

Randomized exams: leveraging uncertainty to improve student learning

Nurassyl Toltay
Nazarbayev University

May 2025

Acknowledgements

I wish to express my sincere gratitude to my thesis advisor, Alejandro Melo Ponce, for his guidance, constructive feedback, and support throughout this research. His consistent encouragement and enthusiasm kept my research spark alive throughout my Master's program. Being not just my supervisor but also a life advisor, I will be forever grateful to him.

Additionally, I am thankful to Vladyslav Nora for his unexpected yet extremely valuable contributions to my thesis. His harsh but insightful comments consistently pushed me to improve my work.

I thank my family and friends for their jokes about how I am about to ruin students' lives, as they made this research journey an unexpectedly fun experience.

Lastly, I express my love and gratitude to my girlfriend, Dariya Karaguzhina, for her invaluable support and patience, which proved so helpful throughout this process.

All errors are, of course, my own.

Contents

1	Introduction	4
2	Related Literature	5
3	Model Framework	7
3.1	Main Problem	7
3.2	Example	11
3.3	Main Result	12
4	Conclusion	15
5	Appendix: Proofs	17
5.1	Proof of Proposition 1	17
5.2	Proof of Proposition 2	18
5.3	Proof of Proposition 3	20

Abstract

This thesis explores the impact of randomization of exam timing on student effort in a dynamic principal-agent model. The model compares between deterministic exams, which have predetermined dates, and randomized exams, where students are uncertain about when they will be tested. Based on standard assumptions regarding the shape of students' cost functions, the analysis finds that randomized tests result in significantly more effort than deterministic tests. This outcome arises from the need for the student to maintain some preparedness under uncertainty. Additionally, the study outlines optimal grading functions, demonstrating that a simple pass/fail exam successfully leverages the benefits of random timing, while a two-threshold method can optimize performance by encouraging effort more flexibly throughout the periods. The findings presented contribute to the existing literature on monitoring, contract design, and information asymmetry by demonstrating that uncertainty over timing - when combined with well-designed grading standards - can be used as a means to improve agent performance in an educational environment.

1 Introduction

This paper develops a theoretical framework to examine the role of randomized examination timing in shaping student incentives. Traditional assessment structures typically rely on fixed exam dates announced well in advance. Such setups encourage students to concentrate effort close to the exam, a pattern that may be suboptimal from a learning perspective, particularly if consistent effort throughout the academic term is more desirable [8]. Introducing uncertainty into the timing of assessment may shift this behavior, incentivizing students to distribute effort more evenly across periods.

To study this idea formally, we construct a discrete-time principal-agent model, where a professor (the principal) designs a grading policy, and a student (the agent) responds by choosing effort over two periods. The student, aware of the policy, decides how much effort to allocate in the first period and then observes whether the exam has already taken place before choosing effort in the second period. The student incurs convex costs for effort, while the professor derives payoff from the total effort exerted across periods.

The model is deliberately simple, as it is used to deliver clean and transparent results about how randomization affects incentives. The simplicity

allows for sharper intuition about the effects of uncertainty in evaluation timing without unnecessary institutional detail.

For a proper analysis, we introduce two regimes: a deterministic setting, where the student knows exactly when the exam will occur, and a randomized setting, where the exam occurs in either period with known probability α . We show that under general conditions, randomized observation strictly dominates deterministic timing in terms of total effort induced from the student. This finding hinges on the student’s strategic response to uncertainty: when the exam may occur in either period, the student finds it optimal to exert meaningful effort in both periods, as opposed to concentrating it in just one.

A key result of the paper is that a simple pass/fail grading rule - where the student passes if effort exceeds a threshold - suffices to achieve this improvement under randomization. However, further gains are possible. By introducing a two-threshold grading function that rewards different levels of effort in different periods, the professor can better align incentives with the timing structure, extracting even higher total effort.

Overall, this study contributes to the theoretical literature on educational incentives, or, in broader context, to the literature on incentive design, by identifying conditions under which randomized exam timing, combined with appropriately chosen grading rules, can be an effective tool to increase student effort. While our model abstracts from many institutional features of real-world courses, it offers insight into how uncertainty and information design can be leveraged to improve learning outcomes through incentive alignment.

2 Related Literature

The analysis builds on foundational results in contract theory, reviewed in Georgiadis [5], while introducing a new approach in principal-agent modeling: using uncertainty about exam timing as a tool to influence agent behavior.

Uncertainty has been studied before in a related context. Dye analyzes a setting where the principal chooses when to use a costly monitoring device, which gives additional information beyond observing output [3]. Dye finds that deterministic monitoring is optimal compared to the randomized. One reason why our results differ is that we assume there is no background signal of performance - no output is observed unless the exam (monitoring) takes

place. This creates a different situation both for the principal and the agent.

A related model is developed by Georgiadis and Szentes, who study a situation where the principal also sees no output unless a costly monitoring action is taken [6]. However, their model is static: the timing of monitoring is not a choice variable, so they cannot study whether randomized observation can help. Still, their setup is important because it shares the same limited information structure used in our model.

Other work has looked at how feedback and beliefs affect effort. Ely, Georgiadis and Rayo show that a principal can sometimes increase effort by withholding performance feedback [4]. In their model, feedback is only given when the agent reaches a certain threshold, and the prior uncertainty keeps the agent working harder. While our model does not involve thresholds or feedback rules, both models rely on the idea that limiting what the agent knows can improve outcomes.

A more recent and especially relevant paper is by Castiglioni, Marchesi, and Gatti [1]. They consider a setting, where the agent has private information about their type, which affects effort cost and success probability. The principal must design a menu of contracts to get the agent to reveal their type. They show that adding randomization - specifically, letting contracts be lotteries over payments - makes the problem easier to solve and can lead to better outcomes. While their model focuses on randomizing payments and ours focuses on randomizing exam timing, both use randomness as a way to align incentives when the principal has limited tools. Another key difference is that their setting is static and involves private types, while ours is dynamic and assumes the cost structure is known. Despite these differences, their paper shows how useful randomization can be in complex incentive problems.

Harbaugh and Rasmusen demonstrate that coarse grading schemes (such as pass/fail or honors/no honors) can outperform finer, fully informative systems in motivating behavior [7]. They argue that reducing informational precision can increase effort by mitigating strategic avoidance or gaming, particularly in environments where disclosure affects downstream participation. While their model centers on signaling and certification to outside observers, our findings point to a similar conclusion: that simplicity in grading (e.g., pass/fail or two-threshold rules) can be an effective tool for raising effort. In our model, these coarse schemes interact with uncertainty over exam timing to incentivize early and sustained preparation. Thus, despite different institutional contexts and mechanisms, both papers highlight how reducing informational granularity can produce higher performance through

behavioral channels.

Dubey and Geanakoplos provide another rationale for coarse grading, showing that broader grade categories can amplify effort incentives in environments with heterogeneous student abilities and social comparison [2]. By clustering students into fewer grading bins, professors can intensify competition for top designations, which in turn motivates stronger performance. While our model abstracts from peer dynamics or status concerns, the insight that effort responds non-monotonically to the fineness of evaluation carries over. In particular, our findings suggest that relatively simple, non-continuous grading functions - combined with timing uncertainty - are not only analytically tractable but also incentive-compatible. Both papers, therefore, support the idea that coarse evaluations, when paired with uncertainty or strategic ambiguity, may be more effective than elaborate point-based grading systems in motivating agents.

Together, these papers help explain the modeling choices in this thesis. Each of them explores important ideas: how monitoring works, when timing matters, how feedback shapes behavior, and how randomization can improve design. Each one helps motivate the framework used in our model of randomized exams and student effort.

3 Model Framework

3.1 Main Problem

This section introduces the formal model used to analyze the relationship between randomized examinations and student effort. The model is structured as a discrete-time principal-agent problem involving two agents:

- The *student*, who selects effort levels to maximize utility derived from the grading outcome, subject to the disutility of exerting effort.
- The *professor*, who aims to maximize total effort exerted by the student, treated as a proxy for learning outcomes.

The interaction unfolds over two periods. The sequence of events proceeds as follows:

1. In period 0, the professor selects the grading function $g : \mathbb{R}_+ \rightarrow [0, 1]$ and the binary random variable D where

$$D(\alpha) = \begin{cases} 1 & \text{with probability } \alpha, \\ 0 & \text{with probability } 1 - \alpha, \end{cases}$$

which represents the professor's choice of an exam path: $D = 1$ corresponds to holding the exam in period 1, while $D = 0$ places the exam in period 2. The professor formally selects the distribution $D(\alpha)$, that is, the probability α with which $D = 1$ occurs. While this ultimately amounts to choosing α , specifying the random variable D explicitly validates the use of expectations over its realizations in both the professor's and student's objective functions. For ease of notation, when not needed, I remove the dependence of D on its "parameter" α . Thus, the professor's strategy is described as a *pair* (D, g) .

2. In period 1, the student chooses an effort level $e_1 \in \mathbb{R}_+$, observing the grading function and examination probability announced by the professor (as the realization of D is not yet announced). The student incurs effort cost $c(e_1)$ where $c : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is assumed to be an increasing and convex function, which is separable across periods, with cost in each period depending solely on the effort exerted in that period.
3. At the end of period 1, the student observes whether an examination has occurred, or formally, observes the realization of D .
4. In period 2, conditional on the information observed, the student chooses a second-period effort level e_2 , which will become a random variable, whose realization is observed by the student and also the professor. I will thus introduce two possibilities for period 2 effort:

$$e_2(D = 1) = e_2^I \in \mathbb{R}_+, \quad e_2(D = 0) = e_2^{II} \in \mathbb{R}_+.$$

Thus, the student's strategy is a *triple*: $e = (e_1, e_2^I, e_2^{II})$.

The appropriate equilibrium concept for this dynamic interaction is the Subgame Perfect Nash Equilibrium (SPNE) which is a strategy profile

$$((D^*, g^*), (e_1^*, e_2^{I*}, e_2^{II*}))$$

such that:

1. Given the professor's announced (D^*, g^*) , the student's strategy $(e_1^*, e_2^I, e_2^{II*})$ maximizes the student's expected utility,

$$(e_1^*, e_2^I, e_2^{II*}) \in \arg \max_e \mathbb{E} [\pi_S((D^*, g^*), (e_1, e_2^I, e_2^{II*}))].$$

2. Anticipating the student's best response, the professor chooses (D^*, g^*) to maximize their expected payoff,

$$(D^*, g^*) \in \arg \max_{(D, g)} \mathbb{E} [\pi_P((D, g), (e_1^*, e_2^I, e_2^{II*}))].$$

The professor's payoff is then given by:

$$\pi_P((D, g), e) = \begin{cases} e_1 + e_2^I & \text{if } D = 1, \\ e_1 + e_2^{II} & \text{if } D = 0. \end{cases}$$

Similarly, the student's payoff is specified as:

$$\pi_S((D, g), e) = \begin{cases} g(e_1) - c(e_1) - c(e_2^I) & \text{if } D = 1, \\ g(e_2^{II}) - c(e_1) - c(e_2^{II}) & \text{if } D = 0. \end{cases}$$

The objective of each agent is to maximize expected payoffs, but with the definitions we introduced above, now it is clear that the expectations are taken with respect to the probability distribution of the random variable D .

The functional structure of the agents' payoffs is supported by several modeling assumptions, some of which were previously introduced. In particular, we assume that the student's cost of effort is increasing and convex, and that costs are period-specific, that is, the cost incurred in one period depends solely on the effort exerted in that period. To further justify the payoff formulation, we now introduce additional assumptions. First, we assume perfect observability: the professor accurately observes the student's effort whenever an exam takes place. Second, we impose a commitment assumption: once the professor announces the grading function and the exam probability, these choices are binding. Third, payoffs are temporally separable for both agents. The professor simply values the sum of efforts across the two periods with no additional weighting, while the student's cost is additively separable over time. We also assume that the cost function is common knowledge, known by both players at the outset. Finally, we treat effort as disposable - useful only in the period in which it is exerted. This reflects

a stylized student who is forgetful (benefiting from no carryover of learning across periods) and primarily motivated by grades rather than intrinsic learning. While these assumptions serve to isolate the mechanisms of interest and deliver clean theoretical results, relaxing some of them - such as allowing for noisy observability of effort or heterogeneity of cost functions - may yield interesting directions for future research.

The model compares two observation regimes:

- **Deterministic observation:** The examination date is known ex ante to the student. Formally, $\alpha \in \{0, 1\}$, resulting in the following optimization problem:

$$\begin{aligned}
 V_{det} = \max_{\alpha \in \{0,1\}, g} & \quad \mathbb{E} [(e_1 + e_2^I) \cdot D(\alpha) + (e_1 + e_2^{II}) \cdot (1 - D(\alpha))] \\
 \text{s.t. } e \in \arg \max_e & \quad \mathbb{E} \left[[g(e_1) - c(e_1) - c(e_2^I)] \cdot D(\alpha) \right. \\
 & \quad \left. + [g(e_2^{II}) - c(e_1) - c(e_2^{II})] \cdot (1 - D(\alpha)) \right] \quad (\text{IC})
 \end{aligned}$$

where V_{det} is the optimal value of the problem.

- **Randomized observation:** The examination date is uncertain, $\alpha \in (0, 1)$, resulting in:

$$\begin{aligned}
 V_{rand} = \max_{\alpha \in (0,1), g} & \quad \mathbb{E} [(e_1 + e_2^I) \cdot D(\alpha) + (e_1 + e_2^{II}) \cdot (1 - D(\alpha))] \\
 \text{s.t. } e \in \arg \max_e & \quad \mathbb{E} \left[[g(e_1) - c(e_1) - c(e_2^I)] \cdot D(\alpha) \right. \\
 & \quad \left. + [g(e_2^{II}) - c(e_1) - c(e_2^{II})] \cdot (1 - D(\alpha)) \right] \quad (\text{IC})
 \end{aligned}$$

where V_{rand} is the optimal value of the problem.

As you can see, the difference in the formal specification of the two problems lies only in the difference of the support of α .

3.2 Example

To illustrate the mechanics of the model, consider a specific functional form of the cost and grading functions. Let $c(e) = e^3/3$ and $g(e) = \min\{e, 1\}$. Under these assumptions, the professor's problem simplifies to:

$$\begin{aligned} \max_{\alpha} \quad & \mathbb{E}[e_1 + e_2] \\ \text{s.t.} \quad & e \in \arg \max_e \mathbb{E} \left[\min\{e_1, 1\} - e_1^3/3 - (e_2^I)^3/3 \cdot D \right. \\ & \left. + [\min\{e_2^{II}, 1\} - e_1^3/3 - (e_2^{II})^3/3] \cdot (1 - D) \right] \quad (\text{IC}). \end{aligned}$$

Lemma 1 *Under deterministic observation, the student's effort is characterized by:*

$$e = \begin{cases} \arg \max [g(e) - c(e)], & \text{if observed,} \\ 0, & \text{otherwise.} \end{cases}$$

The logic is immediate: in the unobserved period, effort yields no benefit and is therefore zero.

Under deterministic observation, our problem yields $V = 1$.

Notice that in the randomized observation regime, the indicator function is only realized after the first period ends. Hence, the student has to take the expectation of the exam in the first period, reducing the problem in the first period to:

$$\max_{e_1} \left[\alpha \min\{e_1, 1\} - \frac{e_1^3}{3} \right],$$

with solution:

$$e_1^* = \alpha^{1/2}.$$

In period 2, the student observes whether an exam has occurred and chooses effort accordingly:

$$e_2^* = \begin{cases} 0, & \text{with probability } \alpha, \\ 1, & \text{with probability } 1 - \alpha. \end{cases}$$

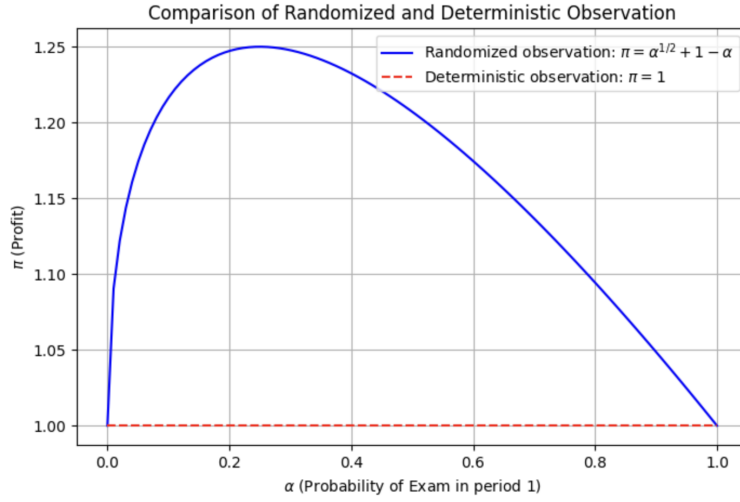
The resulting values are:

- $V = 1$ under deterministic observation.
- $V = \alpha^{1/2} + (1 - \alpha)$ under randomized observation.

It follows immediately that $V_{rand} > 1$ for all $\alpha \in (0, 1)$. The optimal examination probability is:

$$\alpha^* = 0.25.$$

Similarly, you can see the dominance of the randomized observation graphically:



Additionally, consider an alternative specification of the cost function, $c(e_t) = e_t^2/2$. In this case, the professor's payoff under both deterministic and randomized observation is identical:

$$V_{det} = 1 = V_{rand}.$$

This illustrates that the advantage of randomized observation depends on the specific properties of the cost function.

3.3 Main Result

To generalize the analysis, we consider broader functional forms of the cost and grading functions, going back to the original unspecified formulation of the model:

$$\begin{aligned}
& \max_{\alpha \in (0,1), g} \mathbb{E}[e_1 + e_2] \\
& \text{s.t. } e \in \arg \max_e \mathbb{E} \left[[g(e_1) - c(e_1) - c(e_2^I)] \cdot D \right. \\
& \quad \left. + [g(e_2^{II}) - c(e_1) - c(e_2^{II})] \cdot (1 - D) \right] \quad (\text{IC})
\end{aligned}$$

Since the professor selects the grading function, but takes the student's cost structure as given, it is natural to analyze which grading functions can optimally induce effort.

We begin with a benchmark result under deterministic observation:

Lemma 2 *In a deterministic examination setting, for any increasing cost function $c(\cdot)$, there exists a grading function $g(\cdot)$ that can induce the highest possible effort level $c^{-1}(1)$. This function takes the form:*

$$g(e) = \begin{cases} 1, & e \geq c^{-1}(1), \\ 0, & \text{otherwise.} \end{cases}$$

Effort beyond $c^{-1}(1)$ leads to negative utility for the student, whose optimal strategy would be to put the effort level of 0. Hence, higher effort levels are not inducible.

In the randomized observation regime, the student's optimization problem in period 1 is:

$$\max_{e_1} g(e_1) \cdot \alpha - c(e_1),$$

while effort in period 2 follows Lemma 1.

The following proposition establishes the central result:

Proposition 1 *Let \bar{V}_{det} denote the maximum value achievable under a deterministic grading scheme, given a cost function. Then, for any convex and increasing cost function $c(\cdot)$, there exists a grading function $g(\cdot)$ that induces $V' > \bar{V}_{det}$. This function is given by:*

$$g_{p/f}(e) = \begin{cases} 1, & e \geq c^{-1}(\alpha^*), \\ 0, & \text{otherwise} \end{cases},$$

where

$$\alpha^* = \operatorname{argmax} \mathbb{E} [\pi_P((D(\alpha), g_{p/f}), e)]$$

The proof, provided in the appendix, demonstrates that randomization strictly improves effort relative to deterministic observation, for any increasing and convex cost function.

This insight motivates consideration of more flexible grading schemes. In particular, the pass/fail structure imposes the same threshold across periods, which may limit second-period effort. A refinement involves introducing distinct thresholds for each period.

Proposition 2 *Given any convex and increasing cost function of the student $c(\cdot)$, there exists a grading function $g(\cdot)$ that achieves the highest possible optimal value \bar{V} . The function is given by:*

$$g_{t/t}(e) = \begin{cases} 1, & e \geq c^{-1}(1 + \alpha^*\theta^* - \theta^*), \\ \theta^*, & e \geq c^{-1}(\alpha^*\theta^*), \\ 0, & \text{otherwise.} \end{cases}$$

where

$$(\alpha^*, \theta^*) \in \arg \max \quad \mathbb{E} [\pi_P((D(\alpha), g_{t/t}), e)]$$

yielding maximum optimal value:

$$\bar{V} = c^{-1}(\alpha^*\theta^*) + (1 - \alpha^*)c^{-1}(1 + \alpha^*\theta^* - \theta^*)$$

(Proof moved to appendix.)

This two-threshold grading scheme enables the professor to fine-tune incentives across periods, maximizing cumulative effort.

Hence, the model I've created produces SPNE:

$$((D^*, g^*), (e_1^*, e_2^{I^*}, e_2^{II^*})) = ((D(\alpha^*), g_{t/t}), (c^{-1}(\alpha^*\theta^*), 0, c^{-1}(1 + \alpha^*\theta^* - \theta^*)))$$

where

$$(\alpha^*, \theta^*) \in \arg \max \quad \mathbb{E} [\pi_P((D(\alpha), g_{t/t}), e)]$$

Finally, we establish a specification under which the professor strictly prefers $g_{t/t}$ and where they are indifferent between the two functions:

Proposition 3 *Given any convex and increasing cost function of the student $c(\cdot)$, the professor should utilize the following grading function:*

$$g^*(e) = \begin{cases} g_{t/t}, & \alpha^* \leq \frac{3-\sqrt{5}}{2}, \\ g_{t/t} \text{ or } g_{p/f}, & \alpha^* > \frac{3-\sqrt{5}}{2}, \end{cases}$$

where

$$\alpha^* = \arg \max \mathbb{E} [\pi_P((D(\alpha), g_{p/f}), e)]$$

The intuition behind this result lies in how the two grading schemes allocate incentives across periods. Under the pass/fail structure, the equilibrium effort in the first period tends to be higher, while the effort in the second period is comparatively lower than under the two-threshold alternative. When the examination is likely to occur early - that is, when α is sufficiently large - the principal benefits more from the elevated first-period effort induced by the pass/fail rule. In such cases, the gain from better aligning incentives in the second period under the two-threshold function does not outweigh the simplicity and performance of the pass/fail design. Although the two-threshold rule weakly dominates in terms of V_{rand} , the pass/fail function can, in some instances, replicate its performance when the optimal effort is concentrated earlier in time. This observation clarifies the boundary at which a simpler grading scheme becomes sufficient for maximizing total student effort.

This proposition might be useful for future research when cost of exam complexity is introduced, so that the professor optimally chooses the simpler exam when presented with a choice between two exams producing similar expected total effort.

4 Conclusion

This paper studies the effect of randomizing examination dates in a discrete-time principal-agent framework. Introducing uncertainty into the timing of observation alters students' incentives, leading to higher cumulative effort under appropriate conditions. The model shows that, relative to deterministic exams, randomized examinations generate superior outcomes in terms of total student effort.

The paper also highlights the role of grading function design. While a simple pass/fail grading structure is sufficient to improve effort under randomization, further gains are possible through the implementation of a two-threshold grading function. This structure allows the principal to provide different incentives across periods, yielding higher total effort.

The analysis identifies conditions under which the principal prefers the pass/fail scheme to the two-threshold scheme, balancing complexity and effectiveness.

Several avenues for future research are suggested. Extending the model to T periods would allow for a richer dynamic analysis of effort allocation over time. Incorporating heterogeneous student populations with varying cost functions would yield additional insights into the optimal design of grading schemes and examination timing, while also allowing for potential screening analysis for the professor.

5 Appendix: Proofs

5.1 Proof of Proposition 1

The student's optimization problem in the two-period model requires solving the following problem in period 1:

$$\max_{e_1} g(e_1) \cdot \alpha - c_1(e_1).$$

To ensure participation while maximizing effort, the professor sets a threshold:

$$\bar{e} = c^{-1}(\alpha).$$

Since the exam remains identical across periods, the total expected effort is given by:

$$V_{rand} = c^{-1}(\alpha)(2 - \alpha).$$

Comparing this to the deterministic case, where it is $V_{det} = c^{-1}(1)$, we analyze whether randomized observation yields greater effort:

$$c^{-1}(\alpha)(2 - \alpha) \stackrel{\geq}{\leq} c^{-1}(1).$$

Rearranging:

$$f(\alpha) = c^{-1}(\alpha)(2 - \alpha) - c^{-1}(1).$$

If $f(\alpha) > 0$ for some $\alpha \in (0, 1)$, then randomized observation induces greater effort. Since $f(1) = 0$, we examine the slope at $\alpha = 1$:

$$f'(\alpha) = c^{-1'}(\alpha)(2 - \alpha) - c^{-1}(\alpha).$$

Evaluating at $\alpha = 1$:

$$f'(1) = c^{-1'}(1) - c^{-1}(1).$$

Rewriting in terms of the cost function:

$$\frac{1}{c'(c^{-1}(1))} - c^{-1}(1) \stackrel{\geq}{\leq} 0.$$

Rearranging:

$$\frac{1}{c^{-1}(1)} \stackrel{\geq}{\leq} c'(c^{-1}(1))$$

Given that $\frac{1}{c^{-1}(1)}$ represents the average slope over $e \in [0, c^{-1}(1)]$ and $c'(c^{-1}(1))$ is the endpoint slope, convexity implies $c'(c^{-1}(1))$ is greater. Thus, randomized exams induce greater effort than deterministic ones.

5.2 Proof of Proposition 2

Formally, we define:

$$g(e) = \begin{cases} 1, & e \geq \gamma_2, \\ \theta, & e \geq \gamma_1, \\ 0, & \text{otherwise.} \end{cases}$$

where $\theta < 1$.

The objective is to determine whether this structure can induce greater total effort than the single-threshold pass/fail system and, if so, whether it represents the optimal grading function.

Period 1: Student Optimization In the first period, the student faces uncertainty regarding whether they will be observed. Their problem is:

$$\max_{e_1} \alpha \cdot g(e_1) - c(e_1).$$

For the threshold γ_1 , participation requires that expected utility be non-negative:

$$\alpha\theta - c(\gamma_1) \geq 0.$$

Setting this expression to equality at the optimum gives:

$$\gamma_1 = c^{-1}(\alpha\theta).$$

Period 2: Student Optimization In the second period, students know with certainty that they will be evaluated if not previously observed. The professor, anticipating this, can set a stricter threshold γ_2 to induce higher effort. The threshold must satisfy two constraints:

$$\gamma_2 > \gamma_1 = c^{-1}(\alpha\theta), \quad 1 - c(\gamma_2) \geq \theta - \alpha\theta,$$

meaning that the threshold should be higher and that the student should benefit more from putting higher effort.

Solving these constraints while maximizing effort leads to:

$$\gamma_2 = c^{-1}(1 + \alpha\theta - \theta).$$

Comparing Total Effort We now compare the total expected effort under the two-threshold scheme to the effort induced by the pass/fail system:

$$c^{-1}(\alpha\theta) + (1 - \alpha)c^{-1}(1 + \alpha\theta - \theta) \gtrless c^{-1}(\alpha)(2 - \alpha).$$

To analyze the effectiveness of the two-threshold scheme, we compute the partial derivative with respect to θ :

$$\left. \frac{\partial f}{\partial \theta} \right|_{\alpha,1} = c^{-1'}(\alpha)(3\alpha - 1 - \alpha^2).$$

The term $c^{-1'}(\alpha)$ is always positive, while $(3\alpha - 1 - \alpha^2)$ changes sign at $\alpha = \frac{3-\sqrt{5}}{2}$. This implies that for lower values of α , the two-threshold method outperforms the single-threshold approach, whereas for higher α , the standard pass/fail grading scheme is optimal. However, the two-threshold method can replicate any outcome of Pass/Fail (by limiting θ to 1), while the range where two-threshold function is of higher performance cannot be replicated by simple pass/fail. Hence, two-threshold grading overall produces higher performance than Pass/Fail.

Optimality of the Two-Threshold Structure We now assess whether further specifications of $g(e)$ could improve performance. Additional discrete thresholds beyond two are redundant in a two-period model since only two thresholds are utilized. Hence, any scheme with more than two discrete thresholds collapses to the two-threshold method.

As for the continuously increasing grading functions, suppose they induce some equilibrium efforts e_1^* and e_2^* that yield grades g_1 and g_2 .

For period 1, e_1^* is necessarily at least under the same restrictions as γ_1 (which is the equilibrium first-period effort under two-threshold grading):

$$\alpha * g_1 - c(e_1^*) \geq 0$$

For period 2, g_2 and e_2^* are necessarily at least under the same constraints as θ and γ_2 (which is the equilibrium second-period effort under two-threshold grading):

$$g_2 > g_1, \quad g_2 - c(e_2^*) \geq g_1 - c(e_1^*).$$

Given the maximum possible value for g_2 is 1, and acknowledging that $g_1 = \theta$, putting these values back into the constraints and solving for e_1^* and

e_2^* yields the same results as for the two-threshold grading, where producing $e_1^* = c^{-1}(\alpha\theta) = \gamma_1$, and $e_2^* = c^{-1}(1 + \alpha\theta - \theta) = \gamma_2$. Hence, a continuously increasing grading function, having the least number of constraints, can at most replicate the result of the two-threshold method but never surpass it. Thus, we conclude that the grading function:

$$g_{t/t}(e) = \begin{cases} 1, & e \geq c^{-1}(1 + \alpha^*\theta^* - \theta^*), \\ \theta^*, & e \geq c^{-1}(\alpha^*\theta^*), \\ 0, & \text{otherwise.} \end{cases}$$

where

$$(\alpha^*, \theta^*) \in \arg \max \quad \mathbb{E} [\pi_P((D(\alpha), g_{t/t}), e)]$$

achieves the highest possible effort level in this model.

5.3 Proof of Proposition 3

Assume that α^* is the optimal alpha of the Pass/Fail exam, achieving total effort V^* .

We can replicate V^* using the 2-threshold exam by having $\theta = 1$.

Hence, examine $\frac{\partial f}{\partial \theta}$ at $(\alpha, \theta) = (\alpha^*, 1)$ to see its sign.

If the sign is negative (which is true for $\alpha < \frac{3-\sqrt{5}}{2}$), decreasing θ achieves $V' > V^*$, implying that the highest effort can be achieved through 2-threshold grading.

If the sign is positive, decreasing θ achieves $V' < V^*$, implying that the highest effort can be achieved through Pass/Fail grading.

References

- [1] Mauro Castiglioni, Giulia Marchesi, and Nicola Gatti. Designing menus of contracts efficiently: The power of randomization. In *Proceedings of the 2022 ACM Conference on Economics and Computation*, pages 570–571. Association for Computing Machinery, 2022.
- [2] Pradeep Dubey and John Geanakoplos. How to pay workers when wages also confer status. Technical Report 1714, Cowles Foundation Discussion Paper, 2009. Yale University.
- [3] Ronald A. Dye. Optimal monitoring policies in agencies. *The RAND Journal of Economics*, 17(3):339–350, 1986.
- [4] Jeffrey Ely, George Georgiadis, and Luis Rayo. Belief distortion and performance feedback. *Forthcoming, Theoretical Economics*, 2024. Available at: <https://faculty.chicagobooth.edu/luis-rayo/research>.
- [5] George Georgiadis. Contracting with moral hazard: A review of theory and empirics. *Working Paper*, 2022. Kellogg School of Management, Northwestern University. Available online: <https://georgiadis.northwestern.edu/research/>.
- [6] George Georgiadis and Balázs Szentes. Optimal monitoring design. *Econometrica*, 88(6):2473–2502, 2020.
- [7] Rick Harbaugh and Eric Rasmusen. Coarse grades: Informing the public by withholding information. *American Economic Journal: Microeconomics*, 10(1):210–235, 2018.
- [8] Kyung Ryun Kim and Eun Hee Seo. The relationship between procrastination and academic performance: A meta-analysis. *Personality and Individual Differences*, 82:26–33, 2015.