Market Signaling with Grades

Brendan Daley

Fuqua School of Business, Duke University, Durham, NC

Brett Green[1,*]

Haas School of Business, University of California, Berkeley, CA

Abstract

We consider a 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. We derive a necessary and sufficient condition—loosely, that the grade is sufficiently informative relative to the dispersion of (marginal) signaling costs across types—under which the presence of grades substantively alters the equilibrium predictions. Specifically, separating equilibria do not survive stability-based refinements. Instead, the prediction depends on the prior distribution over the sender's type. For example, with two types it involves full pooling when the distribution places sufficient weight on the high type and partial pooling otherwise. Finally, the equilibrium converges to the complete-information outcome as the distribution 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

1. Introduction

Signaling models typically assume that observable, costly actions are the only channels that can convey information about the sender. These models implicity overlook the existence

[*]Corresponding Author

Email addresses: bd28@duke.edu (Brendan Daley), bgreen@haas.berkeley.edu (Brett Green)

[1]Address: 2200 Piedmont Avenue, Berkeley, CA 94720. Phone: 510.575.9980

of *grades*.[2] A *grade* refers to any imperfect public message about the sender's type (e.g., test scores, analyst ratings, product reviews). Grades are prevalent in many environments that have been modeled as signaling games, such as education (Spence, 1973), financial markets (Leland and Pyle, 1977), advertising (Kihlstrom and Riordan, 1984), and warranties (Gal-Or, 1989).

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 observe a stochastic grade in addition to a costly action chosen by the sender. The likelihood of each grade depends on the sender's privately-known type and (potentially) on his chosen action. After observing the chosen action and the realized grade, receivers take actions in response. The purpose of this exercise is to understand how the presence of grades affects equilibrium behavior in signaling games as well as the implications for policy and applied work.

To fix ideas, consider an amended version of Spence's job-market signaling model: a worker, who is privately informed whether his productivity is high or low, chooses an education level (his action); potential employers observe the worker's education level and, in addition, his performance on a test (his grade) prior to making offers.[3] Our first insight is that, in the presence of an informative test, some degree of pooling on the education level is more plausible than the widely adopted prediction of separation.[4] The economic intuition for this result is that if the worker were to choose the separating level of education, attempting to convince the market that he is the high type, employers would infer that he is eager to de-emphasize the results of the test. The low type has more incentive to do this since he expects a worse outcome on the test. Choosing a very high level of education would therefore backfire, leading employers to infer that the worker is more likely to be the low type. Clearly, then, the worker will not choose the separating level of education.

[2]See 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 (2013), Alos-Ferrer and Prat (2012), and Daley and Green (2012), see Section 2 for further discussion.

[3]More generally, employers may observe the worker's entire transcript and any other academic distinctions.

[4]Without grades, the least cost separating equilibrium is the unique stable outcome in this model and its well-known generalizations (Cho and Kreps, 1987; Cho and Sobel, 1990; Ramey, 1996).

Instead, the high type must resolve the tradeoff between how much to exploit his cost advantage (i.e., that education is less costly to him) 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*) that states precisely when the test is informative enough relative to the cost advantage to induce some reliance on it. More specifically, we show that separating equilibria do not survive stability-based refinements when the test is RC-Informative. In addition, the presence of grades changes the relative payoffs in pooling equilibria. When types pool on the costly action, the grade contains additional information about the worker's ability. Therefore, in contrast to pooling equilibria in the model without grades, a high type expects to earn a higher wage than a low type who obtains the same level of education.

Our second key insight is that the addition of grades can resolve the discontinuity of the equilibrium prediction as the prior converges to degeneracy. Regardless of whether grades are available, in the complete-information game where the worker's type is common knowledge, the unique equilibrium outcome involves no education. However, by introducing even the *slightest* possibility that the worker is a lower type, the unique stable outcome of the gradeless model involves the high-type worker choosing a non-trivial level of education to distinguish himself from this unlikely possibility. This implies that both the high type's strategy and payoff are discontinuous in the receivers' prior belief; a prediction that has been a source of criticism for signaling models. In the presence of an RC-Informative test, the model is not susceptible to the same criticism; as the prior becomes degenerate, all stable equilibria converge to the complete-information outcome.

The presence of an RC-Informative test also has implications for Pareto efficiency; we characterize the set of perfect bayesian equilibria (PBE) payoffs and find that all Pareto efficient equilibria involve some degree of pooling (and therefore reliance on grades) when the test is RC-Informative. Generically, there is a unique PBE satisfying the D1 refinement (Cho and Kreps, 1987; Banks and Sobel, 1987), which depends both on the informativeness of the test as well as on the market's prior belief about the worker. With an RC-Informative test, the equilibrium involves partial pooling when the prior assigns a low probability to the high type: the high type chooses an education level that is less costly than the least cost separating equilibrium (LCSE) level, while the low type mixes between revealing himself by forgoing education and imitating the high type. Both the education level 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 an education level that is lower still. As a result, the expected utility of both types increases as the prior belief about the sender becomes more favorable.

Focusing on the equilibrium satisfying D1, we explore the welfare implications of grades. The more informative the test, the more the high type relies on his grade to convey information to employers and the less he relies on costly education. The high type's welfare increases with test informativeness while the low type can be made better or worse off. Overall, as the test becomes more informative, the amount of resources devoted to costly (and inefficient) signaling activities decreases and efficiency improves.

We generalize our main results by extending our analysis to a more general class of preferences and by allowing the accuracy of the test to vary with the action—encapsulating a broad range of applications, including those mentioned at the outset. In this environment, the sender decides how informative a test to subject himself to, knowing receivers 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 also provide evidence consistent with our findings and discuss the implications of our results for empirical work. In a supplementary appendix, we consider a model with N types and obtain similar results; the presence of an informative test leads to reliance on the grade as well as convergence to the complete-information outcome.

The remainder of the paper is organized as follows. In Section 2, we review the related literature. In Section 3, we introduce grades to the job-market signaling model and present the main findings. We generalize the model and our results in Section 4. Section 5 discusses additional applications, implications and extensions. Section 6 concludes. Proofs are located in Appendix A.

2. Related Literature

2.1. Noisy Signaling

It is important to distinguish our approach from the literature on "noisy signaling" (Matthews and Mirman, 1983; Carlsson and Dasgupta, 1997). In a noisy signaling model, receivers do not perfectly observe the sender's action, but rather a noisy signal whose distribution depends on it (for example, a random shock is added to the sender's action 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 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, rather than a single observation encompassing both aspects. Inherently absent in a noisy signaling model, our work illustrates the trade-off between these two channels and explains the importance of both in delivering our predictions.

Some settings may be accurately modeled by combining elements of both our model and noisy signaling models. Senders decide both how costly an observable action to undertake and how much to invest, unobservably, toward influencing the stochastic measure (i.e., the grade). We discuss this extension in Section 5.3.

2.2. Signaling and Grades

Ours is not the first paper to introduce additional information 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 values to employers. 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 action 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 (2013) 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 action level, while the medium type perfectly separates by incurring strictly positive signaling costs. In Section 5.2, we discuss how our results provide an alternative explanation for some of the peculiarities that countersignaling seeks to explain. Alos-Ferrer and Prat (2012) 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 response 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. We also consider a more general structure for the external information, including allowing the quality of information to vary with the sender's costly action. Finally, on the more technical side, we show that the double-crossing property (Matthews and Moore, 1987) arises naturally in our model for the relevant indifference curves and facilitates a tractable equilibrium analysis.

3. Job-Market Signaling with Grades

In this section, we illustrate the role of grades and the main insights of the paper within the canonical signaling example of Spence (1973). There is a worker (the *sender*) and multiple competing firms (*receivers* or *the market*). The worker is privately informed about his ability (*type*) and chooses a level of education (*action*). Firms observe both the worker's level of education and his GPA (*grade*), then compete in Bertrand fashion to hire the worker.

Let $t \in \{L, H\}$ denote the sender's type and $x \in \mathbb{R}_+$ denote his action. 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 $x \in S_t$. After the sender chooses an action, the grade $g \in \mathbb{R}$ is then realized. The action, x, and the grade, g, are both publicly observed. Next, each receiver i simultaneously makes an offer $W_i(x, g)$ to the sender, who then decides which offer to accept, if any.

If a type-t sender chooses an action x and accepts an offer of W, his utility is $W - C_t x$, where $0 < C_H < C_L$ (i.e., the single-crossing property holds). If he rejects all offers, his utility is $-C_t x$. The utility of receiver i is zero if her offer is rejected and is $V_t - W_i$ if her offer is accepted, where $0 < V_L < V_H$.

3.1. Grades and Tests

Given $t \in \{L, H\}$, the grade is a real-valued random variable, G, with density function f_t. We refer to the pair of probability density functions $\{f_L, f_H\}$ as a *test*. Let $R(g) \equiv f_L(g)/f_H(g)$.[5] Note that $R(g)$ measures the informativeness of the grade g. If $R(g) = 1$, then the grade offers no information about the sender's type. Alternatively, if $R(g) > 1$ (< 1), then the grade causes a Bayesian to decrease (increase) the probability assigned to the high type. A test is *statistically informative* if there exists a measurable 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 $t \in \{L, H\}$. 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). For technical convenience, we will require the test to satisfy the following conditions.

T.1 For both types t, f_t is continuous almost everywhere.

T.2 Grades are boundedly-informative: $\inf_g R(g) > 0$ and $\sup_g R(g) < \infty$.

T.3 The Monotone Likelihood Ratio Property (MLRP) holds with R weakly decreasing over the common support of f_L and f_H.

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

Remark 3.1. *For many tests, the set of grades is finite: pass/fail, letter grades A to F, etc. This can easily be accommodated.*[6] *For expositional purposes, we will often use a* ***symmetric binary test****, in which there are two outcomes,* $\{pass, fail\}$*, with* $p = \Pr(pass|t{=}H) = \Pr(fail|t{=}L) \in (\frac{1}{2}, 1)$.

3.1.1. RC-Informative Tests

While $R(g)$ measures the informativeness of a *grade*, we are also interested in measuring the informativeness of *tests*. Blackwell (1951) and Lehmann (1988) provide the two predominant notions for what it means for one test to be more informative than another. The crucial notion of informativeness in our analysis is the low type's expected likelihood ratio, $\mathbb{E}[R(g)|t{=}L]$. 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.[7]

Fact 3.2. *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* $\mathbb{E}[R(g)|L] \geq \mathbb{E}[\hat{R}(g)|L]$.

Intuitively, what matters for equilibrium analysis is not just the test informativeness, but rather the test informativeness *relative* to the high type's advantage in taking the costly action. Specifically, we will show that it is the following condition that is both necessary and sufficient for substantively affecting the equilibrium predictions.

Definition 3.3. *The test is* ***RC-Informative*** *if and only if* $\mathbb{E}[R(g)|L] > \frac{C_L}{C_H}$.

While $\mathbb{E}[R(g)|L]$ measures the informativeness of the test, $\frac{C_L}{C_H}$ is a measure of the high type's *cost advantage*. Hence, RC-Informativeness is simply that the test is informative enough relative to the cost advantage.

Example 3.4. *A symmetric binary test is RC-Informative if and only if:*

$$p > \frac{1}{2}\left(1 + \frac{\sqrt{\left(\frac{C_L}{C_H}\right)^2 + 2\frac{C_L}{C_H} - 3}}{\frac{C_L}{C_H} + 3}\right),$$

where the right-hand side is increasing in the cost advantage. For instance, if the cost advantage is $\frac{C_L}{C_H} = \frac{3}{2}$*, then the test is RC-Informative if and only if* $p > \frac{2}{3}$.

[6] To encompasses 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.

[7] Unlike the partial orderings endowed by Blackwell and Lehman, the notion used here endows a complete ordering of all tests.

Remark 3.5. *Our assumption that $C_L > C_H$ corresponds to the standard single-crossing property, facilitating comparison to the standard gradeless model. However, our analysis does not rely on this assumption. When $C_L = C_H$, any statistically informative test is RC-Informative, and all of our results under RC-Informativeness hold in this environment.*

3.2. Solution Concept and Preliminary Analysis

We use perfect bayesian equilibrium (PBE) as our solution concept.[8] After observing x and g, receivers update to some final belief $\mu_f(x, g) \equiv \Pr(t = H|x, g)$. Receivers' 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 observed action. Along the equilibrium path, the interim belief is pinned down 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).

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

Let $W(x, g)$ be the highest offer from the receivers after observing x and g. Since receivers compete in Bertrand fashion, in any PBE, $W(x, g)$ must be equal to the expected value of V_t given the receivers' final belief.

$$W(x, g) = \mu_f(x, g)V_H + (1 - \mu_f(x, g))V_L.$$

Therefore, the interim belief is sufficient to compute type t's (highest) expected offer:

$$w_t(\mu) \equiv \int \frac{\mu V_H + (1 - \mu)V_L}{\mu + (1 - \mu)R(g)} f_t(g)dg. \tag{2}$$

Thus, the expected utility of a type-t worker depends only on his chosen action and the resultant interim belief: $u_t(x, \mu) = w_t(\mu) - C_t x$. Given a schedule of interim beliefs, $\mu(x)$, an action, x_0, is optimal for a type-t sender if and only if it satisfies:

$$x_0 \in \arg\max_x u_t(x, \mu(x)).$$

[8]The notion of PBE that we employ requires that receivers hold identical beliefs off the equilibrium path and that these beliefs are updated using Bayes rule after any history for which it is possible to do so. Both of these requirements are analogous to ones in Fudenberg and Tirole (1991, pp. 331-333).

3.3. Belief Indifference Curves

Without grades, the indifference curves of interest are those over the space of actions and offers (or, equivalently, actions and final beliefs). With grades, based on the analysis conducted in Section 3.2, it is indifference curves over the space of actions 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 actions to interim beliefs parameterized by utility levels. Provided such a belief exists, let $b_t(x|\hat{u})$ be the interim belief such that $u_t(x, b_t(x|\hat{u})) = \hat{u}$.

A comparison to the gradeless model may be 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. Throughout the paper x^{LCSE} denotes the high type's LCSE action.

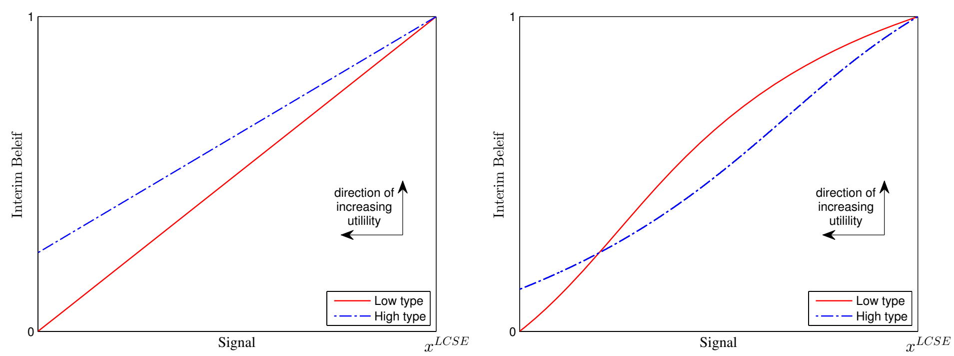

(a) The Gradeless Model (b) The Model with Grades

Figure 1 – BICs for the LCSE payoffs

Without the test, the low type's curve is steeper than and below the high type's for all $x < x^{LCSE}$; the difference in the slope of the two types' curves derives only from their differing costs. With the test, BICs acquire curvature according to each type's expectation 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 expected 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.

BICs have a useful feature, termed the *monotone relative slope property (MRSP)*, that is highlighted in Figure 2. Specifically, 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 $\mu > \mu^*$. Notice that as a special case of MRSP, BICs satisfy the single-crossing property if and only if $\mu^* = 1$.

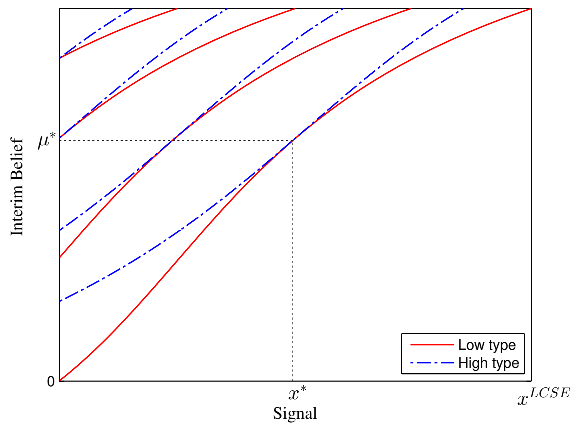

FIGURE 2 – Illustration of the monotone relative slope property

Intuitively, the high type has two advantages over the low type; a *cost advantage* and a *grade advantage*. The cost advantage is independent of the interim belief and can be measured by C_L/C_H. The grade advantage, on the other hand, depends on receivers' interim belief. Starting from μ close to zero, an increase in μ makes the test more important; receivers rely more on the grade when their prior is intermediate and therefore, offers are more sensitive to grades. This benefits the high type relatively more than the low type (since the high type expects to get a higher grade) and hence the high type's grade advantage is increasing in this region. In contrast, for sufficiently large μ, further increases in μ decrease the importance of the test and benefit the low type more than the high type. Since the relative slopes of the BICs tell us which type needs more compensation in interim belief to incur a marginally more expensive action, MRSP results from this relationship between interim belief and test importance.

The relevance of RC-Informativeness (Definition 3.3) will now become apparent.

Lemma 3.6. *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 MRSP. (All formal proofs are found in Appendix A.) One interpretation of the lemma is as follows. In the absence of grades, there is no difference between interim and final belief, so BICs inherit the single-crossing property. Further, the addition of a test that is *not* RC-Informative alters the BICs, but not substantially enough to overturn single-crossing. Only when the test is RC-Informative, is $\mu^* < 1$ and do BICs no longer satisfy single-crossing.

3.4. 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 action perfectly identifies the sender's type, and differ from full pooling equilibria in that multiple actions are on the equilibrium path, each leading to a different interim belief, which also differs from the prior. MRSP 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.

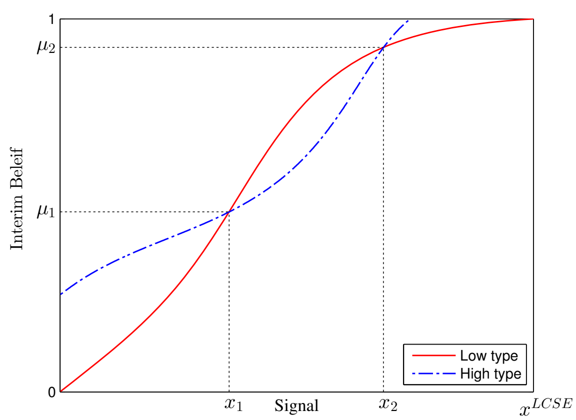

Figure 3 – BICs for a common support equilibrium

The set of PBE depends on the prior. Describing the set of equilibria for every prior is a straightforward but tedious exercise. Instead, we turn to characterizing the set of PBE payoffs when the test is RC-Informative and illustrate how the set changes with the prior.[9] Since each receiver's expected utility is zero in all equilibria, it suffices to characterize the set for the sender. The results are depicted in Figure 4. Two belief levels play a key role. The first is μ^*. Let x^* be the unique action 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 $\underline{\mu} < \mu^*$.

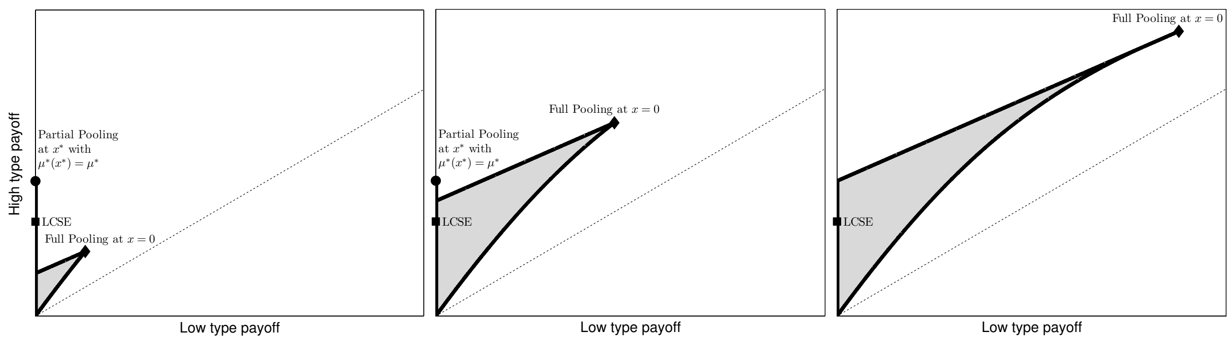

(a) For $\mu_0 < \underline{\mu}$ (b) For $\mu_0 \in (\underline{\mu}, \mu^*)$ (c) For $\mu_0 > \mu^*$

Figure 4 – The set of PBE payoffs for various priors. In all panels, the dotted black line indicates the 45-degree line. Payoffs produced by separating equilibria are located on the vertical axis between the origin and the LCSE payoffs. The remainder of the feasible payoffs on the vertical axis can be achieved by partial pooling equilibria in which $0 \in S_L, \notin S_H$. The upper linear boundary of the shaded area has slope $\frac{C_H}{C_L}$ and corresponds to full-pooling equilibria.

From Figure 4, it is clear that there are only two candidates for Pareto efficient equilibria.[10] 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 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 action. Hence, when the test is RC-Informative, there is always a way to utilize it that Pareto dominates

[9]Characterizing the set when the test is not RC-Informative produces little insight beyond that obtained in the gradeless model. Further, it is straightforward to extend the analysis here to the general model studied in Section 4.

[10]Again, receivers are indifferent between all equilibria. Our notion of Pareto efficiency is *ex interim* in that one equilibrium Pareto dominates another if and only if it is weakly better for *both* types of the sender and strictly greater for at least one of them, as in Riley (2001). This usage is sometimes motivated by noting that it is equivalent to model the environment with a continuum of senders, a fraction μ_0 of which are high types, all participating in the market simultaneously.

separating outcomes.

Proposition 3.7. *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 $\mu_0 \leq \underline{\mu}$.*

Along with the results in Sections 3.5 and 3.6, this indicates that RC-Informativeness is precisely the condition that characterizes when the presence of grades substantively alters our findings compared to those in the gradeless model.

3.5. Equilibrium under D1: Senders Rely on Informative Tests

As in most signaling models, the set of PBE is large because of the flexibility afforded to off-equilibrium-path beliefs. In order to produce further insights and predictions, the set of equilibria must be refined. We focus our attention on PBE 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 it uniquely selects the LCSE (Cho and Kreps, 1987; Cho and Sobel, 1990; Ramey, 1996).[11]

In our model, the D1 refinement can be stated as follows.[12] Fix an equilibrium endowing expected utilities $\{u_L^*, u_H^*\}$. Consider an action 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). If $B_H(x, u_H^*) \subset B_L(x, u_L^*)$, then D1 requires that $\mu(x) = 0$.

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

[11] While 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.

[12] In 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).

(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."

The following proposition demonstrates that, generically, there is a unique D1 equilibrium and describes how it varies with the prior. Recall that $u_L(x^*, \mu^*) = V_L$.

Proposition 3.8. *If the test is RC-Informative, then*

- *if $\mu_0 > \mu^*$, the unique D1 equilibrium is full pooling at $x = 0$.*
- *if $\mu_0 < \mu^*$, the unique D1 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 D1 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 D1 equilibrium is the LCSE.

This result says that if the test is too weak (or the cost advantage is too great), the predictions match those of the gradeless model. In fact, the underlying analysis matches as well—when the test is not RC-Informative, $\mu^* = 1$ (Lemma 3.6) and BICs satisfy the single-crossing property, just as in the gradeless model. More importantly, Proposition 3.8 says that for more informative tests (or lesser cost advantages), the predictions change—at least some degree of reliance on the test (because types no longer separate on x) will take place at all priors, and full reliance on the test occurs for sufficiently high priors.

To see why some degree of pooling must occur when the test is RC-Informative, consider the LCSE and suppose receivers observe an off-path action of $x^{LCSE} - \epsilon$. Clearly, $\mu(x^{LCSE} - \epsilon)$ must be close to one in order for this to be a profitable deviation for either type. Recall that MRSP implies the high type's BIC is steeper at all $\mu > \mu^*$. In words, when μ is near one, a marginal decrease in the interim belief is relatively less costly to the high type than to the low because it increases the importance of the test. Therefore, unlike in the gradeless model, receivers do not interpret this deviation as a negative signal because it indicates a willingness to put more emphasis on the test. In accordance with D1, receivers will assign probability one to $t = H$ after observing the deviation, making it profitable for both types, and breaking the candidate separating equilibrium.

In the D1 equilibrium, the high type uses the costly action to influence the receivers' interim belief only as long as he has a relative advantage in doing so (i.e., as long as his BIC is flatter than the low type's). When the test is RC-Informative, $\mu^* < 1$ (Lemma 3.6) and therefore the high type does not fully separate through the costly action. Instead, when the prior is below μ^*, he uses the costly action to the point where the interim belief reaches μ^*. When the prior is above μ^*, any attempt to influence the interim belief through the costly action is viewed as an attempt to de-emphasize the test. Neither type wants to incur costs if it will decrease the receivers' interim belief, and hence both types fully pool on the costly action at $x = 0$.

3.6. Equilibrium Convergence

Consider the complete-information game where $t = H$ is common knowledge; the unique equilibrium outcome involves the sender choosing $x = 0$ and obtaining an offer (and utility) of V_H. By introducing even the slightest probability that the sender is the low type, the stable equilibrium of the gradeless model predicts that the high type will fully separate by choosing x^{LCSE}. This prediction has been a source of criticism for the gradeless model. One key insight of this paper is that the availability of other sources of information (i.e., grades) can resolve this discontinuity. Within the job-market signaling setting, it is clear from Proposition 3.8 that an RC-Informative test is both necessary and sufficient for convergence of the D1 equilibrium to the full-information outcome as $\mu_0 \to 1$. In Section 4, we generalize this result by providing necessary and sufficient conditions for this same result to obtain in a richer environment.

To formalize these ideas, we will need to delve a bit deeper into the notion of convergence. We introduce the following two notions of convergence.

Definition 3.9. *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 a distribution 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 a distribution 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 LCSE of the gradeless model does not converge to the complete-information outcome as $\mu_0 \to 1$ by either metric.[13] We can now formalize the convergence result within this canonical setting.

[13] When $\mu_0 \to 1$ ($\mu_0 \to 0$), convergence in total mass to Υ is equivalent to convergence of Υ_H (Υ_L) to Υ.

Corollary 3.10. *As $\mu_0 \to 1$, by either notion of convergence (type-by-type or in total mass), the D1 equilibrium converges to the complete-information outcome if and only if the test is RC-Informative.*

As mentioned earlier, RC-Informativeness is the precise condition under which both (i) the predictions of the gradeless model change and (ii) convergence to the full-information outcome is obtained.[14]

3.7. Comparative Statics and Welfare

In this section we investigate comparative statics of the D1 equilibrium by varying the test informativeness and the cost advantage. To do so, we use symmetric binary tests (see Remark 3.1), parameterized by $p \in (\frac{1}{2}, 1)$, where higher p corresponds to a more informative test. In addition, we fix C_L (and therefore x^{LCSE}) and vary the cost advantage by changing C_H; we include the case of $C_H = C_L$ (see Remark 3.5). Below, we classify the main findings into three categories.

Reliance on the Test: *As the informativeness of the test increases, the high type relies more on it, and less on the costly action. Analogously, as the high type's cost advantage increases, he relies more on the costly action, and less on the test.*

The claims are easiest to see in the case of low priors, where (from Proposition 3.8) the high type's reliance on the costly action is captured by x^* (the action he chooses), and his reliance on the test is decreasing in μ^* (the resultant interim belief).[15] Figure 5 illustrates how x^* and μ^* vary with p for several different cost advantages. Modulo boundary solutions, as p increases, x^* decreases and μ^* strictly (weakly) decreases if the cost advantage is strict

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.

[14]Notice that because the high type chooses x^* with probability one for all $\mu_0 < \mu^*$, the D1 equilibrium of the model with grades converges to the complete-information outcome only in total mass as $\mu_0 \to 0$. Convergence 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.

[15]To see that his reliance on the test is decreasing in μ^*, recall that the importance of the test depends on the market's interim belief. Test importance is highest when the belief is intermediate, and decreases as the belief increases from there. Since the high type has a cost advantage, μ^* is always at least as large as the level that maximizes test importance (see Daley and Green, 2013, pg. 14). Additionally, an increase in μ^* enlarges the set of priors for which the D1 equilibrium is partial pooling, which features more reliance on the costly action and less reliance on the test than do full pooling equilibria.

(weak). As $p \to 1$, $x^* \to 0$. That is, as the test becomes completely informative, the high type relies only on the test. As $p \to \frac{1}{2}$, whether 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 $(x^{LCSE}, 1)$. This is simply the recovery of the result in the gradeless model: when there is a cost advantage and no meaningful test, the unique D1 equilibrium is the LCSE. On the other hand, if $C_H = C_L$, $\mu^* = \frac{1}{2}$ for all p and $x^* \to \frac{1}{2}x^{LCSE}$. No matter how weak the test is, the high type relies on it as heavily as possible since he has no cost advantage. Finally, for any fixed p, both x^* and μ^* are increasing in the cost advantage.

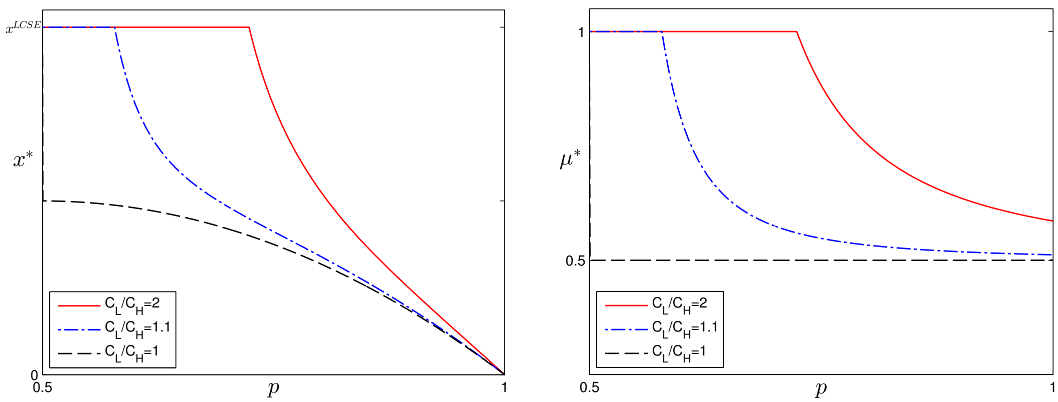

(a) (b)

FIGURE 5 – Comparative statics for x^* and μ^* with symmetric binary tests.

Payoffs: *The high type's welfare increases with the test informativeness, but may increase or decrease with his cost advantage. The low type's welfare decreases with the cost advantage, but may increase or decrease with the test informativeness.*

For a fixed cost advantage, Figure 6(a,b) plots the sender's type-dependent expected payoffs, u_t^*. The key difference between the two panels is that in (a) $\mu_0 = \frac{1}{4} < \mu^*$ for all $p \in (\frac{1}{2}, 1)$, meaning the D1 equilibrium is either separating or partial pooling, whereas in (b) because $\mu_0 = \frac{3}{4}$ the D1 equilibrium switches from partial to full pooling when p crosses 0.61. In both cases, the high type's payoff is increasing in p. Notice from (b) that, despite enjoying the benefit of decreased signaling costs for $p > 0.61$, the low type's welfare decreases as the test continues to get better at identifying his type.

Not surprisingly, the low type's payoff is decreasing in the cost advantage. To see that the high type can also be worse off with a larger cost advantage, consider a game where

$\mu^* < 1$ and μ_0 is just above μ^*, so the D1 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 D1 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.

Efficiency: *Efficiency increases with the test informativeness, but in a way that depends on the prior. If the prior is low, efficiency increases continuously with test informativeness and approaches first-best only as the test becomes perfect. If the prior is high, the first-best outcome can be achieved with an imperfect test.*

Recall that the market is competitive, so the sender captures all the surplus. In Figure 6(c), efficiency is measured by the percentage of the potential surplus the sender attains: $\mathbb{E}[u_t^*|\mu_0]/\mathbb{E}[V_t|\mu_0]$.[16] Not surprisingly, efficiency increases with p and the full surplus is realized as $p \to 1$. However, the two priors illustrate two different avenues by which this is achieved. For $\mu_0 = \frac{3}{4}$, the D1 equilibrium is fully efficient for all $p > 0.61$, since the equilibrium is full pooling at $x = 0$. For $\mu_0 = \frac{1}{4}$, the D1 equilibrium involves inefficient expenditure on the costly action for every $p < 1$. Only as $p \to 1$ does this expenditure decrease to zero (see Figure 5(a)).

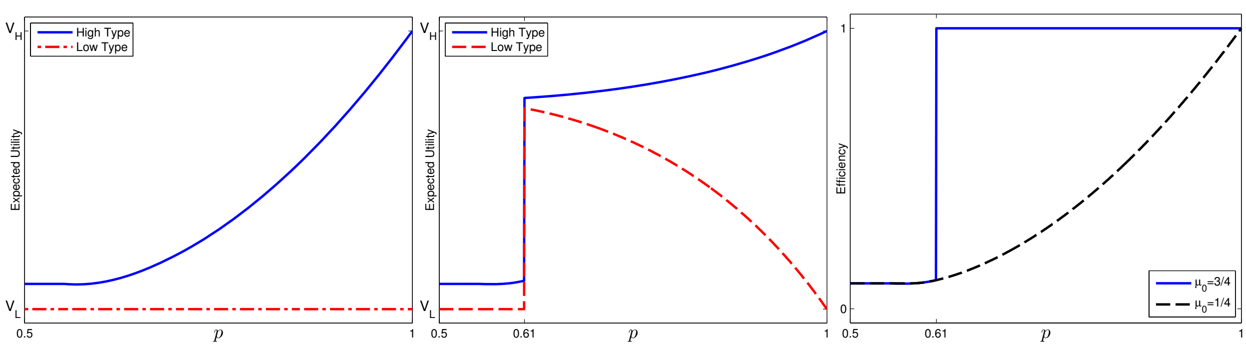

(a) Sender payoff given $\mu_0 = 1/4$ (b) Sender payoff given $\mu_0 = 3/4$ (c) Efficiency

FIGURE 6 – Sender welfare (a,b) and efficiency (c) as they depend on the informativeness of the symmetric binary test for two different priors. All panels use $C_L/C_H = 1.1$.

[16] The denominator also corresponds to the surplus attained in the complete-information outcome. Using a percentage measure for efficiency allows us to meaningfully compare efficiency for different priors, even though the potential surplus, $E[V_t|\mu_0]$, varies with the prior.

4. Generalized Preferences and Testing Technologies

In this section we generalize both the set of allowable sender utility functions and the relationship between the costly action and the test. In addition to demonstrating their robustness, this broadens the scope of our results to a variety of other applications (see Section 5.1 for examples). Further, it allows the statistical informativeness of the test to be determined endogenously by the sender's action. To do so, it is convenient to model the response from receivers in reduced form, as in Mailath (1987). As before, receivers observe the sender's costly action, x, and realized grade, g, and update their (common) belief about the sender's type to a final belief $\mu_f(x, g)$. The payoff of a type-t sender is now $U_t(x, \mu_f)$.[17]

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).

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 action 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 action chosen by the sender in the complete-information setting. A.2 and A.3 also imply that there exists a unique $x^{LCSE} > x^*_L$, such that $U_L(x^*_L, 0) = U_L(x^{LCSE}, 1)$. Finally, we assume that

A.4 $x^*_H < x^{LCSE}$.

[17]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 Section 3, receivers make simultaneous wage offers to the sender, who decides which offer (if any) to accept. In this case $U_t(x, \mu_f) = V_L + \mu_f(V_H - V_L) - C_t x$.

In combination with A.1, A.4 implies that (i) the least-cost-separating strategy profile $\sigma_L(x_L^*) = 1, \sigma_H(x^{LCSE}) = 1$ is part of a 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.

In addition, the accuracy of the test is now permitted to depend on the sender's choice of action. To do so, define a *testing technology* to be a family of tests indexed by x and denoted by $\{f_L(\cdot|x), f_H(\cdot|x)\}$, each satisfying T.1-T.3, and where $f_t(\cdot|x)$ is continuously differentiable in x. Let $R(g|x) = f_L(g|x)/f_H(g|x)$.

Finally, as in the analysis in Section 3, the receivers' interim belief, μ, and the sender's expected utility as a function of his costly action and the interim belief, $u_t(x, \mu)$, will play a key role. In this more general environment,

$$u_t(x, \mu) = \int U_t(x, \mu_f(x, g)) f_t(g|x) dg \tag{3}$$

4.1. Generalized RC-Informativeness and its Consequences

Because both the informativeness of the test and any advantage the high type enjoys in taking the costly action may vary with the chosen level of the action, we need to generalize the notion of RC-Informativeness.

Definition 4.1. *For any $x \in [x_H^*, x^{LCSE}]$, the test is **RC-Informative at $\boldsymbol{x}$** if and only if*

$$\mathbb{E}[R(g|x)|L, x] > \frac{U_{L,1}(x, \mu_f)}{U_{L,2}(x, \mu_f)} \Big/ \frac{U_{H,1}(x, \mu_f)}{U_{H,2}(x, \mu_f)} \Bigg|_{\mu_f = 1} \tag{4}$$

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' action-μ_f indifference curves at $(x, 1)$. Notice that in the model from Section 3, the right-hand-side of (4) is C_L/C_H. In terms of results, RC-Informativeness at x^{LCSE} gives a sufficient condition for separation to be eliminated and all D1 equilibria to rely on grades.

Proposition 4.2. *If the test is RC-Informative at x^{LCSE}, then for any $\mu_0 \in (0, 1)$, all D1 equilibria involve some degree of pooling.*

A sufficient condition for convergence to the complete-information outcome also relies on RC-Informativeness.

Proposition 4.3. *If the test is RC-Informative at all $x \in (x_H^*, x^{LCSE}]$, then as $\mu_0 \to 1$, by either notion of convergence (type-by-type or in total mass), the set of D1 equilibria converges to the complete-information outcome.*

4.2. The Double-Crossing Property

In Section 3, BICs satisfy a key property, MRSP, which stems from the difference in how each type is affected by the test, depending on the market's interim belief (see Section 3.3). In richer settings, this property may fail to hold. However, the key economic implications are restored by an appropriate generalization, termed the *double-crossing property* (DCP), which we now introduce.

For any utility level $\hat{u}$, and either type sender, there exist actions large enough, such that no $\mu \in [0,1]$ can deliver expected utility $\hat{u}$ when they are chosen. Define $\bar{x}(\hat{u}_L) \equiv \max\{x : b_L(x, \hat{u}_L) = 1\}$.[18] Notice that $x^{LCSE} \equiv \bar{x}(U_L(x_L^*, 0))$.

The Double-Crossing Property (DCP). *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 $x < x_0$.*

DCP 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 a point of intersection, then it lies everywhere below to the right.

DCP is implied by MRSP, but not the converse. The property arises naturally on BICs in many signaling models with grades. Clearly, it is satisfied in the job-market signaling model of Section 3. In addition, each of the applications discussed in Section 5.1 satisfy DCP.[19] We maintain that DCP holds for the remainder of this section, in which case the sufficient conditions of Propositions 4.2 and 4.3 are also necessary.

Proposition 4.4. *Under DCP, if the test is not RC-Informative at x^{LCSE}, then for all $\mu_0 \in (0,1)$, the LCSE is a D1 equilibrium.*

Proposition 4.5. *Under DCP, if there exists an $x \in (x_H^*, x^{LCSE}]$ such that the test is not RC-Informative at x, then as $\mu_0 \to 1$, the set of D1 equilibria does not converge to the complete-information outcome by either notion of convergence (type-by-type or in total mass).*

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

[19] It is possible to construct utility functions and testing technologies such that DCP fails by drastically changing the terms in (4) over a small interval of action levels. Preventing such rapid changes is sufficient to ensure the property holds.

In many environments the informativeness of the test is likely to be increasing in x. For example, 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. More relevant for equilibrium analysis is whether the informativeness is increasing *relative* to utility differences. If the difference between the left- and right-hand sides of (4) is increasing in x, results can be both strengthened and simplified. In this case, if the test fails RC-Informativeness at x^{LCSE}, then the test is not RC-Informative at any $x < x^{LCSE}$. Under this stronger hypothesis, Proposition 4.4 can be strengthened to: the LCSE is the *unique* D1 equilibrium for all priors. Further, convergence results hinge only on the simplified condition of RC-Informativeness at $x = x_H^*$.

4.2.1. Full Characterization

In Appendix A, we completely characterize the set of D1 equilibria under DCP (see Proposition A.4). Here we briefly discuss the highlights. Let $\underline{u}_L \equiv U_L(x_L^*, 0)$ and $\bar{u}_L \equiv \max_x U_L(x, 1)$, which bound the payoff the low type can acheive. Now, for any $\hat{u}_L \in [\underline{u}_L, \bar{u}_L]$, let $x_H(\hat{u}_L)$ denote the maximizer of $u_H(x, b_L(x|\hat{u}_L))$. That is, $x_H(\hat{u}_L)$ is the optimal action for the high type to choose if the schedule of interim beliefs is $\mu(x) = b_L(x|\hat{u}_L)$. That the maximizer is unique follows from DCP. 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*.

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 by relying at least 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 by completely separating from the low type using the costly action.

Figure 7 illustrates an example. For each of the low type's plotted BICs, the dotted curves are the BICs corresponding to the high type's constrained-maximal expected utility: $u_H(x_H(\hat{u}_L), \mu_H(\hat{u}_L))$. The dark line running through the tangency points is the solution locus. Notice that $\mu_H(\hat{u}_L)$ is non-decreasing in Figure 7. This *non-decreasing locus* property is necessary and sufficient to ensure that the D1 equilibrium is generically unique. Like DCP, this property also arises naturally in many signaling models with grades. It is satisfied in the job-market signaling model of Section 3 and in each of the applications discussed in Section 5.1. For each candidate utility level of the low type, the solution locus identifies

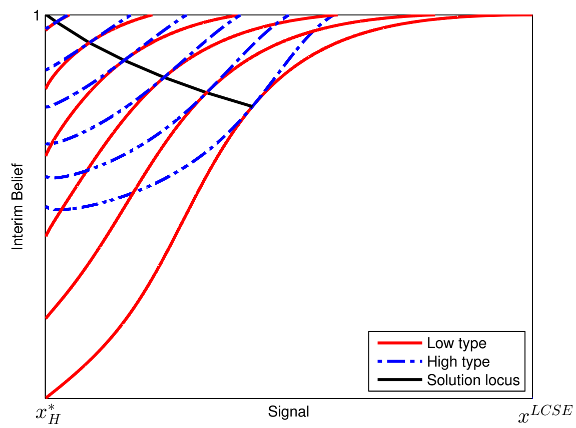

Figure 7 – The Solution Locus

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 illustrates.

Proposition 4.6. *Suppose that DCP holds and $\mu_H(\hat{u}_L)$ is non-decreasing. Then, there exists a unique D1 equilibrium for almost all priors, $\mu_0 \in (0,1)$. Generically, if $\mu_0 \leq \mu_H(\underline{u}_L)$ the D1 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* $\sigma_L^*(x_L^*) = 1 - \sigma_L^*(x_H(\underline{u}_L))$
- $\mu(x_H(\underline{u}_L)) = \mu_H(\underline{u}_L)$, *and* $\mu(x) = 0$ *if* $x \neq x_H(\underline{u}_L)$

If $\mu_0 > \mu_H(\underline{u}_L)$ the D1 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 x^{LCSE} then $\mu_H(\underline{u}_L) < 1$ and the D1 equilibrium involves partial pooling for priors below $\mu_H(\underline{u}_L)$ and full pooling for priors above. On the other hand, if $x_H(\underline{u}_L) = x^{LCSE}$, then $\mu_H(\underline{u}_L) = 1$ and the LCSE is the unique D1 equilibrium for all $\mu_0 \in (0,1)$.

A connection can be made to equilibrium selection in the gradeless model.[20] In the standard gradeless model D1 selects the LCSE. That is, it selects the equilibrium that is optimal for the high type *conditional* on the low type getting his full-information payoff. However, when the test is RC-Informative at x^{LCSE}, the equilibrium with this same property

[20] We are grateful to an anonymous referee for articulating this connection.

is the partial pooling equilibrium in which the high type chooses $x_H(\underline{u}_L)$ and the resultant interim belief is $\mu_H(\underline{u}_L) < 1$. This equilibrium exists only when $\mu_0 \leq \mu_H(\underline{u}_L)$, and continues to be selected by D1 in this case. When $\mu_0 > \mu_H(\underline{u}_L)$, in every D1 equilibrium the low type attains a payoff $u_L^* > \underline{u}_L$, but the high type continues to achieve the maximum payoff consistent with the low type attaining u_L^*.

5. Applications, Evidence, and Extensions

Signaling theory has been applied to a broad array of economic environments (see Riley (2001) for a survey). We believe that grades are an important feature of many of these environments. The model presented in Section 4 is general enough to capture a broad set of applications. In this section we discuss several examples (Section 5.1), summarize the main empirical implications of our results (Section 5.2), and consider two natural extensions (Section 5.3).[21]

5.1. Applications

The following examples fit within the framework of Section 4. In addition, it is straightforward to verify that each satisfies both DCP and the non-decreasing locus property (Section 4.2). To ensure these two properties hold, some structure on preferences and the testing technology is needed (see footnote 19). As in Section 3, the examples below include considerably more structure than needed in an attempt at retaining parsimony. Throughout, we maintain that $0 < V_L < V_H$.

1. **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. Demand for the product is increasing in the expected quality of the good, captured by the demand curve: $Q(P, \mu_f) = a\mathbb{E}[V_t|\mu_f] - bP$. 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

[21]In addition, a number of applied-theory working papers have a signaling-with-grades component (Garfagnini, 2012; Gervais and Strobl, 2012; Li and Li, 2012; Lee, 2013; Taneva, 2013). Each of these papers seeks to answer questions tailored to a single application (such as electoral campaigning or incentives for money managers), for which the theory of signaling in the presence of grades is the necessary first step that the analysis builds on. We believe that the present paper can provide the foundation of a unifying framework for this type of work, as already seen by its adoption in some of them.

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 test in this example represents product reviews, such as those provided by Yelp, CNET, Zagat, Angie's List, etc., which does not depend on x.

2. **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 $U_t(x, \mu_f) = \mathbb{E}[V_t|\mu_f] + x - \frac{1}{2}C_t x^2$, where $0 < C_L < C_H$. 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 test in this example is conducted by third-party reviewers such as Consumer Reports, J.D. Power and Associates, etc., which does not depend on x.

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 this objective as the sender's expected payoff prior to the realization of the grade.[22]

3. **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 and variance $\sigma^2 > 0$. The entrepreneur chooses the fraction of the company he retains,

[22]When 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 4). Since mean-variance preferences are generally not consistent with an expected-utility representation, one should interpret the mean-variance representation in Applications 3-4 as a substitute for expected utility. Finally, to retain the key tradeoff facing the sender, one must assume the sender's risk aversion does not dominate his preference for being seen as a high type (γ is less than some appropriately chosen upper bound $\bar{\gamma}$).

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. The grade is an analyst's recommendation, *buy* or *sell*, prior to the issuance date, which is the outcome of a symmetric binary test (see Remark 3.1). The expected payoff of a type-t entrepreneur before the analyst's recommendation is announced is given by:

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

4. **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 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 cost of an auditor with quality x is given by $c(x)$, where $c' > 0$, $c'' \geq 0$. 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.[23] The sender's expected payoff before the grade is realized is

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

5.2. Implications and Evidence

Below, we summarize our main findings and provide examples of their connections to empirically observed phenomenon, policy implications, and testable predictions.

The Relevance of Grades: *The presence of an informative test implies that grades convey meaningful information in equilibrium because the predicted behavior changes from separation to pooling on the level of costly signaling.*

[23] This testing technology involves a set of grades that perfectly distinguish the high type, violating T.2. The main issue this raises is whether such a grade can overturn a degenerate belief. 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 example (in part because the grade is completely uninformative at $x = 0$). Therefore we do not take a position on this matter, and our results immediately extend to cover this example.

If types are 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 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—controlling for ratings. Our model examines the interaction between these two potential sources of information and, consistent with this evidence, explains the importance of both in market outcomes.[24]

Test Precision: *Less informative grades increase the amount of resources devoted to inefficient signaling activities.*

This result has 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.[25] 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.

[24] An alternative explanation for the relevance of grades is that costly actions 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, 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' response. However, unlike in our model, the D1 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 D1 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 action and a grade to convey that type to the market.

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

Since the 1980's, the average G.P.A. 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 anecdotal observation that extra-curricular involvement has become an increasingly essential part of competitive applications at elite universities.

Outside of education, (inefficient) retention of equity in a venture capital or IPO setting by an entrepreneur has been cited as a signal of his firm's quality (Leland and Pyle, 1977). Our model predicts that better analyst reports (i.e., grades) should decrease the share the entrepreneur retains. If one believes that analyst reports have become more prevalent and (cumulatively) more informative over time, then (controlling for other factors) we should see less entrepreneur equity retention in the time series.

Reputation Matters for Signaling: *Signaling behavior depends on the initial market belief about the sender; a sender with a better reputation incurs lower signaling costs.*

This can explain why established firms, with strong reputations for quality, expend less on signaling than do upstart firms with lesser reputations, but similar quality. Warranties for new cars provide evidence that signaling behavior varies with the manufacturer's reputation in manner consistent with our model and inconsistent with a gradeless model (see Application 2 above). Standard theory suggests that 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, despite both having very similar reliability?[26]

The heuristic explanation for the differing warranties is that Honda 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

[26]Reliability according to data collected by *Consumer Reports*. Summaries 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.

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 explaining the comprehensiveness of their warranty, Hyundai CEO, John Krafcik said, "...we were seeing indications that [Hyundai's] quality was really much better than the perception."[27] Correspondingly, the market has a more favorable prior for a new-issue Civic than a new-issue Elantra. While the standard theory predicts that this should be irrelevant, the model with grades matches the observed facts; holding type constant, a firm with a better reputation (as measured by the prior) incurs smaller signaling costs.

By definition, a sender with a stronger reputation (e.g., Honda) is distinguishable from other, less reputable, senders (e.g., Hyundai) *prior* to choosing their action. Hence, this mechanism differs from countersignaling (Feltovich et al., 2002), wherein, among *ex-ante indistinguishable* agents, higher types signal less vigorously than their medium type counterparts. Determining which theory may apply is then context specific. In some instances, it may be more reasonable to think of agents as being ex-ante indistinguishable. Certainly, when there are differences in reputation across senders, it is not. In such cases, the reputation mechanism provides a more plausible explanation for why highly-regarded agents signal less vigorously.

Signaling without Single-Crossing: *In the presence of grades, the sender can (imperfectly) signal his type through money burning activities.*

In many signaling models, the prediction 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 signaling action corresponds to a monetary expenditure. Because our analysis does not rely on (strict) single-crossing, our results apply equally well to such settings and provide 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 (see Application 1 above). The use of advertising campaigns containing little or no obvious informational content (i.e., money

[27] From interview in *Newsweek* March 24, 2010.

burning) 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 an 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).

5.3. Extensions

In many relevant applications, the sender may have channels through which he has influence over the outcome of the test. Such channels may affect the outcome of the test *directly* by undertaking some costly action (e.g., a student chooses how much effort to exert studying prior to an exam), or *indirectly* by improving the distribution of the quality being tested: the sender's type (e.g., a firm chooses how much to invest in product quality). In a previous working version of this paper (Daley and Green, 2013), we provide a formal, but preliminary, analysis of both extensions. The key findings are as follows.

Indirect Influence: Ex-Ante Investment

Suppose that, before learning his type, the sender chooses how much to "invest" in his type: the more he invests, the more likely he is to become the high type. Our statements below pertain to the equilibria in which continuation play in the signaling stage satisfies D1.

There are two important benchmarks: the investment level that arises 1) if the type is perfectly revealed to receivers and 2) if there is no grade (or equivalently, if the test is not RC-Informative, since the LCSE remains the unique continuation play satisfying D1). The first benchmark corresponds to the first-best efficient investment level if receivers compete away all of their surplus (as in the job-market signaling model of Section 3), which we maintain here for ease of exposition. A key insight from Spence (1973) was that private information can lead to an inefficient allocation of resources devoted to conveying that information. In the second benchmark, the presence of an information asymmetry creates another inefficiency in the ex-ante stage, *underinvestment*. This is because being the high type is less valuable if one must expend resources to demonstrate it.

Returning to the setting with grades, an important consideration is whether the sender's ex-ante investment is observable.[28] When investment is unobservable, the sender always invests less than in the complete-information/first-best benchmark. Interestingly, investment can be higher or lower than the no-grade benchmark, despite the fact that the continuation play is always (weakly) more efficient with grades than without. This is because the presence of grades can improve the payoff of *both* types to a point that sufficiently weakens the incentive to become the high type.

Observability strengthens the incentives to invest. This is because the sender considers not only the benefit of increasing the probability of becoming the high type given a fixed continuation play, but also internalizes that increasing investment changes the continuation play since it directly determines the receivers' "prior" in the signaling stage. Because of this, the investment level can be even higher than the complete-information/first-best benchmark. The intuition is that, unlike in the complete-information benchmark, increasing investment can have the extra benefit of decreasing the expected amount of costly signaling in the continuation game. In fact, this "overinvestment" is constrained-socially-optimal if a social planner is unable to influence behavior in the signaling stage.

Direct Influence: Ex-Interim Hidden Effort

In this extension, after the sender chooses x, but before the grade is realized, the sender chooses a level of costly, unobservable effort toward improving his grade on the test. The game can be thought of as a combination of our model and a noisy signaling model (see Section 2). Further, each component can be analyzed in sequence. That is, any choice of x leads to some interim belief μ that together "endow" a noisy signaling continuation game.

Despite the fact that test informativeness depends on the effort level, the appropriate generalization of RC-informativeness compares the cost advantage to the informativeness of the test when the sender puts in zero effort. This is because, as $\mu \to 1$ the sender's effort in the noisy-signaling game converges zero; there is no reason to exert costly (hidden) effort if the market is already convinced of your type. Hence, the slope of the adapted BICs at $\mu = 1$ will be just *as if* the sender was unable to exert (useful) effort, much like in the original model. With these generalization in hand, Propositions 4.2-4.6, and their proofs, generalize

[28] This makes no difference in either benchmark because, unlike in the model with grades, continuation play does not depend on the receivers' post-investment belief.

in the obvious way.[29]

This extension serves largely as a check that the model is robust to a realistic feature of some environments. The main substantive implication it generates is that the observable action and hidden effort are, to a degree, substitutes. If hidden effort is less costly for the high type, then he will engage in more of it compared to the low type, which increases the informativeness of the test. The high type will then rely more on the test and less on the observable action, similar to the effect of increasing the informativeness of the (exogenous) test in a setting without hidden effort (see Section 3.7). On the other hand, the inverse statements hold if hidden effort is less costly for the low type.

6. Conclusion

Grades are a prevalent force in many incomplete information environments and often convey meaningful information. In a strategic setting, there is a subtle interaction between the information conveyed by costly observable actions and information conveyed by grades. In equilibria that are robust to standard refinements, the sender resolves the trade-off between how much to rely on each of the two transmission mechanisms. If the test is sufficiently informative, the high type relies on its ability to convey information, and the stable equilibrium involves some degree of pooling. Further, the addition of grades yields a more intuitive outcome as the prior puts greater weight on the high type. Both types gain in utility, and the stable equilibrium converges to full pooling at the efficient signaling level—long thought to be the appropriate convergence, but not achieved in gradeless signaling models. These results extend to a model with N types.

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). It may be worth noting that the stronger refinements within this class, such as NWBR and Universal Divinity, yield the same results as D1 in our model.

A number of non-stability based refinements have been developed to refine the set of

[29] Propositions 4.4-4.6 rely on DCP, which in this extended environment becomes intractable to analytically verify for many examples. However, it is not difficult to solve examples numerically and check whether DCP and non-decreasing locus hold.

equilibria in signaling games.[30] Many of these notions eliminate strategy profiles 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 then to alter a somewhat undesirable feature prevalent in signaling models. It is appealing that the predictions along this dimension are aligned, but largely coincidental.

In many examples, such as education, the costly signaling action involves waiting to trade. Static signaling models, such as the one presented here, 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 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 and show that some degree of signaling occurs in equilibrium. Daley and Green (2012) 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 predictions 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

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

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.

Acknowledgements

We are grateful to Alessandro Pavan, an anonymous associate editor and an anonymous referee for their helpful feedback. We are also indebted to Andrzej Skrzypacz, Paul Milgrom, and Jeremy Bulow for their guidance. In addition, we 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, Washington University and Yahoo! Research for useful conversations and comments.

J. G. Riley, Silver signals: Twenty-five years of screening and signaling, Journal of Economic Literature 39 (2001) 432–478.

A. Weiss, A sorting-cum-learning model of education, The Journal of Political Economy 91 (1983) 420–442.

H. Fang, Social culture and economic performance, The American Economic Review 91 (2001) 924–937.

N. Feltovich, R. Harbaugh, T. To, Too cool for school? signalling and countersignalling, RAND Journal of Economics 33 (2002) 630–349.

I. Kremer, A. Skrzypacz, Dynamic signaling and market breakdown, Journal of Economic Theory 133 (2007) 58–82.

G.-M. Angeletos, C. Hellwig, A. Pavan, Signaling in a global game: Coordination and policy traps, Journal of Political Economy 114(3) (2006) 452–485.

G.-M. Angeletos, A. Pavan, Selection-free predictions in global games with endogenous information and multiple equilibria, Theoretical Economics (forthcoming) (2013).

C. Alos-Ferrer, J. Prat, Job market signaling and employer learning, Journal of Economic Theory 147(5) (2012) 1787–1817.

B. Daley, B. Green, Waiting for news in the market for lemons, Econometrica 80(4) (2012) 1433–1504.

M. A. Spence, Job market signaling, Quarterly Journal of Economics 90 (1973) 225–243.

H. E. Leland, D. H. Pyle, Informational asymmetries, financial structure, and financial intermediation, The Journal of Finance 32 (1977) 371–387.

R. E. Kihlstrom, M. H. Riordan, Advertising as a signal, Journal of Political Economy 92 (1984) 427–450.

E. Gal-Or, Warranties as a signal of quality, Canadian Journal of Economics 22 (1989) 50–61.

I. K. Cho, D. Kreps, Signaling games and stable equilibria, Quarterly Journal of Economics 102 (1987) 179–221.

I. K. Cho, J. Sobel, Strategic stability and uniqueness in signaling games, Journal of Economic Theory 50 (1990) 381–413.

G. Ramey, D1 signaling equilibria with multiple signals and a continuum of types, Journal of Economic Theory 69 (1996) 508–531.

J. Banks, J. Sobel, Equilibrium selection in signaling games, Econometrica 55 (1987) 647–662.

S. A. Matthews, L. J. Mirman, Equilibrium limit pricing: The effects of private information and stochastic demand, Econometrica 51 (1983) 981–996.

H. Carlsson, S. Dasgupta, Noise-proof equilibria in two-action signaling games, Journal of Economic Theory 77 (1997) 432–460.

S. Matthews, J. Moore, Monopoly provision of quality and warranties: An exploration in the theory of multidimensional screening, Econometrica 55(2) (1987) 441–467.

D. Blackwell, Comparison of experiments, Proceedings of the Second Berkeley Symposium on Mathematics, Statistics and Probability, (1950) (1951) 93–102.

E. Lehmann, Comparing location experiments, Annals of Statistics 16 (1988) 521–523.

D. Fudenberg, J. Tirole, Game Theory, MIT Press, 1991.

B. Daley, B. Green, Market signaling with grades, 2013. *Mimeo.* Available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1532846.

G. Mailath, Incentive compatibility in signaling games with a continuum of types, Econometrica 55 (1987) 1349–1365.

U. Garfagnini, Escalation of 'commitment' and the value of monitoring, Working Paper (2012).

S. Gervais, G. Strobl, The industrial organization of money management, Working Paper (2012).

H. Li, W. Li, Misinformation, Working Paper (2012).

S. Lee, Plea bargaining: On the selection of jury trials, Working Paper (2013).

I. Taneva, Optimal precision of information in auctions, Working Paper (2013).

S. Titman, B. Trueman, Information quality and the valuation of new issues, Journal of Accounting and Economics 8 (1986) 159–172.

E. B. Jones, J. D. Jackson, College grades and labor market rewards, The Journal of Human Resources 25 (1990) 253–266.

P. L. Hsueh, D. S. Kidwell, Bond ratings: Are two better than one?, Financial Management 17 (1988) 46–53.

S. Rojstaczer, 2009. http://gradeinflation.com/(accessed September 20, 2009).

P. Nelson, Advertising as information, The Journal of Political Economy 82 (1974) 729–754.

P. Milgrom, J. Roberts, Price and advertising signals of product quality, The Journal of Political Economy 94 (1986) 796–821.

L. Thomas, S. Shane, K. Weigelt, An empirical examination of advertising as a signal of product quality, Journal of Economic Behavior & Organization 37 (1998) 415–430.

E. Kohlberg, J.-F. Mertens, On the strategic stability of equilibria, Econometrica 54 (1986) 1003–1037.

G. Mailath, M. Okuno-Fujiwara, A. Postlewaite, Belief-based refinements in signalling games, Journal of Economic Theory 60 (1993) 241–276.

S. Grossman, M. Perry, Perfect sequential equilibrium, Journal of Economic Theory 39 (1986) 97–119.

J. Hillas, Sequential equilibria and stable sets of beliefs, Journal of Economic Theory 64 (1994) 78–102.

J. M. Swinkels, Education sginalling with preemptive offers, Review of Economic Studies 66 (1999) 949–970.

S. Karlin, Total Positivity, volume 1, Stanford University Press, 1968.

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, receivers' best responses are summarized by a function of $\mu_f(x, g)$, (i.e., the highest wage offer in the model of Section 3 and the identity function in the model of Section 4). 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, $R(g|x)$. 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.[31]

[31] In addition, let $B_t^0(x, u_t^*)$ 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.

A.2. Proofs

We begin by stating and proving that MRSP holds for the job-market signaling model, a result that will be used in several of the proofs that follow.

Lemma A.1. *The BICs of the job-market signaling model of Section 3 satisfy MRSP.*

Proof. 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 $\frac{\partial b_t}{\partial x} = -\frac{\partial u_t}{\partial x} / \frac{\partial u_t}{\partial \mu} = \frac{C_t}{w_t'(\mu)}$, and therefore $\frac{\partial b_H}{\partial x} \le \frac{\partial b_L}{\partial x} \Leftrightarrow \frac{w_L'(\mu)}{w_H'(\mu)} \le \frac{C_L}{C_H}$. To verify MRSP it is sufficient to show that $w_L'(\mu)/w_H'(\mu)$ is strictly increasing (i.e., $w_t'(\mu)$ is strictly log-supermodular).[32] Re-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 $g' > g$,

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

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 $g \to g'$, $g' \to g$. By doing so, the above can be written as

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

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

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

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). □

Proof of Lemma 3.6. 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 MRSP. □

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

Proof of Proposition 3.7. 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 3.6, RC-Informativeness implies that $\frac{\partial}{\partial x} b_H(x^{LCSE}|\underline{u}_H) > \frac{\partial}{\partial x} b_L(x^{LCSE}|V_L)$. Hence, there exists an $\epsilon > 0$ such that $b_L\left(x^{LCSE} - \epsilon|V_L\right) > \mu_0$ and $b_L\left(x^{LCSE} - \epsilon|V_L\right) \in B_H\left(x^{LCSE} - \epsilon, \underline{u}_H\right)$. The following equilibrium Pareto dominates the LCSE: $\sigma_H(x^{LCSE} - \epsilon) = 1$, $\sigma_L(x^{LCSE} - \epsilon) = \frac{1-b_L(x^{LCSE}-\epsilon|V_L)}{b_L(x^{LCSE}-\epsilon|V_L)} \frac{\mu_0}{1-\mu_0}$ and $\sigma_L(0) = 1 - \sigma_L(x^{LCSE} - \epsilon)$, with $\mu(x^{LCSE} - \epsilon) = b_L(x^{LCSE} - \epsilon|V_L)$ and $\mu(x) = 0$ for all $x \neq x^{LCSE} - \epsilon$.

Now suppose the test is not RC-Informative. MRSP and Lemma 3.6 imply that BICs satisfy the single-crossing property. The result is well-known for this case (see Mailath et al. (1993) for instance). □

A.2.1. Full Characterization of the Set of D1 Equilibria in the General Model

We now turn to characterizing the set of D1 equilibria in the general two-type model (Section 4) under DCP, with Proposition 3.8 following as a special case. We start by proving Lemma A.2, which does not rely on DCP.

Lemma A.2. *If the payoffs $\{u^*_H, u^*_L\}$ are supported by a D1 equilibrium, then there does not exist an x' such that $b_L(x'|u^*_L) > b_H(x'|u^*_H)$.*

Proof of Lemma A.2. 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. □

In Section 4.2, we introduced the solution locus and stated several of its properties. We formalize them here. To do so, 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 $\mu(x) = b_L(x|\hat{u}_L)$. The high type's maximization problem can be written as

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

Let $x_H(\hat{u}_L)$ be the solution to ($\star$) for a given $\hat{u}_L$, and let $\mu_H(\hat{u}_L) \equiv b_L(x_H(\hat{u}_L)|\hat{u}_L)$.

Lemma A.3. *For each $\hat{u}_L \in [\underline{u}_L, \bar{u}_L]$, $x_H(\hat{u}_L)$ exists. Under DCP, the mappings $\hat{u}_L \to x_H(\hat{u}_L)$ and $\hat{u}_L \to \mu_H(\hat{u}_L)$ on the domain $[\underline{u}_L, \bar{u}_L]$ are continuous functions.*

Proof. 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. DCP ensures that the solution is unique. To see this, fix a $\hat{u}_L \in [\underline{u}_L, \bar{u}_L]$. 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 DCP 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, DCP implies that the corresponding high type's BIC lies everywhere else above the low type's, making the solution unique, and hence the mapping $\hat{u}_L \to x_H(\hat{u}_L)$ is a function. From Berge's Theorem of the Maximum, x_H is continuous. Because b_L is also continuous, $\hat{u}_L \to \mu_H(\hat{u}_L)$ is continuous. □

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

Proposition A.4. *Under DCP, 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 4.6 for this case satisfies D1.*

In all D1 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 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 DCP, $B_H(x, u_H^*) \subseteq B_L(x, u_L^*)$ 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 to satisfy D1, 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 A.2 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 $x \in S_H^*$, $b_L(x|u_L^*) = b_H(x|u_H^*) = \mu(x)$. DCP 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. □

Proof of Proposition 3.8. Follows directly from Proposition A.4 and the following lemma, which characterizes the solution locus in the job-market signaling model. □

Lemma A.5. *In the job-market signaling model of 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 MRSP, the only candidate for a point of tangency (and thus an interior 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. □

Proof of Proposition 4.2. Consider any separating candidate equilibrium profile where the high type chooses $x^s > x^{LCSE}$ and garners utility u_H^s. Let $x' \in (x^{LCSE}, x^s)$. $B_L(x', \underline{u}_L) = \emptyset$ 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 $\mathbb{E}[R(g|x)|L, x] > \frac{U_{L,1}(x,1)}{U_{L,2}(x,1)} \Big/ \frac{U_{H,1}(x,1)}{U_{H,2}(x,1)}$. Therefore, if the test is RC-Informative at x^{LCSE}, then $\frac{\partial}{\partial x} b_H(x^{LCSE}|\underline{u}_H) > \frac{\partial}{\partial x} b_L(x^{LCSE}|\underline{u}_L)$. Hence, there exists $\epsilon > 0$ such that $b_H(x^{LCSE} - \epsilon|\underline{u}_H) < b_L(x^{LCSE} - \epsilon|\underline{u}_L)$ contradicting Lemma A.2. Hence all D1 equilibria involve pooling. □

Proof of Proposition 4.3. Let $\{\mu_0^k\}$ be any sequence of priors that converges to 1, and $(\Upsilon_L^{*,k}, \Upsilon_H^{*,k})$ be a D1 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 $\delta \to 0$ as $\epsilon \to 0$.

These follow 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 $\epsilon \to 0$, $X^k \to \{x_H^*\}$. Now fix any $x' \in (x_H^*, x^{LCSE}]$. 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 4.2). 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 $\epsilon \to 0$, Lemma A.2 implies that $x' \notin X^k$ for all k large enough, giving the result. □

Proof of Proposition 4.4. Again, letting $\underline{u}_H$ be the LCSE utility for the high type, if the test is not RC-Informative at x^{LCSE}, then $\frac{\partial}{\partial x} b_H(x^{LCSE}|\underline{u}_H) \leq \frac{\partial}{\partial x} b_L(x^{LCSE}|\underline{u}_L)$ (see proof of Proposition 4.2). By DCP, $b_H(x'|\underline{u}_H) > b_L(x'|\underline{u}_L)$ for all $x' < x^{LCSE}$, implying $x_H(\underline{u}_L) = x^{LCSE}$. The result follows from Proposition A.4. □

Proof of Proposition 4.5. Suppose there exists an $x' \in (x_H^*, x^{LCSE}]$ such that the test is not RC-Informative at x'. If $x' = x^{LCSE}$, then the LCSE satisfies D1 for any prior (Proposition 4.4), and the set of equilibria does not convergence to the complete-information outcome. If $x' \in (x_H^*, x^{LCSE})$, 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 4.4 for the case where $x' = x^{LCSE}$. By continuity of the solution locus (Lemma A.3) and Proposition A.4, 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 $\mu_0 \to 1$, $\epsilon \to 0$, and hence the set of D1 equilibria does not converge to the complete-information outcome. □

Proof of Proposition 4.6. Note that since $\mu_H(\cdot)$ is non-decreasing, $\mu_H(\underline{u}_L)$ corresponds to the lower bound of $\mu_H(\cdot)$. Further, by definition of $\bar{u}_L$, $\mu_H(\bar{u}_L) = 1$. Thus, $\mu_H : [\underline{u}_L, \bar{u}_L] \to [\mu_H(\underline{u}_L), 1]$.

Let $M \subseteq [\mu_H(\underline{u}_L), 1]$ denote the set of m such that 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.4.[33] □

[33]For 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 action levels $x \in [x_{\min}(\mu_0), x_{\max}(\mu_0)]$.

Supplement to "Market Signaling with Grades"

Brendan Daley and Brett Green

November 5, 2013

B. Supplementary Appendix: More than Two Types

Let us expand the type space to be $\{1, ..., N\}$, $N > 2$. 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_f^t. The sender's utility function continues to satisfy A.1–A.2.[34] 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 $t' \neq t$, $U_{t,2}(x, a) = U_{t',2}(x, a)$.

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.[35] 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

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

[35] That is, U_t specifies both type t's utility tradeoff between the costly action and the market response 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.

informative grades eliminate separation in D1 equilibria.[36] (Recall that S_t denotes the support of Υ_t.)

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

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

then there does not exist a D1 equilibrium in which type t separates; for any action $x \in S_t^*$ 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 actions. The result is then similar to Proposition 4.2: if grades are informative enough, then in equilibrium only the lowest type may assign positive probability to an action 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 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 stable equilibria.[37,38]

Our second result generalizes the convergence established in Proposition 4.3.

Proposition B.2. *Fix a type $t > 1$. If RC_t holds, then as $\mu_0 \to \mu_t$, the set of D1 equilibria converges in total mass to the complete-information outcome. Further, as $\mu_0 \to \mu_N$, 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 D1 equilibrium, and the convergence properties from Proposition B.2 hold.

Propositions B.1 and B.2 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 signaling 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

[36] The D1 refinement for the N-type model can be stated as follows. Fix an equilibrium endowing expected utilities $\{u_t^*\}_{t\in\{1,\ldots,N\}}$. Consider an action 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).

[37] Existence of countersignaling equilibria requires that grades must be relatively poor at differentiating the *low* type from the *medium* type even without refinements. As suggested by Proposition B.1, there exists generic parameters for which D1 eliminates otherwise tenable countersignaling strategy profiles.

[38] Proposition B.1 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 actions greater than the one chosen by the highest type, which are strictly inferior to 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.

D1 uniquely selects the LCSE (Cho and Kreps, 1987; Cho and Sobel, 1990). With grades, D1 may fail to select a unique equilibrium. 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 scaler (unless $N = 2$). For instance, in a three-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.[39]

B.1. Proofs

The following lemma is used in the proof of Proposition B.1.

Lemma B.3. *Under the setup and assumptions of Appendix B, 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

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

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 $f_{t'}(\cdot|x)$). □

[39] When $N = 2$ 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.

Proof of Proposition B.1. 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,

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

Also, by Lemma B.3,

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

Rearranging yields

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

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 B.3 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 $B^0_{t'}(x^*, u^*_{t'}) \cap B^0_t(x^*, u^*_t) = \{\mu_t\}$. Consider any $\delta > 0$, and define $D_\delta \equiv \{\mu : ||\mu - \mu_t|| < \delta\}$ and D^c_δ 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 $\{B_{t'}(x^* - \epsilon, u^*_{t'}) \cap D^c_\delta\} \subseteq \{B_t(x^* - \epsilon, u^*_t) \cap D^c_\delta\}$ 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 action 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 4.2. Implicit differentiation gives that, at point (x, μ), $x = \chi_j(\mu, \hat{u})$, for some $\hat{u}$:

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

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

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

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 $\epsilon \in (0, \epsilon_{t'})$, $B_{t'}(x^* - \epsilon, u^*_{\tilde{t}}) \subset B_t(x^* - \epsilon, u^*_t)$. 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. □

Proof of Proposition B.2. Let $\{\mu_{0,k}\}$ be any sequence of priors that converges to μ_t, and $(\Upsilon_1^{*,k}, ..., \Upsilon_N^{*,k})$ be a D1 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 $\delta \to 0$ as $\epsilon \to 0$.

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 $\epsilon \to 0$, $X^k \to \{0\}$. For any $x > 0$, by choosing ϵ small enough relative to ϵ_m in the proof of Proposition B.1, the argument given there to show that $x^* \not> 0$ 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 $\mu_0 \to \mu_N$. □