{f}_H\}$ in the sense of either Blackwell or Lehmann, then $\mathbb{E}[R(g)|L] \geq \mathbb{E}[\hat{R}(g)|L]$.*

As we have alluded to, the accuracy of the test the sender is subjected to may depend on his chosen signal. For example, a firm issuing new stock hires an auditor. More reputable and higher quality auditors are both more expensive and more accurate. The story is similar for a firm hiring credit rating agencies before issuing bonds. In education, more years of schooling produces a longer and more informative transcript, or a more difficult course of study is better at distinguishing students of varying abilities. To make explicit this relationship, define a *testing technology* to be a family of tests indexed by x and denoted by $\{f_L(\cdot|x), f_H(\cdot|x)\}$, where $f_t(\cdot|x)$ is continuously differentiable in x . Let $R(g|x) = f_L(g|x)/f_H(g|x)$ denote the likelihood ratio of the grade g given the signal x .

In other applications, such as advertising and product warranties (see below), it is more natural to assume that the distribution of the grade is independent of the signal.

Definition 1. *A Constant Testing Technology satisfies $f_t(g|x) = f_t(g|x')$ for all t, g, x, x' .*

2.2 Applications

Signaling theory has been applied to a broad array of economic environments (see Riley (2001) for a survey). As motivated in the Introduction, we believe that grades are a prominent feature of many of these environments and have kept our model quite general in order

¹²Formally, the test $\{f_L, f_H\}$ is weakly more informative than test $\{\hat{f}_L, \hat{f}_H\}$ if for every convex function φ

$$\int \varphi \left(\frac{f_t(g)}{f_{t'}(g)} \right) f_t(g) dg \geq \int \varphi \left(\frac{\hat{f}_t(g)}{\hat{f}_{t'}(g)} \right) \hat{f}_t(g) dg$$

for each $t, t' \in \{L, H\}$. See Blackwell and Girshick (1954) for equivalent conditions and Jewitt (2007) for a more recent discussion and proof of the Blackwell/Lehmann equivalence in dichotomies.

to capture a broad class of such settings. Below we discuss several specific applications that fit within our framework.

1. **Job Market Signaling**: As in Spence (1973), the sender is a student who desires employment, and the receivers are potential employers. Employers value a type- t worker at V_t , where $V_L < V_H$. Given μ_f , the employers make simultaneous wage offers to the sender. Because of competition, the highest wage offered in equilibrium must be the expected value of the worker: $\mathbb{E}[V_t|\mu_f] \equiv \mu_f V_H + (1 - \mu_f)V_L$. The student likes higher wages, but dislikes schooling: $U_t(x, \mu_f) = \mathbb{E}[V_t|\mu_f] - C_t k(x)$, where $k' > 0$, $k'' \geq 0$ and $C_L \geq C_H > 0$. The grade in this example has the quite literal interpretation as the student's transcript (as well as any other statistically differentiating academic honors).

2. **Advertising**: As in Kihlstrom and Riordan (1984), the sender is a firm offering a good of uncertain quality t , and the receivers are potential customers. The demand for the product is increasing in the expected quality of the good, captured by the demand curve: $Q(P, \mu_f) = a\mathbb{E}[V_t|\mu_f] - bP$. For simplicity, the firm's marginal cost of production is zero, regardless of t , so its profit is $\Pi(\mu_f) \equiv (a\mathbb{E}[V_t|\mu_f])^2/4b$. Prior to bringing its good to market, the firm can engage in non-informative advertising (i.e., money burning). Therefore, $U_t(x, \mu_f) = \Pi(\mu_f) - x$. The grade in this example represents reviews of the good, such as those provided by Yelp, CNET, Zagat, Angie's List, etc.

3. **Warranties**: As in Gal-Or (1989), the sender is a firm offering a good of uncertain durability t , and the receivers are a unit mass of potential customers. The firm can offer a warranty policy that replaces the good in the event of failure before time x . Higher quality goods breakdown less frequently. Therefore, customers' willingness to pay is increasing in μ_f , and the firm's (expected) cost for its warranty policy is decreasing in t . This is represented explicitly by $U_t(x, \mu_f) = \mathbb{E}[V_t|\mu_f] + x - \frac{1}{2}C_t x^2$, for some $C_L \geq C_H > 0$.¹³ Notice, that U_t can be non-monotonic in x because, regardless of μ_f and t , it may be profit maximizing to offer a warranty. The grade in this example represents the results of third-party tests such as those provided by Consumer Reports, J.D. Power and Associates, etc.

¹³This simplifies the model of Gal-Or (1989) by abstracting from the cost of initial production of the good and assuming that (i) consumers are homogenous (i.e., $b = 0$ in her notation) and (ii) V_L is large enough relative to C_L . The more general version of the model could be incorporated into our framework, but is much more cumbersome to work with and interpret.

In the following two (simplified) applications from the finance literature, the sender has mean-variance preferences over his final wealth level (\tilde{W}) and maximizes $\mathbb{E}[\tilde{W}] - \frac{\gamma}{2}\text{Var}[\tilde{W}]$, where $\text{Var}[\tilde{W}]$ is the variance of final wealth and γ is a measure of the sender's risk aversion. These applications can be analyzed within our framework by interpreting these preferences to be the sender's *expected* payoff, where the expectation is taken over all possible realizations of the grade given his choice of signal.¹⁴ To retain the key tradeoff facing the sender, we will assume the sender's risk aversion does not dominate his preference for being seen as a high type ($\gamma \leq \bar{\gamma}$).

4. **Financial Structure and Inside Information**: As in Leland and Pyle (1977), the sender is an entrepreneur looking to sell a portion of his company to a market of investors (the receivers). The future returns for the company are random with mean V_t , ($V_L < V_H$) and variance $\sigma^2 > 0$. The entrepreneur chooses the fraction of the company he retains, which serves as the signal. Because the market is competitive, the equity offering will yield a price of $P(\mu_f) \equiv \mathbb{E}[V_t|\mu_f]$ per unit offered. Thus, the expected payoff of a type- t entrepreneur *before* the grade is realized is:

$$x \left(V_t - \frac{\gamma}{2} x \sigma^2 \right) + (1 - x) \left(\mathbb{E}[P(\mu_f)|t, x] - \frac{\gamma}{2} (1 - x) \text{Var}[P(\mu_f)|t, x] \right)$$

The grade in this example can be interpreted as an analyst's recommendation prior to the issuance date. For this application, we assume that the recommendation is binary (e.g., *buy* ($g = 1$) or *sell* ($g = 0$)) and symmetric so that $\Pr(g = 1|H) = \Pr(g = 0|L) = p \in (0.5, 1)$.¹⁵

5. **Auditors and Equity Issuance**: As in Titman and Trueman (1986), a firm plans to issue equity to raise funds for a project. As in the previous example, the future returns for the company are random with mean V_t , ($V_L < V_H$) and variance $\sigma^2 > 0$. The percentage of the firm the company retains is fixed at some $\alpha \in (0, 1)$. Prior to the issuance, the firm chooses an auditor, whose quality is observable and thus serves as the signal. The auditor then prepares a statement, which serves as the grade. Higher quality auditors provide more informative statements, but are also more expensive; the

¹⁴When choosing a signal, the sender will maximize $\mathbb{E}[U_t(x, \mu_f(x, g))]$, where the expectation is taken over realizations of the grade. Hence, it is the sender's expected utility function that plays the crucial role for analysis. For expected-utility maximizers, the construction of the expected utility function follows easily from U_t and the testing technology (see Section 2.3). Since mean-variance preferences are generally not consistent with an expected-utility representation, one should interpret the mean-variance representation in Applications 4-5 as a substitute for expected utility.

¹⁵A binary symmetric test is not a necessary restriction for our results, but it facilitates a tractable analysis of the environment.

cost of an auditor with quality x is given by $c(x)$, where $c' > 0$, $c'' \geq 0$. The sender's expected payoff before the grade is realized is

$$\alpha \left(V_t - \frac{\gamma}{2} \alpha \sigma^2 \right) + (1 - \alpha) \left(\mathbb{E}[P(\mu_f)|t, x] - \frac{\gamma}{2} (1 - \alpha) \text{Var}[P(\mu_f)|t, x] \right) - c(x)$$

To fix ideas and facilitate a tractable analysis, we specify that conditional on (x, t) , $G \sim U[0, 1 + V_t \cdot x]$. Thus a higher quality auditor is more likely to distinguish a high-type firm from a low one.¹⁶

2.3 Solution Concept and Preliminary Analysis

We use Perfect Bayesian equilibrium (PBE) as our solution concept with the additional requirement that receivers hold identical beliefs off the equilibrium path. After observing x and g , receivers update to some final belief $\mu_f(x, g) \equiv \Pr(t = H|x, g)$. Because receivers are Bayesian, the updating can be decomposed into a first update based on x and a second update based on g . The first update results in an *interim belief*. Let μ denote an arbitrary interim belief and $\mu(x)$ be the interim belief as a function of the signal observed. Along the equilibrium path, the interim belief is determined by Υ_L , Υ_H , and the belief consistency requirement of PBE. The second update is purely statistical—receivers update from their interim belief based on the observation of the grade via Bayes rule as given by (1).

$$\mu_f(x, g) = \frac{\mu(x)}{\mu(x) + (1 - \mu(x))R(g|x)} \quad (1)$$

Therefore, given x , the interim belief is sufficient for a type- t sender to compute his expected utility, denoted by u_t .

$$u_t(x, \mu) \equiv \int U_t(x, \mu_f(x, g)) f_t(g|x) dg \quad (2)$$

Equilibrium Selection

As in most signaling models, the set of PBE is large because of the flexibility afforded to off-equilibrium-path beliefs. In Section 3.2, we describe the different forms of PBE that can exist within an application of the model and characterize the set of achievable equilibrium payoffs. In order to produce further insights and predictions, the set of equilibria must be refined.

¹⁶Notice that this testing technology involves a set of grades that perfectly distinguish the high type, which raises the question of whether such a grade can overturn a degenerate prior. In general, the set of equilibria can depend on the way in which this issue is addressed. However, it has no equilibrium implications for this particular application (in part because the grade is completely uninformative at $x = 0$), and therefore we do not take a position on this matter.

Starting in Section 3.3, we focus our attention on equilibria satisfying the D1 refinement (Banks and Sobel, 1987; Cho and Kreps, 1987). Using D1 leads to sharp predictions and also has the advantage of facilitating comparison with the gradeless model where, if single-crossing holds strictly, it uniquely selects the LCSE for any number of types.^{17,18}

3 Job Market Signaling with Grades

In this section, we analyze a simple application, which distills much of the intuition for the more general results to come in Sections 4 and 5. In particular, we focus on the job-market signaling application with linear cost function ($k(x) = x$) and constant testing technology.

Given any interim belief, μ , the highest expected offer received by a type- t worker is independent of x and given by:

$$w_t(\mu) \equiv \int \frac{\mu V_H + (1 - \mu)V_L}{\mu + (1 - \mu)R(g)} f_t(g) dg \quad (3)$$

Thus, the expected utility of a type- t worker is $u_t(x, \mu) = w_t(\mu) - C_t x$. Our analysis begins by investigating the indifference curves for both types over the space of signals and interim beliefs.

3.1 Belief Indifference Curves

Without grades, the indifference curves of interest are those over the space of signals and offers. Further, the sequential rationality of receivers implies that these are equivalent to indifference curves over the space of signals and final beliefs. Based on the analysis conducted in Section 2.3, we see that *with* grades it is indifference curves over the space of signals and *interim* beliefs that are crucial for analysis. We refer to these as *Belief Indifference Curves (BICs)*. It will be useful to think of BICs as functions from signals to interim beliefs parameterized by utility levels. Let $b_t(x|\hat{u})$ be the interim belief such that $u_t(x, b_t(x|\hat{u})) = \hat{u}$.¹⁹

A comparison to the gradeless model is helpful. In the gradeless model, there is no distinction between final and interim beliefs, meaning BICs align with standard indifference curves. Figure 1 illustrates the BICs for the LCSE utility levels in both settings for an example where $C_L > C_H$. Throughout the paper \bar{x} denotes the high type's LCSE signal. Without the test, the low type's curve is steeper than and below the high type's for all $x < \bar{x}$. The difference in the slope of the two types' curves derives only from their differing costs. With the test, the indifference curves acquire curvature according to each type's expectation

¹⁷Cho and Kreps (1987); Cho and Sobel (1990); Ramey (1996)

¹⁸While the set of equilibria is often larger, the main economic insights produced by our model hold under the more mild Divinity refinement (Banks and Sobel, 1987). Results available upon request.

¹⁹If the equation cannot be satisfied for an interim belief in $[0, 1]$, then $b_t(x|\hat{u}) = \emptyset$.

regarding his grade on the test. The crucial observation is that because the input is *interim* belief, the shapes of the indifference curves for the two types are different. Because the low type is more likely to receive a lower grade, his indifference curve in the interior lies everywhere above where it did in the gradeless model. To maintain the same utility when grades are available, the low type needs more favorable beliefs to offset the outcome he expects on the test. The opposite is true for the high type.

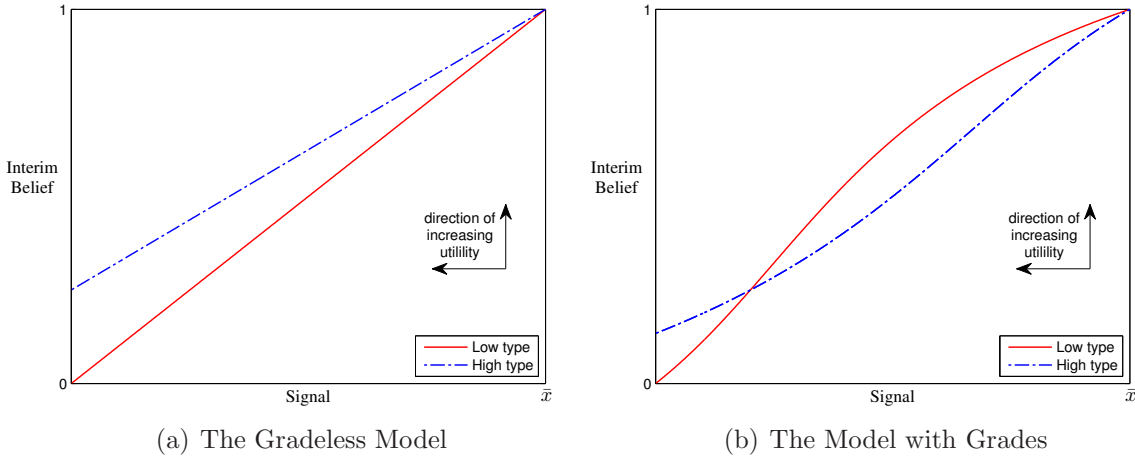


FIGURE 1 – BICs for the LCSE payoffs

3.1.1 The Strong Crossing Property

In this application, BICs have a useful feature, termed the Strong Crossing Property, that is highlighted in Figure 2. Notice first that for each type, all BICs for a given type are horizontal translations of one another—a more costly signal decreases utility, but does not affect the utility tradeoff between x and μ .

Definition 2. *BICs satisfy the **Strong Crossing Property (SCP)** if there exists a unique $\mu^* \in (0, 1]$ such that the slope of the high type’s BICs is less than the low type’s for all $\mu < \mu^*$ and the slope of the high type’s BICs is greater than the low type’s for all $\mu > \mu^*$.*

Lemma 1. *For any constant testing technology, the Job-Market Signaling Application satisfies SCP.*

The slope of the sender’s BIC is $-\frac{\partial u_t}{\partial x} / \frac{\partial u_t}{\partial \mu} = C_t / \frac{dw_t}{d\mu}$. To understand *SCP*, we start with another lemma that shows how the interim belief drives w_t .

Lemma 2. *For any test, there exists a unique $\mu_{\max} \in (0, 1)$ such that $\frac{dw_L}{d\mu} < \frac{dw_H}{d\mu}$ for all $\mu < \mu_{\max}$, and $\frac{dw_L}{d\mu} > \frac{dw_H}{d\mu}$ for all $\mu > \mu_{\max}$.*




FIGURE 2 – Illustration of the Strong Crossing Property

The reasoning behind the lemma is as follows. In the extremes of $\mu = 0, 1$, the Bayesian receivers disregard the result of the test. As the belief becomes more intermediate, the sender's grade begins to matter for his expected offer. Further, there is a unique belief (μ_{\max}) where the test is maximally important, as measured by $w_H - w_L$. Hence, increasing the belief from 0 toward μ_{\max} increases the importance of the test, and is therefore relatively better for the high type than it is for the low type. Similarly, increasing the belief from μ_{\max} toward 1 decreases the importance of the test and benefits the low type more than the high type.

To move from Lemma 2 to *SCP*, we must also incorporate any cost advantage. The slope of the BICs tells us which type needs more compensation from increased interim belief to incur a marginally more expensive signal. If there is no cost advantage, the belief where the low type begins needing less μ -compensation than the high type is when the test starts becoming less important. That is, if $C_H = C_L$ then $\mu^* = \mu_{\max}$. Further, μ^* is increasing in the cost advantage as measured by $\frac{C_L}{C_H}$. If the cost advantage is large enough then $\mu^* = 1$ and the single-crossing condition holds: at every point, the high type's BIC is flatter than the low type's. In other words, when the cost advantage is large relative to the grade advantage (in the manner made precise by Definition 3), the BICs recover the same essential property as in the gradeless model.

3.1.2 RC-Informative Tests

The following condition will play a major role in the analysis of the model.

Definition 3. *The test is **RC-Informative** if $\mathbb{E}[R(g)|L] > \frac{C_L}{C_H}$.*

Recall that $\mathbb{E}[R(g)|L]$ is one way to measure the informativeness of the test, while $\frac{C_L}{C_H}$ is a measure of the cost advantage. Hence, RC-Informativeness is simply that the test is

informative enough relative to the cost advantage. Notice that if $C_H = C_L$, any statistically informative test is RC-Informative.

Lemma 3. *The following statements are equivalent.*

1. *The test is RC-Informative.*
2. *At $\mu = 1$, the high type's BIC is steeper than the low type's BIC.*
3. $\mu^* < 1$.

That statements 1 and 2 are equivalent is a direct calculation. The equivalence of 2 and 3 follows immediately from *SCP*.

3.2 The Set of PBE

There are many Perfect Bayesian Equilibria of the game. As in the gradeless model, there exist separating, full pooling, and partial pooling equilibria. In addition, there exists a new form of equilibria we designate *common support* equilibria. In a common support equilibrium $S_L = S_H$, but $\Upsilon_L \neq \Upsilon_H$. Common support equilibria differ from separating or partial pooling equilibria in that no on-path signal perfectly identifies the sender's type, and differ from full pooling equilibria in that multiple signals are on the equilibrium path, each leading to a different interim belief, which also differs from the prior. *SCP* implies that in any common support equilibrium $S_L = S_H = \{x_1, x_2\}$, where $x_1 < x_2$ and $\mu(x_1) < \mu_0 < \mu(x_2)$. Figure 3 depicts BICs for payoffs which can be supported by a common support equilibrium given any $\mu_0 \in (\mu_1, \mu_2)$ as labeled.²⁰




FIGURE 3 – BICs for a common support equilibrium

²⁰Common support equilibria are not robust to stability-based refinements such as D1 and Divinity (see Section 3.3). In addition, there exist partial pooling equilibria in which $S_L \cap S_H$ is a pair, rather than a singleton (in contrast to the gradeless model with strict single-crossing). Such equilibria fail D1 as well.

The set of PBE depends on the prior. Enumerating each equilibrium for every prior is a straightforward but tedious exercise. Instead, we turn to characterizing the set of equilibrium payoffs when the test is RC-Informative and illustrate how the set changes with the prior.²¹ This will also allow us to discern the set of Pareto efficient equilibria for each prior. Two belief levels will play a key role. The first is μ^* . Let x^* be the unique signal satisfying $u_L(x^*, \mu^*) = V_L$. The second key belief level is $\underline{\mu} \equiv b_H(0|u_H(x^*, \mu^*))$: the belief μ such that the high type is indifferent between $(0, \mu)$ and (x^*, μ^*) . Note that $\underline{\mu} < \mu^*$.




FIGURE 4 – The set of PBE payoffs for various priors

Referring to Figure 4, payoffs produced by separating equilibria are located on the vertical axis between $u_L = u_H = V_L$ and the LCSE payoffs. The remainder of the feasible payoffs on the vertical axis can be achieved by partial pooling equilibria in which $0 \in S_L, \notin S_H$. Regardless of the prior, there exist equilibria such that the low type's payoff is V_L , but the high type achieves a payoff greater than in any separating equilibrium. Next, notice that there exists a linear boundary of the equilibrium payoff set connecting the vertical axis and the point corresponding to full pooling at $x = 0$. This linear boundary represents the payoffs produced by full pooling equilibria. The left endpoint of this boundary corresponds to payoffs for the full pooling equilibrium at the signal x satisfying $u_L(x, \mu_0) = V_L$: the most expensive signal consistent with the low type weakly preferring to remain in the pool. Because the interim belief is the same as the prior belief in a full pooling equilibrium, each type gains by pooling at a less expensive signal, and the slope of the linear boundary is $\frac{C_H}{C_L}$.

We can see that there are only two candidates for Pareto efficient equilibria. The first is full pooling at $x = 0$. For any prior, this is the payoff-maximal equilibrium for the low type, and hence is always Pareto efficient. The second is partial pooling at x^* in which the

²¹Characterizing the set when the test is not RC-Informative is no more difficult than in, and produces little difference from, the gradeless model. Furthermore, the analysis conducted in this section can easily be extended to the general two-type model of Section 4.

high type chooses x^* with probability one and the low type mixes between 0 and x^* such that $\mu(x^*) = \mu^*$. This equilibrium exists if and only if $\mu_0 \leq \mu^*$ and is Pareto efficient if and only if $u_H(x^*, \mu^*) \geq u_H(0, \mu_0)$, which is equivalent to $\mu_0 \leq \underline{\mu}$. Grades convey information *in equilibrium* if and only if types do not separate by choice of signal. Hence, when the test is RC-Informative, there is always a way to utilize it that Pareto dominates separating outcomes.

Proposition 1. *For any μ_0 , if the test is RC-Informative, then grades convey payoff-relevant information in every Pareto efficient equilibrium. If the test is not RC-Informative, then the LCSE is a Pareto efficient equilibrium if and only if $\mu_0 \leq \underline{\mu}$.*

The above conclusion when the test is not RC-Informative necessarily holds in the gradeless model (recall that $\mu^* = 1$ in this case). Along with the results in Section 3.3, this indicates that RC-Informativeness is precisely the criterion that characterizes when the presence of grades substantively alters the predictions from the gradeless model.

3.3 Equilibrium under D1: Senders Rely on Informative Tests

In the two-type model, the D1 refinement can be stated as follows.²² Fix an equilibrium endowing expected utilities $\{u_L^*, u_H^*\}$. Consider a signal x that is not in the support of either type's strategy. Define $B_t(x, u_t^*) \equiv \{\mu : u_t(x, \mu) > u_t^*\}$. If $B_L(x, u_L^*) \subset B_H(x, u_H^*)$, then D1 requires that $\mu(x) = 1$ (where \subset denotes strict inclusion). Conversely, if $B_H(x, u_H^*) \subset B_L(x, u_L^*)$, then D1 requires that $\mu(x) = 0$.²³

The refinement can be interpreted as follows. Suppose that x is not in the support of either type's equilibrium strategy, but it is observed nonetheless. A receiver uses the following reasoning, "The sender must have misunderstood my beliefs, otherwise he would not have chosen this deviation. Of all the possible beliefs that he could (mistakenly) think I will have after seeing x , for which subset of these beliefs would the low type prefer this deviation? For which subset would the high type prefer this deviation? If there are beliefs such that the high type would prefer this deviation and the low type would not, and there are no beliefs such that the low type would prefer the deviation and the high type would not (that is, if $B_L(x, u_L^*) \subset B_H(x, u_H^*)$), then given the deviation to x , I should believe that the sender is of type H ."

²²In Appendix A, we explain the equivalence, in our model, between this definition and D1's original definition (Banks and Sobel, 1987; Cho and Kreps, 1987). See footnote 39 for a formal statement of the D1 refinement in the model with N types.

²³Let $B_t^0(x, u_t^*)$ be the set of interim beliefs μ such that $u_t(x, \mu) = u_t^*$. In Banks and Sobel (1987) and Cho and Kreps (1987), D1 requires that if $B_{t'}^0(x, u_{t'}^*) \cup B_{t'}(x, u_{t'}^*) \subset B_t(x, u_t^*)$ then the receivers assign zero probability to the sender being type t' . The continuity of the preferences and action spaces in our model make the two statements equivalent.

Remark 1. For the remainder of the paper and unless otherwise qualified, *an equilibrium* refers to a PBE (with common off-path receiver beliefs) that satisfies D1.

The following proposition characterizes the set of equilibria. Recall that $u_L(x^*, \mu^*) = V_L$.

Proposition 2. *If the test is RC-Informative, then*

- *if $\mu_0 > \mu^*$, the unique equilibrium is full pooling at $x = 0$.*
- *if $\mu_0 < \mu^*$, the unique equilibrium is partial pooling. The high type chooses x^* with probability 1, and the low type mixes over $x = 0$ and x^* such that $\mu(x^*) = \mu^*$ (i.e., $\sigma_L^*(x^*) = \frac{\mu_0}{1-\mu_0} \frac{1-\mu^*}{\mu^*}$ and $\sigma_L^*(0) = 1 - \sigma_L^*(x^*)$).*
- *if $\mu_0 = \mu^*$, all equilibria are full pooling and can be supported at x iff $x \in [0, x^*]$.*

If the test is not RC-Informative, then for all priors the unique equilibrium is the LCSE.

In essence, the proposition states that if the test is too weak (or the cost advantage is too great), the predictions match those of the gradeless model. The underlying analysis matches as well—when the test is not RC-Informative, $\mu^* = 1$ and the single-crossing condition holds just as in the gradeless model. However, for more informative tests (or lesser cost advantages), the equilibrium predictions change—at least some degree of pooling (and therefore reliance on the test) will take place at all priors, and full pooling will take place for sufficiently high priors. Notice that despite fully pooling, the expected offers and utility levels are different for the two types because offers are conditioned on the realized grade. An immediate implication is that if $C_H = C_L$, then the equilibrium will involve pooling; if the high type has no cost advantage, he relies at least partially on his grade advantage.

We also have the following connection between D1 and Pareto efficiency.

Corollary 1. *If the test is RC-Informative, then*

- *if $\mu_0 > \mu^*$, the unique equilibrium is the only Pareto efficient PBE.*
- *if $\mu_0 \leq \underline{\mu}$, the unique equilibrium is one of the two Pareto efficient PBE.*

Only when $\mu_0 \in (\underline{\mu}, \mu^*)$ do the two criteria have no PBE in common under RC-Informativeness. Even in this case, though, the unique equilibrium Pareto dominates the LCSE.

To sketch the argument for the proposition, we start with a useful lemma.

Lemma 4. *If the payoffs $\{u_H^*, u_L^*\}$ are supported by an equilibrium, then there does not exist a signal x' such that $b_L(x'|u_L^*) > b_H(x'|u_H^*)$.*

We can use Figure 3 to illustrate the lemma. The figure depicts BICs for a candidate u_L^* and u_H^* . Consider a signal $x' \in (x_1, x_2)$. Notice that $b_L(x'|u_L^*) > b_H(x'|u_H^*)$ for all $x' \in (x_1, x_2)$. The lemma claims that this is inconsistent with u_L^* and u_H^* being supported in equilibrium. To see this we ask, for what values of $\mu(x')$ would a sender of type t prefer to deviate to x' ? The set is simply $\{\mu : b_t(x'|u_t^*) < \mu \leq 1\}$. Since $b_L(x'|u_L^*) > b_H(x'|u_H^*)$, the set of beliefs is strictly larger for the high type. D1 then requires that $\mu(x') = 1$. With these off-path beliefs, either type would prefer to deviate to x' and receive payoff $V_H - C_t x' > u_t^*$.

One interpretation of this analysis is that neither at x_1 nor at x_2 has the high type found the optimal balance between relying on the test and relying on the costly signal. At x_1 he relies too much on the test and should increase his costly signal to dissuade the low type from imitating him with such high probability. At x_2 he relies too little on the test and should decrease his signal to make better use of the free information the test provides.

From Lemma 4, we immediately see that the LCSE will fail under RC-Informativeness. RC-Informativeness implies that the high type's BIC is steeper than the low type's at $\mu = 1$ (Lemma 3). Therefore, for the LCSE payoffs, which by construction have the corresponding BICs for each type intersect at $(\bar{x}, 1)$, the high type's BIC falls below the low type's for a deviation slightly below \bar{x} (Figure 1(b)).

Lemma 4 can also be used to derive further implications. Consider any candidate equilibrium utility level for the low type, u_L^* , and its corresponding BIC. In equilibrium, if the low type receives u_L^* , the high type must receive the maximal utility from over (x, μ) pairs on the low type's BIC, $b_L(\cdot|u_L^*)$. Lemma 4 implies that he must receive at least this much. However, if he received a higher utility level, then the high type's equilibrium BIC would lie everywhere above the low type's—meaning the low type would prefer to deviate and imitate the high type. Returning to our interpretation, in equilibrium the high type must achieve his optimal reliance on the two information channels.

SCP implies that when $\mu^* < 1$, the high type's utility-maximizing (x, μ) pair on the low type's BIC will either be at μ^* (a point of tangency) or at $x = 0$ (if μ^* does not reside on the low type's BIC). Hence, each candidate equilibrium utility level for the low type implies a unique equilibrium strategy and utility level for the high type. The final step is, for each prior, to determine which utility levels for the low type can be made consistent with equilibrium strategies/beliefs. When $\mu_0 < \mu^*$ only $u_L^* = V_L$ can simultaneously satisfy low-type incentive compatibility and be consistent with receiver beliefs on the equilibrium path. Only the full-pooling utility level can do so for $\mu_0 > \mu^*$ (see Figure 2).

3.3.1 Equilibrium Convergence

Consider the complete-information game where $t = H$ is common knowledge. The sender chooses $x = 0$ and receives an offer (and utility) V_H . By introducing the slightest probability that the sender is the low type, the gradeless model predicts that the high type will choose his LCSE signal. When the test is RC-Informative, the model with grades does not have this discontinuity. As the prior increases, the equilibrium moves to full pooling at $x = 0$.²⁴ To delve a bit deeper into convergence properties, we need to consider two types of convergence.

Definition 4. Let $\{\mu_0^n\}$ be a sequence of priors converging to μ_0 , and $\{\Upsilon_L^n, \Upsilon_H^n\}$ be any sequence of strategy profiles such that $\Upsilon_L^n, \Upsilon_H^n$ is an equilibrium when the prior is μ_0^n .

- The set of equilibrium strategy profiles **converges type-by-type** to Υ if for every sequence $\{\Upsilon_L^n, \Upsilon_H^n\}$, Υ_t^n converges in distribution to Υ for all t .
- The set of equilibrium strategy profiles **converges in total mass** to Υ if for every sequence $\{\Upsilon_L^n, \Upsilon_H^n\}$, $\mu_0^n \cdot \Upsilon_H^n + (1 - \mu_0^n) \cdot \Upsilon_L^n$ converges in distribution to Υ .

Clearly, equilibrium convergence type-by-type implies convergence in total mass. The gradeless model has been criticized because the LCSE does not converge, by either metric, to the complete-information outcome as $\mu_0 \rightarrow 1$.²⁵ An immediate corollary of Proposition 2 is that the model with grades is not susceptible to the same criticism.

Corollary 2. If the test is RC-Informative, the equilibrium converges to the complete-information outcome type-by-type as $\mu_0 \rightarrow 1$.

Notice that the equilibrium of the model with grades converges to the complete-information outcome only in total mass as $\mu_0 \rightarrow 0$. As discussed in Section 6.1, this is in contrast to noisy signaling models, where convergence is type-by-type for both $\mu_0 \rightarrow 0$ and $\mu_0 \rightarrow 1$.²⁶

3.4 Comparative Statics

To aid understanding, we present a few remarks on comparative statics. As the high type's cost advantage increases, he relies less heavily on grades. Formally, as $\frac{C_L}{C_H}$ increases, μ^*

²⁴Notice that the equilibrium does not discontinuously jump from partial pooling to full pooling at $\mu_0 = \mu^*$ either. This is because of the multiplicity of equilibria at $\mu_0 = \mu^*$.

²⁵When $\mu_0 \rightarrow 1$ ($\mu_0 \rightarrow 0$), convergence in total mass to Υ is equivalent to convergence of Υ_H (Υ_L) to Υ . Looking at convergence in total mass, however, guarantees that from an uninformed party's perspective (i.e., that of a modeler, econometrician, or even a receiver) the distribution of x limits to the desired distribution as the prior tends to its limit.

²⁶Convergence in total mass to the complete-information outcomes as the prior moves all of its weight to the lowest type is a straightforward result for *all* PBE, independent of the informativeness of grades and the number of types. We therefore do not repeat the result when examining convergence in future sections, and instead focus on convergence properties when the prior's mass moves to any other type.

also increases. Because $\mu^* \geq \mu_{\max}$, the test becomes less important in equilibrium as $\frac{C_L}{C_H}$ increases.²⁷

Regarding payoffs, it is worth noting that decreasing C_H can adversely affect the high types equilibrium payoff. To see this, consider a game where $\mu^* < 1$ and μ_0 is just above μ^* , so the equilibrium is full pooling at $x = 0$. As C_H decreases, μ^* increases above μ_0 . The high type's cost advantage is now too great to sustain full pooling, shifting the equilibrium to partial pooling at x^* . This carries a discrete increase in signaling costs but an arbitrarily small change in the high type's expected offer in equilibrium. The high type is forced to exercise his cost advantage—even if it leads to a decrease in his own expected utility—a feature familiar from the gradeless model.

Because of the flexibility afforded to the structure of our tests, general comparative statics on the accuracy of the test are cumbersome. To simplify, let us consider the family of symmetric binary tests (see Application 4), which are parameterized by $p \in (\frac{1}{2}, 1)$. Clearly, as p tends to either $\frac{1}{2}$ or 1, the test becomes completely uninformative or completely informative respectively. As p increases, the test becomes more informative according to both Blackwell and Lehmann, so $\mathbb{E}[R(g)|L]$ increases as well. For such tests, $\mu_{\max} = \frac{1}{2}$ for all p .

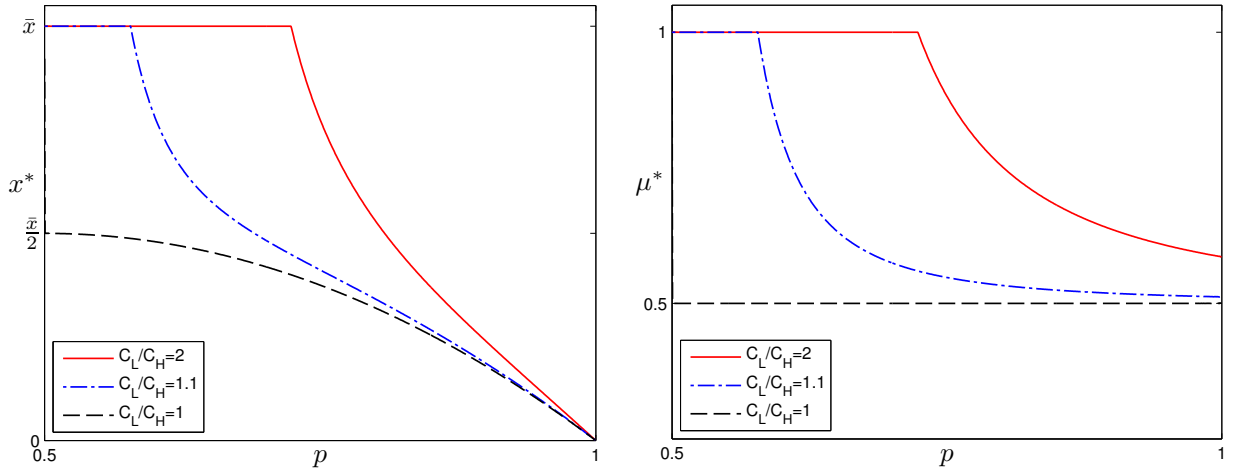


FIGURE 5 – Comparative statics for symmetric binary tests

Figure 5 illustrates how x^* and μ^* vary with p for several different cost advantages.²⁸ Modulo boundary solutions, as p increases, x^* decreases and μ^* strictly (weakly) decreases if the cost advantage is strict (weak). As $p \rightarrow 1$, $x^* \rightarrow 0$ and $\mu^* \rightarrow \frac{1}{2} = \mu_{\max}$. That is, as the test becomes completely informative, there is complete reliance on it. As $p \rightarrow \frac{1}{2}$, whether

²⁷Notice that the change in x^* given a change in $\frac{C_L}{C_H}$ is ambiguous because it will depend on the absolute, as well as relative, changes in costs.

²⁸To avoid scaling problems, C_L and (therefore) \bar{x} are fixed, while C_H varies. See footnote 27.

or not the high type has a cost advantage, no matter how slight, makes a difference. When $C_H < C_L$, (x^*, μ^*) limits to $(\bar{x}, 1)$. This is simply the recovery of the result in the gradeless model: when there is a cost advantage and no meaningful test, the unique equilibrium is the LCSE. On the other hand, if $C_H = C_L$, $\mu^* = \mu_{\max}$ for all p and $x^* \rightarrow \frac{\bar{x}}{2}$. No matter how weak the test is, the high type relies on it as heavily as possible since he has no cost advantage. Finally, unlike increasing the high type's cost advantage, increasing p is always at least weakly beneficial to the high type.

4 Analysis of the General Two-Type Model

Building upon the intuition from the previous section, we now turn to analysis of the general two-type model. To do so, first define

$$Z(x, \mu_f) \equiv \frac{U_{L,1}(x, \mu_f)}{U_{L,2}(x, \mu_f)} \bigg/ \frac{U_{H,1}(x, \mu_f)}{U_{H,2}(x, \mu_f)}$$

Notice that $Z(x, \mu_f)$ relates the slopes of the two types' indifference curves over the familiar space of signals and *final belief*, as opposed to the space of signals and interim beliefs where BICs reside.²⁹ For instance, if $U_{t,1} < 0$ for all (x, μ_f, t) , then the standard Spence-Mirrlees condition is equivalent to $Z(x, \mu_f) > 1$ for all (x, μ_f) . In the application from the previous section, $Z(x, \mu_f) = \frac{C_L}{C_H}$.

Because we allow the test to vary with the signal, we will need to evaluate the informativeness of the test at signaling levels relevant for equilibria.

Definition 5. *For any $x \in [x_H^*, \bar{x}]$, we say that the test is RC-Informative at x if $\mathbb{E}[R(g|x)|L, x] > Z(x, 1)$.*

The generalization of RC-Informativeness compares the same measure of the statistical informativeness of the testing technology at x to the ratio of the slopes of the two types' signal- μ_f indifference curves at $(x, 1)$. Before stating results, two comments are in order. First, similar to Lemma 3, for any x , the high type's BIC is steeper than the low type's at $\mu = 1$ if and only if the test is RC-Informative at x . Second, Lemma 4 applies to the general model as well, leading to the following.

Proposition 3. *If the test is RC-Informative at \bar{x} , then for any $\mu_0 \in (0, 1)$, all equilibria involve some degree of pooling.*

RC-Informativeness gives the precise condition needed on the testing technology for separation to be eliminated and all equilibria to rely on grades. A sufficient condition for

²⁹To accommodate mean-variance preferences, let $U_t(x, \mu_f)$ denote the expected payoff for a fixed (deterministic) μ_f . For example, in Application 4, $U_t(x, \mu_f) = x(V_t - \frac{\gamma}{2}x\sigma^2) + (1-x)P(\mu_f)$.

convergence to the complete-information outcome also relies on RC-Informativeness. Recall that x_H^* denotes the optimal signal for a type- H sender when receivers assign probability one to type H .

Proposition 4. *If the test is RC-Informative at all $x \in (x_H^*, \bar{x}]$, then as $\mu_0 \rightarrow 1$, by either notion of convergence (type-by-type or in total mass), the set of equilibria converges to the complete-information outcome.*

It is possible that all equilibria involve pooling but do not converge to the complete-information outcome. In this case, as $\mu_0 \rightarrow 1$, there exists a sequence of equilibria converging to full pooling at x' , for some $x' \in (x_H^*, \bar{x})$ at which the test is not RC-Informative.

4.1 The Crossing Property

Characterizing the set of equilibria for different specifications of utility functions and testing technologies poses no serious difficulty, but would result in an arduous list of various cases. Instead of embarking on this exercise, we identify a condition, deemed the Crossing Property (CP), that both generates a clean characterization and distills the critical intuition seen in the application from Section 3. For this reason, one can view the role of CP in the model with grades as analogous to the role of the single-crossing property in the gradeless model—in fact, the single-crossing property for BICs is a special case of CP that results when signal- μ_f indifference curves satisfy single-crossing and the test is sufficiently uninformative.³⁰

Like the Strong Crossing Property, CP is a condition on BICs. For any utility level \hat{u} , and either type sender, there exist signals large enough, such that no $\mu \in [0, 1]$ can deliver expected utility \hat{u} when they are chosen. Define $\bar{x}(\hat{u}_L) \equiv \max\{x : b_L(\bar{x}(\hat{u}_L), \hat{u}_L) = 1\}$.³¹ We abuse notation slightly by continuing to let $\bar{x} \equiv \bar{x}(U_L(x_L^*, 0))$ denote the high type's LCSE signal.

The Crossing Property (CP). *Consider any feasible \hat{u}_L, \hat{u}_H such that $b_L(x_0|\hat{u}_L) = b_H(x_0|\hat{u}_H)$ for some $x_0 \in [0, \bar{x}(\hat{u}_L)]$. If $\frac{\partial}{\partial x} b_H(x_0|\hat{u}_H) \leq \frac{\partial}{\partial x} b_L(x_0|\hat{u}_L)$, then $b_H(x|\hat{u}_H) > b_L(x|\hat{u}_L)$ at all $x < x_0$.*

CP generalizes SCP , the key feature of the application from Section 3. It is implied by the SCP , but not the converse. CP says that if the high type's BIC is flatter than the low type's at a point of intersection, then the high type's BIC lies everywhere above the low type's at all points to the left. Further, it implies that if the low type's indifference curve is flatter at

³⁰For BICs to satisfy single-crossing, it is necessary that the test is not RC-Informative at any x and sufficient that for all x, g , $R(g|x) < \bar{R}$, for $\bar{R} > 1$ small enough.

³¹For a given \hat{u}_L in the feasible set of utility levels, there can exist two signaling levels $\bar{x}' < \bar{x}$ such that $b_L(\bar{x}', \hat{u}_L) = b_L(\bar{x}, \hat{u}_L) = 1$. By A.2, it must be that $\bar{x}' < x_L^* < x_H^*$.

a point of intersection, then it lies everywhere below to the right. In Appendix B, we verify that CP arises naturally in signaling models with grades. In particular, we demonstrate that for any constant testing technology, Applications 1-3 discussed in Section 2.2 satisfy CP and that Applications 4-5 satisfy CP for the specified testing technologies.³²

It is possible to construct utility functions and testing technologies such that CP fails by drastically changing Z or the rate at which the informativeness of the test changes over a small interval of signaling levels. Preventing such rapid changes is sufficient to ensure the property holds. We assume that CP holds for the remainder of this section.

Under CP , the sufficient conditions of Propositions 3 and 4 are also necessary for their results.

Proposition 5. *If the test is not RC-Informative at \bar{x} , then for all $\mu_0 \in (0, 1)$, the LCSE is an equilibrium.*

Proposition 6. *If there exists an $x \in (x_H^*, \bar{x}]$ such that the test is not RC-Informative at x , then as $\mu_0 \rightarrow 1$, the set of equilibria does not converge to the complete-information outcome by either notion of convergence (type-by-type or in total mass).*

As evidenced by our earlier discussion, in many environments the informativeness of the test may be *increasing* in x . More importantly for equilibrium analysis is whether the informativeness is increasing *relative* to utility differences. If $\mathbb{E}[R(g|x)|L, x] - Z(x, 1)$ is increasing in x , results can be strengthened and simplified. In this case, if the test fails RC-Informativeness at \bar{x} , then the test is not RC-Informative at any $x < \bar{x}$. Under this stronger hypothesis, Proposition 5 can be strengthened to: the LCSE is the *unique* equilibrium for all priors. Further, convergence results hinge only on the simplified condition of RC-Informativeness at $x = x_H^*$.

Full Characterization

We now turn to characterizing the set of equilibria under CP . The procedure involves two steps and provides a generalization of the analysis from Section 3. The first step is to identify the high type's optimal (x, μ) pair for each candidate low-type utility level. Such pairs were straightforward to identify in the application studied in Section 3 (either $\mu = \mu^*$ or $x = 0$). A more general approach involves solving a parameterized maximization problem and is employed here. Fix a candidate equilibrium utility level for the low type, \hat{u}_L , and consider the high type seeking to maximize his expected utility given the schedule of interim beliefs

³²A constant testing technology is sufficient, but not necessary, to ensure that CP holds within Applications 1-3 (see Section 4.2).

$b_L(x|\hat{u}_L)$. The high type's (parameterized) problem is

$$\max_{x \in [0, \bar{x}(\hat{u}_L)]} u_H(x, b_L(x|\hat{u}_L)) \quad (\star)$$

Recall that in any equilibrium, the low type must obtain an expected utility of at least $\underline{u}_L \equiv U_L(x_L^*, 0)$. Let $\bar{u}_L \equiv \max_x U_L(x, 1)$ denote the maximum possible expected utility attainable by the low type. For each $\hat{u}_L \in [\underline{u}_L, \bar{u}_L]$, (\star) maximizes a continuous function over a compact domain. Existence of a solution is immediate. *CP* ensures that the solution is unique. Let $x_H(\hat{u}_L)$ be the solution to (\star) for a given \hat{u}_L , and let $\mu_H(\hat{u}_L) \equiv b_L(x_H(\hat{u}_L)|\hat{u}_L)$. We will refer to the curve $\{x_H(\hat{u}_L), \mu_H(\hat{u}_L)\}$ for $\hat{u}_L \in [\underline{u}_L, \bar{u}_L]$ as the *solution locus*.³³

Properties of the solution locus depend on the informativeness of the testing technology. If the test is RC-Informative at $\bar{x}(\hat{u}_L)$, then $\mu_H(\hat{u}_L) \in (0, 1)$: the high type maximizes expected utility in (\star) by relying partially on the outcome of the test. If the test is not RC-Informative at $\bar{x}(\hat{u}_L)$, then the locus lies along the upper boundary, $x_H(\hat{u}_L) = \bar{x}(\hat{u}_L)$ and $\mu_H(\hat{u}_L) = 1$: the high type maximizes his expected utility in (\star) by completely separating from the low type using the costly signal.

Figure 6 illustrates an example of the procedure. For each of the low type's plotted BICs, the dotted curves are the BICs for the high type corresponding to interior solutions of (\star) . The dark line running through the solution points is the solution locus. Notice that $\mu_H(\hat{u}_L)$ is non-decreasing in Figure 6. This property is necessary and sufficient to ensure that the equilibrium is generically unique.³⁴ The non-decreasing locus property also arises naturally in signaling models with grades. In Appendix B, we demonstrate that Applications 1-3 discussed in Section 2.2 (for any constant testing technology) and Applications 4-5 (for the specified testing technologies) satisfy the non-decreasing locus property.

For each candidate utility level of the low type, the solution locus identifies the corresponding candidate strategy ($\sigma_H(x_H(\hat{u}_L)) = 1$) and utility level for the high type. The final step is to link candidates to equilibria. The link is made via the equilibrium belief consistency condition as the following proposition demonstrates.

³³To see that the solution to (\star) is unique, fix a \hat{u}_L . If there exists an x' and \hat{u}_H such that $b_L(x'|\hat{u}_L)$ and $b_H(x'|\hat{u}_H)$ are tangent, then *CP* implies that $b_H(x|\hat{u}_H) > b_L(x|\hat{u}_L)$ for all $x \neq x'$. The tangency point is the unique solution to (\star) . If there is not a tangency point for any (x, \hat{u}_H) , then one of the boundaries is the solution. Again, *CP* implies that the corresponding high type's BIC lies everywhere else above the low type's, making the solution unique. From Berge's Theorem of the Maximum, x_H is continuous. Because b_L is also continuous, μ_H is continuous.

³⁴Proposition A.1 in Appendix A characterizes the set of equilibria for arbitrary solution loci under *CP*. Proposition 2 illustrates how multiplicity arises if $\mu_H(\hat{u}_L)$ is constant over an interval, which can obtain only for (at most) a countable set priors.

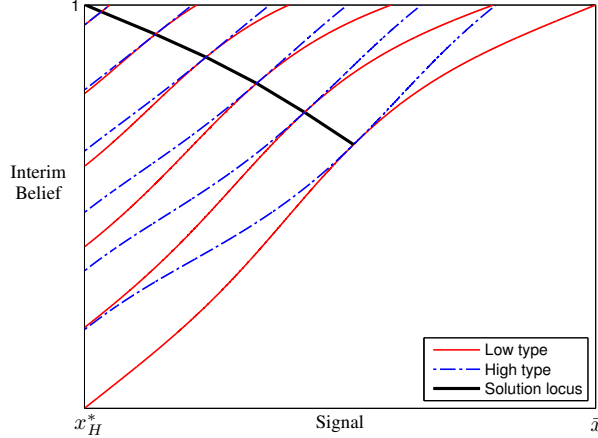


FIGURE 6 – Solving (\star) for each \hat{u}_L

Proposition 7. *Suppose that $\mu_H(\hat{u}_L)$ is non-decreasing. Then, there exists a unique equilibrium for almost all priors, $\mu_0 \in (0, 1)$. Generically, if $\mu_0 \leq \mu_H(\underline{u}_L)$ the equilibrium is:*

- $\sigma_H^*(x_H(\underline{u}_L)) = 1$, $\sigma_L^*(x_H(\underline{u}_L)) = \frac{1 - \mu_H(\underline{u}_L)}{\mu_H(\underline{u}_L)} \frac{\mu_0}{1 - \mu_0}$, and $\sigma_L^*(0) = 1 - \sigma_L^*(x_H(\underline{u}_L))$
- $\mu(x_H(\underline{u}_L)) = \mu_H(\underline{u}_L)$, and $\mu(x) = 0$ if $x \neq x_H(\underline{u}_L)$

If $\mu_0 > \mu_H(\underline{u}_L)$ the equilibrium is full pooling at $x_H(u_L^*)$, where u_L^* satisfies $\mu_H(u_L^*) = \mu_0$.

If the test is RC-Informative at \bar{x} then $\mu_H(\underline{u}_L) < 1$ and the equilibrium involves partial pooling for priors below $\mu_H(\underline{u}_L)$ and full pooling for priors above. On the other hand, if $x_H(\underline{u}_L) = \bar{x}$, then $\mu_H(\underline{u}_L) = 1$, and the LCSE is the unique equilibrium for all $\mu_0 \in (0, 1)$.

4.2 Illustration of the Algorithm with Non-Constant Testing Technology

We illustrate our algorithmic approach for constructing the equilibrium within the job-market signaling environment with a testing technology that varies with the signal. As in Application 5, define the testing technology such that conditional on (x, t) , $G \sim U[0, 1 + V_t \cdot x]$ and fix $V_L = 0$ and $V_H = 1$. This specification implies that $\mathbb{E}[R(g|x)|L, x] = \int_0^1 \frac{1}{1/(1+x)} dg = 1 + x$, and hence the informativeness is increasing in the signal.

We first determine under what conditions the LCSE is an equilibrium. To do so, recall that $U_t(x, \mu_f) = \mathbb{E}[V_t | \mu_f] - C_t k(x)$ and so $Z(x, \mu_f) = \frac{C_L k'(x)}{C_H k'(x)} = \frac{C_L}{C_H}$. Therefore, the test is RC-Informative at x if and only if $\mathbb{E}[R(g|x)|L, x] = 1 + x > \frac{C_L}{C_H}$. Proposition 5 and $\mathbb{E}[R(g|x)|L, x]$ increasing then tell us that if $\frac{C_L}{C_H} \geq 1 + \bar{x} = 1 + k^{-1}(1/C_L)$, then the LCSE is the unique equilibrium for all priors.³⁵ On the other hand if $\frac{C_L}{C_H} < 1 + \bar{x}$, then the LCSE is not an equilibrium (for any prior). We will focus attention on this case.

³⁵Proposition 5 uses *CP*, which can be verified for this specification by following closely the methods in Appendix B.

For a given (x, μ) , the low type's expected offer is

$$w_L(x, \mu) = \int_0^1 \frac{\mu}{\mu + (1 - \mu)(1 + x)} dg = \frac{\mu}{1 + (1 - \mu)x}$$

Therefore, the BIC for the low type is given by

$$b_L(x|\hat{u}_L) = \frac{(1 + x)(\hat{u}_L + C_L k(x))}{1 + x(\hat{u}_L + C_L k(x))}$$

and thus $w_H(x, b_L(x|\hat{u}_L)) = \frac{\hat{u}_L + x - C_L k(x)}{1 + x}$. Noting that $\bar{x}(\hat{u}_L) = k^{-1}((1 - \hat{u}_L)/C_L)$, the maximization problem for the high type (\star) becomes

$$\begin{aligned} \max_x & \frac{\hat{u}_L + x - C_L k(x)}{1 + x} - C_H k'(x) \\ \text{s.t. } & 0 \leq x \leq k^{-1}((1 - \hat{u}_L)/C_L) \end{aligned}$$

The first-order condition for an interior solution requires that

$$1 - \hat{u}_L - C_L k(x) - (1 + x) \left(1 + x - \frac{C_L}{C_H}\right) C_H k'(x) = 0 \quad (4)$$

for some $x \in [0, \bar{x}(\hat{u}_L)]$. Let x^c be the minimal signal such that the test is RC-Informative for all $x > x^c$ (i.e., $x^c \equiv \frac{C_L}{C_H} - 1$), and let u_L^c be defined implicitly by $k^{-1}((1 - u_L^c)/C_L) = x^c$. For any feasible $\hat{u}_L \geq \hat{u}_L^c$ and all $x < \bar{x}(\hat{u}_L)$, the left-hand side of (4) is strictly positive and a corner solution obtains: $x_H(\hat{u}_L) = \bar{x}(\hat{u}_L)$, $\mu_H(\hat{u}_L) = 1$. For any feasible $\hat{u}_L < u_L^c$, there exists a unique $x_H(\hat{u}_L) < \bar{x}(\hat{u}_L)$, defined implicitly as the solution to (4), which solves (\star).³⁶ By inspection, the function $x_H(\hat{u}_L)$ is strictly decreasing. Inserting the inverse function $x_H^{-1}(x)$ for \hat{u}_L into b_L gives

$$b_L(x|x_H^{-1}(x)) = \frac{(1 + x)(1 + x - \frac{C_L}{C_H})C_H k'(x) - 1}{x(1 + x - \frac{C_L}{C_H})C_H k'(x) - 1}$$

which provides a representation for the interior portion of the solution locus (i.e., for $x \in (x^c, \bar{x}]$). The above expression is strictly decreasing in x , which then implies the solution locus is strictly increasing for $\hat{u}_L < u_L^c$. Therefore, for each prior, there exists a unique equilibrium.

Specifying a linear cost function ($C(x) = x$) and the same cost parameter for both types

³⁶To see this, note that the left-hand side of (4) evaluated at $x < x^c$ is strictly positive and strictly negative evaluated at $\bar{x}(\hat{u}_L)$ implying the existence of a solution. Further, the problem is strictly concave on $x \in [x^c, \bar{x}(\hat{u}_L)]$ implying uniqueness.

($C_L = C_H = c$) permits simple closed form expressions for the solution to the algorithm. In particular,

$$x_H(\hat{u}_L) = \sqrt{\frac{1 - \hat{u}_L + c}{c}} - 1, \quad \mu_H(\hat{u}_L) = \frac{\hat{u}_L^2 - 2c - \hat{u}_L \left(c + \sqrt{c(1 - \hat{u}_L + c)} \right)}{\hat{u}_L^2 - 4c}$$

Notice that $\mu_H(\underline{u}_L) = \mu_H(0) = \frac{1}{2}$ and hence by Theorem 7, for $\mu_0 < \frac{1}{2}$, the equilibrium involves partial pooling at the signal $x_H(\underline{u}_L) = \sqrt{1 + \frac{1}{c}} - 1$. For $\mu_0 \geq \frac{1}{2}$, the equilibrium involves full pooling at a signal that is decreasing in the prior. Since $x_H(\bar{u}_L) = x_H(1) = 0$, the equilibrium converges to the complete-information outcome as the prior goes to one. It should be noted that convergence does not obtain if $C_H < C_L$, in which case $x_H(\bar{u}_L) = \sqrt{\frac{C_L}{C_H}} - 1 > 0$; since the test is completely uninformative at $x = 0$, the high type will always rely (in part) on his cost advantage even as the prior goes to one.

5 More than Two Types

Let us expand the type space to be $\{1, \dots, N\}$, $N > 2$. The receivers' prior, μ_0 , is a probability distribution with full support over the type space. When describing beliefs (be they prior, interim or final) we use superscripts to denote the probability assigned to the various types (e.g., μ_0^t is the probability the prior assigns to type t). We now assume that after observing x and g , the receivers each choose some action, and that all payoff-relevant information (from the sender's perspective) from a profile of receivers' actions can be summarized by a scalar $a \in \mathbb{R}_+$. For example, in equity issuance, a represents the market clearing price of the sale—the values of other bids have no direct bearing on the sender's payoff. We therefore, continue to represent the sender's utility, U_t , as a function of two arguments, with the first argument remaining x and the second argument now a instead of μ_f . U_t remains differentiable, with $U_{t,2} > 0$.

Let $a^*(\mu_f, x)$ denote the unique value of a that results when each receiver is playing a best response in the continuation game that follows the sender's choice of x and the common final belief μ_f . For all x and t , a^* is differentiable in x and μ^t . The sender's utility functions continues to satisfy A.1–A.2.³⁷ To simplify analysis, we replace A.3 and A.4 with the following

A.3' For any t and fixed μ_f , $U_t(x, a^*(\mu_f, x))$ is strictly decreasing in x .

A.4' For any x, a and $t' \neq t$, $U_{t,2}(x, a) = U_{t',2}(x, a)$.

³⁷A.1, the weak Spence-Mirrlees condition, is generalized in the standard way. In A.2, $\mu_{t'}$ is replaced with $a^*(\mu_{t'}, x)$.

A.3' says that the signaling action is wasteful. This implies that, in equilibrium, the sender never *directly* gains from signaling and that $x_t^* = 0$ for all t . A.4' implies that types do not differ in their risk preference regarding lotteries over values of a .³⁸ Finally, in line with our applications, we assume that

A.5 For all x , if μ_f first-order stochastically dominates μ'_f , then $a^*(\mu_f, x) > a^*(\mu'_f, x)$

That is, when the final belief puts unambiguously more weight on higher types, the equilibrium response from the receivers is more favorable to the sender.

Let

$$Z_{t',t}(x, a^*) = \frac{\frac{d}{dx}U_{t'}(x, a^*)}{U_{t',2}(x, a^*)} \bigg/ \frac{\frac{d}{dx}U_t(x, a^*)}{U_{t,2}(x, a^*)}$$

A testing technology is now a collection of densities $\{f_t(\cdot|x)\}_{t=1}^N$, which we assume satisfies strict MLRP for all x . Let $R_{t',t}(g|x) \equiv f_{t'}(g|x)/f_t(g|x)$. Our first result generalizes the insight that informative grades eliminate separation in equilibrium.³⁹ (Recall that S_t denotes the support of Υ_t .)

Proposition 8. *Fix a type $t > 1$. If for all x and $t' < t$,*

$$\begin{aligned} \mathbb{E}[R_{t',t}(g|x)|t'] &> Z_{t',t}(x, a^*(\mu_t, x)) \text{ for all } t' < t, \text{ and} \\ \mathbb{E}[R_{t',t}(g|x)|t'] &< Z_{t',t}(x, a^*(\mu_t, x)) \text{ for all } t' > t \end{aligned} \tag{RC}_t$$

then there does not exist an equilibrium in which type t separates; for any signal x such that $x \in S_t^$ there exists a type $\tilde{t} \neq t$ such that $x \in S_{\tilde{t}}^*$.*

RC_t is a generalization of RC-Informativeness and can be loosely paraphrased as: for all x , strict MLRP must not only hold, but hold “strongly” enough relative to differences in preferences over signals. The result is then similar to Proposition 3: if grades are informative enough, then in equilibrium only the lowest type may assign positive probability to a signal that perfectly identifies him—which, of course, eliminates full separation, as well as other strategy profiles such as countersignaling from Feltovich et al. (2002), where the medium

³⁸That is, U_t specifies both type t 's utility tradeoff between the costly signal and market action at any given (x, a) , as well as his risk preferences on lotteries over a for any given x . A.4' maintains that any difference between the preferences of different types is found in the former, not the latter. Notice that without grades A.4' is without loss of generality in the following sense: for any $\{U_t\}_{t=1}^N$ not satisfying A.4', there exists $\{\tilde{U}_t\}_{t=1}^N$ satisfying A.4' such that both collections produce the same indifference curves and set of equilibria. Without grades, the market response, $a^*(\mu_f, x)$, is deterministic, and hence the type-varying risk preferences of the sender have no bearing on behavior.

³⁹The D1 refinement for the N -type model can be stated as follows. Fix an equilibrium endowing expected utilities $\{u_t^*\}_{t \in \{1, \dots, N\}}$. Consider a signal x that is not in the support of any type's strategy. If there exists t, t' such that $B_{t'}(x, u_{t'}^*) \subset B_t(x, u_t^*)$, then D1 requires that the interim belief following x assigns zero weight to t' (where \subset denotes strict inclusion).

type separates while low and high types pool. Once again, sufficiently informative tests (relative to cost advantages) necessitate that grades will convey meaningful information *in equilibrium*.^{40,41}

Our second result generalizes the convergence established in Proposition 4.

Proposition 9. *Fix a type $t > 1$. If RC_t holds, then as $\mu_0 \rightarrow \mu_t$, the set of equilibria converges in total mass to the complete-information outcome. Further, as $\mu_0 \rightarrow \mu_N$, the convergence holds type-by-type.*

Again we find that RC-Informativeness is sufficient to both eliminate separation and ensure convergence to the complete-information outcome. Notice that if all types have the same utility function, then strict MLRP implies that RC_t holds for all $t > 1$. That is, in an environment without cost advantages, no type $t > 1$ separates in any equilibrium, and the convergence properties from Proposition 9 hold.

Propositions 8 and 9 show that the major economic insights demonstrated in the two-type model extend to the larger type space. In general, D1 and stronger stability-based refinements do not select a unique equilibrium with the larger type space. It is possible to construct examples in which both full pooling and a generalized version of the partial-pooling equilibria from the two-type model persist.

To see why this is, consider the job-market signal model—in which V_t is increasing in t and $a^*(\mu_f, x) = \mathbb{E}[V_t|\mu_f]$. Recall that without grades, if the strict Spence-Mirrlees condition holds, then D1 uniquely selects the LCSE (Cho and Kreps, 1987; Cho and Sobel, 1990). This result fails when grades are present. The key difference is in how the market interim belief, μ , maps into best responses. In the gradeless model, the sender will be offered his expected value given μ . Hence, for any N , $\mathbb{E}[V_t|\mu]$ is a sufficient statistic for any μ . That is, any μ can be reduced to a scalar—the expected market value given that belief. Comparing the sets of interim beliefs that make a deviation profitable for each type is therefore reduced to ascertaining which type needs a higher value of $\mathbb{E}[V_t|\mu]$ to make the deviation in question profitable.

When grades are present, this is no longer true. Because the interim belief will be updated based on the grade, it cannot be reduced to a scalar (unless $N = 2$). For instance, in a three-

⁴⁰Countersignaling requires that grades must be relatively poor at differentiating the *low* type from the *medium* type even without refinements. However, there exists generic parameters for which D1 eliminates otherwise tenable countersignaling strategy profiles.

⁴¹Proposition 8 does not hold for a continuum of types. For example, the LCSE survives any stability-based refinement for the simple reason that there is nothing off-path to deviate to except signals greater than the one chosen by the highest type, which are strictly dominated by mimicking the highest type. However, unlike the gradeless model under single-crossing, with a continuum of types (Mailath, 1987; Ramey, 1996), the presence of grades allows outcomes other than the LCSE to survive refinements, lending credibility to their study.

type example with $V_2 = 0.5V_1 + 0.5V_3$, the interim belief $\mu = (\mu^1, \mu^2, \mu^3) = (0, 1, 0)$ has very different implications for each type’s expected offer from the interim belief $\mu' = (0.5, 0, 0.5)$. This is despite $\mathbb{E}[V_t|\mu] = \mathbb{E}[V_t|\mu']$. Under the first interim belief, the realization of the grade will have no effect on the final belief (and offer), while the grade will be quite important under the second. Without this equivalence, refinements based on comparing the relevant belief sets lose some of their bite.⁴²

6 Conclusion

Grades are a prevalent force in many incomplete information environments, and we believe that they often convey meaningful information. The insight of Spence (1973) is that even without grades, senders can signal valuable information to the market. However, in a strategic setting there is a subtle interaction between the two channels. To understand such environments we have characterized the equilibria for a class of two-type signaling games.

In equilibria that are robust to standard refinements, the high type resolves the trade-off between how much to rely on each of the two transmission mechanisms. If the test is sufficiently informative relative to the signal, the high type relies at least in part on the information contained in the test, and the equilibrium involves some degree of pooling. Further, the addition of grades yields a more intuitive outcome as the prior puts higher weight on the high type. Both type’s gain in utility, and if the grade remains sufficiently informative relative to the cost advantage, the equilibrium converges to full pooling at the efficient signaling level—long thought to be the appropriate convergence, but not achieved in gradeless signaling models. Finally, we extended these results to a model with N types.

6.1 Related Literature

Noisy Signaling

It is important to distinguish our approach from the work done on “noisy signaling” (Matthews and Mirman, 1983; Carlsson and Dasgupta, 1997). In a noisy signaling model, receivers do not perfectly observe the sender’s choice, but rather a noisy “signal” whose distribution depends on it (for example, a random shock is added to the sender’s choice). This specification is well-suited to environments where receivers cannot observe the sender’s decision and must make inferences based on resultant data—for example, in Matthews and Mirman (1983) a

⁴²When $N = 2$ a given increase in the interim belief has two effects: *i*) raising the expected μ_f of both types, and *ii*) increasing or decreasing the importance of the grade (depending on from where and to where the belief is increased). When $N > 2$ we can decouple these two effects. For a given set of types, it is now often possible to find changes to a given belief that increase each type’s expected offer and increase the importance of the grade, as well as changes to the same belief that still increase expected offers but decrease the importance of the grade. Intuitively, higher types value the first kind of change more than lower types do and vice versa, reducing the power of stability-based refinements.

potential entrant firm observes only a market price which depends on both the incumbent’s decision and a stochastic market demand. In other economic settings, especially those highlighted above, the model with grades is a more accurate description of the strategic situation: employers observe *both* years of schooling *and* grades on transcripts, consumers observe *both* new car warranties *and* ratings by *Consumer Reports*, etc., rather than a single observation encompassing both aspects. In addition, exogenous information is often informative about the sender’s type, not simply white noise (while the noisy signaling approach does not preclude this feature, it is usually assumed that the distribution of the noise is independent of type). Thus, noisy signaling models do not illustrate the sender’s trade-off between relying on his grade advantage versus his cost advantage.

In noisy signaling models, as in our model, equilibrium behavior depends on the prior. With additive shocks the costly action chosen by each type will tend to the complete-information outcome, as the prior shifts all of its weight toward the high type. However, unlike the model with grades, this is also true as the prior shifts all of its weight toward the *low* type. This illustrates the fundamental difference as to how the equilibrium of each model depends on the prior as discussed in Section 3.3.1. In a noisy signaling model, the reliance is almost entirely dictated by the properties of Bayesian updating—when the prior is extreme, it *cannot* be moved by any reasonable amount of expenditure. In the model with grades, because the sender’s choice is perfectly observable, the sender has the potential to significantly alter the market beliefs starting from any prior—the question is whether he will want to. Whether or not we should expect the high type to choose signals which are vanishingly small when receivers believe his existence extremely unlikely should therefore depend on the setting being modeled.⁴³

Some settings may be accurately modeled by combining elements of both our model and noisy signaling models. Senders have two decisions to make: how costly an observable signal to send (as in our model) and how much unobservable investment to make toward influencing the stochastic measure (as in a noisy signaling model). These models can have many equilibria, including equilibria qualitatively similar to ours and equilibria that rely only on the noisy measure as in a noisy signaling model. The equilibrium selection issue becomes a difficult one because deviations now have both an observable and unobservable component, rendering standard refinements ineffective.

⁴³Finally, in both models, taking the appropriate limit (taking grades to be completely uninformative and taking the variance of the exogenous shock to zero) returns us to the gradeless/noise-free model. In the model with grades, if single-crossing is strict then, for every prior, the equilibrium converges to the LCSE—matching the prediction of the gradeless/noise-free model. In a noisy signaling model, the limit of the equilibrium will depend on the distribution of the shock. In many common setups, such as normally-distributed additive shocks, it can be shown that for every prior the equilibrium converges to the full-pooling outcome—in contrast to the prediction of the gradeless/noise-free model.

Signaling and Grades

Ours is not the first paper to introduce grades in a signaling framework. Weiss (1983) considers a model of education in which students are tested. Unlike our model, he assumes that passing grades are productive in themselves: if two students of the same type receive different grades, they have different market values. The test reveals something that the sender did not already know about himself. Weiss argues that only separating equilibria and full-pooling equilibria at the cheapest signal are reasonable (in a specific formal sense). Fang (2001) demonstrates that in the presence of noisy information, signaling through a seemingly irrelevant activity can be supported in equilibrium. He interprets this activity as “social culture.” In other related work, Angeletos et al. (2006); Angeletos and Pavan (2011) study signaling through policy interventions in a global games setting with exogenous information revelation. One key difference in their model is that each receiver *privately* observes a different piece of information about the sender’s type, whereas the information in our model (i.e., the grade) is *publicly* observed.

Feltovich et al. (2002) examine a three-type signaling model in which the market observes additional information correlated with the sender’s type. The authors identify conditions for countersignaling equilibria to exist. In a countersignaling equilibrium, the high type pools with the low type at the least costly signaling level, while the medium type perfectly separates by incurring strictly positive signaling costs. Alos-Ferrer and Prat (2011) amend the canonical two-type job-market signaling model to one in which after being hired, the sender’s type is gradually revealed via on-the-job employer learning. Market forces raise or lower the sender’s wage as the belief evolves. They show that equilibria involving pooling can survive the Intuitive Criterion when on-the-job revelation of type is fast enough. The assumption that information and wages gradually evolve after the sender has been hired can be mapped into our model by considering: *i*) each realized path of the gradual process as a *grade*, and *ii*) the present value of the sender’s wage stream as the action from receivers in our model.

In contrast to these two papers, we fully characterize the set of equilibria in a two-type model and show that the main economic insights extend to a model with more types. In addition, we allow a more general structure for the external information, including allowing the quality of information to vary with the sender’s costly signal. Finally, we identify a key property (the Crossing Property) that arises naturally in a setting with grades, and—like the single-crossing property in gradeless signaling models—facilitates a tractable analysis and sharp predictions. We demonstrate that this property holds in a variety of economic applications and connect our results to such settings. This provides a theoretical framework for understanding the empirical findings discussed in Section 1.

6.2 Final Remarks

We have used D1 to refine the set of PBE within our model. D1 belongs to a class of refinements derived from the notion of strategic stability (Kohlberg and Mertens, 1986).⁴⁴ It is worth noting that the stronger refinements within this class, such as NWBR and Universal Divinity, yield the same results as D1 in our model (for both the two-type and N -type specifications).

A number of non-stability based refinements have been developed to refine the set of equilibria in signaling games.⁴⁵ Many of these notions eliminate equilibria based on other potential equilibria in order to select more “reasonable” equilibria (for specific games) than their stability-based counterparts. Among these refinements, the one that is perhaps the most relevant to this work is the concept of undefeated equilibria (Mailath et al., 1993). In the gradeless job-market signaling model with two types, the undefeated criterion (uniquely) selects the LCSE when $\mu_0 < \underline{\mu}$ and the efficient full-pooling outcome when $\mu_0 > \underline{\mu}$. Like our model, Mailath et al. (1993) predict the complete-information outcome when the prior is sufficiently large. In fact, part of their motivation for introducing the refinement is to yield a more satisfying prediction in the gradeless job-market signaling model.

Our motivation is quite different. We are interested in understanding signaling environments with grades and how their presence affects equilibrium predictions. Our use of refinements is to gain traction rather than to alter a somewhat undesirable feature prevalent in signaling models. It is appealing that the predictions along this dimension are aligned, but largely coincidental.⁴⁶

In many examples, such as education, the costly signal involves waiting to trade. Static signaling models ignore the dynamic aspects of such an environment. Swinkels (1999) demonstrates that all trade takes place immediately—there is no signaling through delay—in the

⁴⁴Cho and Sobel (1990) show that D1 is equivalent to strategic stability in two-player monotonic signaling games with finite strategy spaces. Our game is monotonic, but we have multiple receivers and continuous strategy spaces. Consider the following alterations to the model in Section 3. First, as suggested by Fudenberg and Tirole (1991), there is a single receiver whose objective is to make an offer as close as possible to the sender’s expected value. Second, each player can choose only from a finite grid of actions. In this game, D1 and stability are equivalent. It can be shown that the D1 equilibria of this game converge to the D1 equilibrium of our model, as the grid approaches our continuous action space.

⁴⁵See for example Mailath et al. (1993), Grossman and Perry (1986), and Hillas (1994).

⁴⁶Mailath et al. (1993) also introduce a related notion, *lexicographically maximum sequential equilibrium* (LMSE), and show that the LMSE is undefeated under four regularity assumptions. While these concepts are not formally defined to refine mixed-strategy equilibria, which are focal in our analysis, they can be adapted in a natural way in order to study their implications when applied to our model. For simplicity, consider the setup in Section 3 and assume the test is RC-Informative. For all $\mu_0 < \underline{\mu}$, a continuum of partial-pooling equilibria are undefeated, while for $\mu_0 > \underline{\mu}$ the efficient full-pooling outcome is the unique undefeated equilibrium. Further, the LMSE corresponds to the D1 equilibrium for priors below $\underline{\mu}$ and to the efficient full-pooling equilibrium for priors above (which also agrees with the D1 equilibrium for priors above μ^*).

gradeless job-market signaling model when preemptive private offers can be made frequently by the market. Kremer and Skrzypacz (2007) amend this analysis by having a grade revealed at a commonly known fixed future date. Signaling occurs as trade does not always happen immediately. Daley and Green (2011) study a dynamic setting, analogous to the model in this paper, in which information is revealed gradually. One interpretation is that this dynamic model relaxes the assumption of a seller's ability to commit to delay trade until a fixed date x . From this standpoint, one could investigate how the timing of information revelation interacts with commitment power to impact trading patterns and welfare. Despite their differences, a few similarities in the equilibria of the two models emerge. In each there is cutoff prior below which the low type mixes between delaying trade and not, while the high type always chooses to delay. The payoffs to each type are constant over priors below the cutoff, with the low type's equal to his market value. Delay is decreasing in the prior, while the expected value to both types is increasing in the prior, once it is above the cutoff.

References

- ALOS-FERRER, C. AND J. PRAT (2011): “Job Market Signaling and Employer Learning,” Mimeo.
- ANGELETOS, G.-M., C. HELLWIG, AND A. PAVAN (2006): “Signaling in a Global Game: Coordination and Policy Traps,” *Journal of Political Economy*, 114(3), 452–485.
- ANGELETOS, G.-M. AND A. PAVAN (2011): “Preempting Speculative Currency Attacks: Robust Predictions in a Global Game with Multiple Equilibria,” Mimeo.
- BANKS, J. AND J. SOBEL (1987): “Equilibrium Selection in Signaling Games,” *Econometrica*, 55, 647–662.
- BLACKWELL, D. (1951): “Comparison of Experiments,” *Proceedings of the Second Berkeley Symposium on Mathematics, Statistics and Probability*, (1950), 93–102.
- BLACKWELL, D. AND M. GIRSHICK (1954): *Theory of Games and Statistical Decisions*, John Wiley & Sons, Inc.
- CARLSSON, H. AND S. DASGUPTA (1997): “Noise-Proof Equilibria in Two-Action Signaling Games,” *Journal of Economic Theory*, 77, 432–460.
- CHO, I. K. AND D. KREPS (1987): “Signaling Games and Stable Equilibria,” *Quarterly Journal of Economics*, 102, 179–221.
- CHO, I. K. AND J. SOBEL (1990): “Strategic Stability and Uniqueness in Signaling Games,” *Journal of Economic Theory*, 50, 381–413.
- DALEY, B. AND B. GREEN (2011): “Waiting for News in the Market for Lemons,” *Econometrica* (Forthcoming).
- FANG, H. (2001): “Social Culture and Economic Performance,” *The American Economic Review*, 91, 924–937.
- FELTOVICH, N., R. HARBAUGH, AND T. TO (2002): “Too Cool for School? Signalling and Countersignalling,” *RAND Journal of Economics*, 33, 630–349.
- FUDENBERG, D. AND J. TIROLE (1991): *Game Theory*, MIT Press.
- GAL-OR, E. (1989): “Warranties as a Signal of Quality,” *Canadian Journal of Economics*, 22, 50–61.
- GROSSMAN, S. AND M. PERRY (1986): “Perfect Sequential Equilibrium,” *Journal of Economic Theory*, 39, 97–119.
- GROSSMAN, S. J. (1981): “The Informational Role of Warranties and Private Disclosure about Product Quality,” *The Journal of Law and Economics*, 24, 461–483.
- HILLAS, J. (1994): “Sequential Equilibria and Stable Sets of Beliefs,” *Journal of Economic Theory*, 64, 78–102.
- HSUEH, P. L. AND D. S. KIDWELL (1988): “Bond Ratings: Are Two Better than One?” *Financial Management*, 17, 46–53.
- JEWITT, I. (2007): “Information Order in Decision and Agency Problems,” <http://individual.utoronto.ca/xianwenshi/reading/jewitt07.pdf>.
- JONES, E. B. AND J. D. JACKSON (1990): “College Grades and Labor Market Rewards,” *The Journal of Human Resources*, 25, 253–266.
- KARLIN, S. (1968): *Total Positivity*, vol. 1, Stanford University Press.
- KIHLSTROM, R. E. AND M. H. RIORDAN (1984): “Advertising as a Signal,” *Journal of Political Economy*, 92, 427–450.

- KOHLBERG, E. AND J.-F. MERTENS (1986): "On the Strategic Stability of Equilibria," *Econometrica*, 54, 1003–1037.
- KREMER, I. AND A. SKRZYPACZ (2007): "Dynamic Signaling and Market Breakdown," *Journal of Economic Theory*, 133, 58–82.
- LEHMANN, E. (1988): "Comparing Location Experiments," *Annals of Statistics*, 16, 521–523.
- LELAND, H. E. AND D. H. PYLE (1977): "Informational Asymmetries, Financial Structure, and Financial Intermediation," *The Journal of Finance*, 32, 371–387.
- LUTZ, N. A. (1989): "Warranties as Signals under Consumer Moral Hazard," *RAND Journal of Economics*, 20, 239–255.
- MAILATH, G. (1987): "Incentive Compatibility in Signaling Games with a Continuum of Types," *Econometrica*, 55, 1349–1365.
- MAILATH, G., M. OKUNO-FUJIWARA, AND A. POSTLEWAITE (1993): "Belief-Based Refinements in Signalling Games," *Journal of Economic Theory*, 60, 241–276.
- MATTHEWS, S. A. AND L. J. MIRMAN (1983): "Equilibrium Limit Pricing: The Effects of Private Information and Stochastic Demand," *Econometrica*, 51, 981–996.
- MILGROM, P. AND J. ROBERTS (1986): "Price and Advertising Signals of Product Quality," *The Journal of Political Economy*, 94, 796–821.
- NELSON, P. (1974): "Advertising as Information," *The Journal of Political Economy*, 82, 729–754.
- RAMEY, G. (1996): "D1 Signaling Equilibria with Multiple Signals and a Continuum of Types," *Journal of Economic Theory*, 69, 508–531.
- RILEY, J. G. (2001): "Silver Signals: Twenty-Five Years of Screening and Signaling," *Journal of Economic Literature*, 39, 432–478.
- ROJSTACZER, S. (2009): <http://gradeinflation.com/> (accessed September 20, 2009).
- SPENCE, M. A. (1973): "Job Market Signaling," *Quarterly Journal of Economics*, 90, 225–243.
- SWINKELS, J. M. (1999): "Education Signalling with Preemptive Offers," *Review of Economic Studies*, 66, 949–970.
- THOMAS, L., S. SHANE, AND K. WEIGELT (1998): "An Empirical Examination of Advertising as a Signal of Product Quality," *Journal of Economic Behavior & Organization*, 37, 415–430.
- TITMAN, S. AND B. TRUEMAN (1986): "Information Quality and The Valuation of New Issues," *Journal of Accounting and Economics*, 8, 159–172.
- WEISS, A. (1983): "A Sorting-cum-Learning Model of Education," *The Journal of Political Economy*, 91, 420–442.

A Appendix

A.1 D1 Equivalence

The way we define D1 is slightly different than how it is defined in Banks and Sobel (1987) and Cho and Kreps (1987). Specifically, we use receiver *interim beliefs* in the definition of $B_t(x, u_t^*)$ rather than [rationalizeable] *receiver best response* profiles. As a result, any non-empty $B_t(x, u_t^*)$ is a subset of the $N-1$ dimensional unit simplex (e.g., $[0, 1]$ in the two-type model), rather than a much more complicated subset of all possible final belief schedules, $\mu_f(\cdot, \cdot)$. This simplification is crucial for making our analysis tractable and, as we now demonstrate, is without loss of generality.

In our model, the receivers' best response is a function $a^*(\mu_f(x, g))$, (where a^* is the identify function in the two-type model). Bayesian updating from the prior μ_0 to the final belief $\mu_f(x, g)$ can be decomposed into a first update from μ_0 to an interim belief $\mu(x)$, then a second update from $\mu(x)$ to $\mu_f(x, g)$ based on the realization of the grade. The second update is purely statistical—it is based only on the commonly-known likelihood ratios, $R_{v',t}(g|x)$. Therefore, given any receiver best response profile, the interim belief is sufficient to compute the sender's expected payoff, allowing us to pose the refinement using interim beliefs.

A.2 Proofs

Proof of Lemma 2. First note that w_t is twice continuously differentiable and meets the criteria for exchanging the order of integration and differentiation by the functional form of the integrand in (3) and assumptions on f_t . It is immediate that $w_H(\mu) = w_L(\mu)$ at $\mu = 0, 1$ and hence $\int_0^1 (w'_H(\mu) - w'_L(\mu))d\mu = 0$. Because the test is statistically informative, $w_H(\mu) > w_L(\mu)$ for all $\mu \in (0, 1)$. This immediately implies that $w'_H(0) \geq w'_L(0)$. Moreover, $w'_H(\mu) - w'_L(\mu)$ is strictly decreasing, as shown below.

$$\begin{aligned} \frac{d}{d\mu}(w'_H(\mu) - w'_L(\mu)) &= (V_H - V_L) \int \frac{d^2}{d\mu^2} \left(\frac{\mu}{\mu + (1 - \mu)R(g)} \right) (f_H(g) - f_L(g))dg \\ &= -2(V_H - V_L) \int \frac{(1 - R(g))^2}{(\mu + (1 - \mu)R(g))^3} f_L(g)dg < 0 \end{aligned}$$

Since $w'_H - w'_L$ is decreasing and integrates to zero, it must be strictly positive at $\mu = 0$. Therefore, it crosses zero exactly once and from above. \square

Proof of Lemma 1. Fix \hat{u}_t . By definition, $u_t(x, b_t(x|\hat{u}_t)) = \hat{u}_t$. Total differentiation of both sides with respect to x gives $\frac{\partial u_t}{\partial x} + \frac{\partial u_t}{\partial \mu} \frac{\partial b_t}{\partial x} = 0$, hence $\frac{\partial b_t}{\partial x} = -\frac{\partial u_t}{\partial x} / \frac{\partial u_t}{\partial \mu} = \frac{C_t k'(x)}{w'_t(\mu)}$, and therefore $\frac{\partial b_H}{\partial x} \leq \frac{\partial b_L}{\partial x} \Leftrightarrow \frac{w'_L(\mu)}{w'_H(\mu)} \leq \frac{C_L}{C_H}$. To verify *SCP* it is sufficient to show that $w'_L(\mu)/w'_H(\mu)$ is strictly increasing (i.e., $w'_t(\mu)$ is strictly log-supermodular).⁴⁷ Order the grade space so that $R(g) = f_L(g)/f_H(g)$ is weakly increasing over the common support. Define $h(\mu, g) \equiv R(g)/(\mu + (1 - \mu)R(g))^2$. Note that $w'_t(\mu) = \int h(\mu, g)f_t(g)dg$. Moreover, $h(\mu, g)$ is log-supermodular since for any $g' > g$,

$$\frac{d}{d\mu} \left(\frac{h(\mu, g')}{h(\mu, g)} \right) = 2 \frac{R(g')}{R(g)} (R(g') - R(g)) \frac{(\mu + (1 - \mu)R(g))}{(\mu + (1 - \mu)R(g'))^3} \geq 0$$

⁴⁷The proof method used here is adapted from Karlin (1968) (Chapter 3, Proposition 5.1).

Since $w'_t > 0$, it is enough to show that for any $\mu' > \mu$, $w'_L(\mu')w'_H(\mu) - w'_L(\mu)w'_H(\mu') > 0$:

$$\begin{aligned}
& w'_L(\mu')w'_H(\mu) - w'_L(\mu)w'_H(\mu') \\
&= \iint h(\mu', g')f_L(g')h(\mu, g)f_H(g)dgdg' - \iint h(\mu, g)f_L(g)h(\mu', g')f_H(g')dg'dg \\
&= \iint \frac{f_L(g')}{f_H(g')}h(\mu', g')h(\mu, g)f_H(g)f_H(g')dgdg' - \iint \frac{f_L(g)}{f_H(g)}h(\mu, g)h(\mu', g')f_H(g)f_H(g')dgdg'
\end{aligned}$$

Decompose the region of integration into two sets, $g' > g$ and $g' < g$. Then convert the second region into the first by the change of variable $g \rightarrow g'$, $g' \rightarrow g$. By doing so, the above can be written as

$$\iint_{g' > g} \left[\frac{f_L(g')}{f_H(g')} - \frac{f_L(g)}{f_H(g)} \right] [h(\mu', g')h(\mu, g) - h(\mu', g)h(\mu, g')] f_H(g)f_H(g')dgdg' \quad (5)$$

Both bracketed terms of the integrand are non-negative for all $g' > g$. Hence, we can bound (5) from below by integrating over a subset of the region. Let $\mathcal{G}_- \equiv \{g : f_H(g) > f_L(g)\}$ and $\mathcal{G}_+ \equiv \{g : f_H(g) < f_L(g)\}$. Then (5) is bounded below by

$$\int_{g' \in \mathcal{G}_+} \int_{g \in \mathcal{G}_-} \left[\frac{f_L(g')}{f_H(g')} - \frac{f_L(g)}{f_H(g)} \right] [h(\mu', g')h(\mu, g) - h(\mu', g)h(\mu, g')] f_H(g)f_H(g')dgdg' \quad (6)$$

Both bracketed terms are strictly positive over the region of integration. Since the test is informative, the region has strictly positive measure implying that (6) is strictly positive and therefore so too is (5). \square

Proof of Lemma 3. At $\mu = 1$, by direct calculation the slope of the BIC is $\frac{C_H}{(V_H - V_L)}$ for the high type and $\frac{C_L}{(V_H - V_L) \int R(g)f_L(g)dg} = \frac{C_L}{(V_H - V_L)\mathbb{E}[R(g)|L]}$ for the low type. The high type's BIC is steeper if and only if $\frac{C_H}{(V_H - V_L)} > \frac{C_L}{(V_H - V_L)\mathbb{E}[R(g)|L]} \Leftrightarrow \mathbb{E}[R(g)|L] > \frac{C_L}{C_H}$, which establishes the equivalence of 1 and 2. The equivalence of 2 and 3 is immediate from *SCP*. \square

Proof of Proposition 1. Fix a prior μ_0 . Suppose that the test is RC-Informative. Grades convey information in equilibrium if and only if there is pooling. Therefore, it is sufficient to show that all separating equilibria are Pareto dominated. First, notice that, by definition, the LCSE Pareto dominates all other separating equilibria. Second, let \underline{u}_H be the LCSE utility for the high type. By Lemma 3, the test being RC-Informative implies that $\frac{\partial}{\partial x} b_H(\bar{x}|\underline{u}_H) > \frac{\partial}{\partial x} b_L(\bar{x}|V_L)$. Hence, there exists an $\epsilon > 0$ such that $b_L(\bar{x} - \epsilon|V_L) > \mu_0$ and $b_L(\bar{x} - \epsilon|V_L) \in B_H(\bar{x} - \epsilon, \underline{u}_H)$. The following equilibrium Pareto dominates the LCSE: $\sigma_H(\bar{x} - \epsilon) = 1$, $\sigma_L(\bar{x} - \epsilon) = \frac{1 - b_L(\bar{x} - \epsilon|V_L)}{b_L(\bar{x} - \epsilon|V_L)} \frac{\mu_0}{1 - \mu_0}$ and $\sigma_L(0) = 1 - \sigma_L(\bar{x} - \epsilon)$, with $\mu(\bar{x} - \epsilon) = b_L(\bar{x} - \epsilon|V_L)$ and $\mu(x) = 0$ for all $x \neq \bar{x} - \epsilon$.

Now suppose the test is not RC-Informative. *SCP* and Lemma 3 imply that BICs satisfy the single-crossing property. The result is well-known for this case (see Mailath et al. (1993) for instance). \square

We now turn to characterizing the set of D1 equilibria in the general two-type model (Section 4) under *CP*, with the characterization in Section 3 following as a special case. We start by proving Lemma 4, which does not rely on *CP*.

Proof of Lemma 4. Fix a payoff vector $\{\hat{u}_H, \hat{u}_L\}$, with BICs $b_H(x|\hat{u}_H)$, $b_L(x|\hat{u}_L)$ such that $\exists x'$ at which $b_L(x'|\hat{u}_L) > b_H(x'|\hat{u}_H)$. Suppose S_H^* , S_L^* are the supports of equilibrium strategies endowing payoffs \hat{u}_H, \hat{u}_L . If $x' \in S_H^*$ and $x' \in S_L^*$, then belief consistency requires $\mu(x') = b_H(x'|\hat{u}_H) = b_L(x'|\hat{u}_L)$ contradicting the premise that $b_L(x'|\hat{u}_L) > b_H(x'|\hat{u}_H)$. If $x' \in S_H^*$ and $x' \notin S_L^*$, then belief consistency requires $\mu(x') = b_H(x'|\hat{u}_H) = 1 \geq b_L(x'|\hat{u}_L)$, again contradicting the premise. If $x' \notin S_H^*$ and $x' \in S_L^*$, then belief consistency requires $\mu(x') = b_L(x'|\hat{u}_L) = 0 \leq b_H(x'|\hat{u}_H)$, contradicting the premise. Finally, if $x' \notin S_L^* \cup S_H^*$ and $b_L(x'|\hat{u}_L) > b_H(x'|\hat{u}_H)$, then D1 requires $\mu(x') = 1 > b_H(x'|\hat{u}_H)$, implying the high type can profitably deviate to x' . Hence, $\{\hat{u}_H, \hat{u}_L\}$ cannot be supported by any D1 equilibrium. \square

The following proposition characterizes the set of all D1 equilibria in the two-type model under the Crossing Property and for an arbitrary solution locus.

Proposition A.1. *Under CP, the set of D1 Equilibria is as follows:*

- Full Pooling Equilibria: Consider a point on the solution locus $(x_H(\hat{u}_L), \mu_H(\hat{u}_L))$, $\hat{u}_L \in [\underline{u}_L, \bar{u}_L]$. Full pooling at $x_H(\hat{u}_L)$ is a D1 equilibrium if and only if $\mu_0 = \mu_H(\hat{u}_L)$.
- Partial or Full Separation: For any given prior $\mu_0 \leq \mu_H(\underline{u}_L)$, the equilibrium identified in Proposition 7 for this case satisfies D1.

In all equilibria, $\mu(x) = 0$ for all x off the equilibrium path. There are no other D1 equilibria.

Proof. Verifying that the proposed strategies and beliefs constitute a PBE is routine. To see that the off-path beliefs satisfy D1, notice that in all equilibria $S_H^* = \{x_H(u_L^*)\}$. Therefore, by CP, $B_H(x, u_H^*) \subseteq B_L(x, u_L^*)$ for all x off path. Hence, the off-the-path beliefs satisfy D1.

We now eliminate all candidate equilibria not identified in the proposition. First, we demonstrate that if u_L^* is the low type's payoff in equilibrium, then it must be that $\sigma_H^*(x_H(u_L^*)) = 1$ and $\mu(x_H(u_L^*)) = \mu_H(u_L^*)$. Second, we show that, for each prior, only the equilibria identified in the proposition are compatible with the first claim and the belief-consistency condition for equilibrium.

To see the first claim, fix an equilibrium endowing payoffs u_L^*, u_H^* . Lemma 4 implies that $b_L(x|u_L^*) \leq b_H(x|u_H^*)$ for all $x < \bar{x}(u_L^*)$. In addition, for all $x \in S_H^*$, $\mu(x) = b_H(x|u_H^*) \leq b_L(x|u_L^*)$ to ensure the appropriate payoff for the high type as well as incentive compatibility for the low type. Hence, for all $x \in S_H^*$, $b_L(x|u_L^*) = b_H(x|u_H^*) = \mu(x)$. CP implies that $S_H^* = \{x_H(u_L^*)\}$.

For the second claim, fix a μ_0 . If there exists an equilibrium such that $u_L^* > \underline{u}_L$, it must be that $S_L^* \subseteq S_H^* = \{x_H(u_L^*)\}$. Because S_L^* must be non-empty, the equilibrium must involve full pooling at $x_H(u_L^*)$. This will satisfy $\mu(x_H(u_L^*)) = \mu_H(u_L^*)$ and belief consistency if and only if $\mu_H(u_L^*) = \mu_0$. Finally, if there exists an equilibrium such that $u_L^* = \underline{u}_L$, then $S_H^* = \{x_H(\underline{u}_L)\}$, $\mu(x_H(\underline{u}_L)) = \mu_H(\underline{u}_L)$, and $S_L^* \subseteq \{0, x_H(\underline{u}_L)\}$. Belief consistency then requires that $\mu_H(\underline{u}_L) \geq \mu_0$. It is immediate that the (mixed) strategy identified in the proposition is the unique one consistent with Bayesian updating by the receivers. This eliminates all equilibria but those put forth in the proposition. \square

Proof of Proposition 2. Follows directly from Proposition A.1 and the following lemma, which characterizes the solution locus in the job-market signaling application from Section 3. \square

Lemma A.1. *In the job-market signaling application from Section 3, the solution locus is as follows:*

$$\mu_H(\hat{u}_L) = \begin{cases} \mu^* & \hat{u}_L \leq u_L(0, \mu^*) \\ b_L(0|\hat{u}_L) & \hat{u}_L > u_L(0, \mu^*) \end{cases}, \quad x_H(\hat{u}_L) = \begin{cases} b_L^{-1}(\mu^*|\hat{u}_L) & \hat{u}_L \leq u_L(0, \mu^*) \\ 0 & \hat{u}_L > u_L(0, \mu^*) \end{cases}$$

Proof. Consider first the case in which the test is RC-Informative. By the Strong Crossing Property, the only candidate for a point of tangency (and thus a maximizer of (\star) for some \hat{u}_L) is at $\mu^* < 1$. If $\hat{u}_L \leq u_L(0, \mu^*)$, then the point of tangency is achieved at $(x, \mu) = (b_L^{-1}(\mu^*|\hat{u}_L), \mu^*)$. If $\hat{u}_L > u_L(0, \mu^*)$, then tangency cannot be achieved since the high type's BICs are steeper than the low type's at all (x, μ) that deliver \hat{u}_L to the low type. In this case, the only possible solution is at $x = 0$, which therefore necessitates $\mu_H(\hat{u}_L) = b_L(0|\hat{u}_L)$. If the test is not RC-Informative then $\mu^* = 1$ and the high type's BICs are everywhere flatter. Hence, $\mu_H(\hat{u}_L) = 1$ for all $\hat{u}_L \in [V_L, V_H]$, which necessitates that $x_H(\hat{u}_L) = b_L^{-1}(1|\hat{u}_L)$ as implied by the lemma. \square

Proof of Proposition 3. Consider any separating candidate equilibrium profile where the high type chooses $x^s > \bar{x}$ and garners utility u_H^s . Let $x' \in (\bar{x}, x^s)$. $B_L(x', \underline{u}_L) = \emptyset$ and $B_H(x', u_H^s) \neq \emptyset$. D1 mandates that $\mu(x') = 1$, making a deviation to x' profitable for the high type, breaking the equilibrium.

Consider now the LCSE as a candidate (D1) equilibrium, letting \underline{u}_H be the LCSE utility for the high type. At $\mu = 1$, by direct calculation the slope of the BIC is $-\frac{U_{H,1}(x,1)}{U_{H,2}(x,1)}$ for the high type and $-\frac{U_{L,1}(x,1)}{U_{L,2}(x,1)} \cdot \frac{1}{\mathbb{E}[R(g|x)|L,x]}$ for the low type. The high type's BIC is steeper at $(x, 1)$ if and only if $\mathbb{E}[R(g|x)|L, x] > Z(x, 1)$. Therefore, if the test is RC-Informative at \bar{x} , then $\frac{\partial}{\partial x} b_H(\bar{x}|\underline{u}_H) > \frac{\partial}{\partial x} b_L(\bar{x}|\underline{u}_L)$. Hence, there exists $\epsilon > 0$ such that $b_H(\bar{x} - \epsilon|\underline{u}_H) < b_L(\bar{x} - \epsilon|\underline{u}_L)$ contradicting Lemma 4. Hence all D1 equilibria involve pooling. \square

Proof of Proposition 4. Let $\{\mu_0^k\}$ be any sequence of priors that converges to 1, and $(\Upsilon_L^{*,k}, \Upsilon_H^{*,k})$ be an equilibrium strategy profile for prior μ_0^k . Then for any $\epsilon > 0$ there exists an n such that for all $k > n$

- There exists a $X^k \subseteq S_H^{*,k}$ such that, for all $x \in X^k$: $1 - \mu^{*,k}(x) < \epsilon$.
- The total mass attributed to $\{x : x \notin X^k\}$ by $\Upsilon_H^{*,k}$ is less than δ , with $\delta \rightarrow 0$ as $\epsilon \rightarrow 0$.

These follows easily from the fact that the μ_0^k assigns vanishingly small weight to the low type. It is therefore sufficient to show that as $\epsilon \rightarrow 0$, $X^k \rightarrow \{x_H^*\}$. Now fix any $x' > x_H^*$. RC-Informativeness establishes that $\frac{\partial}{\partial x} b_H(x'|u_H(x', 1)) > \frac{\partial}{\partial x} b_L(x'|u_L(x', 1))$ (see proof of Proposition 3). U_L and U_H differentiable implies that for η small enough $\frac{\partial}{\partial x} b_H(x'|u_H(x', 1 - \eta)) > \frac{\partial}{\partial x} b_L(x'|u_L(x', 1 - \eta))$. Therefore, as $\epsilon \rightarrow 0$, Lemma 4 implies that $x' \notin X^k$ for all k large enough, giving the result. \square

Proof of Proposition 5. Again, letting \underline{u}_H be the LCSE utility for the high type, if the test is not RC-Informative at \bar{x} , then $\frac{\partial}{\partial x} b_H(\bar{x}|\underline{u}_H) \leq \frac{\partial}{\partial x} b_L(\bar{x}|\underline{u}_L)$ (see proof of Proposition 3). By CP, $b_H(x'|\underline{u}_H) > b_L(x'|\underline{u}_L)$ for all $x' < \bar{x}$, implying $x_H(\underline{u}_L) = \bar{x}$. Result follows from Proposition A.1. \square

Proof of Proposition 6. Suppose there exists an $x' \in (x_H^*, \bar{x}]$ such that the test is not RC-Informative at x' . If $x' = \bar{x}$, then the LCSE satisfies D1 for any prior (Proposition 5), and the set of equilibria does not convergence to the complete-information outcome. If $x' \in (x_H^*, \bar{x})$, then $\mu_H(u_L(x', 1)) = 1$ and $x_H(u_L(x', 1)) = x'$ by the same argument given in the proof of Proposition 5 for the case where $x' = \bar{x}$. By continuity of the solution locus and Proposition A.1, for μ_0 arbitrarily close to 1, there exists an $x'' \in (x' - \epsilon, x' + \epsilon)$, such that full-pooling at x'' is a D1 equilibrium. As $\mu_0 \rightarrow 1$, $\epsilon \rightarrow 0$, and hence the set of equilibria does not converge to the complete-information outcome. \square

Proof of Proposition 7. Note that since $\mu_H(\cdot)$ is non-decreasing, $\mu_H(\underline{u}_L)$ corresponds to the lower bound of $\mu_H(\cdot)$, further the set of (x, μ) such that $u_L(x, \mu) = \bar{u}_L$ is the singleton $\{(x_L^*, 1)\}$ and therefore by definition of (\star) , it must be that $\mu_H(\bar{u}_L) = 1$. Thus, $\mu_H : [\underline{u}_L, \bar{u}_L] \rightarrow [\mu_H(\underline{u}_L), 1]$.

Let $M \subseteq [\mu_H(\underline{u}_L), 1]$ denote the set of m such that there exists $\hat{u}_L < \hat{u}'_L$ with $m = \mu_H(\hat{u}_L) = \mu_H(\hat{u}'_L)$. Consider the preimage $\mu_H^{-1}(m) \equiv \{\hat{u}_L \in [\underline{u}_L, \bar{u}_L] : \mu_H(\hat{u}_L) = m\}$. By continuity of μ_H , μ_H^{-1} is non-empty on $[\mu_H(\underline{u}_L), 1]$. Let $f(m) \equiv \min\{\mu_H^{-1}(m)\}$ and note that f is a strictly increasing (left-continuous) function from $[\mu_H(\underline{u}_L), 1]$ to $[\underline{u}_L, \bar{u}_L]$. Further, $M = \{m : f(m^-) \neq f(m+)\}$ where $f(m^-)$ and $f(m+)$ denote the left and right limits respectively. By Froda's theorem, M is (at most) a countable set and hence the set of such points is non-generic.

Now μ_H weakly increasing implies: (i) $\mu_0 < \mu_H(\underline{u}_L) \implies \mu_0 < \mu_H(\hat{u}_L)$ for all $\hat{u}_L \geq \underline{u}_L$ and (ii) for any $\mu_0 \geq \mu_H(\underline{u}_L)$ and $\mu_0 \notin M$, there exists a unique \hat{u}_L such that $\mu_0 = \mu_H(\hat{u}_L)$. The rest follows from Proposition A.1.⁴⁸ \square

The following lemma is used in the proof of Proposition 8.

Lemma A.2. *Under the setup and assumptions of Section 5, fix any x and $t > t'$. Then*

- *There exists a constant K such that $U_t(x, a) - U_{t'}(x, a) = K$ for all a .*
- *For any non-degenerate interim belief μ , $u_t(x, \mu) - u_{t'}(x, \mu) > K$.*

Proof. The first statement is equivalent to assumption A.4'. Now, fix any x , $t > t'$, and non-degenerate interim belief μ . Let $U_t(x, a) - U_{t'}(x, a) = K$. Then

$$u_t(x, \mu) - u_{t'}(x, \mu) = \int U_t(x, a^*(\mu_f(x, g|\mu), x))f_t(g|x)dg - \int U_{t'}(x, a^*(\mu_f(x, g|\mu), x))f_{t'}(g|x)dg$$

Suppressing some of the dependencies to simplify notation, rewrite this as

$$\begin{aligned} u_t(x, \mu) - u_{t'}(x, \mu) &= \int U_t(x, a^*)f_t(g|x)dg - \int U_{t'}(x, a^*)f_{t'}(g|x)dg \\ &= \int [U_t(x, a^*) - U_{t'}(x, a^*)]f_t(g|x)dg + \int U_{t'}(x, a^*)[f_t(g|x) - f_{t'}(g|x)]dg \\ &= K + \int U_{t'}(x, a^*)f_t(g|x)dg - \int U_{t'}(x, a^*)f_{t'}(g|x)dg > K \end{aligned}$$

⁴⁸For each $\mu_0 \in M$, let $x_{\min}(\mu_0) = \min\{x : \mu_H(\hat{u}_L) = \mu_0 \text{ for some } \hat{u}_L \in [\underline{u}_L, \bar{u}_L]\}$ and define $x_{\max}(\mu_0)$ analogously. For any such μ_0 , the set of equilibrium consists of a continuum of full-pooling equilibria at signaling levels $x \in [x_{\min}(\mu_0), x_{\max}(\mu_0)]$.

where the final inequality follows from assumption A.5 (because strict MLRP implies that the distribution of $\mu_f(x, g|\mu)$ when the density of G is $f_t(\cdot|x)$ strictly first-order stochastically dominates the distribution of $\mu_f(x, g|\mu)$ when the density of G is $f_{t'}(\cdot|x)$). \square

Proof of Proposition 8. For the purpose of contradiction, fix a candidate equilibrium with x^* and $t > 1$ such that $x^* \in S_t$ and $x^* \notin S_{\tilde{t}}$ for all $\tilde{t} \neq t$. We will first show that x^* cannot be 0. We will then demonstrate that, under RC_t , for any $x^* > 0$, there exists an $\epsilon > 0$ such that choosing $x^* - \epsilon$ leads to a higher payoff than choosing x^* for type t under any off-path receiver interim beliefs satisfying D1, implying the candidate equilibrium fails the criterion.

Suppose that $x^* = 0$. For arbitrary $t' < t$, for any $x \in S_{t'}$, both $x > 0$ and

$$u_{t'}^* = u_{t'}(x, \mu(x)) \geq u_{t'}(x^*, \mu_t) = U_{t'}(x^*, a^*(\mu_t, x^*))$$

hold. Further, assumptions A.3', A.5, and $U_{t,2} > 0$ imply that for such x , $\mu(x) \neq \mu_{t'}$. Therefore, given that the type space is finite, there exists $x' \in S_{t'}$ such that $\mu(x')$ is non-degenerate. Let a' be the certainty equivalent for type t' when choosing x' :

$$U_{t'}(x^*, a') = u_{t'}(x', \mu(x')) \geq U_{t'}(x^*, a^*(\mu_t, x^*))$$

By assumption A.1,

$$U_t(x', a') - U_t(x^*, a^*(\mu_t, x^*)) \geq U_{t'}(x', a') - U_{t'}(x^*, a^*(\mu_t, x^*)) \geq 0 \quad (7)$$

Also, by Lemma A.2,

$$u_t(x', \mu(x')) - u_{t'}(x', \mu(x')) > U_t(x', a') - U_{t'}(x', a')$$

Rearranging yields

$$u_t(x', \mu(x')) - U_t(x', a') > u_{t'}(x', \mu(x')) - U_{t'}(x', a') = 0 \quad (8)$$

From (7) and (8),

$$\begin{aligned} u_t(x', \mu(x')) - U_t(x', a') + U_t(x', a') - U_t(x^*, a^*(\mu_t, x^*)) &> 0 \\ u_t(x', \mu(x')) - U_t(x^*, a^*(\mu_t, x^*)) &> 0 \end{aligned}$$

implying that type t garners a strictly higher payoff by choosing x' instead of x^* , in violation of the hypothesis. Hence, $x^* \neq 0$.

Fix now $x^* > 0$ and a type $t' < t$. We wish to show that there exists an $\epsilon_{t'} > 0$ such that, for all $\epsilon \in (0, \epsilon_{t'})$, $B_{t'}(x^* - \epsilon, u_{t'}^*) \subset B_t(x^* - \epsilon, u_t^*)$. There are two cases to cover: 1) $u_{t'}^* > U_{t'}(x^*, a^*(\mu_t, x^*))$, or 2) $u_{t'}^* = U_{t'}(x^*, a^*(\mu_t, x^*))$. Define $B_t^0(x, \hat{u})$ to be the set $\{\mu : u_t(x, \mu) = \hat{u}\}$.

Case 1: First, if $t = N$, then by assumption A.5 $u_{t'}^* > U_{t'}(x^*, a^*(\mu_t, x^*))$ implies that $B_{t'}(x^*, u_{t'}^*) = B_{t'}^0(x^*, u_{t'}^*) = \emptyset$. If $t < N$, then $u_{t'}^* > U_{t'}(x^*, a^*(\mu_t, x^*))$ implies that $B_{t'}(x^*, u_{t'}^*) \subset B_t(x^*, u_t^*)$ with $\inf\{\|\mu - \mu'\| : \mu \in B_t^0(x^*, u_t^*), \mu' \in B_{t'}^0(x^*, u_{t'}^*)\} > \alpha$, for some $\alpha > 0$. To see this, let $\tilde{\mu}$ be an element of $B_t^0(x^*, u_t^*)$ not equal to μ_t . By A.5, $\tilde{\mu}$ is non-degenerate.

Therefore, Lemma A.2 implies that

$$u_t(x^*, \tilde{\mu}) - u_{t'}(x^*, \tilde{\mu}) > U_t(x^*, a^*(\mu_t, x^*)) - U_{t'}(x^*, a^*(\mu_t, x^*))$$

Rearranging this gives

$$\begin{aligned} u_t(x^*, \tilde{\mu}) - U_t(x^*, a^*(\mu_t, x^*)) &> u_{t'}(x^*, \tilde{\mu}) - U_{t'}(x^*, a^*(\mu_t, x^*)) \\ 0 &> u_{t'}(x^*, \tilde{\mu}) - U_{t'}(x^*, a^*(\mu_t, x^*)) \end{aligned}$$

Hence, $u_{t'}^* > U_{t'}(x^*, a^*(\mu_t, x^*)) > u_{t'}(x^*, \tilde{\mu})$, establishing the claim—given that $U_{t'}$ (and therefore $u_{t'}$) and a^* are continuous. Therefore, again by continuity of $U_{t'}$ and a^* , (whether $t = N$ or not) there exists an $\epsilon_{t'} > 0$ such that $B_{t'}(x^* - \epsilon, u_{t'}^*) \subset B_t(x^* - \epsilon, u_t^*)$ for all $\epsilon \in (0, \epsilon_{t'})$.

Case 2: The same argument just given for Case 1 shows that $u_{t'}^* = U_{t'}(x^*, a^*(\mu_t, x^*))$ implies that $B_{t'}(x^*, u_{t'}^*) \subset B_t(x^*, u_t^*)$ and that $B_{t'}^0(x^*, u_{t'}^*) \cap B_t^0(x^*, u_t^*) = \{\mu_t\}$. Consider any $\delta > 0$, and define $D_\delta \equiv \{\mu : \|\mu - \mu_t\| < \delta\}$ and D_δ^c to be the complement of D_δ . Continuity of $U_t, U_{t'}$ and a^* implies then that there exists an $\gamma_{t'}(\delta) > 0$ such that $\{B_{t'}(x^* - \epsilon, u_{t'}^*) \cap D_\delta^c\} \subseteq \{B_t(x^* - \epsilon, u_t^*) \cap D_\delta^c\}$ for all $\epsilon \in (0, \gamma_{t'}(\delta))$.

We now need to show that there exists a $\delta > 0$ and $\lambda_{t'}(\delta) > 0$ such that $\{B_{t'}(x^* - \epsilon, u_{t'}^*) \cap D_\delta\} \subset \{B_t(x^* - \epsilon, u_t^*) \cap D_\delta\}$ for all $\epsilon \in (0, \lambda_{t'}(\delta))$. To do this, for any type j , define $\chi_j(\mu, \hat{u})$ to be the signal x that gives type j utility \hat{u} when choosing x leads to interim belief μ (if no such x in \mathbb{R}_+ exists, then $\chi_j(\mu, \hat{u}) = \emptyset$). It is immediate that χ_j is differentiable where it is strictly positive. We can proceed analogously to the proof of Proposition 3. Implicit differentiation gives that, at point (x, μ) , $x = \chi_j(\mu, \hat{u})$, for some \hat{u} :

$$\begin{aligned} \frac{d\chi_j}{d\mu^k} &= -\frac{du_j}{d\mu^k} / \frac{du_j}{dx} \\ &= \frac{-\int_{\mathbb{R}} U_{j,2} \left[\frac{da^*}{d\mu_f^1} \cdot \frac{d\mu_f^1}{d\mu^k} + \dots + \frac{da^*}{d\mu_f^N} \cdot \frac{d\mu_f^N}{d\mu^k} \right] f_j(g|x) dg}{du_j/dx} \end{aligned} \quad (9)$$

Of use here will be (9) evaluated at a degenerate belief, μ_l for $l \neq k$.

$$\begin{aligned} \left. \frac{d\chi_j}{d\mu^k} \right|_{\mu=\mu_l} &= \frac{-\int_{\mathbb{R}} U_{j,2} \left[\frac{da^*}{d\mu_f^k} \cdot \frac{f_k}{f_l} - \frac{da^*}{d\mu_f^l} \cdot \frac{f_k}{f_l} \right] f_j(g|x) dg}{dU_j(x, a^*(\mu_l, x))/dx} \\ &= \left[\frac{da^*(\mu_l, x)}{d\mu_f^k} - \frac{da^*(\mu_l, x)}{d\mu_f^l} \right] \mathbb{E}[R_{k,l}(g|x)|j] \left(\frac{U_{j,2}(x, a^*(\mu_l, x))}{dU_j(x, a^*(\mu_l, x))/dx} \right) \end{aligned}$$

Because we are in Case 2, $\chi_t(\mu_t, u_t^*) = \chi_{t'}(\mu_t, u_{t'}^*) = x^*$. Therefore, by RC_t , there exists a $\delta > 0$, such that $\chi_t(\mu, u_t^*) > \chi_{t'}(\mu, u_{t'}^*)$ for all μ such that both $\mu \in D_\delta$ and $\mu^i > 0$ for a unique type $i \neq t$. Further, $\chi_t, \chi_{t'}$ differentiable imply that, locally, any directional derivative is the convex combinations of the partial derivatives, which extends the result to: there exists a $\delta > 0$, such that $\chi_t(\mu, u_t^*) > \chi_{t'}(\mu, u_{t'}^*)$ for all $\mu \in D_\delta$. Because u_t and $u_{t'}$ are decreasing in x , it follows that for any such δ there exists an $\lambda_{t'}(\delta) > 0$ such

that $\{B_{t'}(x^* - \epsilon, u_{t'}^*) \cap D_\delta\} \subset \{B_t(x^* - \epsilon, u_t^*) \cap D_\delta\}$ for all $\epsilon \in (0, \lambda_{t'}(\delta))$. Finally, let $\epsilon_{t'} = \min\{\lambda_{t'}(\delta), \gamma_{t'}(\delta)\}$, and we have that $B_{t'}(x^* - \epsilon, u_{t'}^*) \subset B_t(x^* - \epsilon, u_t^*)$ for all $\epsilon \in (0, \epsilon_{t'})$.

We have shown that for each type $t' < t$ there exists $\epsilon_{t'} > 0$ such that, for all $\epsilon \in (0, \epsilon_{t'})$, $B_{t'}(x^* - \epsilon, u_{t'}^*) \subset B_t(x^* - \epsilon, u_t^*)$. Let $\epsilon_m = \min\{\epsilon_{t'} : t' < t\}$. Then for any $\epsilon \in (0, \epsilon_m)$, $B_{t'}(x^* - \epsilon, u_{t'}^*) \subset B_t(x^* - \epsilon, u_t^*)$ for every $t' < t$. For $\epsilon \in (0, \epsilon_m)$, if $x^* - \epsilon$ is off the equilibrium path, D1 prescribes that the interim belief puts zero weight on any type $t' < t$, following a deviation to $x^* - \epsilon$. So by assumptions A.3' and A.5, $u_t(x^* - \epsilon, \mu(x^* - \epsilon)) > u_t(x^*, \mu(x^*)) = U_t(x^*, a^*(\mu_t, x^*))$ for any μ that is D1 admissible, making the deviation profitable for type t and breaking the equilibrium.

If $x^* - \epsilon$ is on-path, then either $\exists t' < t$ such that $(x^* - \epsilon) \in S_{t'}$ or there does not. If there does not, then the same argument from the previous paragraph establishes that the deviation is profitable for type t . If there does exist such a $t' < t$, then $u_{t'}(x^* - \epsilon, \mu(x^* - \epsilon)) = u_{t'}^*$. However, we have just demonstrated that under RC_t , $B_{t'}(x^* - \epsilon, u_{t'}^*) \subset B_t(x^* - \epsilon, u_t^*)$, implying, by continuity of u , that $(x^* - \epsilon, \mu(x^* - \epsilon)) \in B_t(x^* - \epsilon, u_t^*)$. Hence, type t receives a strictly higher payoff at $x^* - \epsilon$ than at x^* , contradicting the equilibrium hypothesis and establishing the proposition. \square

Proof of Proposition 9. Let $\{\mu_{0,k}\}$ be any sequence of priors that converges to μ_t , and $(\Upsilon_1^{*,k}, \dots, \Upsilon_N^{*,k})$ be an equilibrium sender-strategy profile for prior $\mu_{0,k}$. Then for any $\epsilon > 0$ there exists an n such that for all $k > n$

- There exists a $X^k \subseteq S_t^{*,k}$ such that, for all $x \in X^k$, $\|\mu^{*,k}(x) - \mu_t\| < \epsilon$.
- The total mass attributed to $\{x : x \notin X^k\}$ by $\Upsilon_t^{*,k}$ is less than δ , with $\delta \rightarrow 0$ as $\epsilon \rightarrow 0$.

These follows easily from the fact that the $\mu_{0,k}$ assigns vanishingly small weight to all other types. It is therefore sufficient to show that as $\epsilon \rightarrow 0$, $X^k \rightarrow \{0\}$. For any $x > 0$, by choosing ϵ small enough relative to ϵ_m in the proof of Proposition 8, the argument given there to show that $x^* \not\approx 0$ and the differentiability of the utility function for each type establish that for k large enough $x \notin X^k$. This establishes the convergence in total mass. Finally, if $t = N$, then every type $t' < N$ receives payoff approaching $U_{t'}(0, a^*(\mu_N, 0))$, his maximum feasible payoff, from imitating type N and a strictly lower payoff from not doing so, establishing that convergence will be type-by-type when $\mu_0 \rightarrow \mu_N$. \square

B Verification of the Crossing and Non-Decreasing Locus

Here we verify both the Crossing Property and the non-decreasing locus property (i.e., $\mu_H(\cdot)$ non-decreasing) for the applications discussed in Section 2.2. For Applications 1-3, we allow any constant testing technology, and for Applications 4-5, we consider the testing technologies specified therein. For notational convenience, fix $V_H = 1$ and $V_L = 0$.

1. *Job-Market Signaling*: This application satisfies the stronger property of *SCP* (see Lemma 1). For non-decreasing locus, see Lemma A.1.⁴⁹
2. *Advertising*: This application also satisfies *SCP*. Recall that $\Pi(\mu_f) = (a\mathbb{E}[V_t|\mu_f])^2/4b$. Denoting expected profit by π , we have that

$$\pi'_t(\mu) = \frac{a^2}{2b} \int \frac{\mu R(g)}{(\mu + (1 - \mu)R(g))^3} f_t(g) dg$$

As in the proof of Lemma 1, it suffices to show that $\frac{\pi'_L(\mu)}{\pi'_H(\mu)}$ is strictly increasing in μ . Following the same steps as that proof, order the grade space so that R is increasing in g and let $h(\mu, g) = \frac{\mu R(g)}{(\mu + (1 - \mu)R(g))^3}$. Notice that $h(\mu, g)$ is log-supermodular since for any $g_2 > g_1$,

$$\frac{\partial}{\partial \mu} \left(\frac{h(\mu, g_2)}{h(\mu, g_1)} \right) = 3 \frac{R(g_2)}{R(g_1)} (R(g_2) - R(g_1)) \left(\frac{\mu + (1 - \mu)R(g_1)}{(\mu + (1 - \mu)R(g_2))^2} \right)^2 \geq 0$$

The rest of the proof follows the same steps as the proof of Lemma 1. For non-decreasing locus, see Lemma A.1.⁵⁰

3. *Warranties*: First derive that $\frac{\partial}{\partial x} b_t(x|\hat{u}_t) = \frac{C_t x - 1}{w'_t(\mu)}$. Note that b_L is increasing, and b_H is decreasing, for all $x \in [0, V_H/C_H]$. Hence *CP* holds trivially at all (x, μ) such that $x \leq V_H/C_H$, and so we restrict attention to $x > V_H/C_H$. Consider (x_0, μ_0) such that $b_H(x_0|\hat{u}_H) = b_L(x_0|\hat{u}_L) = \mu_0$, and $\frac{\partial}{\partial x} b_H(x_0|\hat{u}_H) \leq \frac{\partial}{\partial x} b_L(x_0|\hat{u}_L)$. Therefore

$$\begin{aligned} \frac{C_H x_0 - 1}{w'_H(\mu_0)} &\leq \frac{C_L x_0 - 1}{w'_L(\mu_0)} \\ \implies \frac{w'_L(\mu_0)}{w'_H(\mu_0)} &\leq \frac{C_L x_0 - 1}{C_H x_0 - 1} \end{aligned} \quad (10)$$

Again, from the proof of Lemma 1, the left-hand side of (10) is strictly increasing in μ and constant in x , while the right-hand side is strictly decreasing in x and constant in μ ,

⁴⁹More generally for any cost functions, $C_t(x)$, let $f(x) = \frac{C'_t(x)}{C'_H(x)}$. If $f(x)$ is (weakly) decreasing then there is a unique optimizer for each \hat{u}_L and the solution locus is non-decreasing. Note that f weakly decreasing is sufficient, but not necessary, for these results. The proof of this follows immediately from $\frac{w'_L(\mu)}{w'_H(\mu)}$ strictly increasing (see proof of Lemma 1).

⁵⁰It is possible to show that *CP* holds for a broader class of Π functions: those that retain the log-supermodularity of the product function $f(\mu, g) \equiv \Pi'(\mu_f(\mu, g))h(\mu, g)$. For example, if Π is either convex or concave, and $\Pi''' \leq 0$, *CP* is satisfied.

immediately implying a non-decreasing locus and also the bound $w'_L(\mu) < \frac{C_L x_0 - 1}{C_H x_0 - 1} w'_H(\mu)$ for all $\mu < \mu_0$. Now consider any $x < x_0$, by definition of b_t

$$w_t(\mu_0) + x_0 - \frac{1}{2} C_t x_0^2 = w_t(b_t(x|\hat{u}_t)) + x - \frac{1}{2} C_t x^2$$

And therefore

$$w_t(\mu_0) - w_t(b_t(x|\hat{u}_t)) = \int_{b_t(x|\hat{u}_t)}^{\mu_0} w'_t(s) ds = (x_0 - x) \left(\frac{C_t}{2} (x_0 + x) - 1 \right) \quad (11)$$

Taking the rightmost expression from (11) for $t = L$, dividing by the same term for $t = H$, and using the bound on $w'_L(\mu)$, for $\mu < \mu_0$, obtained from (10), we get that

$$\frac{C_L \frac{(x_0+x)}{2} - 1}{C_H \frac{(x_0+x)}{2} - 1} = \frac{\int_{b_L(x|\hat{u}_L)}^{\mu_0} w'_L(s) ds}{\int_{b_H(x|\hat{u}_H)}^{\mu_0} w'_H(s) ds} < \left(\frac{C_L x_0 - 1}{C_H x_0 - 1} \right) \frac{\int_{b_L(x|\hat{u}_L)}^{\mu_0} w'_H(s) ds}{\int_{b_H(x|\hat{u}_H)}^{\mu_0} w'_H(s) ds}$$

To see that $b_H(x|\hat{u}_H) > b_L(x|\hat{u}_L)$, note that $\frac{C_L \frac{(x_0+x)}{2} - 1}{C_H \frac{(x_0+x)}{2} - 1} > \frac{C_L x_0 - 1}{C_H x_0 - 1}$ and therefore $\frac{\int_{b_L(x|\hat{u}_L)}^{\mu_0} w'_H(s) ds}{\int_{b_H(x|\hat{u}_H)}^{\mu_0} w'_H(s) ds} > 1$, implying CP holds for this application.

4. *Financial Structure:* We verify the properties for all $\gamma \leq \bar{\gamma} \equiv 1/(V_H - V_L)$. This parametric restriction implies that the sender's expected payoff is increasing in the interim belief, i.e., that his risk aversion does not dominate the benefit from being seen as a high type. Recall that

$$u_t(x, \mu) = x(V_t - \frac{\gamma}{2} \sigma^2 x) + (1-x)(w_t(\mu) - \frac{\gamma}{2}(1-x)v_t(\mu))$$

where $v_t(\mu) = Var[P(\mu_f)|t]$. Since the test is binary and symmetric, we have that $v_H(\mu) = v_L(\mu)$ for all μ and hence we drop the subscript on v_t . The premise of CP , $b_L(x_0|\hat{u}_L) = b_H(x_0|\hat{u}_H) = \mu_0$ and $\frac{\partial}{\partial x} b_H(x_0|\hat{u}_H) \leq \frac{\partial}{\partial x} b_L(x_0|\hat{u}_L)$, for some (x_0, μ_0) , is equivalent to

$$\frac{w_H(\mu_0) - 1 + \gamma \sigma^2 x_0 - \gamma(1-x_0)v(\mu_0)}{w'_H(\mu_0) - \frac{\gamma}{2}(1-x_0)v'(\mu_0)} \leq \frac{w_L(\mu_0) + \gamma \sigma^2 x_0 - \gamma(1-x_0)v(\mu_0)}{w'_L(\mu_0) - \frac{\gamma}{2}(1-x_0)v'(\mu_0)} \quad (12)$$

Consider any $x' < x_0$. There are two cases: either $b_H(x'|\hat{u}_L) \geq \mu_0$ or not. Since $b_L(\cdot|\hat{u}_L)$ is strictly increasing for all $\hat{u}_L \in [\underline{u}_L, \bar{u}_L]$, $b_L(x'|\hat{u}_L) < b_L(x_0|\hat{u}_L) = \mu_0$. Combining this with the first case immediately gives $b_H(x'|\hat{u}_L) > b_L(x'|\hat{u}_L)$. Verifying CP for the remaining case is broken into the following two subcases:

- (i) If $\mu_0 \leq 1/2$, then because $w'_H(\mu) > w'_L(\mu)$ for all $\mu < 1/2$ and the numerator on the left-hand side of (12) is strictly less than the numerator on the right-hand side for all (x, μ) , it follows that (12) holds strictly for all (x, μ) such that $\mu < \mu_0$ and hence $\frac{d}{dx} b_H < \frac{d}{dx} b_L$ for all such points. Since $b_H(x_0|\hat{u}_H) = b_L(x_0|\hat{u}_L)$, it follows immediately that $b_H(x'|\hat{u}_H) > b_L(x'|\hat{u}_L)$.

(ii) If $\mu_0 > 1/2$, rearrange (12) to get

$$\frac{w_H(\mu_0) - 1 + \gamma\sigma^2x_0 - \gamma(1-x_0)v(\mu_0)}{w_L(\mu_0) + \gamma\sigma^2x_0 - \gamma(1-x_0)v(\mu_0)} \leq \frac{w'_H(\mu_0) - \frac{\gamma}{2}(1-x_0)v'(\mu_0)}{w'_L(\mu_0) - \frac{\gamma}{2}(1-x_0)v'(\mu_0)} \quad (13)$$

Taking the partial derivative of the right-hand side of (13) with respect to the interim belief and using the fact that i) $w''_L > w''_H$ and (see Proof of Lemma 2) and ii) $w'_L > w'_H$, $v' < 0$ for all $\mu > 1/2$, shows that the right-hand side is strictly decreasing in μ . Letting $Q(x_0, \mu_0)$ denote the term on the left-hand side of (13), we have that for all $\mu < \mu_0$

$$Q(x_0, \mu_0) \left(w'_L(\mu) - \frac{\gamma}{2}(1-x_0)v'(\mu) \right) < w'_H(\mu) - \frac{\gamma}{2}(1-x_0)v'(\mu) \quad (14)$$

By definition of b_t

$$\int_{b_t(x'|\hat{u}_t)}^{\mu_0} \left(w'_t(s) - \frac{\gamma}{2}(1-x')v'(s) \right) ds = \frac{(x_0 - x')}{(1-x')} \left(w_t(\mu_0) - V_t + \gamma \left(\frac{x' + x_0}{2} \sigma^2 - \left(1 - \frac{x' + x_0}{2} \right) v(\mu_0) \right) \right) \quad (15)$$

Let T_t denote the expression on the RHS of (15). Integrating both sides of (14) over $[b_L(x'|\hat{u}_L), \mu_0]$ and substituting in T_L gives

$$Q(x_0, \mu_0) \cdot T_L < \int_{b_L(x'|\hat{u}_L)}^{\mu_0} \left(w'_H(s) - \frac{\gamma}{2}(1-x)v'(s) \right) ds \quad (16)$$

Note that $T_H/T_L = Q(\frac{x'+x_0}{2}, \mu_0) < Q(x_0, \mu_0)$, since Q is strictly increasing in its first argument and $x' < x_0$. Therefore, $T_H < Q(x_0, \mu_0)T_L$, which from (16) implies that $T_H < \int_{b_L(x'|\hat{u}_L)}^{\mu_0} (w'_H(s) - \frac{\gamma}{2}(1-x)v'(s)) ds$. The last inequality holds if and only if $b_H(x'|\hat{u}_H) > b_L(x'|\hat{u}_L)$, since the integrand is strictly positive.

For non-decreasing locus, first note that if the test is not RC-Informative at x_0 then it is not RC-Informative for all $x' < x_0$. Combined with CP , this implies that if $\mu_H(\hat{u}_L) = 1$, then $\mu_H(\hat{u}'_L) = 1$ for all $\hat{u}'_L > \hat{u}_L$. Next, let (x_0, μ_0) denote an arbitrary interior solution (i.e., $\mu_0 < 1$) at which point (13) must hold with equality, which requires $\mu_0 > 1/2$. Evaluated at (x_0, μ_0) , the right-hand side of (13) is decreasing in both x and μ , while the left-hand side is increasing in both x and μ . Hence, for any $x' < x_0$, in order to have (13) hold with equality at (x', μ') , it must be that $\mu' > \mu_0$. Now, consider any $\hat{u}_L \in (u_L(x_0, \mu_0), \bar{u}_L]$. Since $u_{L,1} < 0$, $u_{L,2} > 0$, it must be that either $x_H(\hat{u}_L) < x_0$, $\mu_H(\hat{u}_L) > \mu_0$ (or both). The latter case gives the desired result. In the first case, if a corner solution obtains, then $\mu_H(\hat{u}_L) = 1 > \mu_0$. If the solution is interior, then (13) must hold with equality at $(x_H(\hat{u}_L), \mu_H(\hat{u}_L))$ and therefore, using the argument provided above, it must be that $\mu_H(\hat{u}_L) > \mu_0$.

5. *Auditing*: Notice first that $Var[P(\mu_f)|L] = 0$ regardless of x . Furthermore $[\underline{u}_L, \bar{u}_H] =$

$[-\frac{\gamma}{2}\alpha^2\sigma^2, (1-\alpha) - \frac{\gamma}{2}\alpha^2\sigma^2]$. Using the functional form of the testing technology, we get

$$\begin{aligned} u_L(x, \mu) &= (1-\alpha)\frac{\mu}{1+(1-\mu)x} - c(x) - \frac{\gamma}{2}\alpha^2\sigma^2 \\ u_H(x, \mu) &= \alpha + (1-\alpha)\frac{\mu + (1-\mu)x}{1+(1-\mu)x} - c(x) - \frac{\gamma}{2}(1-\alpha)^2v_H(x, \mu) - \frac{\gamma}{2}\alpha^2\sigma^2 \end{aligned}$$

where $v_H(x, \mu) = Var[P(\mu_f)|t, x]$ is bounded and twice-differentiable. We proceed by verifying the properties for the case that $\gamma = 0$, for which we obtain analytic solutions. By continuity of preferences and the testing technology, and because $v_H(x, \mu)$ is bounded and twice continuously differentiable, there exists a $\bar{\gamma} > 0$ such that both properties continue to hold for all $\gamma < \bar{\gamma}$.

Solving $u_L(x, \mu) = \hat{u}_L$ for μ gives $b_L(x|\hat{u}_L) = \frac{(1+x)(\hat{u}_L+c(x))}{(1-\alpha)+x(\hat{u}_L+c(x))}$, plugging this into (\star) , $u_H(x, b_L(x|\hat{u}_L)) = \frac{\alpha+\hat{u}_L+x(1-c(x))}{1+x}$. The first-order condition for an interior solution requires that

$$\hat{u}_L = (1-\alpha) - c(x) - x(1+x)c'(x) \quad (17)$$

Clearly, the right-hand side is (i) strictly decreasing and continuous in x , (ii) greater than \bar{u}_L at $x = 0$, and (iii) less than \underline{u}_L for x large enough. Hence the solution to (17), denoted by $x_A(\hat{u}_L)$, exists, is unique for all $\hat{u}_L \in [\underline{u}_L, \bar{u}_L]$ and is non-negative. Checking the second-order condition at an interior solution is routine. Furthermore, note that $x_A(\hat{u}_L) \leq \bar{x}(\hat{u}_L)$ (which satisfies $\hat{u}_L = (1-\alpha) - c(x)$) Thus, we have that $x_H(\hat{u}_L) = x_A(\hat{u}_L)$, which is strictly decreasing and therefore invertible. In addition, CP must also hold (if CP failed, there would exist utility pairs $\{u_L, u_H\}$ whose BICs had points of tangency, i.e., solutions to (17), that fail the second-order condition, which violates the argument above). For increasing locus, it suffices to show that $b_L(x|x_H^{-1}(x))$ is decreasing in x on $[0, \bar{x}]$. Plugging in the inverse function to b_L and taking the derivative we get that

$$\frac{d}{dx}b_L(x|x_H^{-1}(x)) = - \left(\frac{(1-\alpha)c'(x) + x^2c'(x) + (1-\alpha)xc''(x)}{(1-\alpha + x^2c'(x))^2} \right) < 0$$

which completes the verification of both properties.