Discussion Paper Series – CRC TR 224

Discussion Paper No. 221
Project B 03

Inducing Effort Through Grades

David Rodina [1]
John Farragut [2]

October 2020

[1] University of Bonn. E-mail: drodina1988@gmail.com
[2] Northwestern University. E-mail: johndfarragut@gmail.com

Funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 is gratefully acknowledged.

Inducing Effort through Grades*

David Rodina[†], John Farragut[‡]

October 17, 2018

Abstract

We study the problem of a principal who wants to incentivize an agent's investment in productivity through an information disclosure policy. The agent participates in a market and benefits from an improved perception of his productivity. Under asymmetric information about the agent's ability we explore how qualitative features of the optimal deterministic grading scheme depend on the distribution of the agent's ability. Perhaps surprisingly, random grades can induce higher investment. When the effect of the agent's investment is subject to exogenous shocks and there is no asymmetric information, in a wide variety of circumstances the optimal disclosure policy has a relatively simple threshold form.

*We are particularly grateful to Eddie Dekel and Asher Wolinsky for valuable insights and continual encouragement. We also wish to thank Simon Board, Johannes Hörner, Wojciech Olszewski, Bruno Strulovici, and three anonymous referees for helpful comments.

[†] University of Bonn, Contact> drodina1988@gmail.com. Funding by the Deutsche Forschungsgemeinschaft (DFG) through CRC TR 224 (Project B03) is gratefully acknowledged.

[‡] Northwestern University. Contact: johndfarragut@gmail.com

1 Introduction

An agent (he) takes costly effort to improve the quality of a good. The good will be purchased by a market, and the price it is willing to pay is increasing in the quality. However, the market does not observe the quality directly. Instead, a principal (she) discloses information about the quality to the market. The objective of the principal is to maximize the transaction price. The only instrument that is available to the principal is the information disclosure policy; transfers are not available to her.

An example of such a situation is the problem of a school trying to make its students work hard. Students increase their productive value through studying, and the wage they will receive on the labor market is increasing in their perceived value. The labor market learns about the value through transcripts, and the school can design how much information they contain. The objective of the school is to maximize the expected wage of its graduates, say because schools get ranked according to this criterion. Another application is to credit rating agencies and the bond market; we describe in Section 5 how our model applies to such an environment.

In particular, one can ask what qualitative form the optimal disclosure policy should take: Should the grades be fine and accurately reflect productivity, or be coarse and lump together many different outcomes? Can it be optimal to induce a random grade? An answer to the latter question could also shed some light on the related question of what a good test looks like. Conventional wisdom says that a good test should have a high signal to noise ratio (see Lazear (2006)). This is true if one wants to learn about fixed characteristics of an agent. But that this also provides good incentives for the agent to exert effort is less clear. Certainly if the grade is pure noise then it provides no incentives to spend effort, as the market draws no inference from the grade. Yet it is not clear why effort should be "monotone in the amount of noise" if one allows for very general ways of revealing the performance.

We consider two specifications of the model. In the first part, the productive value is a deterministic function of the agent's effort, and the agent's ability to raise the quality is private information. The nature of the optimal grading policy depends on the distribution of abilities. If it is sufficiently convex (that is, there are relatively many high ability agents), then the quality is revealed in a very coarse way: there is only a small number of "letter grades". On the other hand, if the distribution is sufficiently concave, then it is optimal to reveal the quality when it exceeds a threshold. The basic trade-off is that a coarser disclosure policy discourages lower ability agents, but encourages higher ability agents to exert more effort. In the second part, we consider a setting where there is no asymmetric information, but the outcome can be only imperfectly influenced by effort. In a wide variety of cases, the optimal policy has a threshold form. Fully revealing the quality above the threshold raises incentives to exert effort for sufficiently high effort levels, and garbling below the threshold deters the agent from shirking because the perceived quality of an agent who doesn't meet the threshold is low.

It is worth remembering the following two features of the standard principal agent model where the principal observes some signal about effort, and ex-ante commits to a wage schedule (see e.g. Grossman and Hart (1983)):

1. A more informative performance signal is always preferable.

2. The wage schedule rewards signal realizations that are indicative of higher effort.

In our model, the principal cannot commit to the wage schedule; it is determined by the market belief about the agent's performance. It is precisely the statistical inference which is absent in the standard model that makes coarsening the signal potentially worthwhile, as the incentive effect of the wage schedule does not figure in the market's consideration.

We present a unified model in Section 2. In Sections 3 and 4, we analyze specialized versions of the model. In Section 5 we discuss how our work relates to the literature and explain how our setting fits the motivating example of credit rating agencies. Section 6 concludes, and an Appendix contains proofs not found in the main body.

2 The model

There are two strategic players, a principal and an agent, and one non-strategic player, the market.

The agent has a type θ in a bounded interval $\Theta = [\underline{\theta}, \overline{\theta}] \subset \mathbb{R}$ drawn according to a cdf Φ with full support, and continuous density ϕ. The agent chooses effort e from a finite set $E \subset \mathbb{R}_+$, and receives a wage $w \in \mathbb{R}_+$. The type, effort, and wage together determine the agent's utility $u(w, e, \theta)$. u is strictly increasing in w for all $(e, \theta) \in E \times \Theta$, and is strictly decreasing in e for all $(w, \theta) \in \mathbb{R}_+ \times \Theta$. We denote by $\underline{e}$ the least element of E.

The principal does not observe the agent's type or effort directly. Instead he observes a *score* q. A score is a random variable that takes values in a given finite set $Q \subset \mathbb{R}_+$, and the distribution of which depends upon the agent's effort. $F(\cdot|e)$ denotes the cdf over scores $q \in Q$ given effort e, and $f(\cdot|e)$ the probability mass function. We use $\underline{q}$ to denote the smallest element in Q.

Rather than being allowed to commit to an arbitrary wage schedule (a mapping from scores to wages), the principal commits to a mapping from scores to "grades." Formally, a *grading scheme* is a function $H : Q \to \Delta(\mathbb{R})$ which maps each score q to a finite support distribution $H(\cdot|q)$ on $\mathbb{R}$. H is a *deterministic grading scheme* or *grading function* if it maps each score to a degenerate distribution. In this case we can alternatively write the mapping as $g : Q \to \mathbb{R}$, where $g(q) = g_0$ means q is mapped to the point mass at g_0. The grading scheme is *fully revealing* if it is deterministic and $g(\cdot)$ is injective. For a given grading scheme H, the set of *grades* is $G = \cup_{q\in Q} \operatorname{supp} H(\cdot|q) \subset \mathbb{R}$. This implies that for g a grading function, $G = \cup_{q\in Q}\{g(q)\}$.

The grade can be interpreted as a signal released to the market, which the market uses to determine a wage. Given a grading scheme H and its associated set of grades G, a *market wage function* (or simply *wage function*) is a function $w : G \to \mathbb{R}$ such that if the agent receives grade g, then he is paid wage $w(g)$. We make the important assumption that the wage is equal to the expectation of the score conditional on the grade. Formally, $w(g) = \mathbb{E}[q|g]$, where the expectation takes account of the distribution of types, the grading scheme, and equilibrium play of the agent. We note here that under a fully revealing grading scheme $g(\cdot)$, if q is chosen with positive probability (so that Bayes' rule is unambiguous) then $w(g(q)) = q$. Later assumptions will imply that $w(g(q)) = q$ for q off the path as well.

The timing is as follows. First, the principal publicly commits to a grading scheme H. Then the agent privately chooses effort e, which generates a score q that is observed by the principal but not the market. The market observes a grade g generated according to the score and the grading scheme. Finally the agent is paid the market wage $w(g)$ corresponding to his grade.

The principal's goal is to maximize the expected score of the agent.[1] We take the perspective of mechanism design. For each grading scheme H there is a set of equilibria that may be played by the agent and market (the equilibrium notion will be made precise), and we assume that among these the principal-preferred equilibrium is selected. The principal then chooses a grading scheme that maximizes his payoff from the corresponding equilibrium.

The basic equilibrium notion that we use is weak perfect Bayesian equilibrium (WPBE). Let H be any grading scheme. A *WPBE for H* consists of: for each $\theta \in \Theta$ a distribution over efforts $e^*(\theta) \in \Delta(E)$, and a market wage schedule $w : G \to \mathbb{R}$, such that, (1) $e^*(\theta)$ maximizes the expected payoff of type θ given H, F, and $w(\cdot)$, and (2) for all $g \in G$ that occur with positive probability it holds that $w(g) = \mathbb{E}[q|g]$, where the joint distribution of (g, q) is derived from Φ, e^*, F, and H according to Bayes' rule.

We will consider two versions of the model in Sections 3 and 4.

[1] Under our assumptions, maximizing the expected score is equivalent to maximizing the expected wage.

3 Asymmetric information

In this section we examine a version of the model we call the *asymmetric information* environment. Its main features are that the agent has private information, and the score is a deterministic function of effort.

Definition 3.1 *An asymmetric information environment satisfies:*

1. *Asymmetric information:* $\underline{\theta} < \overline{\theta}$.
2. $Q = E$ *and* $F(\cdot|e)$ *is the point mass at* e.[2]
3. $v(q, \theta) := u(q, q, \theta)$ *is continuous, single-peaked in* q *for all* θ*, and strictly supermodular in* (q, θ).[3]

Part 2 states the notational convention of identifying the score q with the effort that leads to that score with probability 1. For example, we may write $u(w, q, \theta)$ to mean $u(w, e, \theta)$ where $q(= e)$ is the score that occurs with probability one given effort choice e.[4] Part 3 defines a shape restriction on the utility function. Notice that $v(q, \theta) = u(q, q, \theta)$ it the agent's utility function under fully revealing grading. For future reference, define the *unrestricted choice* $q^*(\theta) = \arg\max_q v(q, \theta)$ to be type θ's optimal choice of score (= effort) under fully revealing grading.[5] To rule out uninteresting cases, we assume that $\underline{q} < q^*(\underline{\theta}) < q^*(\overline{\theta}) < \overline{q}$, where $\underline{q} = \min Q$ and $\overline{q} = \max Q$. Also, $\overline{q}$ is sufficiently large in that $v(\overline{q}, \overline{\theta}) < v(\underline{q}, \overline{\theta})$.

In this environment, we consider only WPBE that satisfy the following restriction on off-path beliefs: Fix H. For any $g \in G$ that occurs with probability 0 in equilibrium, $w(g) = \min\{q : g \in \operatorname{supp} H(\cdot|q)\}$. This says that upon observing an off-path grade, the market adopts the worst-case belief about the score that is consistent with that grade, given the grading scheme.

First we show that the problem of selecting a grading function is equivalent in a certain sense to selecting a menu of scores from which the agent is allowed to make a choice. Then

[2]As long as $F(\cdot|e)$ is deterministic one can reparametrize efforts and assume $Q = E$ without loss.

[3]Let $X, Y \subset \mathbb{R}$. A function $f : X \to \mathbb{R}$ is *single-peaked* if there exists $\tilde{x} \in X$ such that for any $x, x' \in X$, if $x < x' \leq \tilde{x}$ or $\tilde{x} \leq x' < x$, then $f(x') \geq f(x)$. A function $f : X \times Y \to \mathbb{R}$ is *strictly supermodular* if for any $x, x' \in X$ and $y, y' \in Y$, if $x < x'$ and $y < y'$, then $f(x', y) - f(x, y) < f(x', y') - f(x, y')$.

[4]As written, the market's optimal wage depends only on effort, while the type affects the agent's preferences. Our models covers some cases where the optimal wage depends on ability: suppose that the market's optimal wage is $\hat{q}(e, \theta)$, while the agent's preferences are $\hat{u}(w, e, \theta)$. This can be translated into our setting by defining $u(w, q, \theta) = \hat{u}(w, e(q, \theta), \theta)$, where $e(q, \theta)$ is the cheapest way to produce q for type θ. So the crucial assumption is not that the optimal wage is independent of type, but rather that the grade is a garbling of the optimal wage.

[5]With a continuous type distribution, strict supermodularity and single-peakedness of v imply that $q^*(\theta)$ is a singleton for almost all θ and any selection is nondecreasing in θ.

we derive some properties of the optimal grading function that are true independent of the type distribution, followed by a sharper characterization that depends on the shape of the type distribution. A quadratic example will be used to provide some intuition for the results. In most of the analysis the principal will be restricted to a deterministic grading scheme, but we also briefly discuss random mechanisms at the end.

3.1 Grading functions and score menus

When restricted to deterministic grading schemes, the problem of the principal was described as selecting a grading function. We will now discuss an observation that allows us to reformulate the principal's problem in a way that is more amenable to analysis. Recall that in this environment the mapping from effort to scores is deterministic, so that we can identify e and q, and in particular choosing an effort is the same as choosing a score.

Consider the following game. Instead of a grading function, the principal announces a *menu* (a nonempty set) of scores $M \subset Q$ from which the agent may choose. The agent's choice q from the menu is disclosed to the market, after which he receives the wage q equal to the score that he selected. The agent's and principal's payoffs are the same as in the main model. We call this the *menu game*. Say that a menu M is *nonredundant* if $\underline{q} \in M$ and any score $q \in M - \{\underline{q}\}$ is selected by a positive measure of types. Note that whether a menu is redundant or not is independent of the distribution, as long as it has full support on Θ.

We say that a grading function or menu implements a score schedule $\hat{q} : \Theta \to Q$ if for all θ, $\hat{q}(\theta)$ is the optimal choice for θ when faced with that grading function or menu. Any grading function or menu implements an essentially unique score schedule.[6] The following lemma suggests that when looking for the optimal grading function, one can instead optimize over menus that contain $\underline{q}$ in the menu game if this is convenient, and then translate the optimal menu into its corresponding grading function.

Lemma 3.2 *Take any grading function and its associated score schedule $\hat{q}(\theta)$. Then the menu $M = \{\hat{q}(\theta) : \theta \in \Theta\} \cup \{\underline{q}\}$ is nonredundant and implements $\hat{q}(\theta)$.*

Conversely, take any menu M with $\underline{q} \in M$ and its associated score schedule $\hat{q}(\theta)$. There exists a grading function that implements $\hat{q}(\theta)$.

Proof The outcome of the grading function is implemented by the menu because there are fewer deviations in the menu game. In the other direction, start with a menu and an

[6]Essentially unique means that any two score schedules that can be part of an equilibrium differ on a measure zero set of types. Without loss of optimality, the score schedule is pure.

equilibrium score schedule $\hat{q}(\cdot)$. Define the following set for each q:

$$\hat{Q}(q) = \{\tilde{q} \in Q : \tilde{q} \leq q \text{ and } \exists \theta \in \Theta \text{ s.t. } \hat{q}(\theta) = \tilde{q}\}.$$

This is the set of scores below q that are selected by some type in equilibrium. Now define a grading function by

$$g(q) = \max\{\tilde{q} : \tilde{q} \in \hat{Q}(q)\},$$

and $g(q) = \underline{q}$ if $\hat{Q}(q)$ is empty. If q was selected in the menu game, then $g(q) = q$. Otherwise, $g(q)$ is the largest score $\tilde{q} < q$ such that $\tilde{q}$ was selected (or $\underline{q}$ if none exists). This implements the outcome of the menu game for the following reason. Any score that was selected in equilibrium from the menu is revealed to the market. Any score that was not selected in the menu game is interpreted by the market as the smallest score below it that was selected by some type. If there was no such score and type, the grade is interpreted as the lowest possible score $\underline{q}$. Therefore when faced with the grading function there are two kinds of deviations. The first is a deviation to either $\underline{q}$ or a score chosen by some type in the menu, but either of these was available in the menu and not worthwhile. Any remaining deviation (to a score not chosen in the menu game) is worse than some deviation from the former category. ∎

One might wonder whether the principal can benefit from eliciting information from the agent. In principle one could screen agents by offering different grading functions, and agents pick a function depending on their type. The principal cannot be worse off compared to offering a single grading function, but it turns out that he cannot benefit either, which is not difficult to see given the equivalence between grading schemes and menus.

Lemma 3.3 *There is no benefit from offering a menu of grading functions.*

Proof By Lemma 3.2, every grading function has an equivalent representation as a menu of scores. Similarly, a menu of grading functions is equivalent to a menu of menus of scores. Taking a union over the menus of scores one gets a single menu of scores, which can be translated back into a single grading function that implements the same outcome as the menu of grading functions. ∎

3.2 Deterministic grading schemes

In this section we derive properties of optimal grading functions (i.e. deterministic grading schemes). The potential benefit of random grading schemes is considered in Section 3.3.

3.2.1 Results independent of the distribution

In this part we identify several qualitative features of the optimal grading function that hold for any type distribution. First we introduce terminology for some particular grading functions that will be of interest later.

Recall that a grading function is fully revealing if $q \neq q'$ implies $g(q) \neq g(q')$, and that under such a grading function the agent chooses $q^*(\theta) = \arg\max_q v(q, \theta)$.

A grading function features a "failing grade combined with perfectly revealing grading" if $g(q)$ is constant below score $\hat{q}$ and strictly increasing starting from $\hat{q}$. An example is

$$g(q) = \begin{cases} \underline{q} & \text{if } q < \hat{q} \\ q & \text{if } q \geq \hat{q} \end{cases}$$

Scores below $\hat{q}$ are given the same "failing grade" $\underline{q}$, and it might or might not be that some types pick the failing grade in equilibrium. Scores $q \geq \hat{q}$ receive distinct grades, and hence are perfectly inferred by the market.

A grading function is of the "pass-fail" form if there is some score $\hat{q}$ such that every score below $\hat{q}$ gets mapped into the same grade g_0, and every score weakly above $\hat{q}$ gets mapped into a separate grade g_1. An example is

$$g(q) = \begin{cases} \underline{q} & \text{if } q < \hat{q} \\ \hat{q} & \text{if } q \geq \hat{q} \end{cases}$$

A grading function "with at most n letter grades" is a grading function whose set of grades has size no greater than n, i.e. $|G| = |\cup_{q \in Q} \{g(q)\}| \leq n$. Pass-fail grading is a special case of a grading function with at most two letter grades. We now present two results that are true for any type distribution.

Proposition 3.4 (Pooling at the bottom) *Fully revealing grading is inferior to some grading function featuring a failing grade combined with perfectly revealing grading.*

Proof Start with a fully revealing grading function, which corresponds to the full menu of scores. Since v is single-peaked in q, type $\underline{\theta}$ prefers $q^*(\underline{\theta})$ to the lowest score $\underline{q}$, and there exists a highest $\tilde{q} \geq q^*(\underline{\theta})$ such that $\underline{\theta}$ prefers $\tilde{q}$ to $\underline{q}$, and all other types do so as well by supermodularity. After removing all scores $q \in (\underline{q}, \tilde{q})$, all types choose weakly higher scores: types near $\underline{\theta}$ (for which $q^*(\theta)$ was in the removed interval) select $\tilde{q}$, and the remaining types' choices are unaffected. For a sufficiently fine grid of scores, this leads to a strict improvement. The new menu corresponds to a grading function featuring a failing grade combined with perfectly revealing grading. ∎

The idea is that if the score that the agent would pick from the full menu is not available, he is willing to pick a slightly higher score rather than the lowest possible score. So by removing scores at the bottom of the menu one can induce low types to select a higher score. Note that we have not ruled out that some types "fail" in equilibrium if the optimal grading function features a failing grade combined with perfectly revealing grading (which does not happen in the preceding construction). The result only illustrates that perfectly revealing grading is dominated.

In fact, the preceding argument shows slightly more: for any nonredundant menu that contains $\tilde{q}$ as defined previously, if the menu contains a score q such that $\underline{q} < q < \tilde{q}$, then this menu is dominated.

Proposition 3.5 (Disclosure at the top) *A grading function that is constant above a threshold is weakly worse than one that features revealing grading above the threshold.*

Proof The nonredundant menu corresponding to the initial grading function has some highest score $q_{\max}$. Whenever $q_{\max} < q^*(\overline{\theta})$, adding all the scores between $q_{\max}$ and $q^*(\overline{\theta})$ leads to an improvement. Types close to $\overline{\theta}$ select higher scores (those types for which $q^*(\theta) \in (q_{\max}, q^*(\overline{\theta})])$ while the remaining types' choices are unaffected. The new menu can be represented by a grading function that reveals all scores above $q_{\max}$. ∎

The specification of the grading function for sufficiently high scores is always irrelevant as these scores are never chosen, so it can be optimal to pool all scores above a threshold, but never strictly so. There always exists another grading function that fully reveals at the top.

3.2.2 Results depending on the shape of the type distribution

We have shown that the principal always benefits from pooling scores below some threshold and revealing all scores above some other threshold, regardless of the utility function or type distribution (satisfying our assumptions). To more fully characterize optimal menus, we now consider the effect of various shape assumptions on the type distribution. We start by considering when the principal wants to include or exclude interior intervals of scores.

Consider a nonredundant menu that contains two scores $q_L < q_H$ and contains no scores in between them; otherwise the menu may be arbitrary. Suppose that score $q_L = q^*(\theta_L)$ is the unrestricted choice for some type θ_L, and $q_H = q^*(\theta_H)$ the one for some type $\theta_H > \theta_L$. Then there is an indifferent threshold type $\hat{\theta} \in (\theta_L, \theta_H)$, characterized by $v(q_L, \hat{\theta}) = v(q_H, \hat{\theta})$, such that all types in $[\theta_L, \hat{\theta})$ choose q_L and all types in $[\hat{\theta}, \theta_H]$ choose q_H; existence of $\hat{\theta}$ follows from continuity of v, and the rest from supermodularity and single-peakedness. The payoff to the principal from types in the interval $[\theta_L, \theta_H]$ does not depend upon other details of the menu, and equals

$$\int_{\theta_L}^{\hat{\theta}} q_L \, \mathrm{d}\Phi(\theta) + \int_{\hat{\theta}}^{\theta_H} q_H \, \mathrm{d}\Phi(\theta).$$

Now suppose that we modify the menu by adding all scores in (q_L, q_H). Then each $\theta \in (\theta_L, \theta_H)$ switches to choosing $q^*(\theta) \in (q_L, q_H)$; the behavior of all θ outside the interval is unchanged. The principal gains from types in $(\theta_L, \hat{\theta})$, below the threshold, who choose a higher score; he loses from types in $[\hat{\theta}, \theta_H)$, above the threshold, who choose a lower one. The total change in payoff equals

$$\int_{\theta_L}^{\theta_H} q^*(\theta) \, \mathrm{d}\Phi(\theta) - \left\{ \int_{\theta_L}^{\hat{\theta}} q_L \, \mathrm{d}\Phi(\theta) + \int_{\hat{\theta}}^{\theta_H} q_H \, \mathrm{d}\Phi(\theta) \right\}$$
$$= \int_{\theta_L}^{\hat{\theta}} [q^*(\theta) - q_L] \, \mathrm{d}\Phi(\theta) - \int_{\hat{\theta}}^{\theta_H} [q_H - q^*(\theta)] \, \mathrm{d}\Phi(\theta).$$

In the second line, the first integral is the gain to the principal and the second the loss.

Roughly, whether the above quantity is positive–so that adding the interval benefits the principal–depends upon three factors:

1. the location of the threshold type, determined by the utility function;
2. how much the principal gains from each type below the threshold as compared to what he loses from each type above it, also determined by the utility function;
3. how many types are below the threshold as compared to above it, determined by the distribution.

Before continuing in full generality we present an example that allows us to isolate the effect of the distribution, described in the third factor.

A quadratic example. For exposition we write as if $Q = \mathbb{R}_+$, but any conclusion would be true for a finite grid Q as well. The agent's utility is $u(w, q, \theta) = w - c(q, \theta)$, quasilinear in the wage and with a quadratic cost function that takes the form:

$$c(q, \theta) = cq - a\theta q + \frac{b}{2} q^2,$$

where a, b, c are parameters with $a, b > 0$. Assume the parameters are such that $q^*(\theta) \geq 0$ for all θ. The agent's choice under fully revealing grading is

$$q^*(\theta) = \frac{1}{b}[1 + a\theta - c].$$

Proposition 3.6 *For a quadratic cost function, an optimal grading function features*

- *a failing grade combined with revealing grading if Φ is concave (ϕ is decreasing),*
- *at most three letter grades if Φ is convex (ϕ is increasing),*
- *at most three letter grades followed by revealing grading if ϕ is single peaked.*

More generally, if the type distribution is convex on some interval $[\theta_1, \theta_2]$, then the optimal menu can contain at most two scores in $[q^*(\theta_1), q^*(\theta_2)]$. On the other hand, if the optimal menu contains an interval of scores $[q^*(\theta_1), q^*(\theta_2)]$ for some θ_1, θ_2, then necessarily the type distribution has to be concave on $[\theta_1, \theta_2]$.

Proposition 3.6 will be explained by looking at the menu version of the problem (we work with closed subsets of Q to ensure that every type has an optimal choice).

Consider a menu that contains two scores $q_L < q_H$, preferred by θ_L, θ_H, respectively, and no score in between them. Suppose that we add to the menu the interval of scores (q_L, q_H). The quasilinear-quadratic utility implies that the indifferent type is equal to the midpoint: $\hat{\theta} = \frac{1}{2}\theta_L + \frac{1}{2}\theta_H$. By linearity we can write $q^*(\theta) = q_L + \beta(\theta - \theta_L) = q_H - \beta(\theta_H - \theta)$ where $\beta = \frac{a}{b} > 0$. Using this, the change in the principal's payoff is

$$\begin{aligned}
&\int_{\theta_L}^{\hat{\theta}} [q^*(\theta) - q_L] \, \mathrm{d}\Phi(\theta) - \int_{\hat{\theta}}^{\theta_H} [q_H - q^*(\theta)] \, \mathrm{d}\Phi(\theta) \\
&= \beta \int_{\theta_L}^{\hat{\theta}} [\theta - \theta_L] \, \mathrm{d}\Phi(\theta) - \beta \int_{\hat{\theta}}^{\theta_H} [\theta_H - \theta] \, \mathrm{d}\Phi(\theta).
\end{aligned}$$

Now integrate by parts and manipulate; one finds that the expression is nonnegative if and only if

$$\Phi\left(\frac{1}{2}\theta_L + \frac{1}{2}\theta_H\right) \geq \frac{1}{\theta_H - \theta_L} \int_{\theta_L}^{\theta_H} \Phi(\theta) \, \mathrm{d}\theta. \tag{1}$$

A sufficient condition for the inequality to hold is that Φ is concave on $[\theta_L, \theta_H]$. If Φ is convex on $[\theta_L, \theta_H]$ then the opposite inequality holds. Therefore if Φ is concave on $[\theta_L, \theta_H]$ and the menu contains q_L, q_H, the principal benefits from adding the entire interval of scores (q_L, q_H) to the menu. If Φ is convex he is better off excluding them.

This suggests the following results about the optimal nonredundant menu M^* (remember that $[q^*(\underline{\theta}), q^*(\overline{\theta})]$ is the set of scores selected under fully revealing grading):

1. When Φ is concave, there can be no "holes" in $M^* \cap [q^*(\underline{\theta}), q^*(\overline{\theta})]$.
2. When Φ is convex, there can be no intervals in $M^* \cap [q^*(\underline{\theta}), q^*(\overline{\theta})]$.

3. In $M^* \setminus [q^*(\underline{\theta}), q^*(\overline{\theta})]$, there are at most two scores selected by some type: one below $q^*(\underline{\theta})$, and one above $q^*(\overline{\theta})$.

In the proof of Proposition 3.6, these observations are used to show that in the case of a concave Φ, the optimal nonredundant menu has the form $\{\underline{q}\} \cup [q_L, q_H]$. This can be implemented by a grading function which features a failing grade combined with perfectly revealing grading. Moreover, the threshold is such that no type selects the failing grade.

In the case of a convex Φ, we have shown that an optimal menu contains no interval of scores (that are unrestricted choices for some types), but we have not ruled out for example the possibility that the menu includes many scores spaced closely together. Such menus turn out to be not optimal, as one can show that at most two scores are selected from the optimal menu.[7] This implies that the optimal grading function features at most three letter grades.

Another interesting case is when the density ϕ is single peaked, so that there is an interior type $\hat{\theta}$ such that to the left (right) of $\hat{\theta}$, ϕ is increasing (decreasing). Combining previous arguments, one can show that all scores above a certain threshold are included in the optimal menu, and at most two scores are selected below the threshold. therefore, the grading function coarsens low scores into at most three letter grades, and fully reveals all scores above the threshold.

We now return to the general case.

General preferences. Apart from the example, whenever removing the interior of an interval of scores, then any type that initially selected a score from the interval will select one of the endpoints. A threshold type will be indifferent between the endpoints. All higher types will select the upper endpoint and hence select a higher score, and all lower types will select the lower endpoint and score. The quadratic setup simplified the issues of the location of the threshold type and how much the principal gained or lost from each type: for each low type $\theta_L + \Delta \in [\theta_L, \hat{\theta})$ for which there was a loss from removing the interval, there was a "mirror" high type $\theta_H - \Delta \in (\hat{\theta}, \theta_H]$ the same distance from the other endpoint and for which the gain exactly offset the loss, and furthermore the gain/loss was linear in type. This allowed us to isolate the effect of the shape of the distribution. Concavity meant there were relatively more low types as compared to high types, and so upon removing an interval the total loss from the low types outweighed the total gain from the high types. So the principal preferred to include intervals. With a convex distribution there were relatively more high types than low types, so that removing intervals helped the principal.

To gain some intuition for the general case, consider an environment the same as in the previous paragraph except that for each low type $\theta_L + \Delta$ that causes a gain when adding the

[7] One can directly show if three scores are selected, the principal can benefit by removing the intermediate one.

interval, the loss from the "mirror" high type $\theta_H - \Delta$ is twice as large. (The threshold type is the same as before.) When should the principal include the interior of the interval? Now it is not sufficient that "there are more low types than high types," i.e. that the distribution is concave, because the principal loses twice as much from each high type by adding the interval as he gains from the mirrored low type; if there are "almost" as many high types as low types then he is better off not adding it. However if the number of low types sufficiently outweighs the number high types–intuitively, if the distribution is concave enough–then the principal is better served to include the interval.

This line of reasoning suggests that if the relative gain/loss to the principal is uniformly bounded across types and the distribution of types is concave enough (or convex or single-peaked enough), then optimal menus should have the same qualitative properties as those in the example. To make this precise we turn to our formal results.

We say that "condition X is true for a sufficiently C type distribution" if there exists a type distribution Φ that has a continuous density and full support such that X is true for Φ, and X remains true for any transformation $\Gamma \circ \Phi$ such that $\Gamma \circ \Phi$ is a cdf with a continuous density and full support and both $\Gamma \circ \Phi$ and Γ have property C.

Proposition 3.7 *Under the "asymmetric information" assumptions, the following are true.*

- *If Φ is sufficiently convex, an optimal menu has at most three scores from which at most two are selected in equilibrium. It can be implemented by a grading function that has two or three letter grades.*

- *If Φ is sufficiently concave, an optimal menu includes all scores above some cutoff. It can be implemented by a grading function that has a failing grade combined with perfectly revealing grading.*

- *If $\phi = \Phi'$ is sufficiently single peaked at a type $\theta_0 \in \Theta$, an optimal menu features a cutoff. Above the cutoff all scores are included, and up to the cutoff at most three scores are included from which at most two are selected in equilibrium. It can be implemented by a grading function that has at most three letter grades combined with perfectly revealing grading.*

The proof appears in the appendix. The approach is quite different from the one used in the quadratic example and must address more subtleties. The key step is a single crossing condition. It states that if we look at the set of scores selected by a nontrivial set of types and take three adjacent scores $q_1 < q_2 < q_3$ from that set, then if it is worthwhile to keep q_2 in the menu (remove q_2 from the menu), then this remains true if we take a concave (convex) transformation of the type distribution. However it is not necessarily true that the payoff difference is monotone in the relevant transformation.

3.3 Random grading schemes

So far we considered deterministic grading schemes, which we saw was equivalent to offering a simple menu of scores and then perfectly revealing the choice. Can one implement more outcomes using random grading schemes? If yes, can the principal exploit them to increase the expected score? As an example, suppose the grade is pure noise. Since it carries no information about the score, the wage will be a constant independent of the grade. Therefore every type of the agent will spend no effort. If one could somehow show that expected effort is monotone in the amount of noise, one could conclude that deterministic grading functions are optimal. An example in Appendix B shows that this is true in a setting where realized score and noise follow a normal distribution.

Nevertheless, the following reasoning indicates why more general randomizations can raise the expected score. Start with a deterministic grading function where scores q_L and q_H, with $q_L < q_H$, are selected in equilibrium, resulting in scores g_L and g_H, respectively. All other scores are pooled resulting in the lowest possible wage $\underline{q}$. Moving to a random grading scheme where q_L, q_H lead to respective randomizations $\tilde{g}_L, \tilde{g}_H$ with support $\{g_L, g_H\}$, the lower score q_L becomes relatively more attractive. The principal can fix the type θ_{LH} who is indifferent between q_L and q_H, and adjust these scores to keep him indifferent by raising q_L and reducing q_H after moving to the randomization. Intuitively, this could be beneficial if there are relatively many low ability types, and their marginal cost of increasing the score is low.

However, this intuition has two gaps. First, for a given increase in the expected wage following q_L, the reduction in the expected wage associated with q_H has to be large if there are relatively many low types. Thus, q_H needs to be reduced substantially to keep θ_{LH} indifferent. Second, if low ability types have a low marginal cost of increasing their score, it is not obvious that the original grading scheme which pooled scores in (q_L, q_H) was the best deterministic grading scheme in the first place. Nevertheless, an example in Appendix B formalizes the preceding discussion and shows that a random grading scheme can improve over the best deterministic one.

Finally, it should be noted that the results on menus in Lemmas 3.2 and 3.3 are not true once one allows for random grading schemes. First, it is not true anymore that a grading scheme can be identified with a subset of the 45-degree line in the score-wage space.[8] Second,

[8] When the agent is risk neutral with respect to the wage, lotteries over wages can be identified by their expectation. Looking at the score-*expected wage* space, a random grading scheme can induce a menu that is

it is not without loss of generality to restrict attention to a single grading scheme when the agent is not risk neutral with respect to the wage. A menu of grading schemes allows that a given score leads to a different wage distribution, depending on which grading scheme was selected.

4 Random scores

In this section we consider the *random score* environment where there is no asymmetric information and the score is random conditional on effort. The support of scores is "connected" and effort shifts the distribution "upwards" as defined below.

Definition 4.1 *A random score environment satisfies:*

1. *No asymmetric information:* $\underline{\theta} = \overline{\theta}$.
2. *There exist two functions* $\underline{q}(e)$ *and* $\overline{q}(e)$ *such that*
   - $\text{support}(F(\cdot|e)) = \{q \in Q : \underline{q}(e) \leq q \leq \overline{q}(e)\}$,
   - $\underline{q}(e)$ *and* $\overline{q}(e)$ *are nondecreasing.*
3. $F(q|e)$ *is increasing in the monotone likelihood ratio.*[9]
4. *Quasilinear utility:* $u(w, e, \theta) = w - c(e)$.

We consider only WPBE that satisfy a sequential equilibrium-style refinement, as follows. Fix a grading scheme H and let e^* denote the equilibrium effort strategy of the agent. Then there exists a sequence σ_n of completely mixed effort strategies such that $\sigma_n \to e^*$ and for any grade $g \in G$, $w(g) = \lim_{n\to\infty} \mathbb{E}[q|g, \sigma_n]$. This means that the wage paid by the market is derived from consistent beliefs.[10]

One consequence of the score being random is that the concept of a score menu is not useful, so we work directly with grading schemes. We restrict attention to pure strategy equilibria only (i.e. effort is pure) but allow for arbitrary random grading schemes. In this case, the following can be said about the optimal grading scheme.

Proposition 4.2 *Under the assumptions of the random score environment:*

not a subset of the 45-degree line.

[9]A family of cdfs $F(\cdot|y) : X \to [0, 1]$ parametrized by $y \in Y \subset \mathbb{R}$ with densities $f(\cdot|y)$ is increasing in the *monotone likelihood ratio (MLR)* if for any $x \leq x'$ and $y \leq y'$ such that $x, x' \in \text{supp}\, F(\cdot|y) \cap \text{supp}\, F(\cdot|y')$, it holds that $\frac{f(x|y')}{f(x|y)} \leq \frac{f(x'|y')}{f(x'|y)}$.

[10]If the action set of the market (the set of wages) was finite, our restriction would coincide with sequential equilibrium.

1. *There exists an optimal grading scheme that reveals all scores above a threshold, and every score realized in equilibrium is above the threshold.*

Under the additional assumption that the effort under full revelation e_{FR} is such that the lowest realized score is off-path (i.e., $\underline{q}(e_{FR}) > \underline{q}$):

2. *Full revelation is weakly suboptimal.*
3. *An optimal grading scheme features a failing grade combined with perfectly revealing grading if the minimal effort $e = \underline{e}$ generates the lowest score $q = \underline{q}$ with probability* 1.

This means that instead of optimizing over all possible grading schemes, the principal only needs to optimize over a cutoff above which all scores are revealed. In a second step, he needs to design the grading scheme for scores below the threshold in such a way that the agent is willing to spend enough effort to always make it above the threshold. Starting from any information structure and associated equilibrium effort, we construct a new information structure that is of the form described in the proposition and leads to a weakly higher equilibrium effort.

If one can sustain a sufficiently high effort by some mechanism (e.g. the fully revealing one) so that certain low scores are off-path, then one can always weakly improve upon full revelation. Also, if the lowest effort always generates the lowest possible score, it is optimal to pool all scores below the fully revealing region. We proceed with two examples to provide some insight into the two features of the optimal grading schemes, namely that on-path scores get fully revealed, but off-path scores might lead to a garbling of the grade.

To see why all scores in equilibrium should be revealed, consider the following example. Suppose that the support of the scores is independent of effort. To facilitate a neat exposition, we also let effort be a continuous variable in this example that affects the distribution of scores in a differentiable way. Consider a coarse grading scheme where scores get partitioned into low scores Q_L and high scores Q_H, and equilibrium effort under the coarse grading scheme is $\hat{e}$. When the agent contemplates about deviating, his payoff is

$$\Pr(Q_L|e)\mathbb{E}[q|q \in Q_L, \hat{e}] + \Pr(Q_H|e)\mathbb{E}[q|q \in Q_H, \hat{e}] - c(e).$$

For example if one of the scores in Q_L is realized, then the observer pays $\mathbb{E}[q|q \in Q_L, \hat{e}]$ which is the average score in set Q_L, based on the equilibrium conjecture about the agent's effort. The derivative with respect to e equals

$$\frac{\mathrm{d}\,\Pr(Q_L|e)}{\mathrm{d}e}\mathbb{E}[q|q \in Q_L, \hat{e}] + \frac{\mathrm{d}\,\Pr(Q_H|e)}{\mathrm{d}e}\mathbb{E}[q|q \in Q_H, \hat{e}] - c'(e), \tag{2}$$

and has to be zero at $\hat{e}$. If in turn the score gets perfectly revealed, then using the law of total expectations the agent's objective can be rewritten as

$$\int q \, \mathrm{d}\Phi(q|e) - c(e) = \Pr(Q_L|e)\mathbb{E}[q|q \in Q_L, e] + \Pr(Q_H|e)\mathbb{E}[q|q \in Q_H, e] - c(e).$$

The derivative with respect to e equals

$$\left\{ \frac{\mathrm{d}\Pr(Q_L|e)}{\mathrm{d}e}\mathbb{E}[q|q \in Q_L, e] + \frac{\mathrm{d}\Pr(Q_H|e)}{\mathrm{d}e}\mathbb{E}[q|q \in Q_H, e] - c'(e) \right\} + \sum_{i \in \{L,H\}} \Pr(Q_i|e)\frac{\mathrm{d}\mathbb{E}[q|q \in Q_i, e]}{\mathrm{d}e}.$$

Evaluating this at $\hat{e}$, the equilibrium effort under the coarse grading scheme, the term in curly brackets is zero as this is the condition determining $\hat{e}$ as can be seen in equation (2). Now under the MLR condition, $\mathrm{d}\mathbb{E}[q|q \in Q_i, e]/\mathrm{d}e$ is positive, and strictly so in general.[11] This is a partial equilibrium based reasoning starting from a specific coarse grading function, yet the idea why full revelation does better is quite compelling.

When a set of scores gets pooled into grade $\hat{g}$, then the expected score conditional on observing $\hat{g}$ is determined by the market's conjecture about effort. By deviating, the agent can affect the probability generating $\hat{g}$, but not the wage $w(\hat{g})$. Under the MLR property, an upward deviation from equilibrium raises the expected score conditional on $\hat{g}$. But this deviation is unobserved, so unlike in the case of full revelation, the agent does not get rewarded for raising the expected score conditional on $\hat{g}$. Thus, incentives are dampened relative to the case of perfectly revealing grading, where the agent appropriates the "full marginal benefit" of effort. This is suggestive of why the principal would like to reveal scores.

However when the support of scores is allowed to depend on effort, some care needs to be taken as certain grades are off-path, yet can be induced by the agent through a deviation. This can be illustrated when the score is a deterministic function of effort, so that $q(e) = e$. If the score gets fully revealed, the agent maximizes the expected score minus the cost, which is maximized at some interior effort level $\hat{e}$. On the other hand, let e^* be the highest effort that makes the agent indifferent to $e = \underline{e}$. In general, we have $\underline{e} < \hat{e} < e^*$. Then one can find a pass-fail scheme that induces e^* such that the fail grade is associated with an off-path belief of $e = \underline{e}$, and this is better than revealing the score as $e^* > \hat{e}$. This is beneficial since if the high effort level e^* is expected, then low scores are off path and one can pool

[11] It is here that we use the assumption that the support of scores is independent of e, as even under the sequential equilibrium type refinement one can assign off-path beliefs that violate the MLR conclusion.

them and assign a low off-path belief to them, making downwards deviations by the agent unprofitable. This is why in general the principal does not want to fully reveal low scores that are not realized in equilibrium.[12] Unless the lowest effort leads to a deterministic score, it will in general be optimal to randomize over grades after off-path score in the principal's most preferred equilibrium.

The reasoning for Proposition 4.2 is similar as that behind Proposition 5.2 in Dewatripont, Jewitt and Tirole (1999) and Proposition 3.4 in Rodina (2017). While they study a career concerns model where the agent wants to maximize the perception of his type, also there the principal wants to maximize effort. By redefining the type in each of these papers in the right way one can embed our model in their setup, yet Proposition 4.2 does not follow from either paper. Dewatripont, Jewitt and Tirole (1999) consider how changes in the grading policy affect the marginal incentives to spend effort (which we did in our reasoning following the proposition), but this local argument is in general not sufficient as shown in Rodina (2017). Even if a given grading policy has lower marginal incentives than an alternative one *at the equilibrium level of effort*, this in general is not enough to compare equilibrium levels. In Appendix A we provide a global argument. The full revelation result in Rodina (2017) does not apply here as there the support of scores is independent of effort.

5 Discussion

Related literature. Our paper belongs to the literature on information design and moral hazard. We distinguish between *investment* and *signaling* models. In either type of model, a principal commits to reveal some information about the agent to the market, who then rewards the agent based on its belief about his productivity. In an investment model, the agent's effort affects his productivity and the principal reveals information about productivity. In a signaling model, productivity depends on the agent's ability (and possibly effort), which is private information. The principal reveals information about the agent's effort, which leads to some inference about his productivity.

Our paper and Boleslavsky and Kim (2018) belong to the class of investment models. For general preferences of the principal and agent, they show how to optimally resolve the tradeoff between inducing a favorable distribution of wages and incentivizing the agent to exert effort. In particular, their paper is related to our random scores section where the agent has no

[12] In the deterministic case a pass-fail scheme is optimal, yet one could also write it as a fully revealing scheme with a failing grade, where the fully revealing part starts at the pass grade. But this type of indeterminacy only arises when the support of scores is degenerate given effort. If the support of θ is non-degenerate, then the best pass-fail scheme is in general worse than the best revealing scheme with a fail grade.

private information. They characterize the optimal disclosure policy for general preferences of the principal and agent when the set of productivities is binary. We consider the case where principal and agent are risk neutral with respect to the market wage, but allow for any number of productivities and relax the full support assumption, which has a qualitative effect on the results.

A separate branch of the literature on information design and moral hazard considers signaling models, such as Board (2009), Rayo (2013), Zubrickas (2015), and Olszewski and Siegel (2016).

Rayo (2013) studies the optimal mechanism to sell a status good that is used by the buyer to signal his type. The seller offers a menu of signals, each of which reveals a subset of types. He derives which types should be pooled into the same signal depending on preferences and the distribution of types. Board (2009) allows for more general payoff effects when a set of agents is pooled compared to Rayo (2013), and explores how these peer effects impact the optimal mechanism. Zubrickas (2015) considers also the case where the market has myopic beliefs about how scores on a test get translated into grades.

The main tradeoff in our asymmetric information section is that when pooling multiple scores into the same grade, lower types tend to spend less effort and higher types more. This is somewhat opposite to Rayo (2013), where the agent wants to maximize the belief about his type. Within a set of types that gets pooled, lower types are willing to pay more since they are perceived as the average which makes them better off, and the opposite effect is true for higher types. There is an additional effect where types that get fully revealed and are just above the pooling region are willing to pay strictly more in order to avoid being perceived as the average type in the pool. In summary, the principal can extract higher payments from high and low types, and less from intermediate ones.

In Rayo (2013), the optimal mechanism switches between separating and pooling regions, and sufficiently high types fall into a separating region. In particular, this rules out mechanisms that have a few coarse categories. It is also not clear whether grading schemes that keep switching between pooling and separating regions are commonly observed in practice. Admittedly, many grading schemes can be optimal in our setting, but they take commonly observed forms in the natural cases we consider. For example, a few coarse categories are sometimes used when evaluating the hygiene of restaurants, a setting where the investment model appears more relevant.

Olszewski and Siegel (2016) consider the welfare effect of different information revelation policies in college admission contests. Lowering the informativeness of test scores through coarse categorizations reduces welfare through a reduced match quality, but nevertheless can lead to a Pareto improvement by reducing incentives to exert wasteful effort.

While the tradeoff under asymmetric information is reasonably straightforward, we do not see an easy way to provide a characterization of the optimal grading function for an arbitrary distribution of types. To appreciate this, a comparison with Olszewski and Siegel (2016) is useful. In their model, wages are given by some distribution and are awarded to agents on an assortative basis. The type with with the p-percentile highest effort receives the p-percentile highest wage. For a given distribution of wages one can use payoff equivalence techniques to solve for the effort allocation. Garbling information leads to a mean preserving contraction of the wage distribution, and this can be used to assess the impact on the effort schedule. On the other hand, in our model the distribution of wages paid is endogenous as it coincides with the distribution of scores generated, and this prevents us from applying payoff equivalence to determine the effort schedule.

In Moldovanu, Sela and Shi (2007) and Dubey and Geanakoplos (2010), agents have status concerns and care how they are perceived relative to their peers. The principal discloses a relative ranking of agents and can choose how informative it is. There is only one agent in our setting, but even with multiple agents there would be no benefit from making the grading scheme not separable across agents.

Related to the literature on information design and moral hazard are career concerns models with information design questions. Some recent examples are Dewatripont, Jewitt and Tirole (1999), Smolin (2015), Rodina (2017), and Hörner and Lambert (2018).

Examples. The story of a university that evaluates its prestige as the average salary of its graduates adapts straightforwardly. Rather than supply details for that, we describe two stylized examples: credit ratings agencies and mechanisms for public good contributions.

Rating agencies. A firm wants to finance a project. The project is successful with probability q and delivers a rate of return k, and with probability $1-q$ it fails and generates no returns. Projects have an inherent quality θ, and the firm can take costly actions that reduce the default probability. It is easier to prevent default for higher quality projects.

The firm needs to raise money on the bond market; in order to do so, the bond must be rated by a rating agency. The agency observes q and publishes a rating g, which is observed by the market. The bond pays \$1 tomorrow, and the market is willing to pay $w = E[q|g]$ for it. The rating agency's payment is proportional to the price of the bond; it receives $\alpha \cdot w$. Therefore, the firm can invest $(1-\alpha)w$. The payoff of the firm is

$$u(w, q, \theta) = q[(1-\alpha)wk - 1] - \alpha w - c(q, \theta).$$

This setup fits the assumptions of the asymmetric information environment from Section 3. The credit rating agency (the principal) commits to a rating scheme in order to maximize the expected payment. The firm's (agent's) payoff function satisfies all the assumptions imposed if $c(q, \theta)$ is continuous, strictly submodular, and sufficiently convex.[13]

Public good contribution mechanism. [14] A charity collects contributions towards a public good, and can specify which contribution levels a citizen can make. The level of the public good is the sum of individual contributions, $Q = \sum_i q_i$, and each citizen has preferences $Q - c(q_i, \theta)$. For example, if q_i is the time citizen i spends to volunteer on a project, then the previous results identify when it can be optimal to require a minimum time commitment, but otherwise flexibility ("failing grade combined with fully revealing grading"), or when there should be just a few options, each specifying how much the citizen needs to volunteer ("letter grades"). This requires a commitment to not accept certain contribution levels. Alternatively, citizens have negligible preferences over the public good (relative to the cost of providing it), but want to induce the belief that their contribution was high.

Commitment. When the principal has commitment power, his problem is the same whether his objective is to maximize the expected score or the expected wage, and so our results are true under either assumption. The second case is relevant if the principal-agent relationship is such that the principal receives a fixed fraction of the agent's wage. For a concrete example consider the residential real estate market. A realtor (principal) markets homes by publicizing information (grades) about them. He is paid a fixed percentage of the sale price (wage) and would like the homeowner (agent) to make costly investments to improve the home's quality (score). Without commitment, the two objectives differ. When the principal cares about the score, he is ex post indifferent about releasing any grade and so is willing to do as prescribed by commitment. When his objective is the wage, he will always release the grade that pays the highest wage and hence only a babbling equilibrium exists and agents exert no effort.

Relation to signaling models. It is important for our results that the principal observes directly the true value of the agent, i.e. that, if known by the market, the score would be sufficient to determine the agent's wage, regardless of the market's beliefs. To appreciate the subtlety, compare to the classic signaling model of Spence (1973). There the agent's value is his unobserved talent (type) and what is observed is his choice of education (score). Even though in equilibrium education choice perfectly reveals the type, this inference hinges upon the equilibrium beliefs and so is not consistent with our assumptions. Ignoring the component of optimal policy design, the Spence model assumes one extreme, that education is a costly signal and per se totally uninformative of market value. We assume the opposite,

[13] Define c on a convex set of q's and suppose it satisfies $2(1-\alpha)k < c_{qq}$.

[14] We would like to thank an anonymous referee for suggesting this example.

that what is observed by the principal is directly sufficient to determine a correct market wage. Other than for a few parametric special cases, relaxing this assumption in general made our problem intractable.

Relation to delegation models. The asymmetric information environment can also be interpreted as a delegation model, in which an agent wants to pick a type dependent action (score) and the principal wants the highest action possible. Furthermore the principal must respect an individual rationality constraint (corresponding to choosing the lowest score). The connection is clear from the equivalence between grading functions and score menus: the score menu plays the role of the permitted set of actions and all related results carry over at once.

6 Conclusion

We study the problem of a principal who wants to raise a productive investment by an agent when the only instrument available is an information disclosure policy about the agent's productivity. The focus is on two distinct cases, and we derive some properties of the optimal disclosure policy. In the first case the agent's incentive to invest is private information, and the main point is that when garbling information, low types get discouraged and the opposite is true for high types. Despite this it appears difficult to provide a full characterization of the optimal mechanism, but some qualitative properties can be elucidated and how they depend on the distribution of types. In the second case there is symmetric information, but the realized productivity is affected by factors beyond the agent's control such as environmental shocks. In a wide class of circumstances the disclosure policy has a threshold form that is fully revealing in equilibrium, which is seemingly at odds with the special case where the agent can perfectly control the productivity.

The tradeoff that the principal faces when garbling information is quite distinct. In the asymmetric information model, it has opposite effects on high and low types, while in the random score model it reduces incentives if information realized on the equilibrium path gets garbled. On the other hand, garbling off-path information can raise incentives so that full disclosure is often not optimal. In an environment that combines asymmetric information and a random productivity we expect both effects to operate. More interesting would be a version of the model where the principal cannot observe productivity perfectly, but a signal of it that can be "gamed" by the agent in a systematic way. In a situation where productivity is a function of multiple factors, some of which are exogenous and others affected by the agent's actions, it could be that the principal designs a test that systematically overemphasizes certain factors over others. This would be at odds with our assumption that the principal directly garbles productivity and would introduce a signaling component to our model, which would necessitate different techniques.

References

Board, Simon. 2009. "Monopolistic Group Design with Peer Effects." *Theoretical Economics*, 4(1): 89–125.

Boleslavsky, Raphael, and Kyungmin Kim. 2018. "Bayesian Persuasion and Moral Hazard." *Mimeo.*

Dewatripont, Mathias, Ian Jewitt, and Jean Tirole. 1999. "The Economics of Career Concerns, Part I: Comparing Information Structures." *The Review of Economic Studies*, 66(1): 183–198.

Dubey, Pradeep, and John Geanakoplos. 2010. "Grading Exams: 100, 99, 98, or A, B, C?" *Games and Economic Behavior*, 69(1): 72–94.

Grossman, Sanford J, and Oliver D Hart. 1983. "An Analysis of the Principal-Agent Problem." *Econometrica*, 7–45.

Hörner, Johannes, and Nicolas Lambert. 2018. "Motivational Ratings." *Mimeo.*

Lazear, Edward P. 2006. "Speeding, Terrorism, and Teaching to the Test." *The Quarterly Journal of Economics*, 1029–1061.

Moldovanu, Benny, Aner Sela, and Xianwen Shi. 2007. "Contests for Status." *Journal of Political Economy*, 115(2): 338–363.

Olszewski, Wojciech, and Ron Siegel. 2016. "Pareto Improvements in the Contest for College Admissions." *Mimeo.*

Rayo, Luis. 2013. "Monopolistic Signal Provision." *The BE Journal of Theoretical Economics*, 13(1): 27–58.

Rodina, David. 2017. "Information Design and Career Concerns." *Mimeo.*

Smolin, Alex. 2015. "Optimal Feedback Design." *Mimeo.*

Spence, Michael. 1973. "Job Market Signaling." *The Quarterly Journal of Economics*, 355–374.

Zubrickas, Robertas. 2015. "Optimal Grading." *International Economic Review*, 56(3): 751–776.

A Proofs

Proof of Proposition 3.6 Since the cost function is quadratic, $q^*(\theta)$ is linear and can be written as $\alpha + \beta\theta$, where $\beta = a/b$, and extended beyond $[\underline{\theta}, \overline{\theta}]$. For any q_i, θ_i will denote the type that satisfies $q^*(\theta_i) = q_i$

Failing grade combined with revealing grading if Φ is concave. A failing grade combined with revealing grading is equivalent to a menu of the form $M = \{\underline{q}\} \cup [\hat{q}, \infty)$. Part 1 below shows that some optimal menu takes the form $\{\underline{q}, q_1\} \cup [q_2, q_3] \cup [q_4, \infty)$, where $[q_2, q_3] \subset [q^*(\underline{\theta}), q^*(\overline{\theta})]$. Part 2 shows that one can further restrict attention to menus of the form $\{\underline{q}, q_1\} \cup [q_2, \infty)$. Given this, part 3 covers the case where $q_2 \leq q^*(\overline{\theta})$, and part 4 the case where $q_2 > q^*(\overline{\theta})$, to show that some optimal menu has the form $M = \{\underline{q}\} \cup [\hat{q}, \infty)$.

1. Take any menu where $q_L, q_H \in [q^*(\underline{\theta}), q^*(\overline{\theta})]$ are selected, yet (q_L, q_H) is excluded. It follows from equation (1) that adding (q_L, q_H) to the menu leads to an improvement. Since to either side of $[q^*(\underline{\theta}), q^*(\overline{\theta})]$ at most one score is selected, we can restrict attention to menus of the form $\{\underline{q}, q_1\} \cup [q_2, q_3] \cup [q_4, \infty)$, where $[q_2, q_3] \subset [q^*(\underline{\theta}), q^*(\overline{\theta})]$.

2. Suppose $q_3 \in [q^*(\underline{\theta}), q^*(\overline{\theta})], q_4 > q^*(\overline{\theta})$ are selected. Let $\hat{\theta} = \frac{\theta_3+\theta_4}{2}$ be the type who is indifferent between q_3 and q_4. Adding (q_3, q_4) to the menu changes the expected score by
$$\int_{\theta_3}^{\hat{\theta}} (q^*(\theta) - q_3)\mathrm{d}\Phi(\theta) - \int_{\hat{\theta}}^{\overline{\theta}} (q_4 - q^*(\theta))\mathrm{d}\Phi(\theta).$$
Straightforward calculations show that the change in the expected score is
$$\beta\left[\Phi(\hat{\theta})(\theta_4 - \theta_3) - \int_{\theta_3}^{\theta_4} \Phi(\theta)\mathrm{d}\theta\right].$$
Since Φ is concave on $[\underline{\theta}, \overline{\theta}]$, it is also concave on $[\underline{\theta}, \theta_4]$, so the change in the expected score is nonnegative. Therefore, we can restrict attention to menus of the form $\{\underline{q}, q_1\} \cup [q_2, \infty)$, where $q_1 < q^*(\underline{\theta})$.

3. Case $q_2 \leq q^*(\overline{\theta})$: Let $\hat{\theta} = \frac{\theta_1+\theta_2}{2}$ be the type who is indifferent between q_1 and q_2. Making q_1 redundant by adding score $q^*(\underline{\theta} + (\underline{\theta} - \theta_1))$ to the menu, types in $[\underline{\theta}, \hat{\theta}]$ increase their score from q_1 to $q^*(\underline{\theta} + (\underline{\theta} - \theta_1))$, gaining $2\beta(\underline{\theta} - \theta_1)$. The indifferent type between

$q^*(\underline{\theta} + (\underline{\theta} - \theta_1))$ and q_2 is $\frac{\underline{\theta}+(\underline{\theta}-\theta_1)+\theta_2}{2} = \theta_2 - (\hat{\theta} - \underline{\theta})$, so types in $(\hat{\theta}, \theta_2 - (\hat{\theta} - \underline{\theta}))$ reduce their score from q_2 to $q^*(\underline{\theta} + (\underline{\theta} - \theta_1))$, losing $\beta(\theta_1 + \theta_2 - 2\underline{\theta})$. The net benefit

$$2\beta(\underline{\theta} - \theta_1)\Phi(\hat{\theta}) - \beta(\theta_1 + \theta_2 - 2\underline{\theta})(\Phi(\theta_2 - (\hat{\theta} - \underline{\theta})) - \Phi(\hat{\theta}))$$

is positive if Φ is concave. Furthermore, adding all scores in $(q^*(\underline{\theta} + (\underline{\theta} - \theta_1)), q_2)$ leads to a further improvement which follows from the first case considered in the proof.

4. Case $q_2 > q^*(\overline{\theta})$: Replacing q_1, q_2 with $q_1 - \varepsilon, q_2 + \varepsilon$ leaves the indifferent type unaffected, and since ε has an unrestricted sign, the expected score has to be independent of ε. Therefore, the original menu can be replaced with $\{\underline{q}\} \cup [q_2', \infty)$, where $q_2' = q_2 + q_1 - \underline{q}$.

At most three letter grades if Φ is convex. The discussion following equation (1) shows that some optimal menu contains no interval of scores that are selected. Now take any menu where score q_1, q_2, q_3 are selected. Denoting by θ_{ij} the type who is indifferent between q_i and q_j, removing q_2 changes the expected score by

$$-\int_{\theta_{12}}^{\theta_{13}} (q_2 - q_1)\mathrm{d}\Phi(\theta) + \int_{\theta_{13}}^{\theta_{23}} (q_3 - q_2)\mathrm{d}\Phi(\theta).$$

Straightforward calculations show that this can be written as

$$(q_3 - q_1)\left[\frac{q_2 - q_1}{q_3 - q_1}\Phi(\theta_{12}) + \frac{q_3 - q_2}{q_3 - q_1}\Phi(\theta_{23}) - \Phi(\theta_{13})\right].$$

When Φ is convex this is nonnegative since $\frac{q_2-q_1}{q_3-q_1}\theta_{12} + \frac{q_3-q_2}{q_3-q_1}\theta_{23} = \theta_{13}$.[15]

At most three letter grades followed by revealing grading if ϕ is single peaked. Let $\hat{\theta}$ be the value that maximizes ϕ, and $\hat{q} = q^*(\hat{\theta})$.

Given any menu M, let $\tilde{q} = \min\{q \in M : q \geq \hat{q}\}$. The argument for concave Φ (steps 1 and 2) shows that adding $[\tilde{q}, \infty)$ to the menu leads to an improvement.

From the argument for convex Φ, some optimal menu does not contain any interval of scores below $\hat{q}$, and at most two scores are selected. So there exists a menu of the form $\{\underline{q}, q_1, q_2\} \cup [\tilde{q}, \infty)$ which leads to a weak improvement over the original menu. ∎

Proof of Proposition 3.7 Fix any grid Q. For each of the cases in Proposition 3.7, we will exhibit a type distribution such that the relevant statement is true and show that it remains so under appropriate transformations of the distribution. A generic transformation will be denoted as $\phi \cdot \gamma$ which one transforms density $\phi(\theta)$ into density $\phi(\theta) \cdot \gamma(\theta) / \int_\Theta \phi(x) \cdot \gamma(x)\mathrm{d}x$.

[15] Due to the quadratic cost function, $\theta_{ij} = \frac{q_i+q_j}{2\beta}$.

Note that the set of nonredundant menus is independent of the type distribution as long as its support is kept fixed.

Step 1: Interior scores. Fix any $q_1, q_2, q_3 \in Q$ with $q_1 < q_2 < q_3$. We will be comparing nonredundant menus $M = Q_0 \cup \{q_1, q_2, q_3\}$ and $M' = Q_0 \cup \{q_1, q_3\}$[16], where Q_0 contains no scores between q_1 and q_3 but is otherwise arbitrary, and q_1 is selected by a positive measure of types if $q_1 = \underline{q}$. Let θ_{ij} be the type that is indifferent between q_i and q_j. We have that $\theta_{12} < \theta_{13} < \theta_{23}$, the strict inequalities coming from the fact that M is nonredundant. Let $W(M, \phi)$ be the expected score or welfare under menu M and distribution ϕ, and $\Delta(M, M', \phi)$ the difference in welfare. Since the choice of types outside $[\theta_{12}, \theta_{23}]$ is not affected, we get that

$$\Delta(M, M', \phi) = W(M', \phi) - W(M, \phi) = -\int_{\theta_{12}}^{\theta_{13}} (q_2 - q_1)\ \mathrm{d}\Phi(\theta) + \int_{\theta_{13}}^{\theta_{23}} (q_3 - q_2)\ \mathrm{d}\Phi(\theta).$$

There exists a continuous function $\gamma_1(\theta) > 0$ on $[\underline{\theta}, \overline{\theta}]$ such that $\Delta(M, M', \phi \cdot \gamma_1) > 0$. Also there exists a continuous function $\gamma_2(\theta) > 0$ on $[\underline{\theta}, \overline{\theta}]$ such that $\Delta(M, M', \phi \cdot \gamma_2) < 0$. From now on any function γ, γ_1 or γ_2 will be continuous and strictly positive on $[\underline{\theta}, \overline{\theta}]$.

Step 2: Preserving optimality. If $\Delta(M, M', \phi) > 0$, then $\Delta(M, M', \phi \cdot \gamma_1) > 0$ if γ_1 is increasing. If $\Delta(M, M', \phi) < 0$, then $\Delta(M, M', \phi \cdot \gamma_2) < 0$ if γ_2 is decreasing. This single crossing property says that if it is optimal to add (remove) scores in the interior of the menu, this remains true after a decreasing (increasing) transformation of the density. We will show that $\Delta(M, M', \phi) > 0 \Rightarrow \Delta(M, M', \phi \cdot \gamma_1) > 0$ if γ_1 is increasing, the other case is analogous.

$$\begin{aligned} \Delta(M, M', \phi \cdot \gamma_1) &\overset{s}{=} -\textstyle\int_{\theta_{12}}^{\theta_{13}} (q_2 - q_1)\phi(\theta)\gamma_1(\theta)\ \mathrm{d}\theta + \int_{\theta_{13}}^{\theta_{23}} (q_3 - q_2)\phi(\theta)\gamma_1(\theta)\ \mathrm{d}\theta \\ &= -\gamma_1(\theta') \textstyle\int_{\theta_{12}}^{\theta_{13}} (q_2 - q_1)\phi(\theta)\ \mathrm{d}\theta + \gamma_1(\theta'') \int_{\theta_{13}}^{\theta_{23}} (q_3 - q_2)\phi(\theta)\ \mathrm{d}\theta, \end{aligned}$$

for some $\theta' \in [\theta_{12}, \theta_{13}]$ and $\theta'' \in [\theta_{13}, \theta_{23}]$, so that $\gamma_1(\theta'') \geq \gamma_1(\theta')$ and $\Delta(M, M', \phi \cdot \gamma_1) > 0$.[17]

Step 3: Scores below $q^*(\underline{\theta})$. Fix any $q_1 < q^*(\underline{\theta})$, and let q_1^+ be the next score on the grid to the right of q_1. Compare the nonredundant menu $M = Q_0 \cup \{q_1\}$ (where q_1 is selected by a positive measure of types if $q_1 = q_0$) with the menu $M' = Q \cup \{q_1^+\}$, and let q_2 be the smallest element from Q_0 greater than q_1 (if there is none, then menu M is dominated e.g. by the full menu of all scores). By nonredundancy of M, $q_2 > q^*(\underline{\theta})$, and therefore M' is nonredundant as well. We have that $\underline{\theta} < \theta_{12} < \theta_{1+2}$, where θ_{1+2} is indifferent between q_1^+

[16]If $M = Q_0 \cup \{q_1, q_2, q_3\}$ is nonredundant, then so is $M' = Q_0 \cup \{q_1, q_3\}$

[17]$\overset{s}{=}$ stands for "has the same sign as". the actual expression for $\Delta(M, M', \phi \cdot \gamma_1)$ contains a re-normalization of the density that does not affect the sign.

and q_2. We get that the difference in welfare is

$$\Delta(M, M', \phi) = W(M', \phi) - W(M, \phi) = \int_{\underline{\theta}}^{\theta_{12}} (q_1^+ - q_1)\ \mathrm{d}\Phi(\theta) - \int_{\theta_{12}}^{\theta_{1+2}} (q_2 - q_1^+)\ \mathrm{d}\Phi(\theta).$$

There exists a function γ such that $\Delta(\phi \cdot \gamma) > 0$. Also if $\Delta(\phi) > 0$, then $\Delta(\phi \cdot \gamma) > 0$ if γ is decreasing.

Step 4: Scores above $q^*(\overline{\theta})$. Fix any $q_2 > q^*(\overline{\theta})$. Take any nonredundant menu M with $q_2 \in M$. Since M is nonredundant and $q_2 > q^*(\overline{\theta})$, q_2 has to be the highest score in the menu and the next highest score, call it q_1, has to satisfy $q_1 < q^*(\overline{\theta})$. Let q_1^+ be the next score on the grid to the right of q_1. Consider menu M' that one gets by adding q_1^+ to menu M. We get that the difference in welfare is

$$\Delta(M, M', \phi) = W(M', \phi) - W(M, \phi) = \int_{\theta_{11+}}^{\theta_{12}} (q_1^+ - q_1)\ \mathrm{d}\Phi(\theta) - \int_{\theta_{12}}^{\theta_{1+2}} (q_2 - q_1^+)\ \mathrm{d}\Phi(\theta).$$

There exists a function γ such that $\Delta(\phi \cdot \gamma) > 0$. Also if $\Delta(\phi) > 0$, then $\Delta(\phi \cdot \gamma) > 0$ if γ is decreasing.

Step 5. Now we are ready to wrap up the proof of Proposition 3.7. Because of the grid, there are finitely many menus.

- The case of decreasing ϕ: Take any nonredundant menu that is not of the threshold form. This means that either there is a gap in the menu, or the highest score is strictly smaller than $q^*(\overline{\theta})$. In the latter case the menu is dominated by one that adds all scores up to $q^*(\overline{\theta})$. In the former case, by steps 1-4 one can find a menu of the threshold form that gives a higher expected score under density $\phi \cdot \gamma$, where γ is decreasing. We can repeat this for any menu that is not of the threshold form, and since each time we take a decreasing transformation of the density, eventually all nonredundant menus that are not of the threshold form are suboptimal, and this remains true after a decreasing transformation of the density.

- The case of increasing ϕ: Take any nonredundant menu where three or more scores are selected. By steps 1 and 2, there exists a menu from which only two scores are selected that leads to a higher expected score under density $\phi \cdot \gamma$, where γ is increasing. We can repeat this for any menu from which three or more scores are selected, and since each time we take an increasing transformation of the density, eventually all nonredundant menus from which three or more scores are selected are suboptimal, and this remains true after a increasing transformation of the density.

- The case of single peaked ϕ: This is a combination of the two preceding cases. Let $\hat{\theta}$

be the modal type. It cannot be that a menu has more than two scores weakly smaller than $q^*(\hat{\theta})$ that are selected after sufficiently many single peaked transformations. Then if one looks at scores above $q^*(\hat{\theta})$ that are selected, it cannot be optimal for there to be a gap after sufficiently many single peaked transformations. ∎

Proof of Proposition 4.2 The outline of the proof is as follows. For a given equilibrium $(\tilde{H}, \tilde{e}, \tilde{w})$ we construct a new disclosure policy $\hat{H}$ and describe a wage schedule $\hat{w}$. We then show that given the disclosure policy and wage schedule, the agent would like to deviate to some $\hat{e} \geq \tilde{e}$. Finally, we establish that $(\hat{H}, \hat{e}, \hat{w})$ constitutes an equilibrium by showing that $\hat{w}$ can be derived from consistent beliefs.

Before delving into the proof we define some notation. Starting from any grading policy H with associated set of grades G, for any set of grades $G_1 \subset G$, let $H(G_1|q)$ be the probability that some $g \in G_1$ is realized conditional on score q. Also, $H(\cdot|q, G_1)$ is a distribution over grades that assigns zero probability to $g \notin G_1$, and probability $h(g|q)/H(G_1|q)$ to any $g \in G_1$.

Step 1: Take any (possibly random) grading policy $\tilde{H}$, with associated pure strategy equilibrium effort $\tilde{e}$ and wage schedule $\tilde{w}(g)$. This means that the smallest score that ever gets generated is $\underline{q}(\tilde{e})$. Now move to a disclosure policy that reveals all scores weakly above $\underline{q}(\tilde{e})$, and below $\underline{q}(\tilde{e})$ remains the same. Formally, let G_L be the set of grades that are induced with positive probability by some score $q < \underline{q}(\tilde{e})$, and are off the equilibrium path. For notational convenience, assume from now on that $G_L \cap Q = \emptyset$ (otherwise just relabel the grades in G_L). The new grading scheme has the set of grades $\hat{G} = G_L \cup Q$. The distribution over grades is

$$\hat{H}(\cdot|q) = \begin{cases} \tilde{H}(\cdot|q, G_L) & \text{if } q < \underline{q}(\tilde{e}) \text{ and } \tilde{H}(G_L|q) > 0, \\ \delta_q & \text{if } q \geq \underline{q}(\tilde{e}) \text{ or } q < \underline{q}(\tilde{e}) \text{ and } \tilde{H}(G_L|q) = 0. \end{cases}$$

Independently of the sequence of completely mixed strategies, the wage schedule satisfies $\hat{w}(q) = q$ if $q \geq \underline{q}(\tilde{e})$ or $q < \underline{q}(\tilde{e})$ and $\tilde{H}(G_L|q) = 0$ (any of these scores gets fully revealed). For any $g \in G_L$, let $Q(g)$ be the set of scores that assigns positive probability to g under $\tilde{H}$. While the sequence of fully mixed strategies $\hat{\sigma}_n$ will be specified later, we already state now that it satisfies $\hat{\sigma}_n(e) = \tilde{\sigma}_n(e)$ for $e < \tilde{e}$. We claim that $\hat{w}(g) \leq \tilde{w}(g)$ for $g \in G_L$, since

$$\begin{aligned} \tilde{w}(g) = & \ \lim_{n\to\infty} \ \frac{\sum_{q \in Q(g)} q\tilde{h}(g|q) f(q|\tilde{\sigma}_n)}{\sum_{q \in Q(g)} \tilde{h}(g|q) f(q|\tilde{\sigma}_n)}, \\ \hat{w}(g) = & \ \lim_{n\to\infty} \ \frac{\sum_{q \in Q(g), q < \underline{q}(\tilde{e})} q\tilde{h}(g|q) f(q|\tilde{\sigma}_n)}{\sum_{q \in Q(g), q < \underline{q}(\tilde{e})} \tilde{h}(g|q) f(q|\tilde{\sigma}_n)}. \end{aligned}$$

Step 2: Let $\tilde{t}(q)$ be the expected wage the agent can expect after q is realized under $(\tilde{H}, \tilde{w})$, that is $\tilde{t}(q) = \int \tilde{w}(g)\mathrm{d}\tilde{H}(g|q)$. Similarly, $\hat{t}(q)$ is the expected wage after q under $(\hat{H}, \hat{w})$. Next, define the payoff when facing schedule t as $u(e|t) = \int t(q)\mathrm{d}F(q|e) - c(e)$. We now claim that if we can show that

$$e \leq \tilde{e} \Rightarrow u(e|\hat{t}) \leq u(e|\tilde{t}), \text{with equality at } e = \tilde{e}, \tag{3}$$

then when moving from $\tilde{t}$ to $\hat{t}$ the agent wants to deviate weakly upwards. This is since $\tilde{e}$ maximizes $u(e|\tilde{t})$, so that no $e < \tilde{e}$ can maximize $u(e|\hat{t})$ if equation (3) is true.

To show equation (3), first note that $\hat{t}(q) = q$ for $q \geq \underline{q}(\tilde{e})$. Next, for any score with $q < \underline{q}(\tilde{e})$ and $\tilde{H}(G_L|q) = 0$, we have $\tilde{t}(q) \geq \underline{q}(\tilde{e}) > q = \hat{t}(q)$. This is true since $\tilde{H}(G_L|q) = 0$ implies that q only induces grades that are on-path under $\tilde{H}$, and for any such grade $\tilde{w}(g) \geq \underline{q}(\tilde{e})$. Finally, for any score with $q < \underline{q}(\tilde{e})$ and $\tilde{H}(G_L|q) > 0$, we have $\tilde{t}(q) \geq \hat{t}(q)$. This is true since

$$\begin{aligned} \tilde{t}(q) = &\ \textstyle\sum_{g\in G_L} \tilde{h}(g|q)\tilde{w}(g) + \sum_{g\notin G_L} \tilde{h}(g|q)\tilde{w}(g), \\ \hat{t}(q) = &\ \textstyle\sum_{g\in G_L} \tilde{h}(g|q)\hat{w}(g) / \sum_{g\in G_L} \tilde{h}(g|q), \end{aligned}$$

and the fact that for $g \in G_L$, $\tilde{w}(g) \geq \hat{w}(g)$ and for $g \notin G_L$, $\tilde{w}(g) \geq \underline{q}(\tilde{e}) \geq \hat{t}(q)$.

We can then write

$$\begin{aligned} u(e|\hat{t}) = &\sum_{q<\underline{q}(\tilde{e})} \hat{t}(q)f(q|e) + \Pr(q \geq \underline{q}(\tilde{e}))\mathbb{E}[q|e, q \geq \underline{q}(\tilde{e}), \tilde{e})] - c(e), \\ u(e|\tilde{t}) = &\sum_{q<\underline{q}(\tilde{e})} \tilde{t}(q)f(q|e) + \Pr(q \geq \underline{q}(\tilde{e}))\mathbb{E}[\tilde{t}(q)|e, q \geq \underline{q}(\tilde{e})] - c(e). \end{aligned}$$

Now for any q with $\underline{q}(\tilde{e}) \leq q \leq \overline{q}(\tilde{e})$, $\tilde{t}(q) = \mathbb{E}_{\tilde{H}}[\mathbb{E}[q|g, \tilde{e}]|q]$. So

$$\begin{aligned} \mathbb{E}[\tilde{t}(q)|e, q \geq \underline{q}(\tilde{e})] = &\mathbb{E}[\mathbb{E}_{\tilde{H}}[\mathbb{E}[q|g, \tilde{e}]|q]|e, q \geq \underline{q}(\tilde{e})] \\ = &\mathbb{E}[\mathbb{E}_{\tilde{H}}[\mathbb{E}[q|g, \tilde{e}, q \geq \underline{q}(\tilde{e})]|q]|e, q \geq \underline{q}(\tilde{e})]. \end{aligned}$$

The final equality is true since when taking the expectation over q conditional on $\tilde{e}$, the event $q \geq \underline{q}(\tilde{e})$ is redundant. By the law of iterated expectations,

$$\begin{aligned} \mathbb{E}[q|e, q \geq \underline{q}(\tilde{e})] = &\mathbb{E}_{\tilde{H}}[\mathbb{E}[q|g, e, q \geq \underline{q}(\tilde{e})]|e, q \geq \underline{q}(\tilde{e})] \\ = &\mathbb{E}[\mathbb{E}_{\tilde{H}}[\mathbb{E}[q|g, e, q \geq \underline{q}(\tilde{e})]|q]|e, q \geq \underline{q}(\tilde{e})]. \end{aligned}$$

By the MLR property, since $e \leq \tilde{e}$ we have $\mathbb{E}[q|g, e, q \geq \underline{q}(\tilde{e})] \leq \mathbb{E}[q|g, \tilde{e}, q \geq \underline{q}(\tilde{e})]$. Therefore

$u(e|\hat{t}) \leq u(e|\tilde{t})$. When $e = \tilde{e}$, $F(\underline{q}(\tilde{e})|e) = 0$ so that $u(e|\hat{t}) = u(e|\tilde{t})$.

Step 3: Let $\hat{e} \geq \tilde{e}$ be the maximizer of $u(e|\hat{t})$, which exists by finiteness. The new equilibrium is $(\hat{e}, \hat{H}, \hat{w})$, and it only remains to show that $\hat{w}$ can be derived from consistent beliefs. Under $\hat{H}$, any score is either fully revealed or mapped into a grade in G_L. For any score that is fully revealed, $\hat{H}(\cdot|q) = \delta_q$ and the unique wage derived from consistent beliefs is $w(q) = q$. A grade in G_L can only be generated by scores $q < \underline{q}(\tilde{e}) \leq \underline{q}(\hat{e})$, and is therefore off-path. Any sequence of completely mixed strategies that satisfies $\hat{\sigma}_n(e) = \tilde{\sigma}_n(e)$ for $e < \tilde{e}$ works for our conclusion.

If under full revelation the effort e_{FR} is such that $\underline{q}(e_{FR}) > \underline{q}$, then a grading scheme that pools all scores $q < \underline{q}(e_{FR})$ and fully reveals all scores $q \geq \underline{q}(e_{FR})$ also implements e_{FR}. A completely mixed strategy that features $\hat{\sigma}_n(\underline{e}) = \frac{1}{n}$ and $\hat{\sigma}_n(e) = \frac{1}{n^2}$ for any e with $\underline{e} < e < e_{\tilde{FR}}$ makes downwards deviations less attractive, as low scores are supposed to be generated by $\underline{e}$, minimizing the wage corresponding to the pooling grade. One can then try to implement higher effort levels by raising the threshold above which scores are fully revealed. While such a type of scheme is not necessarily the optimal one, it always does weakly and sometimes strictly better than full revelation. In the special case where $e = \underline{e} \Rightarrow q = \underline{q}$, this provides the strongest punishment for downwards deviations among all grading schemes. ∎

B Examples for Section 3.3

A normal example. The purpose of this example is to illustrate the idea that introducing randomness is detrimental for incentives. To exploit the tractability of the normal distribution we assume here that $Q = \mathbb{R}$. Types are normally distributed on the real line. The agent's utility is as in the quadratic example in 3.2.2, quasilinear in the wage with a quadratic cost function. We restrict attention to grading policies such that the distribution of the score and grade is jointly normal. This is equivalent (up to a normalization on the covariance of q and g) to the collection of policies of the form

$$g_\sigma(q) = q + \sigma\epsilon$$

for $\sigma \in \mathbb{R}$, where ϵ is standard normal and independent of q. Each choice of σ leads to a score schedule $q_\sigma(\theta)$, which the principal wishes to maximize in expectation. Fix σ. By joint normality, the market wage is $w(g) = \mathbb{E}[q_\sigma|g] = \alpha(\sigma) + \beta(\sigma)g$, where $\beta(\sigma) = \frac{\text{cov}(q_\sigma, g_\sigma)}{\text{var}(g_\sigma)} = \frac{\text{var}(q_\sigma)}{\text{var}(q_\sigma)+\sigma^2}$. The expected wage from score q is $\mathbb{E}[w|q] = \alpha(\sigma) + \beta(\sigma)\mathbb{E}[g|q] = \alpha(\sigma) + \beta(\sigma)q$.

Maximizing $\mathbb{E}[w|q] - c(q,\theta)$ over q gives the schedule of scores:

$$q_\sigma(\theta) = \frac{1}{b}\left[\beta(\sigma) + a\theta - c\right].$$

This shows that $\text{var}(q_\sigma)$ does not depend on σ. Therefore setting $\sigma = 0$ maximizes $\beta(\sigma)$ and, since $b > 0$, maximizes $q_\sigma(\theta)$ pointwise and a fortiori in expectation. We conclude that, in the class of policies for which score and grade are jointly normally distributed, a noise-free grading policy is optimal.

A discrete example. Now we present an environment in which a random grading scheme outperforms the best (deterministic) grading function. We use a discrete set of types and a continuous set of scores, but one can approximate our example with a fine grid of scores and continuous type distribution to get an example in line with our assumptions and arrive at the same conclusion.

The set of scores is $Q = \mathbb{R}_+$, and the set of types is $\Theta = \{\theta_1, \theta_2\}$ where type θ_i has mass $\phi_i \in (0,1)$ and $\theta_1 < \theta_2$. Utility is quasilinear: $u(w,q,\theta) = w - c(q,\theta)$, with $c \in C^2$ strictly convex in q, $-c$ strictly supermodular, $c(0,\theta) = 0$, and $c_q(0,\theta) < 1 < c_q(+\infty,\theta)$. We assume that the two types are sufficiently far apart in that $q^*(\theta_2) > \tilde{q}_1$ where $\tilde{q}_1$ is the highest score that θ_1 weakly prefers to $\underline{q} = 0$.

It is straightforward to show that an optimal grading function induces the following behavior. There are two scores $q_1 < q_2$ such that all θ_i choose q_i, θ_1 is indifferent between q_1 and the score $\underline{q} = 0$, and θ_2 is indifferent between q_1 and q_2. Since the grading function is deterministic, the wage for score q_i is equal to q_i. The indifference properties imply that $q_1 - c(q_1,\theta_1) = 0$ and $q_2 - c(q_2,\theta_2) = q_1 - c(q_1,\theta_2)$. These equations admit two solutions. One has the low type excluded ($q_2 > q_1 = 0$), and the other includes both types ($q_2 > q_1 > 0$). There exists $\overline{\phi_1} < 1$ such that for $\phi_1 > \overline{\phi_1}$ the solution with $q_2 > q_1 > 0$ is strictly optimal; assume $\phi_1 > \overline{\phi_1}$ henceforth. One grading function that implements this has set of grades $G = \{0, g_1, g_2\}$ with $g(q) = 0$ for $q \notin \{q_1, q_2\}$ and $g(q_i) = g_i$ for $i = 1, 2$.

Now consider the following condition:

$$c_q(q_2,\theta_2) > c_q(q_1,\theta_1). \tag{$*$}$$

One way to interpret this is that the marginal rate of substitution $-\frac{u_q}{u_w}$ at a type's equilibrium score is increasing in type. We show that if $(*)$ holds, there exists $\overline{\overline{\phi_1}} < 1$ such that if $\phi_1 > \overline{\overline{\phi_1}}$, then the principal strictly benefits from randomization.[18] In what follows q_1, q_2 refer to the

[17] $q_1 < q_2$ follows from the assumption that types are sufficiently far apart.

[18] $\overline{\overline{\phi_1}}$ and $\overline{\phi_1}$ are independent.

quantities from the optimal (deterministic) grading function.

Consider a random grading scheme that implements the following. There are two scores $\hat{q}_1, \hat{q}_2$ with $0 < \hat{q}_1 < \hat{q}_2$. The expected wage from choosing $\hat{q}_i$ is W_i and 0 otherwise, and all θ_i choose $\hat{q}_i$ in equilibrium. Given the wages, the scores are such that θ_1 is indifferent between $\hat{q}_1$ and $\underline{q} = 0$, and θ_2 is indifferent between $\hat{q}_2$ and $\hat{q}_1$. We claim there exists $\overline{\epsilon} > 0$ such that the principal can induce any $W_1, W_2, \hat{q}_1, \hat{q}_2$ such that $W_1 \in [q_1, q_1 + \overline{\epsilon}]$, $\phi_1[W_1 - \hat{q}_1] + \phi_2[W_2 - \hat{q}_2] = 0$, and the conditions above hold. The grading scheme randomizes over two "informative" grades with score-dependent probabilities for scores on the path; off the path it releases a "null" grade.[19]

Now consider the following family of reduced form perturbations starting from the optimal deterministic grading function. For $\epsilon \geq 0$ small, increase W_1 from q_1 to $W_1(\epsilon) \equiv q_1 + \epsilon$, and then adjust the scores and the high wage to maintain equalities in the indifference and wage equations. Let $\hat{q}_i(\epsilon) > 0$ and $W_2(\epsilon)$ be the positive solutions:

$$\begin{aligned} W_1(\epsilon) - c(q_1(\epsilon), \theta_1) &= 0 \\ W_2(\epsilon) - c(q_2(\epsilon), \theta_2) &= W_1(\epsilon) - c(q_1(\epsilon), \theta_2) \\ \phi_1[W_1(\epsilon) - q_1(\epsilon)] + \phi_2[W_2(\epsilon) - q_2(\epsilon)] &= 0. \end{aligned}$$

The principal's payoff from the perturbed grading scheme is $\phi_1 q_1(\epsilon) + \phi_2 q_2(\epsilon)$. The deterministic optimum corresponds to $\epsilon = 0$ and the implicit function theorem applies, so to show that randomization improves the principal's payoff it suffices to show that $\phi_1 q_1'(0) + \phi_2 q_2'(0) > 0$.[20] By differentiating the displayed system of equations and manipulating, one can show that the inequality holds if and only if

$$\phi_1 > \frac{c_q(q_1, \theta_1) - c_q(q_1, \theta_2)}{c_q(q_2, \theta_2) - c_q(q_1, \theta_2)} \equiv \overline{\overline{\phi_1}}.$$

The assumptions on c imply the numerator and denominator are positive, so $(*)$ implies $\overline{\overline{\phi_1}} < 1$. Therefore for probabilities ϕ_1 above the bound, randomization benefits the principal.

[19] Consider the following collection of grading schemes parametrized by $d \in [0, 1]$. The set of grades is $G = \{0, g_1, g_2\}$. Let $h(g|q)$ be the probability of g given q. Put $h(g_1|\hat{q}_1) = 1 - \phi_2 d$ and $h(g_2|\hat{q}_1) = \phi_2 d$, and $h(g_2|\hat{q}_2) = 1 - \phi_1 d$ and $h(g_1|\hat{q}_2) = \phi_1 d$. For $q \notin \{\hat{q}_1, \hat{q}_2\}$, put $h(0|q) = 1$. Since $d, \phi_i \in [0, 1]$, the function $h(\cdot|q)$ is a well-defined probability on G for all q. When $d = 0$, h is equivalent to a deterministic grading function. One can verify that if the indifference conditions hold then for any d the wage equation holds. Then by using the implicit function theorem one can show that for d sufficiently small the indifference conditions also hold, with scores $0 < \hat{q}_1 < \hat{q}_2$. Finally each d pairs to a corresponding ϵ; $d = 0$ to $\epsilon = 0$, and as d increases from 0 the ϵ increases.

[20] Finding a condition such that the inequality goes the other way does not imply that the principal can improve from randomization. This is because randomization can only increase the wage of the low type relative to the deterministic optimum. In other words $\epsilon < 0$ is not possible.

It is interesting to note that the perturbation results in a Pareto improvement: type θ_1 obtains a payoff of 0 in either case, but type θ_2 strictly benefits. The reasoning is simple. The high type is always indifferent to imitating the low type. The perturbation increases the low wage/score to keep the low utility constant, which by strict supermodularity implies the high type strictly prefers the larger values.

LaTeX Souce Files

Click here to download LaTeX Souce Files: Latex source files.zip