

Blind Disclosure

Aaron Kolb Marilyn Pease Daniel W. Sacks Joshua Quick*

February 2021

Abstract

We develop and test a theory of *blind disclosure*. A risk-averse sender chooses whether to disclose information based on a preliminary, private signal. In the unique equilibrium, contrary to the literature's canonical unraveling result, senders disclose only if their preliminary signal exceeds a cutoff. This cutoff rule leads to partial unraveling, with less disclosure in environments with more uncertainty or more risk aversion. Using unique administrative data on disclosed and undisclosed grades in a large university, we find that the model is consistent with student choices during Spring 2020 to conceal letter grades by switching to optional pass-fail grades.

*Kolb: Indiana University, kolba@indiana.edu; Pease: Indiana University, marpease@iu.edu; Sacks: Indiana University, dansacks@indiana.edu; Quick: Indiana University, jdquick@iu.edu. We are grateful to Rick Harbaugh for many detailed conversations about this project, to Ben Motz for assistance with the data, to Aditya Mahendran for excellent research assistance, and to seminar audiences at Indiana University for helpful feedback.

1 Introduction

Many major disclosure decisions must be made under partial uncertainty. For example, when students decide whether to take a course pass/fail, they do not know what grade they will ultimately receive. When homeowners test for lead paint, they are obligated to disclose the outcome, though unremediated paint can reduce house prices (Billings and Schnepel, 2017). When athletes enroll in voluntary drug testing programs, they are likely uncertain about whether they have consumed banned substances.¹ And when firms hire an auditor, they do not necessarily know the outcome of the audit.² Since agents often commit to disclose the results of an evaluation or audit before knowing its outcome, uncertainty looms large in disclosure decisions. These examples contrast with seminal theoretical work, which assumes away uncertainty and yields the stark conclusion that, absent disclosure cost, senders disclose all information, even unfavorable information, out of fear that nondisclosure will be interpreted even less favorably (e.g., Milgrom (1981), Grossman (1981), Viscusi (1978)).

How does this uncertainty affect equilibrium disclosure? In this paper, we develop and test a model of *blind disclosure*, in which a risk-averse sender must choose whether to disclose a piece of information (“grade”) based only on a noisy, private signal of it. We first show that if the sender is at all risk averse and there is any uncertainty over the final grade, then the unraveling logic breaks down. Risk aversion makes senders particularly afraid of disclosing a grade that might turn out to be bad, so they are willing to pool with senders who conceal even if their likely grade is better than the average concealed grade. In the unique equilibrium of our baseline model, there is a cutoff in the sender’s signal of their likely grade above which all senders disclose and below which all senders conceal. For the cutoff type, the negative inference from concealing is equal to the risk premium associated with disclosing the unknown grade. The main testable implication of the model is that the less information the senders have at the time of the disclosure decision, the more likely they are to conceal their grade.

While other factors can also prevent complete unraveling, our examination of uncertainty reveals novel economic forces. Our model generates an endogenous disclosure cost, microfounding previous literature on, and generating forces not present in, models with an exogenous cost. For instance, a less precise signal both increases the risk associated with disclosure and makes the sender’s interim belief less sensitive to that signal, in particular, making senders with low signals more optimistic. In an extension of the model where the sender has private information about both the preliminary signal’s realization and its precision, we show that for some parameter values, students with less precise signals are more likely to disclose despite facing a higher risk premium, highlighting the tension between the above forces. Uncertainty also has implications for the value of public information about the sender’s type; despite reducing the risk premium associated with disclosure, providing a bit of such information can make the sender less likely to disclose, causing the receiver to end up with less precise information overall in equilibrium.

¹The Voluntary Anti-Doping Association manages one such testing program. It is estimated that 6.4-8.8% of doping cases are inadvertent and due to contaminated dietary supplements (Outram and Stewart, 2015).

²A recent, prominent example is the fall of Wirecard, the payment processing company. Facing allegations of financial misconduct, Wirecard hired a new auditor, KPMG, to perform an external review and reassure investors and regulators. The review ended up confirming the misconduct, leading to bankruptcy and criminal investigations (Davies, 2020).

To test the model, we must overcome the challenge of testing any model with private information: the difficulty of observing agents' private information. We address this challenge using unique student data from Spring 2020 at a large public university, exploiting a grading policy change due to the COVID-19 pandemic that allowed students to switch to pass/fail grading mid-Semester. We obtain data on students' disclosed grade from the Registrar, as well as data on students' grades over the course of the semester from Canvas, the university's learning management platform. These data record not only students' grade disclosure decisions and final grades (whether disclosed or not), but also an interim grade, before disclosure decisions were made, and they record grades regardless of students' disclosure decision. We use the interim grade to develop a proxy for students' private signal about their final grade. The predictive power of this signal differs across classes; in some classes it is highly predictive, but in others it is noisier, with substantial residual uncertainty about final grade given the signal. The data show incomplete unravelling. Overall disclosure is 85 percent, but this high average masks the fact that only 15 percent of D grades and 41 percent of C grades are disclosed.

We find clear support for the blind disclosure model as a partial explanation for this incomplete unraveling. Our key empirical result is that students are statistically significantly less likely to disclose a given expected grade if it comes in a class in which signals are noisier. This result holds when we control for observed student or class characteristics and, critically, it also holds in models with student fixed effects. Because we condition on expected grade and student fixed effects, our result is not driven by higher disclosure of higher grades, nor by students with high disclosure rates sorting into classes with more precise signals. Our point estimates indicate that, for a given student, eliminating uncertainty would reduce nondisclosure by about 14 percent. We argue that moral hazard is unlikely to explain this pattern.

Our theoretical and empirical findings can partially address the widespread concerns that appeared in student newspapers and social media in Spring 2020 about how students would be negatively affected by being forced to decide on pass/fail grading before knowing their exact grade. Such concerns focus on the direct effect of increased uncertainty in the presence of risk aversion. But they ignore statistical spillovers and strategic interactions as uncovered by our analysis: the stigma of concealment is reduced once the receiver accounts for the increased uncertainty, and there are strategic complementarities between students' concealment and more favorable interpretation of concealment by the receiver. These equilibrium effects provide a channel through which students might benefit from early deadlines for selecting pass/fail grading.

These results on information and risk also offer new insight into the general question of grading policies under voluntary disclosure. The literature has not previously considered how a grading policy involves a tradeoff between providing more information to receivers at the expense of more risk to the senders, but our application to student grading highlights that this tradeoff can in practice be quite important, and this tradeoff is mediated by the amount of private information available to senders at the time of the disclosure decision. If all weight is put on providing information to receivers then mandating exact grades is best, while if all weight is put on minimizing risk to students, then providing no information at all is best (which is close to mandating pass/fail in our environment where fail is

extremely rare).³ Our results show that any desired level of disclosure can be implemented under voluntary disclosure by adjusting the amount of information provided to students about their grade. Even though blind disclosure is inefficient in the sense that student “mistakes” in guessing their final grade lead to both more risk and less information, the reduced risk from less unraveling can compensate for this inefficiency. Hence, paradoxically, blind disclosure can still be preferable to ex post disclosure when enough weight is put on reducing risk.

2 Contribution and Literature Review

Our analysis of blind disclosure contributes to the empirical and theoretical literature on disclosure. Since unraveling was first analyzed (Viscusi (1978), Milgrom (1981), Grossman (1981)), the theoretical literature has explored the forces that prevent full disclosure. For instance, if there is uncertainty about the sender’s ability to disclose verifiable information—distinct from senders’ uncertainty about what is being disclosed that we study—then types with unfavorable information can pool with those who have no verifiable information to disclose (e.g., Dye (1985), Shin (2003)). Similarly, if disclosure is costly, then disclosure of relatively unfavorable information is not worth it despite the negative inference from concealment (e.g., Viscusi (1978), Jovanovic (1982)).⁴ Market structure can also inhibit disclosure; for example, in some circumstances duopolists will not disclose even absent disclosure costs (Board, 2009). More broadly, in a survey of this literature, Dranove and Jin (2010) note that the theoretical research on the reasons for non-disclosure “has focused on the problems posed by disclosure costs, market structure, and the role of consumers,” while maintaining the assumption that disclosing parties know perfectly the quality that they are disclosing. These factors seems unlikely to fully explain incomplete disclosure in our context, as disclosure is costless and market power seems unimportant. By contrast we argue that uncertainty likely is important in our context. We therefore develop a model of incomplete disclosure that can be applied to situations where quality is uncertain at the time of disclosure, and test the model using unique data, providing the first evidence that uncertainty can reduce disclosure.

Within the large theoretical literature on disclosure, two papers are especially relevant as they also incorporate uncertainty into a disclosure decision. Lubensky and Schmidbauer (2020) analyze the decision of whether to let consumers “test drive” a product when the firm knows product quality and consumers learn this quality and also horizontal match value from trying the product. They find a cutoff equilibrium where senders with (known) private information above a threshold allow

³Separate from risk concerns, the Bayesian persuasion literature investigates how withholding some information can benefit the sender. Chakraborty and Harbaugh (2007) and Ostrovsky and Schwarz (2010) show that elite schools have less incentive to provide information on student ability given the more favorable job market their students face, and argue that grade inflation at elite schools effectively hides information. It may also explain why six of the top ten US News and World Reports schools adopted universal mandatory pass/fail in Spring 2020 while only four other schools in the top 100 did.

⁴Note that our model can be viewed as microfounding the explicit cost of disclosure in these models, as risk aversion creates an implicit cost of disclosure associated with uncertainty. (Moreover, the implicit cost of disclosure for the cutoff type can be an equilibrium object; see Section 3.7.) In addition, because this implicit cost is a function of model parameters, some of our comparative statics involve new forces relative to a model with explicit costs; see Proposition 2 part (vii) and the second part of Proposition 3.

test drives. [Bond and Zeng \(2019\)](#) analyze disclosure of firm news to receivers who have private information about their preferences. Risk aversion leads firms with more extreme good or bad news to be less willing to disclose, so that in equilibrium, only firms with intermediate information reveal. While these papers and the current one demonstrate how the presence of different kinds of uncertainty can lead to partial disclosure, our model differs from these in several respects. Whereas these models feature heterogeneous receivers—and thus can be thought of as reflecting uncertain match quality—we model uncertainty with homogeneous receivers, and thus our model is likely applicable where there is uncertainty in vertical quality. In contrast to both papers above, ours admits a unique equilibrium, which allows us to easily compute comparative statics and generate testable predictions. Furthermore, our model is amenable to empirical analysis in settings where the sender’s information can be measured directly, such as in our grade disclosure setting.

While much of the broad empirical literature on disclosure focuses on its *consequences*,⁵ we contribute to a strand focusing on the *causes* of (non-)disclosure. Especially relevant is work by [Brown et al. \(2012\)](#) and [Bederson et al. \(2018\)](#). Brown et al. document that certain types of movies reliably do not screen for critics, a form of non-disclosure. They argue that receiver naivete is the likely explanation for the failure of unravelling, as some moviegoers do not interpret non-disclosure sufficiently negatively. [Bederson et al. \(2018\)](#) show that not all restaurants disclose their hygiene grades, even restaurants receiving an A grade. As high scoring A’s are especially likely to be concealed, [Bederson et al. \(2018\)](#) argue that this non-disclosure is explained by countersignalling in the sense of [Feltovich et al. \(2002\)](#). Both papers consider but ultimately dismiss uncertainty as a reason for nondisclosure in their contexts, with Brown et al. arguing that professional film studios are unlikely to face much uncertainty about critical response, and [Bederson et al. \(2018\)](#) showing that restaurants can predict the outcome of their inspection, implying limited uncertainty. Thus we find that uncertainty limits disclosure while the prior literature does not. A possible explanation for our different findings is that we study a context where uncertainty is likely to be especially important: the disclosure decision of non-professionals in a new environment.

3 Model

A representative student has an ability $\alpha \in \mathbb{R}$, unknown to both the student and the market. The student and market share a common prior $\alpha \sim N(\mu_\alpha, 1/\tau_\alpha)$, where $\tau_\alpha \in (0, +\infty)$ is the precision of the belief. The student is enrolled in a course which results in a final grade γ that is an unbiased signal of ability, $\gamma|\alpha \sim N(\alpha, 1/\tau_\gamma)$, where $\tau_\gamma \in (0, +\infty) \cup \{+\infty\}$. One interpretation of the prior is the

⁵Seminal papers by [Dranove et al. \(2003\)](#) and [Jin and Leslie \(2003\)](#) show that hospitalizations for food borne illness decline after restaurant hygiene disclosure (because of improved hygiene), but physicians increased cream skimming after their quality was publicized (to manipulate their quality rating). More recent work has shown that posting calories decreases calories of food purchases at Starbucks ([Bollinger et al., 2011](#)), posting HMO plan ratings drive enrollment towards higher quality plans ([Darden and McCarthy, 2015](#)), U.S. News college rankings affects college choice ([Luca and Smith, 2013](#)), as does information on the earnings of college graduates ([Hurwitz and Smith, 2018](#)), and energy efficiency information influences appliance choice ([Houde, 2018](#)).

student’s grade point average (GPA) based on previous courses.⁶

The student privately observes a noisy preliminary score S that is an unbiased signal of the final grade, $S = \gamma + \epsilon$ where $\epsilon \sim N(0, 1/\tau_\epsilon)$ and $\tau_\epsilon \in (0, +\infty)$. The student then chooses an action *Disclose* or *Conceal*.⁷ A (competitive) market observes the student’s choice (and grade, if it is disclosed) and offers the student a wage $w \in \mathbb{R}$ equal to the student’s expected ability. The student’s payoff is given by a CARA utility function $u(w) = 1 - e^{-\lambda w}$, where $\lambda > 0$ is the level of absolute risk aversion.⁸ Except in Section 3.8, we abstract from moral hazard in order to focus our analysis on the strategic disclosure problem.⁹

The solution concept is perfect Bayesian equilibrium (henceforth, *equilibrium*). The market offers a wage equal to its expectation of the student’s type α , determined by Bayes’ rule whenever possible. The student’s action maximizes her payoff given the market wage schedule. In our noisy environment where students make their decision to conceal or disclose their final grade γ without yet knowing it, we will show that all actions are on the equilibrium path.

3.1 Overview of Primary Forces

Proposition 1 below shows that with risk aversion and blind disclosure, the unique equilibrium features incomplete unraveling. To show why some concealment occurs in equilibrium, we identify several key effects of increased uncertainty, that is, of decreased precision in the student’s private signal S . A *risk aversion effect* operates through a student’s greater perceived variance in the final grade. In the presence of risk aversion, increasing this variance (holding fixed the mean belief about the final grade) increases the risk premium the student is willing to pay to avoid the risk inherent in disclosure, leading better performing students to conceal. A *mean-reversion effect* operates as students’ mean beliefs about their grade are now less responsive to the private signal. The direction of this effect is ambiguous; students whose signals are higher than expected are made more pessimistic, and hence more willing to conceal, while the opposite holds when signals are lower than expected. A *sorting effect* operates from the market’s perspective, as high-ability students become more likely to conceal due to receiving a bad signal. Relative to a setting in which disclosure decisions are made with full knowledge of the grade, the market is more optimistic about the ability of students who chose to conceal and offers them a higher wage, fixing the fraction of students who choose to conceal. Finally, two additional indirect effects are at play. In response to a higher wage conditional on conceal, more students choose to conceal (with the marginal students being better performing on average), and as better students begin to conceal, the market raises its wage for those who conceal.

⁶Most simply, if there is no other prior information on student ability and the grades γ_n in $N - 1$ previous courses are i.i.d. $N(\alpha, 1/\tau_\gamma)$ then $\mu_\alpha = \sum^{N-1} \gamma_n / (N - 1)$ and $\tau_\alpha = (N - 1)\tau_\gamma$.

⁷*Conceal* here is a simplification of Pass/Fail options because it eliminates the possibility of failure, but failure is rare in practice.

⁸The advantage of such preferences is that they admit a mean-variance representation, simplifying the analysis: the expected payoff of a normally distributed random wage with mean \bar{w} and variance σ^2 is the same as the payoff from a guaranteed wage of $\bar{w} - \frac{\lambda}{2}\sigma^2$. We explore non-CARA preferences in Section 3.7.

⁹Following most disclosure models, we take student effort as given in our main analysis to focus on the ramifications of information and risk aversion. Moral hazard and schooling are addressed in more depth in the papers of Costrell (1994), Dubey and Geanakoplos (2010), and Boleslavsky and Cotton (2015).

3.2 Equilibrium Analysis

Suppose, as we will confirm in Proposition 1, that there is a cutoff equilibrium where students with a preliminary score above some $S = s^*$ choose to *Disclose* their final grade γ and students with a lower score *Conceal* it. For a student with this score s , the expected payoff from taking a risk and disclosing γ must equal the expected payoff from pooling with types $S < s^*$ who conceal γ ,

$$\mathbb{E}[u(\mathbb{E}[\alpha|\gamma])|S = s^*] = u(\mathbb{E}[\alpha|S < s^*]). \quad (1)$$

First consider the expected *Disclose* payoff on the LHS of indifference condition (1). Notice that upon seeing the disclosed grade γ , the market's expectation $\mathbb{E}[\alpha|\gamma]$ on the LHS above does not further depend on its conjecture about the student's strategy, since S is the student's only source of private information, and it contains no additional information given the grade γ . (In contrast, the market's expectation after observing *Conceal* clearly depends on its conjecture of the student's strategy.) Using the mean-variance representation of expected utility for CARA preferences,

$$\mathbb{E}[u(\mathbb{E}[\alpha|\gamma])|S = s^*] = u\left(\mathbb{E}[\mathbb{E}[\alpha|\gamma]|S = s^*] - \frac{\lambda}{2}\text{Var}[\mathbb{E}[\alpha|\gamma]|S = s^*]\right). \quad (2)$$

We now simplify the mean and variance terms in the argument of u on the right side of (2). Since S is an unbiased signal of α with precision $\tau_S := (1/\tau_\gamma + 1/\tau_\epsilon)^{-1}$, the posterior distribution of α given the signal realization $S = s^*$ is $\alpha|S = s^* \sim N\left(\frac{\tau_\alpha\mu_\alpha + \tau_S s^*}{\tau_\alpha + \tau_S}, (\tau_\alpha + \tau_S)^{-1}\right)$. By the law of iterated expectations,

$$\mathbb{E}[\mathbb{E}[\alpha|\gamma]|S = s^*] = \mathbb{E}[\alpha|S = s^*] = \underbrace{\frac{\tau_\alpha}{\tau_\alpha + \tau_S}}_{\text{mean-reversion effect}} \mu_\alpha + \frac{\tau_S}{\tau_\alpha + \tau_S} s^*. \quad (3)$$

The coefficient on the prior μ_α is decreasing in τ_ϵ through τ_S , and thus increasing in the student's uncertainty, illustrating the mean-reversion effect. Moreover, since γ is an unbiased signal of α with precision τ_γ , the posterior distribution of $\alpha|\gamma$ has precision $\tau_\alpha + \tau_\gamma$. Thus, by the law of total (conditional) variance,

$$\begin{aligned} \text{Var}[\mathbb{E}[\alpha|\gamma]|S = s^*] &= \text{Var}[\alpha|S = s^*] - \mathbb{E}[\text{Var}[\alpha|\gamma]|S = s^*] \\ &= \frac{1}{\tau_\alpha + \tau_S} - \frac{1}{\tau_\alpha + \tau_\gamma} = \underbrace{\frac{\tau_S}{\tau_\alpha + \tau_S} \frac{\tau_\gamma}{\tau_\epsilon(\tau_\alpha + \tau_\gamma)}}_{\text{risk aversion effect}}. \end{aligned} \quad (4)$$

The expression in (4) is decreasing in τ_ϵ , and thus increasing in the student's uncertainty, demonstrating the risk aversion effect. Incorporating (3) and (4) in (2), the LHS of (1) becomes

$$\mathbb{E}[u(\mathbb{E}[\alpha|\gamma])|S = s^*] = u\left(\frac{\tau_\alpha\mu_\alpha + \tau_S s^*}{\tau_\alpha + \tau_S} - \frac{\lambda}{2} \frac{\tau_S}{\tau_\alpha + \tau_S} \frac{\tau_\gamma}{\tau_\epsilon(\tau_\alpha + \tau_\gamma)}\right). \quad (5)$$

Now consider the expected *Conceal* payoff on the RHS of indifference condition (1). Noting that $S \sim N(\mu_\alpha, \nu^2)$ with $\nu^2 := 1/\tau_\alpha + 1/\tau_\epsilon + 1/\tau_\gamma$,¹⁰ by the formula for the expectation of a truncated normal random variable, $\mathbb{E}[S|S < s^*] = \mu_\alpha - \nu\phi(z^*)/\Phi(z^*)$, where $z^* := \left(\frac{s^* - \mu_\alpha}{\nu}\right)$ is the cutoff in terms of the standardized or “curved” score. Therefore, by the law of iterated expectations,

$$\mathbb{E}[\alpha|S < s^*] = \mathbb{E}\left[\underbrace{\frac{\tau_\alpha\mu_\alpha + \tau_S S}{\tau_\alpha + \tau_S}}_{=\mathbb{E}[\alpha|S]} \mid S < s^*\right] = \mu_\alpha - \underbrace{\frac{\tau_S}{\tau_\alpha + \tau_S}}_{\text{sorting effect}} \nu \phi(z^*)/\Phi(z^*). \quad (6)$$

The sorting effect is demonstrated by $\frac{\tau_S}{\tau_\alpha + \tau_S}\nu$ increasing in τ_ϵ : fixing the fraction of students who conceal, when students have more precise information, a *Conceal* decision is a more accurate indicator of a low grade (and in turn, low ability), and the market more strongly revises its belief downward. In other words, with more uncertainty, there is less sorting and a higher expectation of those students who conceal.

Substituting the *Disclose* and *Conceal* payoff derivations (5) and (6) back into (1), writing s^* in terms of z^* , and simplifying, we obtain the equation

$$z^* - \frac{\lambda}{2} \frac{\tau_\gamma}{\nu(\tau_\alpha + \tau_\gamma)\tau_\epsilon} = -\frac{\phi(z^*)}{\Phi(z^*)}. \quad (7)$$

Any equilibrium thus corresponds to a solution to (7). In the appendix, we prove part (i) of the following result, and we prove part (ii) by showing that (7) has a unique solution whenever $\lambda > 0$.

Proposition 1 (Partial Unraveling) *(i) In any equilibrium there exists a cutoff $s^* \in \mathbb{R}$ such that the student plays Disclose if $s \geq s^*$ and Conceal if $s < s^*$. (ii) If the student is risk averse ($\lambda > 0$) such a cutoff equilibrium exists, is unique, and is given by $s^* = z^*\nu + \mu_\alpha$, where z^* is the unique solution to (7).*

The monotonicity of the student’s preferences with respect to her private signal ensures that any equilibrium must involve a cutoff strategy for the student.¹¹ However, the wage after concealment and the student’s cutoff are strategic complements: the wage is increasing in the market’s conjecture of the student’s cutoff, while the student’s cutoff is increasing in the wage. Despite this complementary, there is a unique equilibrium.

3.3 Comparative Statics

The comparative statics of the likelihood that a student discloses, through the cutoff z-score, can be understood through Figure 1, which illustrates the the left and right sides of (7). Since the black line crosses the blue curve from below, changes in parameter values which shift down the left side (in

¹⁰It is convenient to work with the standard deviation ν here rather than the precision, since this allows us to easily express the “z-score” associated with each signal realization, which we identify in our final equation.

¹¹Note that each cutoff determines an essentially unique strategy, with flexibility only in the student’s behavior when the signal is the cutoff itself, which happens with probability zero. Going forward, we drop the qualifier *essentially*, and uniqueness of equilibrium reduces to uniqueness of the equilibrium cutoff.

black) cause the equilibrium cutoff z-score to increase and the probability of disclosure to decrease. In particular, the cutoff z-score is decreasing in τ_ϵ , as in part (i) of Proposition 2 below: more information about the final grade makes it more likely that a student discloses through the risk aversion effect, sorting effect, and strategic effects; this is despite the fact that the mean-reversion effect can work in the opposite direction.¹²




Figure 1: Existence and uniqueness of equilibrium cutoff (z-score). The solid and dashed black lines are the left hand side of (7) evaluated at $\tau_\epsilon = .5$ and $\tau_\epsilon = 1$, respectively, with $\lambda = \tau_\alpha = 4$ and $\tau_\gamma = 3$, while blue curve is the right hand side of (7).

Parts (ii) and (iii) of Proposition 2 show how the probability of disclosure conditional on γ ,¹³ is affected by a change in precision τ_ϵ . On average, higher precision induces a higher probability of disclosure in equilibrium, but there is a heterogeneous effect on students that depends on their grade. As information precision increases, students sort more accurately, and thus students with low (high) grades become more likely to conceal (disclose). Moreover, the students most likely to change their behavior are those whose grade lies in an intermediate interval, as these students are most likely to have a signal between the two values of the cutoff s^* .

Parts (iv)-(vii) of Proposition 2 describe how the probability of disclosure depends on other model parameters. Intuitively, students are less likely to disclose when they are more risk averse. Increasing the precision of the grade τ_γ as a signal of ability has a similar effect, as it raises the stakes for students — the market’s assessment of the student’s ability and therefore the wage is more sensitive to the grade realization. In contrast, increasing the precision of the prior information τ_α eventually leads to more disclosure by reducing the risk premium, a force which arises due to blind disclosure.

¹²We show in Proposition 9 in the appendix that the cutoff s^* is also decreasing in τ_ϵ .

¹³This probability is $\Pr(\epsilon > s^* - \gamma|\gamma) = 1 - \Phi((s^* - \gamma)\sqrt{\tau_\epsilon})$.

However, there exists a region of parameter values for which disclosure is *decreasing* in the precision of the prior. To illustrate, suppose τ_γ is a perfect signal of the student’s ability. Increasing the prior precision then has no effect on the market’s assessment if the student discloses. But the student’s signal becomes relatively less useful in predicting her ability, which dampens the negative inference by the market upon observing *Conceal*. Consequently, the student becomes more willing to conceal. Finally, the fraction of students who conceal is independent of the prior mean about their ability, since higher mean ability results in both higher average performance and a higher expected performance conditional on concealment, two forces which exactly offset each other.

Proposition 2 (Comparative Statics) *Fixing other parameter values, as the precision τ_ϵ of the student’s private signal increases, (i) the equilibrium cutoff z^* decreases and the probability of disclosure increases, with $z^* \rightarrow +\infty$ as $\tau_\epsilon \rightarrow 0$ and $z^* \rightarrow -\infty$ as $\tau_\epsilon \rightarrow +\infty$; (ii) conditional on the grade γ , the probability of disclosure increases (decreases) for students with high (low) grades; and (iii) the students whose probability of disclosure is most affected are those whose grade lies in an intermediate interval. The probability of disclosure is (iv) decreasing in the degree of risk aversion λ , (v) decreasing in the precision τ_γ of the grade as a signal of ability, (vi) independent of the prior mean μ_α of ability, and (vii) nonmonotonic and single-peaked in the prior precision τ_α of ability.*

Propositions 1 and 2 yield testable implications for our grade disclosure setting. From Proposition 1, we predict that a positive fraction of students will choose to conceal their grades in each class. Moreover, from Proposition 2 part (i), we predict that this fraction will be largest in classes in which students have the least precise information about their performance.

3.4 Information Loss

In this section, we quantify the amount of information loss for the market as a result of blind disclosure. We formalize information loss as an increase in the mean squared error (MSE) of the market’s posterior belief of the student’s ability.¹⁴ Letting $\mathbb{E}^m[\alpha]$ denote the market’s posterior expectation of α , the MSE is $\mathbb{E}[(\alpha - \mathbb{E}^m[\alpha])^2]$.

Reducing the precision of the student’s information results in information loss in equilibrium through two channels: there is more concealment of grades, and there is more noise in the sorting of students into those who conceal and those who disclose; the fact that a student concealed becomes less informative of the student’s ability. The first part of Proposition 3 says that any decrease in the precision of the student’s private signal results in information loss in equilibrium.

On the value of public information: If we interpret the common prior $\alpha \sim N(\mu_\alpha, 1/\tau_\alpha)$ as the result of a diffuse prior over student ability followed by a public signal $\mu_\alpha \sim N(\alpha, 1/\tau_\alpha)$, a natural question is whether more public information (i.e., increasing τ_α) benefits the market in the sense of reducing the mean squared error. While a *perfect* public signal reveals the student’s ability and results in zero

¹⁴The notion of mean squared error used in this section to measure uncertainty in the market’s posterior belief is not to be confused with the (root) mean squared error (RMSE) used later in our empirical analysis to measure students’ uncertainty about their performance.

mean squared error, for low values of τ_α , a marginal increase in τ_α can lead to higher mean squared error. While increasing τ_α by definition raises the “baseline” amount of information available to the market, the probability that the student discloses is decreasing in τ_α when the latter is close to zero, from Proposition 2. For some parameter values, this second force dominates the first, and the mean squared error increases with public information; effectively, public information provision can “crowd out” strategic information disclosure. Figure 2 illustrates. The possibility that a marginal increase in public information can be harmful also arises in the coordination game of Morris and Shin (2002) and the cheap talk game of Chen (2012), albeit for different reasons.




Figure 2: Standardized cutoff (z^*) and mean squared error as functions of τ_α for $\tau_\gamma = 5$, $\tau_\epsilon = 1$, $\lambda = 4$.

Proposition 3 *The mean squared error of the market’s assessment of ability is strictly decreasing in τ_ϵ . There exist parameter values for which the mean squared error is nonmonotonic in τ_α .*

An implication of Proposition 3 is that the mean squared error in the market’s estimate of the student’s *grade* is also strictly decreasing in τ_ϵ . In particular, calculating the MSE for the ability under the specification $\tau_\gamma = +\infty$ yields the same result as calculating the MSE for the students grade after relabeling the prior precision over γ as τ_α in the latter calculation.

3.5 Extension: Receiver Uncertainty over Sender’s Information Precision

In many settings of blind disclosure, the receiver might not know the precision of the sender’s private signal. In our grade disclosure application, for example, the market might not know how accurate the student’s preliminary grade estimate was at the time she had to make her disclosure decision in a particular class. Indeed, the market might know that students faced varying deadlines across classes for opting into pass/fail grading without knowing which deadlines occurred in which classes.

In this section, we extend the baseline model to allow for uncertainty by the market about the precision τ_ϵ of the student’s private signal. We show that the existence of equilibrium with incomplete

unraveling is robust to such uncertainty, and it can induce incomplete unraveling even among perfectly informed students. For simplicity, suppose $\tau_\epsilon \in \{\tau_L, \tau_H\}$ with $\tau_L < \tau_H$, and let $q = \Pr(\tau_\epsilon = \tau_H) \in (0, 1)$ be the market’s belief about τ_ϵ . Assume that q , τ_L , and τ_H are common knowledge.

As stated in Proposition 4, when both τ_L and τ_H are positive and finite, for any distribution over these precision types, there exists an equilibrium in cutoff strategies; existence extends to the case where $\tau_H = \infty$ when q is sufficiently small, which we defer to Proposition 5. To characterize cutoff equilibria, we first fix an “input wage” of w given to students who conceal, determine the best-reply cutoff signals for types τ_L and τ_H in response to w , and calculate the “output wage” $f(w)$, i.e., the expected ability of a student who conceals assuming the student follows the best-reply cutoffs. Figure 3 illustrates; note that the output wage $f(w)$ lies between its complete-information counterparts denoted $f(w; H)$ and $f(w; L)$.

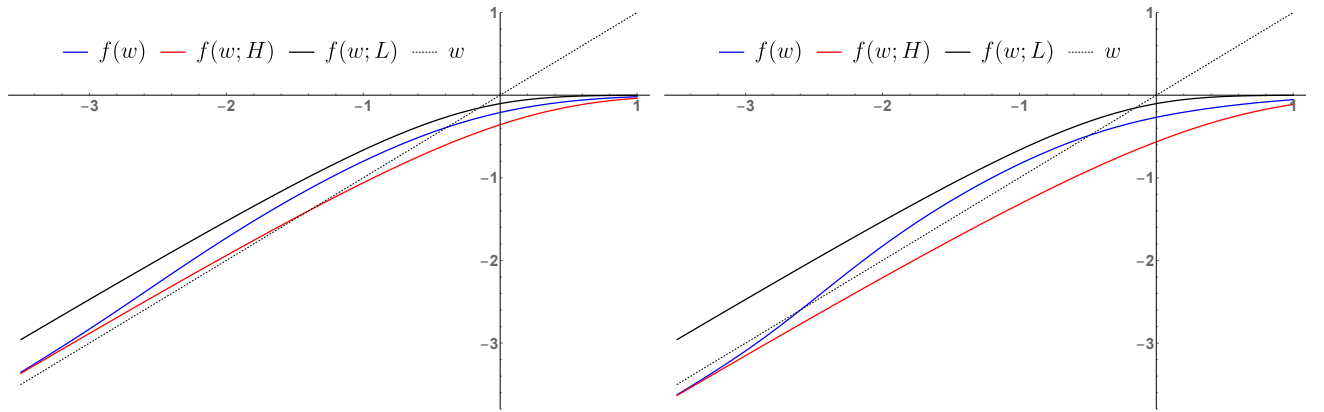


Figure 3: Equilibrium wages correspond to fixed points of the equation $w = f(w)$. Left panel: $\tau_\epsilon^H = 2$; right panel: $\tau_\epsilon^H = \infty$. Other parameter values: $\mu_\alpha = 0$, $\tau_\alpha = \tau_\gamma = 1$, $\tau_\epsilon^L = q = 1/2$, $\lambda = 5$.

When the receiver has uncertainty over the sender’s private information precision, there can be multiplicity of equilibria, as in the right panel of Figure 3.¹⁵ Hence, we focus on the ex ante student-optimal equilibrium. Equilibria can obviously be ranked in terms of the wage after *Conceal*, and since the wage after *Conceal* corresponds to a reservation utility for the student, the ex ante student-optimal equilibrium (and in fact the optimal equilibrium for either precision type) is the one with the highest wage after *Conceal*. Generically, this equilibrium is stable, with the best-response dynamic following a small perturbation in the wage leading back to the equilibrium.

In the student-optimal equilibrium, less-informed students follow higher cutoffs and are more likely to conceal, if and only if there is sufficient risk aversion.¹⁶ To see why, first note that in comparing the decision facing two students with the same signal but different precision types, the risk-aversion and mean-reversion effects are at play, but the sorting effect is absent due to the market’s uncertainty over the student’s information precision. Hence, if type τ_ϵ^L students conceal more, the risk-aversion

¹⁵The right panel of Figure 3 illustrates the case $\tau_\epsilon^H = \infty$, but multiplicity can also arise for sufficiently large $\tau_\epsilon^H < \infty$.

¹⁶Beyond the student-optimal equilibrium, if risk aversion is sufficiently high, then in *any* equilibrium, students of type τ_ϵ^L are more likely to conceal than students of type τ_ϵ^H , and if risk aversion is sufficiently low, they are less likely to conceal in all equilibria.

effect must not be overcome by the mean-reversion effect. Intuitively, when students have low risk aversion, any equilibrium wage for students who conceal is very low, so a student must receive a very low signal to become willing to conceal. Since bad signals induce less pessimism for students of type τ_ϵ^L , those students are less likely than type τ_ϵ^H students to obtain such bad signals that they prefer to conceal. In this case, the mean-reversion effect is strong and acts opposite the risk-aversion effect. When students instead have high risk aversion, the equilibrium wage and thus the student cutoffs are not too low. In turn, this ensures that the mean-reversion effect either reinforces the risk aversion effect, or it acts only weakly in the opposite direction, so that the net effect of increased uncertainty is to reduce disclosure.

Proposition 4 (i) For any receiver uncertainty over precision $\tau_\epsilon \in \{\tau_\epsilon^L, \tau_\epsilon^H\}$ where $0 < \tau_\epsilon^L < \tau_\epsilon^H < \infty$, there exists a cutoff equilibrium where $s^*(\tau_\epsilon^L), s^*(\tau_\epsilon^H) \in \mathbb{R}$, and the wage-maximizing equilibrium is stable. (ii) In the wage-maximizing equilibrium, if and only if λ is sufficiently high, students of type τ_ϵ^L are more likely to conceal. (iii) In any equilibrium, conditional on the grade γ , students of type τ_ϵ^H are more (less) likely than students of type τ_ϵ^L to disclose if γ is high (low), and the difference in disclosure probabilities is largest for grades in an intermediate interval.

As in the baseline model, uncertainty affects students heterogeneously with respect to their true grade. With higher uncertainty, students are less likely to disclose conditional on a high grade and more likely to disclose conditional on a low grade, and the magnitude of the difference vanishes at extreme grades. Figure 4 illustrates.




Figure 4: Disclosure probability conditional on γ for students of type $\tau_\epsilon \in \{1, 3\}$, fixing $q = 1/2$ and $\lambda = \tau_\alpha = \tau_\gamma = 5$.

The next proposition says that uncertainty over the student's information precision can prevent

unraveling even among students who are perfectly informed: provided the fraction of perfectly informed students is not too large, there exists a cutoff equilibrium, with a positive fraction of perfectly informed students choosing to conceal. Intuitively, the presence of students who are less informed (and risk averse) can raise the average quality of the pool of students who conceal, so that a perfectly informed student whose grade is at her respective cutoff has the same expected quality as the entire pool of students who conceal, and therefore she cannot gain by separating.¹⁷ For this quality-raising effect to be sufficiently strong, there must be sufficiently many students who are not perfectly informed. In fact, for high values of q , there is complete unraveling due to a composition-shifting effect: when the wage after *Conceal* is very low, the pool of students who conceal consists mainly of those with high precision signals, and thus $|f(w) - f(w; H)| \rightarrow 0$ as $w \rightarrow -\infty$, but $f(w; H)$ does not have a fixed point when $\tau_\epsilon^H = \infty$.

Proposition 5 *Suppose $\tau_\epsilon \in \{\tau_\epsilon^L, \tau_\epsilon^H\}$, where $\tau_\epsilon^L \in (0, \infty)$ and $\tau_\epsilon^H = \infty$. If q is sufficiently small, there exists a cutoff equilibrium where $s^*(\tau_\epsilon^L), s^*(\tau_\epsilon^H) \in \mathbb{R}$; that is, there is an equilibrium with incomplete unraveling even among students who are perfectly informed.*

3.6 Extension: Receiver Uncertainty over Sender’s Risk Aversion

In our grade disclosure setting, it is reasonable to expect that students have some unobserved heterogeneity in their level of risk aversion; that is, λ , is private information for the student. Here we explore an extension of our baseline model in which the market only knows the distribution over λ . Suppose for simplicity that $\lambda \in \{\lambda_H, \lambda_L\}$, where $0 < \lambda_L < \lambda_H < \infty$, and the market has a known prior belief $q = \Pr(\lambda = \lambda_L) \in (0, 1)$.

Unambiguously, in any equilibrium (and fixing the prior), a student with higher risk aversion has a higher cutoff for disclosure than a student with lower risk aversion, since a student with higher risk aversion is willing to pay a larger risk premium for the same gamble over wages. It follows that students with high risk aversion who conceal have higher average ability than students with low risk aversion who conceal, and thus when risk aversion is unknown, any equilibrium wage after *Conceal* lies between the equilibrium wages when risk aversion is known to be low or high. Hence, pooling students of different levels of risk aversion causes highly risk averse students to “subsidize” less risk averse students: students with low risk aversion end up better off, and those with high risk aversion worse off, than if their risk aversion were known to the receiver.

Proposition 6 *There exists an equilibrium with cutoffs $s^*(\lambda_L), s^*(\lambda_H) \in \mathbb{R}$. In any equilibrium, $s^*(\lambda_H) > s^*(\lambda_L)$, students of type λ_H are more likely to conceal than students of type λ_L , and students of type λ_L (λ_H) obtain a higher (lower) expected payoff than they would in the unique equilibrium of the game in which their type is common knowledge.*

¹⁷This result is similar to the incomplete unraveling result in [Dye \(1985\)](#) which is driven by uncertainty over the sender’s verifiable message set. A key difference is that in our model, the student always has the option to disclose, so all nondisclosure is endogenous.

3.7 Extension: General Preferences

We now demonstrate that the main results can be extended beyond the case of CARA preferences of the student. Let u be any C^2 utility function satisfying $u' > 0$ and $u'' \leq 0$. It continues to hold that any equilibrium is a cutoff equilibrium, and as long as risk aversion does not vanish at low wages, such an equilibrium exists. Intuitively, if risk aversion does not vanish at low wages, then for sufficiently low wages, the risk premium associated with the risk of disclosing an unknown grade exceeds the negative effect of pooling with types below the cutoff, preventing unraveling. Moreover, the equilibrium is unique if risk aversion is nonincreasing (ignoring the degenerate risk neutral case). Nonincreasing risk aversion implies that the risk premium is decreasing in the student's signal, while for the conjectured cutoff student, the negative effect of pooling increases in the cutoff; hence, the indifference condition for the cutoff student can only hold at one signal value. The last part of the proposition generalizes the comparative static with respect to λ in Proposition 2: greater risk aversion implies a lower probability of disclosure, provided that the conditions for uniqueness are satisfied.

Proposition 7 (i) *In any equilibrium, there exists a cutoff $s^* \in \mathbb{R}$ such that the student plays *Disclose* if $s > s^*$ and *Conceal* if $s < s^*$. (ii) If $\liminf_{w \rightarrow -\infty} -\frac{u''(w)}{u'(w)} > 0$, then an equilibrium exists. (iii) If $-\frac{u''(w)}{u'(w)}$ is nonincreasing and not identically zero, the equilibrium is unique. (iv) If u and v satisfy the conditions for uniqueness in part (iii), and if u exhibits greater risk aversion than v in the sense that $-\frac{u''(w)}{u'(w)} \geq -\frac{v''(w)}{v'(w)}$ for all w , then there is less disclosure under u than under v .*

3.8 Extension: Moral Hazard

We now introduce moral hazard, allowing the student to influence her final grade through effort.¹⁸ Suppose that after observing her private signal and choosing whether to disclose her grade, the student privately chooses an effort level $e \geq 0$ which increases her grade by e at a cost $\frac{c}{2}e^2$ subtracted from her wage, where $c > 0$ is common knowledge.¹⁹ There exists an equilibrium in which the student plays *Disclose* if and only if her grade exceeds a cutoff s^* , and in which she chooses effort 0 after *Conceal* and effort $e^* = \frac{\tau_\gamma}{c(\tau_\alpha + \tau_\gamma)}$ after *Disclose*. Here, e^* follows from equating the marginal benefit $\frac{\tau_\gamma}{\tau_\alpha + \tau_\gamma}$ of effort with the marginal cost ce . The indifference condition that pins down s^* now must take into account the student's continuation value from choosing *Disclose*. Since the market fully anticipates the student's effort e^* after *Disclose*, the student's wage distribution after choosing $(\text{Disclose}, e^*)$ conditional on her signal is the same as in the baseline model, but the student's utility from playing *Disclose* is reduced by the effort cost $\frac{c}{2}(e^*)^2$. Hence, adding moral hazard reinforces the incomplete unraveling found in the baseline model. However, adding moral hazard does not change our qualitative prediction that more uncertainty results in less disclosure.²⁰

¹⁸We assume, for simplicity, that effort increases the grade but not the sender's value to the market. We leave exploration of the latter possibility for future work.

¹⁹That is, define S and γ as in the baseline model, but now suppose that the market observes $\gamma + e$ if the student chooses to disclose. For simplicity, we do not model effort chosen before the preliminary signal.

²⁰Through the latter channel, uncertainty reduces effort in an ex ante sense; despite that the equilibrium effort e^* is independent of the precision of the student's preliminary signal, uncertainty makes it more likely that students conceal and choose effort 0.

We find that the probability of disclosure is *increasing* in the cost parameter c . Since the equilibrium effort e^* is inversely proportional to c , the total effort cost $\frac{c}{2}(e^*)^2$ is decreasing in c , and thus the student is more willing to disclose when c increases; taking $c \rightarrow \infty$, equilibrium effort costs tend to 0, and we recover the equilibrium of our baseline model. Intuitively, when c is lower, the student has a larger incentive to exert effort to boost her grade, and is willing to pay a larger total cost; however, in equilibrium, the market raises its expectations of effort accordingly, offsetting the benefit.

Proposition 8 *With moral hazard, there exists an equilibrium characterized by a cutoff s^* and effort e^* after disclosure. The probability of disclosure is increasing in the effort cost parameter c and in the student’s signal precision τ_ϵ .*

4 Empirical Setting and Data

We now turn to our empirical analysis.

4.1 Empirical Setting

We study undergraduate students’ grade disclosure decisions during the Spring 2020 semester at Indiana University, Bloomington. Students left for spring break on March 13, 2020, expecting to return on March 23. As the severity of the COVID-19 pandemic became clear, spring break was extended by a week, and then classes moved exclusively to remote instruction, beginning March 30. As the pandemic itself and the abrupt online transition both presented difficulties for instruction, learning, and assessment, Indiana University implemented a novel pass/fail grading option (“P/F”).²¹ This option was approved on March 24 by the faculty council, which gave instructors authority to award P grades to students who had enrolled in their course for a letter grade, with permission from the Dean of the unit (Anderson, 2020). Campus administrators initially indicated that this option was to be used only in exceptional circumstances when it would not be possible to provide an exact grades. However students interpreted this email to mean that they had the option to request a P/F grade. Following student pressure, the university granted this option. The university let schools and departments determine the deadline to exercise the option. Students therefore had the option to disclose their letter grade, or conceal it by selecting the P/F option. This disclosure was “blind” if students had to make their disclosure decisions before knowing their final grades. Most students learned their final grade some time between May 3, the start of final exam period, and May 11, three days after the end of the exam period, the deadline for posting final grades.

The deadlines for when students had to make their disclosure decision differed across academic units and changed over time. For example, on April 5, Draughon (2020) reported that the policy school and the school of public health initially required that students make their P/F decision prior

²¹This option at Indiana University is called S/F grading, S for Satisfactory. Like P grades, S grades do not affect the university’s GPA calculation. S/F grade differs from traditional P/F grades because S grades count towards progress in major and other requirements. We refer to S/F grades as P/F grades because this terminology is more widely used to describe Spring 2020 grading schemes. In our empirical application we drop students with traditional P/F grades.

to the start of the final exam period on May 3. The business school initially allowed students to request and P/F grade only if they faced “extenuating circumstances” due to the pandemic, and implemented class-specific deadlines which had to be after final exam period began. The College of Arts and Sciences would honor P/F requests at instructors’ discretion. Thus the initial policy required that most students make a decision before seeing their final grades, and gave instructors discretion not to honor these requests. However in the following weeks many (but not all) academic units extended their deadlines for P/F decisions. Most schools instructed faculty to honor all requests. The business school allowed students to make a P/F selection after seeing their final course grade. The College of Arts and Sciences allowed students until May 8 (the last day of the final exam period), to make their selection. For all units, the default policy was to disclose a letter grade unless a student affirmatively asks for a P/F.²²

Despite the fact that some students could in principle have waited until seeing their final grades before making their disclosure decision, this setting fits the blind disclosure model well. We say this for three reasons. First, many schools required students to make their P/F decision before seeing their final grade, and thus they faced some uncertainty over what they were disclosing. Second, in our experience, many students made their P/F requests well before the stated deadlines.²³ These early requests may have been motivated by a concern the university would reverse its P/F policy and not allow such P/F grades at all. In that case, from the student’s perspective, the effective deadline was well before the final grade is revealed. Third, some schools allowed students to request a grade change after their grade was posted, allowing them to retroactively conceal a disclosed grade, or disclose a concealed grade. These grade change requests are extremely rare in our data, despite the fact that many students could have raised their GPA by retroactively concealing low grades and/or disclosing high grades. Thus many students effectively committed to their disclosure decision before seeing their final grade, consistent with our blind disclosure model.

4.2 Data

Data sources: Our data derive from two sources at Indiana University: the registrar’s office and Canvas, an online course management platform. We extract a student-course level data set, drawing records for all Indiana University students enrolled in undergraduate classes at the Bloomington campus, in the Spring of 2020 (Quick et al., 2020). We limit the sample to classes offered for the full semester or the second half of the semester, as first-half classes were not subject to the P/F policy. We exclude observations (student-courses) with non-standard grades such as incomplete, withdrawn, or traditional pass. We drop a small number of students with missing Registrar or Canvas data to obtain our “full” sample, and we limit our “analysis sample” to certain classes as described below. [Appendix Table C.1](#) shows how our sample size changes as we impose this and other restrictions to arrive at our full sample and analysis sample.

²²There were some exceptions for courses which could not give traditional grades, such as performance based courses, which could assign P for the entire course. Such courses are not included in our empirical analysis.

²³Because P/F requests were made through instructors or academic units, rather than the registrar, we do not observe the timing of P/F requests in our data.

From the registrar data we see information on students, their course enrollments, and their grades, if disclosed. We observe student credit hours and cumulative grades earned, from which we derive GPA. These variables are measured as of the beginning of the semester, so they do not reflect Spring 2020 performance; we call our GPA measure “incoming GPA” for this reason. We observe limited course-level characteristics, including course level (100, 200, 300, or 400, indicating increasingly advanced courses), course size (measured as number enrollees), and average incoming GPA of enrolled students, as a proxy for course difficulty. For each student-course, we observe final grades, or a P if students selected P/F and did not fail. We say that a student disclosed their grade if we observe a non-P final grade. We cannot distinguish between students who disclosed a failing grade, or students who elected P/F and failed anyway.

From the Canvas data we observe students scores in their course gradebooks. Canvas, like many web-based platforms, records all user activities, making it possible to reconstruct scores as they stood at any moment. We measure scores at two dates: May 2—just before finals period—and after the semester has ended, the final score. The May 2 score is defined as points awarded as of May 2, divided by total points possible on all assignments due by May 2, on a 0-100 score. The final score measures the percent of total points that a student received at the end of the semester. The final score differs from the May 2 score because of final exams and projects, because of earlier assignments graded after May 2, and because of undated components of the course such as participation. We convert the final score to a final Canvas letter grade using the standard mapping, where for example [80, 83) is a B-, [83, 87) is a B, and [87, 90) is a B+. We convert letter grades to a four-point scale in the standard way, with 2.7 for a B-, 3.0 for a B, 3.3 for an B+, etc.

Derived measures: A key strength of our data is that we can see who discloses in the registrar data, while also seeing non-disclosed grades in the Canvas data. This interpretation of the Canvas data assumes a close concordance between Canvas grades and final grades assigned. We investigate this concordance in [Appendix Figure C.1](#). Canvas and final grades are closely but not perfectly aligned. Overall they agree about three-quarters of the time. However there is considerable heterogeneity across courses. In about 1,400 of the 3,300 courses in our full sample, there is perfect agreement, in the sense that Canvas and registrar grades match exactly for each student in the course. However it is clear that one reason Canvas and registrar grades can disagree is that Canvas scores are rounded up; for example more than half of scores in [89.9, 90) are awarded an A-. Our analysis sample therefore focuses on the 2,519 courses in which the Canvas and Registrar grades never diverge by more than one notch, e.g., B to B+.

We use the Canvas data to develop proxies for students’ signals and the uncertainty of their signal distribution. A natural proxy for the signal would be the May 2 score, as it is known before most P/F deadlines, and predictive of final grades. However the raw May 2 score turns out to be a poor proxy because Canvas scores increase substantially after May 2. (See [Appendix Figure C.3](#).)²⁴ This increase

²⁴This increase could reflect many factors. There may be completed but ungraded assignments—which would count against a student’s May 2 score. There may be undated assignments such as participation which pull up grades. Final assignments might increase grades. Or students may cheat on final exams, which were typically online and not proctored in Spring 2020.

is heterogeneous across classes; in some classes, the May 2 score is a nearly unbiased predictor of final grade, but in others it is severely biased.²⁵ Because students likely know their courses’ specific grading policy—for example, how many points are outstanding, how hard does the final project seem to be, are there ungraded assignments on which they are likely to do well—we can construct a more effective proxy by adjusting the May 2 score for class specific differences. We therefore define students’ signals using a regression framework. To start, we estimate a regression of final Canvas score on May 2 score for student i in course c :

$$FinalScore_{ic} = \beta_{0c} + \beta_{1c}May2_{ic} + e_{ic}. \quad (8)$$

We allow for course-specific slope and intercept because there is likely considerable cross-course heterogeneity in the relationship between May 2 and final scores. Allowing for such heterogeneity increases the predictive power of our signal measure for students’ disclosure decisions; see [Appendix Table C.4](#). To obtain students’ signal, we take the predicted values from this regression and convert them to a four-point scale. For example, if $\hat{\beta}_{0c} = 50$, $\hat{\beta}_{1c} = 0.5$, and $May2_{ic} = 80$, we would have a predicted value of 90, and the signal would be an A- or 3.7. Because this measure is estimated class-by-class, it may not be reliable for small courses. We therefore limit our analysis sample to courses with at least 30 students, our final sample restriction.²⁶

We view this measure as a useful but imperfect proxy for students’ signals. This measure solves the problem of drift, producing unbiased predictions of final score, given May 2 score. It also captures the idea that students know a lot about their courses’ individual grading policies. To establish the validity of our signal measure, we show in [Appendix Table C.4](#) the R^2 from a regression of students’ disclosure decision on their signal, as well as other variables for comparison. The signal strongly predicts disclosure, much more so than student GPA or the raw May 2 score. A downside is that it may contain too much information; it might be unrealistic to assume that students know the exact β s for their courses. On the other hand it also misses some relevant information, as students have private information that may be relevant for forecasting their final performance, for example their overall aptitude. We therefore consider robustness to an alternative proxy which regresses final score on May 2 score as well as student incoming GPA (also with a class-specific coefficient), to capture some student-specific information. Alternatively, we can use students’ actual final canvas grades as proxies for their signals. This proxy clearly contains too much information, but if students have information unobservable to us, then the final grade may be a better proxy for their signal than is their May 2 score. In practice we find that our results are not sensitive to the specific signal proxy that we use.

Our model highlights the uncertainty of the signals as a key force influencing disclosure decisions. We measure uncertainty as the course-specific root mean squared error from [Equation 8](#), which we call $RMSE_c$. Courses will have greater RMSE to the extent that May 2 grades are less predictive of final grades, which reflects the notion of uncertainty in the model. [Appendix Tables C.2 and C.3](#)

²⁵See [Appendix Figure C.2](#), which plots the distribution of course-level prediction errors. There is a spike at zero but a long right-tail, implying that in many classes, May 2 grades are on average equal to final grades, but in some classes, May 2 grades systematically understate final grades.

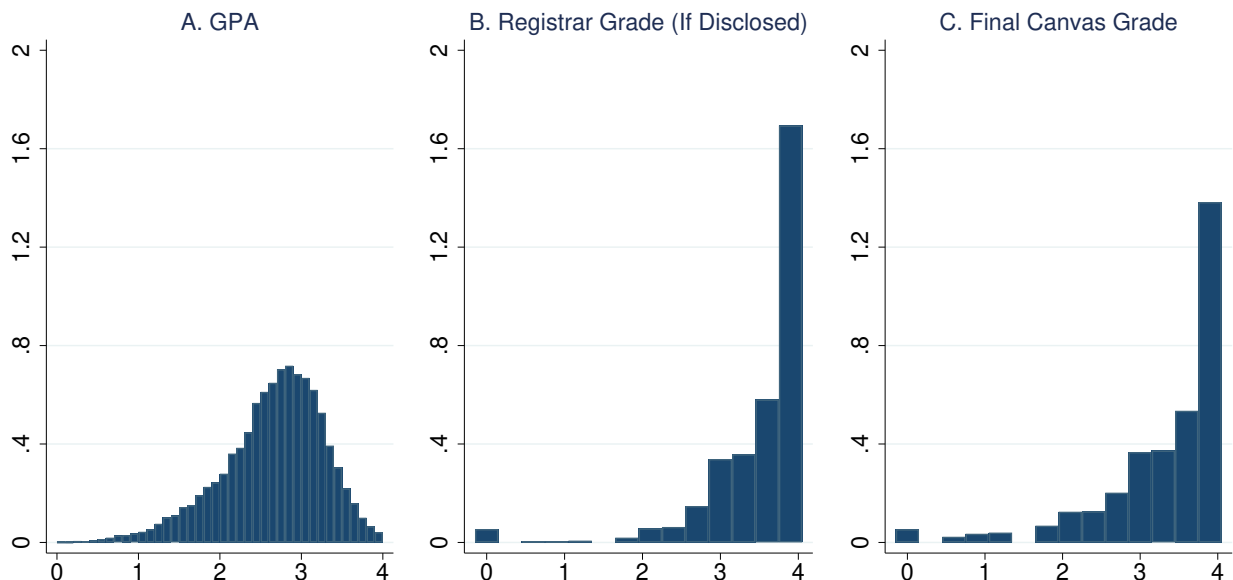
²⁶This sample restriction ends up excluding all the “P/F” only classes, so that our analysis sample consists of classes where students, not instructors, made the P/F decision.

Table 1: Summary statistics

Sample	Full (1)	Analysis (2)
<u>A. Class-level summary statistics</u>		
Course size	33.49 (36.21)	59.79 (41.50)
100-level	0.28 (0.45)	0.22 (0.42)
200-level	0.24 (0.43)	0.28 (0.45)
300-level	0.31 (0.46)	0.37 (0.48)
400-level	0.16 (0.37)	0.13 (0.33)
RMSE	3.60 (3.48)	3.72 (2.20)
# Courses	3,191	778
<u>B. Student-level summary statistics</u>		
GPA	2.63 (0.64)	2.64 (0.64)
# Students	29,315	23,717
<u>C. Student-class-level summary statistics</u>		
Canvas Grade signal	3.28 (0.88)	3.35 (0.82)
Final Canvas Grade	3.28 (0.93)	3.35 (0.86)
Registrar Grade (if disclosed)	3.55 (0.73)	3.55 (0.71)
Disclosed	0.85 (0.36)	0.86 (0.35)
# Observations	106,877	46,515

Notes: Table reports means (standard deviations in parentheses) at the indicated unit of observation. The full sample consists of all enrollments in undergraduate classes in Spring 2020, at Indiana University-Bloomington, in full-term of second-half courses, among students with non-missing GPA and Canvas score information. The analysis sample is further limited students in which Canvas and Registrar grades always agree within one notch, with at least 30 students.

Figure 5: Density of GPA, Registrar Grades, and Final Canvas Grades

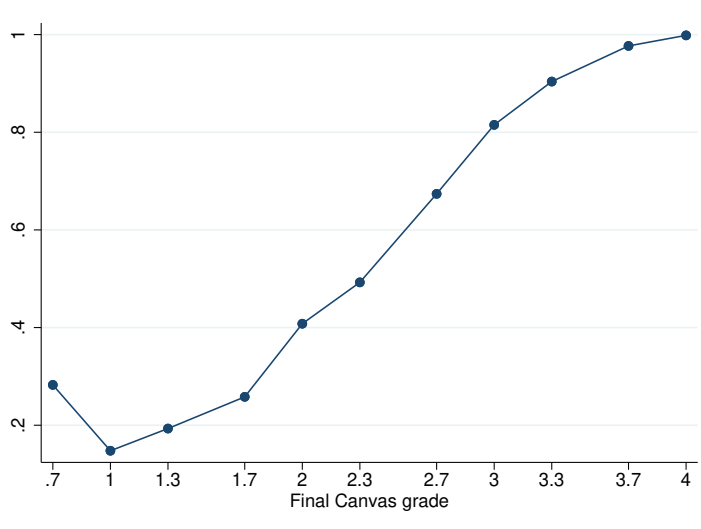


Notes: Figure reports the density of GPA, registrar grades (if disclosed) and final Canvas grades. The sample is the analysis sample as described in the notes to [Table 1](#).

describes the correlates of course-level $RMSE_c$. Uncertainty is greater for classes with lower incoming student GPA, and it varies systematically across schools, with higher uncertainty in the College of Arts and Sciences, and lower uncertainty in the professional schools. However overall we find few strong predictors of $RMSE_c$; over 80 percent of the variation in $RMSE_c$ remains unexplained. We interpret this as indicating that variation in $RMSE_c$ reflects course-specific idiosyncratic factors, such as weight on final exam and the concordance between topics in the first and second half of the course. This idiosyncratic variation is useful for identifying the effect of uncertainty on disclosure. We also consider in robustness checks alternative measure of uncertainty: the course-specific “switch rates,” the fraction of students whose final Canvas grade differs from their signal by at least 1 notch, or at least 2 notches.

Descriptive statistics We report summary statistics for students, courses, and student-course observations in [Table 1](#), for the full sample as well as the analysis sample which is limited to observations in classes with at least 30 students and perfect Canvas-Registrar grade agreement (within one notch). Because our analysis sample focuses on large classes, we end up selecting a quarter of all classes, but about 30 percent of students. The average class in our analysis sample has an RMSE of 3.72. This means that for a given expected final grade—a given signal—final grades vary by about ± 3.72 points, or about one notch. [Figure C.4](#) shows the distribution of RMSE across courses. There is wide variation and two clear modes. In some of our analyses we will divide up the sample into above- and below-median RMSE courses.

Figure 6: Disclosure rates given final grade



Notes: Figure shows the probability of disclosing grades, at each level of final Canvas grade. The sample is the analysis sample as described in the notes to [Table 1](#).

Overall disclosure was high, but this is because final grades were high, not because low grades were always disclosed, as classic models of unravelling predict. Eighty-six percent of grades in our analysis sample were disclosed, and the average disclosed grade is nearly a fully point higher than the average incoming GPA, 3.6 versus 2.6.²⁷ In part these high grades reflect selection—high grades are more likely to be disclosed—but they largely reflect high Canvas grades, as the average final Canvas grade is 3.4. [Figure 5](#) shows the distribution of GPA, Canvas grade, and final grade. Whereas the GPA distribution is centered around its mean, the final Canvas and Registrar grade distributions have a pronounced skew, with a mass at 3.7 and 4 and a long left tail. There is some selection evident in the figures, as A’s and A-’s are over represented among the disclosed grades, relative to the concealed grades. We can see selection more clearly in [Figure 6](#), which shows the disclose rate at each final Canvas grade. Grades of 4.0 and 3.7 are nearly always disclosed, whereas lower grades are much less likely to be disclosed. Importantly, the 86 percent overall disclosure rate masks the fact that lower grades are rare but infrequently disclosed. Our model argues that one reason for this non-disclosure is risk aversion and uncertainty over final grades. We next turn to testing this model.

5 Testing the Blind Disclosure Model

5.1 Approach

We test the prediction of the blind disclosure model that a given signal is less likely to be disclosed if it is drawn from a less precise distribution. To test this prediction we estimate linear probability

²⁷GPA is the average of prior semesters’ grades, among returning student, and so provides a natural benchmark.

models of the following form for the probability that student i discloses her grade in course c :

$$Disclose_{ic} = \alpha + \beta RMSE_c + \sum_s \theta_s 1\{signal_{ic} = s\} + X_{ic}\gamma + \mu_i + \varepsilon + ic. \quad (9)$$

Here our interest is on β , the coefficient on $RMSE_c$, course c 's root mean squared prediction error, which is our proxy for uncertainty, although in robustness tests we consider alternative proxies. The model implies that $\beta < 0$. We control for a set of dummy variables for each signal S , based on May 2 grade, although we consider robustness to alternative signal measures and specifications. We consider a range of specifications. In some specifications we control for course- and student-level observables X_{ic} . These are the course- and student-level variables listed in [Table 1](#), plus a set of dummy variables for the 14 academic groups listed in [Appendix Table C.3](#). Our most tightly controlled specifications account for course-level observables and student fixed effects, μ_i . Because $RMSE_c$ varies across courses, we cluster our standard errors at the course level.

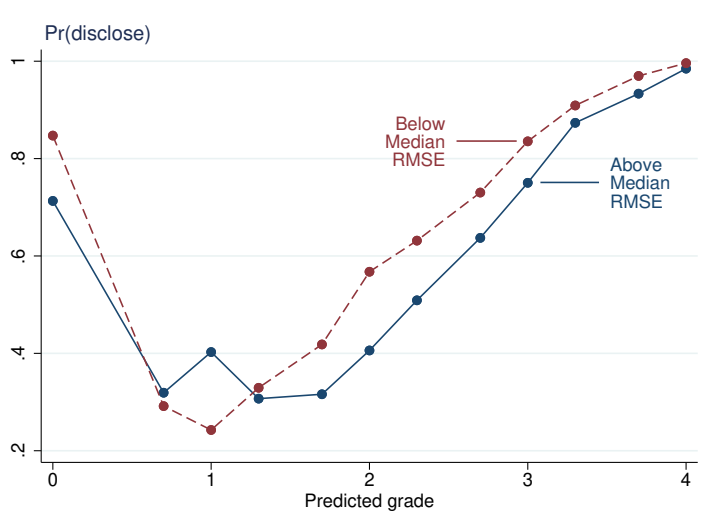
These specifications identify the effect of uncertainty on disclosure decisions by comparing the same student in different classes with different uncertainties, adjusting for the signal they receive and the observable class characteristic. The key identifying assumption is that classes with more uncertainty in their signals do not have other, systematic, unobserved reasons for non-disclosure. This assumption is fairly weak. Although there are likely classes where students would like to disclose doing well (e.g. core classes in their major), these are likely also the classes where students would most like to conceal a poor grade, and hence these classes should not necessarily lead to high disclosure. Indeed it is reassuring to us that few variables strongly predict $RMSE_c$, and controlling for observable course characteristics does not affect our estimates. While it is likely that $RMSE_c$ is measured with error, this should bias us against finding an effect of uncertainty on disclosure. Another possible concern that we discuss and rule out below is reverse causality, that a course's non-disclosure rate may influence its uncertainty.

5.2 Results

Graphical evidence: We begin by showing graphical evidence that uncertainty is associated with an increase in the probability that a given signal is disclosed. [Figure 7](#) shows the probability that a given signal is disclosed, separately for high and low uncertainty classes, with high uncertainty classes ones with above-median $RMSE_c$. At each signal between 1.3 and 3.7, we see clear separation between the high- and low-uncertainty classes, with higher disclosure in the lower uncertainty classes, exactly as the model predicts. Very high signals are nearly always disclosed, regardless of uncertainty. Thus the graphical evidence clearly supports the blind disclosure model except at very low signals. (A signal of 1.3 is about the fourth percentile.) However the higher disclosure rate in uncertain classes of low signals is difficult to interpret because students cannot conceal a final grade of 0 (i.e. an F), and greater uncertainty makes it more likely that a low but passing signal is in fact a final grade F.

Regression results: The graphical evidence does not adjust for observable differences across classes or students that might be correlated with disclosure decisions and signals, nor does it use the

Figure 7: Disclosure rates are lower in high-uncertainty classes



Notes: Figure shows the probability that a given signal (i.e. predicted grade) is disclosed, separately for high uncertainty classes (with above median signal RMSE) and low uncertainty classes. Note that Final grades of F are always disclosed because they cannot be concealed. The sample is the analysis sample as described in the notes to [Table 1](#).

full, continuous variation in $RMSE_c$. We therefore estimate variants of [Equation 9](#), and report the results in [Table 2](#). We begin in Column (1) by showing results that only adjust for signal dummy variables. We estimate that a 1-unit increase in $RMSE$ reduces disclosure probability by about 0.7 percentage points. The estimate is statistically significant. In columns (2)-(4) we add an increasing stringent set of controls. Adding course-level controls in column (2) changes the coefficient only slightly. Further adding student GPA has a slight effect in column (3), and adding student fixed effects (instead of GPA) in column (4) has no further effect.²⁸ We view column (4) as our preferred specification because it is most tightly controlled. Overall, however, our estimates are not sensitive to the controls we use (or indeed whether we control for anything other than the signal itself). We emphasize that these controls are relevant predictors of disclosure decisions; our model R^2 increases substantially when we add them, and our standard errors fall. But they are not correlated with $RMSE_c$. This provides some reassurance (in the spirit of [Altonji et al. \(2005\)](#); [Oster \(2019\)](#)) that unobserved factors are not simultaneously driving $RMSE_c$ and $disclosure_{ic}$. Thus overall we find strong evidence for the prediction that students are less likely to disclose a given signal when it comes from a high-uncertainty distribution.

The model makes a further prediction that the effect of uncertainty on disclosure probabilities is “hump-shaped:” it is greatest for intermediate signals, and lowest for very high and very low signals. To test this prediction, we define an intermediate grade as one between 1.7 and 3.7—roughly one point below and above average GPA. The graphical evidence supports this prediction as well; [Figure 7](#)

²⁸When we include these student fixed effects, we drop students who appear only once in our main sample.

Table 2: Uncertainty reduces disclosure

	(1)	(2)	(3)	(4)	(5)
RMSE	-0.007 (0.002)	-0.006 (0.002)	-0.007 (0.002)	-0.007 (0.001)	-0.002 (0.001)
RMSE \times intermediate signal					-0.008 (0.002)
R^2	0.267	0.292	0.294	0.644	0.644
# Observations	46,514	46,514	46,514	37,208	37,208
# Courses	778	778	778	778	778
Course observables	No	Yes	Yes	Yes	Yes
Student GPA	No	No	Yes	NA	NA
Student fixed effects	No	No	No	Yes	Yes

Notes: Table reports coefficients from a linear probability model of disclosure on course RMSE (Our proxy for uncertainty) as well as the indicated controls. Signal controls are a set of dummy variables for each grade signal (A, A-, B+, etc.). Course observables are dummy variables for course level (200/300/400) and academic unit (Business School, etc.), as well as class size and average incoming GPA of students enrolled in the class. Student GPA is redundant with student fixed effects and hence excluded in columns (4) and (5). The sample is the analysis sample as defined in the notes to [Table 1](#). Robust standard errors, clustered on course, in parentheses.

shows the biggest differences between below-median and above-median $RMSE_c$ courses at intermediate signals. We test this prediction in our regression framework by including an interaction between $RMSE_c$ and “intermediate” signal. We report the results in column (5) of [Table 2](#). We find that the coefficient is significant and positive, implying that indeed uncertainty has the greatest effect on the probability that an intermediate signal is disclosed.

Quantitative interpretation: We find clear qualitative evidence for the blind disclosure model, in the sense that the effect of uncertainty on disclosure is directionally consistent with the model’s implications. This is important because our main empirical goal is to show that blind disclosure is one reason for incomplete unravelling; we do not argue that it is the only or primary reason. However, the estimates also imply a quantitatively important role for uncertainty. Our point estimate of 0.007 implies that, for a given student, moving from average uncertainty of 3.7 to no uncertainty would raise disclosure by about 2.6 percentage points. This is a meaningful share of the overall non-disclosure rate of 14 percent. We interpret this as the direct effect of uncertainty itself—the combined “risk aversion” and “mean-reversion” effects—rather than an equilibrium effect (which also includes the “sorting” effect) where uncertainty also influences receiver’s beliefs about the meaning of a non-disclosed grade. We think the direct effect is a more plausible interpretation for this estimate because we identify the effect of uncertainty by comparing students in different classes but at the same institution (indeed often the same student). It is unlikely that employers know which classes are high or low uncertainty, although it is likely that they know there are differences across classes. Hence our comparison likely

holds fixed beliefs but varies uncertainty in a single equilibrium, and gives us the direct effect. The equilibrium effects of uncertainty are likely to be even greater.

Robustness: Overall we find that uncertainty reduces the probability of disclosing a given signal, supporting the blind disclosure model. This finding is robust to alternative definitions of uncertainty, alternative ways of measuring students’ signals, and alternative samples. We show robustness to alternative uncertainty measures in [Appendix Table C.5](#). Column (1) reports our baseline specification, which measures uncertainty with the course-specific RMSE of May 2 grade as a predictor for final grade. In columns (2) and (3) we use an alternative measure which is based on the proportion of grades which change. Specifically in column (2) our measure of uncertainty is the proportion of students whose final grade differs from their May 2 grade by at least one grade notch, and in column (3) the measure is the proportion that differ by at least two notches. These measures address the concern that our underlying uncertainty measure reflects uncertainty in the continuous grade, but letter grade uncertainty may be especially important for students. In column (4) we measure uncertainty more coarsely, with an indicator for “above-median RMSE” (defined at the course level), as in Figure 7. These alternative measures imply that eliminating uncertainty would increase disclosure 2-4 percentage points (for a single student, holding fixed beliefs), similar to our baseline estimate.

We show robustness to alternative signal measures in [Appendix Table C.6](#). In our baseline specification in column (1), the signal measure is a set of dummy variables for the predicted final letter grade, given May 2 grade, with course-specific coefficients. Discretizing the predicted grade and turning it into a set of dummy variables may exacerbate measurement error problems (as the signal itself is measured with error). In columns (2) and (3) we drop the dummy variables for predicted grade and instead include linear or cubic terms in predicted final grade. A separate issue is that our signal misses some information available to students. We enrich the information set in column (2). To do so, we modify the signal measure so that each students’ prediction depends on their incoming GPA as well as their May 2 grade.²⁹ When we work with this alternative signal, we also use as an RMSE measure the RMSE from the regression of final grade on May 2 grade and student GPA. In column (5) we consider an alternative proxy for students’ signal: their actual final canvas grade, entered as a set of dummy variables. This proxy is surely imperfect but it may be better than our predicted final grade measure (as a proxy for students’ signals) because students may have private information about their final grade. We find a slightly smaller coefficient on uncertainty here. In column (6) we continue to use the final grade as a proxy for students’ signals, but we round up final grades above .5 to the next letter grade (so an 89.5 is a B+ in column (5) but an A- in column (6)). These alternative ways of handling students’ signals measures have very little effect on our estimates.

As a final set of checks, we show in [Appendix Table C.7](#) robustness to alternative samples. Our baseline sample, in column (1), is restricted to classes with at least 30 students, in which canvas and registrar grades agree to within one notch. In column (2), we further restrict the sample to classes in which there is no evidence of “bunching” in the grade distribution. Specifically, we exclude classes in

²⁹That is, we regress final canvas grade on May 2 grade and student-level incoming GPA, with course-specific coefficients.

which at least 1 percent of students have a final Canvas grade that is an exact multiple of 10. In classes with this exact bunching, our signal and precision measures may be less reliable. In column (3) we exclude classes in which a majority of grades are an A. For these classes there may be little uncertainty (despite our measure) as students might be able to forecast their final grade. In column (4) we expand the sample to include classes of all sizes (dropping the requirement that classes contain at least 30 students). In column (5) we tighten this condition to require an exact match for all students, and in column (6) we instead require that either Canvas-Registrar grades match exactly, or that rounded-up Canvas grades match exactly.³⁰ Overall our estimates are not sensitive to the exact sample inclusion criteria. We acknowledge that our estimates would change substantially if we included classes in which the Canvas and Registrar grades disagreed. However in these classes our signal and uncertainty measures are both likely highly error-prone, and so is not particularly concerning to us that our results would change if we included these classes.

Moral hazard: Our analysis so far has abstracted from moral hazard, which in our context refers to the possibility that not disclosing a grade reduces effort incentives, resulting in lower grades. Our personal experience indicates that such moral hazard is likely present in our context. Nonetheless moral hazard is unlikely to explain our main finding, the correlation between class-level uncertainty about final grades and disclosure decisions. Our model, when extended to accommodate moral hazard, implies that moral hazard reduces the level of disclosure and grades, but does not affect the disclosure-uncertainty correlation. Our model thus implicitly controls for moral hazard by including a set of dummy variables for all grades. It is also reassuring that our main results are robust to using realized rather than predicted grades as proxies for student expectations, as realized grades might be especially sensitive to moral hazard, to the extent that students made their final effort choices after making their disclosure choices.

6 Conclusion

We have studied a disclosure model in which the sender must make a decision to disclose information before observing it perfectly. In the presence of risk aversion, such uncertainty leads to incomplete unraveling in equilibrium, with higher uncertainty associated with less disclosure. The model’s main prediction is supported by empirical tests using data on students’ grade disclosure choices in a university setting.

We leave several extensions of the model and empirical questions for future work. One such extension would allow senders to have a additional private information about their ability. In such a model, a student’s decision to disclose a grade would be informative beyond the grade itself, introducing a signaling channel. Another extension would allow senders to simultaneously disclose grades in multiple courses, so that the sender’s decision to disclose one grade would be linked to her private signals about other grades. A further empirical question is how employers and graduate schools react

³⁰So for example if a student has an 89.5 on Canvas and an A- registrar grade, we consider it a match in this specification, but if she has an 89.4 and an A-, we do not.

to undisclosed grades.

References

- Altonji, Joseph G, Todd E Elder, and Christopher R Taber**, “Selection on observed and unobserved variables: Assessing the effectiveness of Catholic schools,” *Journal of political economy*, 2005, *113* (1), 151–184.
- Anderson, Shelby**, “Satisfactory/fail grading now an option for IU-Bloomington spring courses,” *Indiana Daily Student*, March 2020. <https://www.idsnews.com/article/2020/03/satisfactoryfail-grading-now-an-option-for-iu-bloomington-spring-courses>.
- Bederson, Benjamin B., Ginger Zhe Jin, Phillip Leslie, Alexander J. Quinn, and Ben Zou**, “Incomplete Disclosure: Evidence of Signaling and Countersignaling,” *American Economic Journal - Microeconomics*, 2018, *10* (1), 41–66.
- Billings, Stephen B and Kevin T Schnepel**, “The value of a healthy home: Lead paint remediation and housing values,” *Journal of Public Economics*, 2017, *153*, 69–81.
- Board, Oliver**, “Competition and disclosure,” *The Journal of Industrial Economics*, 2009, *57* (1), 197–213.
- Boleslavsky, Raphael and Christopher Cotton**, “Grading standards and education quality,” *American Economic Journal: Microeconomics*, 2015, *7* (2), 248–79.
- Bollinger, Bryan, Phillip Leslie, and Alan Sorensen**, “Calorie posting in chain restaurants,” *American Economic Journal: Economic Policy*, 2011, *3* (1), 91–128.
- Bond, Philip and Yao Zeng**, “Silence is Safest: Non-Disclosure When the Audience’s Preferences are Uncertain,” *American Economic Journal: Microeconomics*, 2019. Working paper.
- Brown, Alexander L., Colin F. Camerer, and Dan Lovallo**, “To Review or Not to Review? Limited Strategic Thinking at the Movie Box Office,” *American Economic Journal: Microeconomics*, 2012, *4*, 1–26.
- Chakraborty, Archishman and Rick Harbaugh**, “Comparative Cheap Talk,” *Journal of Economic Theory*, 2007, *132*, 70–94.
- Chen, Ying**, “Value of public information in sender–receiver games,” *Economics Letters*, 2012, *114* (3), 343–345.
- Costrell, Robert M.**, “A Simple Model of Educational Standards,” *American Economic Review*, 1994, *84*, 956–971.
- Darden, Michael and Ian M McCarthy**, “The star treatment estimating the impact of star ratings on Medicare Advantage enrollments,” *Journal of Human Resources*, 2015, *50* (4), 980–1008.
- Davies, Paul J.**, “How Wirecard went from tech star to bankrupt,” *Wall Street Journal*, July 2020.
- Dranove, David and Ginger Zhe Jin**, “Quality Disclosure and Certification: Theory and Practice,” *Journal of Economic Literature*, December 2010, *48* (4), 935–963.
- , **Daniel Kessler, Mark McClellan, and Mark Satterthwaite**, “Is more information better? The effects of “report cards” on health care providers,” *Journal of political Economy*, 2003, *111* (3), 555–588.

- Draughon, Luzane**, “A guide to S/F grading policies at IU’s various undergraduate schools,” *Indiana Daily Student*, April 2020. <https://www.idsnews.com/article/2020/04/a-guide-to-sf-grading-policies-at-ius-various-undergraduate-schools>.
- Dubey, Pradeep and John Geanakoplos**, “Grading Exams: 100, 99, 98,... or a, b, c?,” *Games and Economic Behavior*, 2010, *69* (1), 72–94.
- Dye, Ronald A.**, “Disclosure of Nonproprietary Information,” *Journal of Accounting Research*, 1985, *23*, 123–145.
- Feltovich, Nick, Richmond Harbaugh, and Ted To**, “Too Cool for School? Signalling and Countersignalling,” *RAND Journal of Economics*, Winter 2002, *33* (4), 630–649.
- Grossman, Sanford J.**, “The informational role of warranties and private disclosure about product quality,” *The Journal of Law and Economics*, 1981, *24* (3), 461–483.
- Houde, Sébastien**, “How consumers respond to product certification and the value of energy information,” *The RAND Journal of Economics*, 2018, *49* (2), 453–477.
- Hurwitz, Michael and Jonathan Smith**, “Student responsiveness to earnings data in the College Scorecard,” *Economic Inquiry*, 2018, *56* (2), 1220–1243.
- Jin, Ginger Zhe and Phillip Leslie**, “The effect of information on product quality: Evidence from restaurant hygiene grade cards,” *The Quarterly Journal of Economics*, 2003, *118* (2), 409–451.
- Jovanovic, Boyan**, “Truthful disclosure of information,” *The Bell Journal of Economics*, 1982, pp. 36–44.
- Lubensky, Dmitry and Eric Schmidbauer**, “Free Product Trials: Disclosing Quality and Match Value,” *Economic Inquiry*, 2020, *58* (4), 1565–1576.
- Luca, Michael and Jonathan Smith**, “Salience in quality disclosure: Evidence from the US News college rankings,” *Journal of Economics & Management Strategy*, 2013, *22* (1), 58–77.
- Milgrom, Paul**, “Good News and Bad News: Representation Theorems and Applications,” *Bell Journal of Economics*, 1981, *12*, 380–391.
- Morris, Stephen and Hyun Song Shin**, “Social value of public information,” *American Economic Review*, 2002, *92* (5), 1521–1534.
- Oster, Emily**, “Unobservable selection and coefficient stability: Theory and evidence,” *Journal of Business & Economic Statistics*, 2019, *37* (2), 187–204.
- Ostrovsky, Michael and Michael Schwarz**, “Information Disclosure and Unraveling in Matching Markets,” *American Economic Journal: Microeconomics*, 2010, *2* (2), 34–63.
- Outram, Simon and Bob Stewart**, “Doping through supplement use: a review of the available empirical data,” *International journal of sport nutrition and exercise metabolism*, 2015, *25* (1), 54–59.
- Quick, Joshua, Ben Motz, and Daniel W. Sacks**, “DATA EXTRACT: Spring 2020 Grade Choice Study,” 2020. <https://doi.org/10.5967/76ek-3d02.IN>.
- Sampford, Michael R.**, “Some inequalities on Mill’s ratio and related functions,” *The Annals of Mathematical Statistics*, 1953, *24* (1), 130–132.

Shin, Hyun Song, “Disclosures and Asset Returns,” *Econometrica*, 2003, *71*, 105–134.

Viscusi, W. Kip, “A Note on ‘Lemons’ Markets with Quality Certification,” *Bell Journal of Economics*, 1978, *9*, 277–279.

A Theory Appendix

Proof of Proposition 1. (i) In any equilibrium, the student receives some known wage $w_c \in \mathbb{R}$ after choosing *Conceal*. Recalling the mean-variance representation for CARA preferences and the simplifications in the main text resulting in (5), the student is indifferent iff

$$\begin{aligned} & \mathbb{E}[\mathbb{E}[\alpha|\gamma]|S = s] - \frac{\lambda}{2} \text{Var}[\mathbb{E}[\alpha|\gamma]|S = s] = w_c \\ \iff & \frac{\tau_\alpha \mu_\alpha + \tau_S s}{\tau_\alpha + \tau_S} - \frac{\lambda}{2} \frac{\tau_S}{\tau_\alpha + \tau_S} \frac{\tau_\gamma}{\tau_\epsilon (\tau_\alpha + \tau_\gamma)} = w_c \end{aligned}$$

The LHS of the second equation above is strictly increasing in s , and hence there is a unique signal value $s^* \in \mathbb{R}$, below (above) which students play *Conceal* (*Disclose*).

(ii) First, observe that the lhs of (7) has limit $+\infty$ as $z \rightarrow +\infty$ while the rhs has limit 0. Second, from known properties (Sampford, 1953) of the function $g(z) := \frac{\phi(z)}{\Phi(z)}$, $0 > g'(z) > -1$, so the lhs and rhs can cross at most once. Third, $z + g(z) \rightarrow 0$ as $z \rightarrow -\infty$,³¹ so the lhs and rhs do cross given that $\lambda > 0$. We conclude that (7) has a unique solution z^* , and hence there is a unique equilibrium with cutoff $s^* = \nu z^* + \mu_\alpha$. ■

Proof of Proposition 2. (i) On the left hand side of (7), $\nu \tau_\epsilon$ (where ν is a function of τ_ϵ) is increasing in τ_ϵ , and hence the left hand side of (7) is increasing in τ_ϵ . Since the left hand side of (7) crosses the right hand side from below at z^* , it follows that z^* is decreasing in τ_ϵ , and thus the probability of disclosure is increasing in τ_ϵ . Now $z + \frac{\phi(z)}{\Phi(z)}$ is positive, strictly increasing, tends to 0 as $z \rightarrow -\infty$, and tends to $+\infty$ as $z \rightarrow +\infty$. As $\tau_\epsilon \rightarrow 0$, $\frac{\lambda}{2} \frac{\tau_\gamma}{\nu(\tau_\alpha + \tau_\gamma)\tau_\epsilon} \rightarrow +\infty$ and hence $z^* \rightarrow +\infty$; similarly, as $\tau_\epsilon \rightarrow +\infty$, $\frac{\lambda}{2} \frac{\tau_\gamma}{\nu(\tau_\alpha + \tau_\gamma)\tau_\epsilon} \rightarrow 0$ and hence $z^* \rightarrow -\infty$.

(ii) Conditional on γ , the probability that a student discloses, parameterized by τ_ϵ , is $\Pr(S > s^*(\tau_\epsilon)|\gamma) = 1 - \Phi((s^*(\tau_\epsilon) - \gamma)\sqrt{\tau_\epsilon})$. Now consider any two values τ_ϵ^L and τ_ϵ^H . It is easy to see that $\gamma \mapsto (s^*(\tau_\epsilon^L) - \gamma)\sqrt{\tau_\epsilon^L}$ and $\gamma \mapsto (s^*(\tau_\epsilon^H) - \gamma)\sqrt{\tau_\epsilon^H}$ are two lines and the first intersects the second once from below. Since Φ is strictly increasing, it follows that $\Pr(S > s^*(\tau_\epsilon^L)|\gamma)$ intersects $\Pr(S > s^*(\tau_\epsilon^H)|\gamma)$ exactly once, from above.

(iii) Specifically, we show that for any two values τ_ϵ^L and τ_ϵ^H , there exist thresholds $\gamma^b < \underline{\gamma} < \bar{\gamma} < \gamma^\sharp$ such that conditional on γ , the change in the probability that the student discloses between $\tau_\epsilon = \tau_\epsilon^L$ and $\tau_\epsilon = \tau_\epsilon^H$ is strictly greater for any $\gamma \in [\underline{\gamma}, \bar{\gamma}]$ than for any $\gamma \notin [\gamma^b, \gamma^\sharp]$.

Consider two values of τ_ϵ : $\tau_\epsilon^L < \tau_\epsilon^H$. The difference in the probability of disclosure (or concealment) conditional on γ is $D(\tau_\epsilon^H, \tau_\epsilon^L; \gamma) := |\Phi((s^*(\tau_\epsilon^H) - \gamma)\sqrt{\tau_\epsilon^H}) - \Phi((s^*(\tau_\epsilon^L) - \gamma)\sqrt{\tau_\epsilon^L})|$. As argued in part (ii), there is exactly one value of γ , call it γ^* , for which $D(\tau_\epsilon^H, \tau_\epsilon^L; \gamma) = 0$. For any $\underline{\gamma}, \bar{\gamma}$ with $\gamma^* < \underline{\gamma} < \bar{\gamma}$, since D is continuous in γ , D is bounded away from 0 on this interval: there exists $\eta > 0$ such that for all $\gamma \in [\underline{\gamma}, \bar{\gamma}]$, $D(\tau_\epsilon^H, \tau_\epsilon^L; \gamma) > \eta > 0$. On the other hand, for any τ_ϵ , $\Phi((s^*(\tau_\epsilon) - \gamma)\sqrt{\tau_\epsilon})$ converges to 0 as $\gamma \rightarrow +\infty$ and converges to 1 as $\gamma \rightarrow -\infty$. It follows that $D(\tau_\epsilon^L, \tau_\epsilon^H; \gamma) \rightarrow 0$ as $\gamma \rightarrow \pm\infty$, and thus there exist $\gamma^b, \gamma^\sharp \in \mathbb{R}$ with $\gamma^b < \underline{\gamma} < \bar{\gamma} < \gamma^\sharp$ such that for all $\gamma \notin [\gamma^b, \gamma^\sharp]$, $D(\tau_\epsilon^L, \tau_\epsilon^H; \gamma) < \eta$. Hence, for all $\gamma' \in [\underline{\gamma}, \bar{\gamma}]$ and $\gamma'' \notin [\gamma^b, \gamma^\sharp]$, $D(\tau_\epsilon^L, \tau_\epsilon^H; \gamma') > D(\tau_\epsilon^L, \tau_\epsilon^H; \gamma'')$.

(iv)-(vii) By straightforward differentiation, the left hand side of (7) is decreasing in λ and τ_γ and independent of μ_α ; then by a similar argument to that in part (i), the probability of disclosure is decreasing in λ and τ_γ and independent of μ_α . Similarly, differentiation wrt τ_α yields that the left hand side of (7) is increasing in τ_α iff $\tau_\gamma \tau_\epsilon (\tau_\alpha - \tau_\gamma) + 2\tau_\alpha^2 (\tau_\gamma + \tau_\epsilon) > 0$. The left hand side of this inequality is strictly increasing, is strictly negative when $\tau_\alpha = 0$ and is strictly positive for sufficiently

³¹This claim can be established by applying L'Hopital's rule twice: $\lim_{z \rightarrow -\infty} (z + g(z)) = \lim_{z \rightarrow -\infty} \frac{z\Phi(z) + \phi(z)}{\Phi(z)} = \lim_{z \rightarrow -\infty} \frac{z\phi(z) + \Phi(z) - z\phi(z)}{\phi(z)} = \lim_{z \rightarrow -\infty} \frac{\phi(z)}{-z\phi(z)} = 0$.

large τ_α . It follows that z^* is nonmonotonic and single-peaked in τ_α . When this inequality holds, the probability of disclosure is increasing in τ_α . ■

Proof of Proposition 3. The mean squared error is

$$\begin{aligned}\mathbb{E}[(\alpha - \mathbb{E}^m[\alpha])^2] &= \Pr(C)\mathbb{E}[(\alpha - \mathbb{E}^m[\alpha])^2|C] + \Pr(D)\mathbb{E}[(\alpha - \mathbb{E}^m[\alpha])^2|D] \\ &= \Pr(C)\mathbb{E}[(\alpha - \mathbb{E}[\alpha|C])^2|C] + \Pr(D)\mathbb{E}[(\alpha - \mathbb{E}[\alpha|\gamma])^2|D] \\ &= \Pr(C)\text{Var}[\alpha|C] + \Pr(D)\mathbb{E}[\text{Var}[\alpha|\gamma]|D].\end{aligned}\tag{10}$$

Recall that $\alpha|S \sim N\left(\frac{\tau_\alpha\mu_\alpha + \tau_S S}{\tau_\alpha + \tau_S}, 1/(\tau_\alpha + \tau_S)\right)$. Hence, by the law of total variance,

$$\begin{aligned}\text{Var}[\alpha|C] &= \text{Var}[\mathbb{E}[\alpha|S]|C] + \mathbb{E}[\text{Var}[\alpha|S]|C] \\ &= \text{Var}\left[\frac{\tau_\alpha\mu_\alpha + \tau_S S}{\tau_\alpha + \tau_S}|C\right] + 1/(\tau_\alpha + \tau_S) \\ &= \left(\frac{\tau_S}{\tau_\alpha + \tau_S}\right)^2 \text{Var}[S|C] + 1/(\tau_\alpha + \tau_S).\end{aligned}$$

Moreover, $\mathbb{E}[\text{Var}[\alpha|\gamma]|D] = \text{Var}[\alpha|\gamma] = 1/(\tau_\alpha + \tau_\gamma)$. Substituting these into (10) yields

$$\begin{aligned}\mathbb{E}[(\alpha - \mathbb{E}^m[\alpha])^2] &= \Pr(C) \left\{ \left(\frac{\tau_S}{\tau_\alpha + \tau_S}\right)^2 \text{Var}[S|C] + 1/(\tau_\alpha + \tau_S) \right\} + (1 - \Pr(C))/(\tau_\alpha + \tau_\gamma) \\ &= \Phi(z^*) \left\{ \left(\frac{\tau_S}{\tau_\alpha + \tau_S}\right)^2 \underbrace{\nu^2[1 - z^*g(z^*) - g(z^*)^2]}_{=:Q(z^*)} + 1/(\tau_\alpha + \tau_S) \right\} \\ &\quad + (1 - \Phi(z^*))/(\tau_\alpha + \tau_\gamma) \\ &=: V(z^*, \tau_\epsilon),\end{aligned}\tag{11}$$

where $g(z^*) := \phi(z^*)/\Phi(z^*)$. Observe that $Q(z^*) = 1 + g'(z^*)$. By standard results, (i) $g'(z^*) \in (-1, 0)$, so $Q(z^*) \in (0, 1)$, and (ii) $g''(z^*) > 0$, so $Q'(z^*) > 0$.³²

Recall from the proof of Proposition 2 that z^* is decreasing in τ_ϵ . Hence, to prove the proposition at hand, it suffices to show that $V(z^*, \tau_\epsilon)$ is increasing in z^* (the effect of increased concealment) and (for fixed z^*) is decreasing in τ_ϵ (sorting is noisier when uncertainty is higher). The term in braces in (11) is greater than $1/(\tau_\alpha + \tau_\gamma)$ since $Q > 0$ and $\tau_\gamma > \tau_S$, and $\Phi(z^*)$ and $Q(z^*)$ are increasing in z^* , so $V(z^*, \tau_\epsilon)$ increasing in z^* . Next, partially differentiating wrt τ_ϵ (and recalling that ν is a function of τ_ϵ) yields $\frac{\partial V}{\partial \tau_\epsilon} = -\frac{\tau_\gamma^2(1-Q(z^*))}{(\tau_\gamma + \tau_\epsilon + \tau_\alpha(\tau_\gamma + \tau_\epsilon))^2} < 0$. We conclude that the mean squared error in equilibrium is decreasing in τ_ϵ .

For the second part of the proposition, we can first calculate the MSE conditional on the signal μ_α and then integrate over μ_α . But since the MSE conditional on μ_α is independent of μ_α , we can simply fix an arbitrary μ_α and calculate MSE as in (11).

Note that $MSE(\tau_\alpha) \leq 1/(\tau_\alpha) \rightarrow 0$ as $\tau_\alpha \rightarrow \infty$, where $1/\tau_\alpha$ is the MSE if the student always conceals. Next, we differentiate (11) with respect to τ_α , using implicit differentiation of (7) with respect to τ_α to eliminate $\frac{d}{d\tau_\alpha} z^*(\tau_\alpha)$. We then use (7) to eliminate λ , yielding an expression for $\frac{d}{d\tau_\alpha} MSE(\tau_\alpha)$ of the form $M(\tau_\alpha, \tau_\epsilon, \tau_\gamma, z^*(\tau_\alpha, \tau_\epsilon, \tau_\gamma, \lambda))$. Note that fixing any $(\tau_\epsilon, \tau_\gamma) \in \mathbb{R}_{++}^2$, for each $\tau_\alpha \in (0, \infty)$ there exists $\lambda(\tau_\alpha) > 0$ such that $z^* = 0$. In the simplified expression for $M(\tau_\alpha, \tau_\epsilon, \tau_\gamma, z^*(\tau_\alpha, \tau_\epsilon, \tau_\gamma, \lambda(\tau_\alpha)))$, the denominator is strictly positive, and the numerator is strictly positive for sufficiently small τ_α .

³²See Sampford (1953) for the properties of $\phi(z)/(1 - \Phi(z)) = g(-z)$.

Hence, there exist parameter values for which $\frac{d}{d\tau_\alpha} MSE(\tau_\alpha) > 0$.³³ For such parameter values, using fact that $MSE(\tau_\alpha) \rightarrow 0$ as $\tau_\alpha \rightarrow \infty$, MSE is nonmonotonic in τ_α . ■

Proposition 9 *The equilibrium cutoff s^* is strictly increasing in τ_ϵ .*

Proof. Write (7) as $z - K/(\nu\tau_\epsilon) = -g(z)$. Let us use prime notation to denote derivatives wrt τ_ϵ . Differentiation yields $z' + K\frac{(\nu\tau_\epsilon)'}{(\nu\tau_\epsilon)^2} = -g'z' = g(g+z)z'$, which implies $z' = -\frac{K\frac{(\nu\tau_\epsilon)'}{(\nu\tau_\epsilon)^2}}{1-g(g+z)}$. Note that $\nu' = -\frac{1}{2\nu\tau_\epsilon^2}$. Using (7) to eliminate K implies $z' = -\frac{(z+g)\frac{(\nu\tau_\epsilon)'}{(\nu\tau_\epsilon)}}{1-g(g+z)}$. Hence

$$\begin{aligned} s' &= z'\nu + \nu'z = -\frac{(z+g)(\nu\tau_\epsilon)'}{\tau_\epsilon(1-g(g+z))} - \frac{z}{2\nu\tau_\epsilon^2} \\ &= -\frac{(z+g)\left[-\frac{1}{2\nu\tau_\epsilon} + \nu\right]}{\tau_\epsilon(1-g(g+z))} - \frac{z}{2\nu\tau_\epsilon^2} \\ &= \frac{(z+g)\left[1-2\nu^2\tau_\epsilon\right] - z\left[1-g(g+z)\right]}{2\nu\tau_\epsilon^2(1-g(g+z))}. \end{aligned}$$

The denominator in the last line above is positive, while the numerator can be expressed as $z+g-(z+g)2\nu^2\tau_\epsilon - z+gz(z+g)$, which is bounded above by $z+g-(z+g)2 - z+gz(z+g) = (z+g)[gz-1] - z$, where we have used that $\nu^2\tau_\epsilon > 1$ and $z+g > 0$. Hence, our goal is to show that for all $z \in \mathbb{R}$, $h(z) := (z+g)[gz-1] - z < 0$. Now with repeated use of L'Hopital's rule, one can show that $\lim_{z \rightarrow -\infty} h(z)/\phi(z) = -\infty$, and hence for sufficiently low z , $h(z) < 0$. Moreover, one can show that if $h(z) = 0$, then $h'(z) < 0$.³⁴ Hence, if h ever crosses 0, it must be strictly positive to the left of the crossing point. We conclude that h never crosses 0 and is strictly negative everywhere. ■

Proof of Proposition 4. Suppose the market pays a wage w if the student plays *Conceal*. We begin by expressing the student's best response cutoffs for each precision level as a function of w and writing the equilibrium existence condition as a fixed point equation in w . Recall from the proof of Proposition 1 the definitions $\nu(\tau_\epsilon) := \sqrt{1/\tau_\alpha + 1/\tau_\gamma + 1/\tau_\epsilon}$ and $\tau_S(\tau_\epsilon) := (1/\tau_\gamma + 1/\tau_\epsilon)^{-1}$, where we now make explicit the dependence on τ_ϵ . For $i \in \{H, L\}$, abbreviate $\nu^i := \nu(\tau_\epsilon^i)$ and $\tau_S^i := \tau_S(\tau_\epsilon^i)$. Recalling (4), the cutoff student $s^*(w; \tau_\epsilon)$ in response to w is determined by the indifference condition

$$w = \frac{\tau_\alpha\mu_\alpha + \tau_S(\tau_\epsilon)s^*}{\tau_\alpha + \tau_S(\tau_\epsilon)} - \frac{\lambda}{2} \frac{\tau_S(\tau_\epsilon)}{\tau_\alpha + \tau_S(\tau_\epsilon)} \frac{\tau_\gamma}{\tau_\epsilon(\tau_\alpha + \tau_\gamma)} \quad (12)$$

$$\implies s^*(w; \tau_\epsilon) = w \left(\frac{\tau_\alpha + \tau_S(\tau_\epsilon)}{\tau_S(\tau_\epsilon)} \right) + \frac{\lambda}{2} \frac{\tau_\gamma}{(\tau_\alpha + \tau_\gamma)\tau_\epsilon} - \frac{\tau_\alpha}{\tau_S(\tau_\epsilon)} \mu_\alpha. \quad (13)$$

Define the standardized cutoff student

$$z^*(w; \tau_\epsilon) := \frac{s^*(w; \tau_\epsilon) - \mu_\alpha}{\nu(\tau_\epsilon)} = (w - \mu_\alpha) \frac{\tau_S(\tau_\epsilon) + \tau_\alpha}{\tau_S(\tau_\epsilon)\nu(\tau_\epsilon)} + \frac{\lambda}{2\nu(\tau_\epsilon)} \frac{\tau_\gamma}{(\tau_\alpha + \tau_\gamma)\tau_\epsilon}.$$

Suppose the market conjectures that the students follow cutoff strategies $s^*(\tau_\epsilon^L), s^*(\tau_\epsilon^H) \in \mathbb{R}$. The market's belief conditional on *Conceal* is

$$\mathbb{E}[\alpha|C] = \frac{q\mathbb{P}(C|\tau_\epsilon = \tau_\epsilon^H)\mathbb{E}[\alpha|C, \tau_\epsilon = \tau_\epsilon^H] + (1-q)\mathbb{P}(C|\tau_\epsilon = \tau_\epsilon^L)\mathbb{E}[\alpha|C, \tau_\epsilon = \tau_\epsilon^L]}{q\mathbb{P}(C|\tau_\epsilon = \tau_\epsilon^H) + (1-q)\mathbb{P}(C|\tau_\epsilon = \tau_\epsilon^L)}.$$

³³A Mathematica notebook file containing the expressions underlying this argument is available from the authors upon request.

³⁴The proofs of these facts about the function h are available from the authors upon request.

Using the expressions from the proof of Proposition 1, define

$$\begin{aligned} M(w; \tau_\epsilon) &:= \mathbb{E}[\alpha|C, \tau_\epsilon] = \mu_\alpha - \frac{\tau_S(\tau_\epsilon)}{\tau_\alpha + \tau_S(\tau_\epsilon)} \nu(\tau_\epsilon) \frac{\phi(z^*(w; \tau_\epsilon))}{\Phi(z^*(w; \tau_\epsilon))} \\ A_H(w) &:= q\mathbb{P}(C|\tau_\epsilon = \tau_\epsilon^H) = q\Phi(z^*(w; \tau_\epsilon^H)) \\ A_L(w) &:= (1-q)\mathbb{P}(C|\tau_\epsilon = \tau_\epsilon^L) = (1-q)\Phi(z^*(w; \tau_\epsilon^L)). \end{aligned}$$

Hence $\mathbb{E}[\alpha|C] = f(w) := \frac{A_H(w)M(w; \tau_\epsilon^H) + A_L(w)M(w; \tau_\epsilon^L)}{A_H(w) + A_L(w)}$, and to establish the existence of an equilibrium in (interior) cutoff strategies, it suffices to find a solution w^* to $w = f(w)$.

By inspection, $f < \mu_\alpha$, so for any $w \geq \mu_\alpha$, $w > f(w)$. We now argue that there exists w such that $w < f(w)$. It suffices to show that $\lim_{w \rightarrow -\infty} [f(w) - w] > 0$. To that end, since for each w , $f(w)$ lies between $M(w; \tau_\epsilon^H)$ and $M(w; \tau_\epsilon^L)$, it suffices to show that for all $\tau_\epsilon \in (0, \infty)$, $\lim_{w \rightarrow -\infty} [M(w; \tau_\epsilon) - w] > 0$.

Using (12) and plugging in $z^*(w; \tau_\epsilon)\nu(\tau_\epsilon) + \mu_\alpha$ for $s^*(w; \tau_\epsilon)$, we have

$$\begin{aligned} M(w; \tau_\epsilon) - w &= \mu_\alpha - \frac{\tau_S(\tau_\epsilon)}{\tau_\alpha + \tau_S(\tau_\epsilon)} \nu(\tau_\epsilon) \frac{\phi(z^*(w; \tau_\epsilon))}{\Phi(z^*(w; \tau_\epsilon))} \\ &\quad - \left[\frac{\tau_\alpha \mu_\alpha + \tau_S(\tau_\epsilon)(z^*(w; \tau_\epsilon)\nu(\tau_\epsilon) + \mu_\alpha)}{\tau_\alpha + \tau_S(\tau_\epsilon)} - \frac{\lambda}{2} \frac{\tau_S(\tau_\epsilon)}{\tau_\alpha + \tau_S(\tau_\epsilon)} \frac{\tau_\gamma}{\tau_\epsilon(\tau_\alpha + \tau_\gamma)} \right] \\ &= \frac{\lambda}{2} \frac{\tau_S(\tau_\epsilon)}{\tau_\alpha + \tau_S(\tau_\epsilon)} \frac{\tau_\gamma}{\tau_\epsilon(\tau_\alpha + \tau_\gamma)} - \frac{\tau_S(\tau_\epsilon)}{\tau_\alpha + \tau_S(\tau_\epsilon)} \nu(\tau_\epsilon) \left[\frac{\phi(z^*(w; \tau_\epsilon))}{\Phi(z^*(w; \tau_\epsilon))} + z^*(w; \tau_\epsilon) \right]. \end{aligned}$$

Now the first term in the last line above is constant in w , and as $w \rightarrow -\infty$, we have $z^*(w; \tau_\epsilon) \rightarrow -\infty$, so using the same fact about the normal distribution as in the proof of Proposition 1, the expression in square brackets vanishes as $w \rightarrow -\infty$. Thus, $\lim_{w \rightarrow -\infty} [M(w; \tau_\epsilon) - w] > 0$ for all τ_ϵ , and $\lim_{w \rightarrow -\infty} [f(w) - w] > 0$. It follows that for sufficiently low w , $f(w) > w$. Since we have already shown there exists w such that $w > f(w)$, by continuity, there exists w^* such that $f(w^*) = w^*$. It follows that the cutoffs $s^*(\tau_\epsilon^L) := s^*(w^*; \tau_\epsilon^L) \in \mathbb{R}$ and $s^*(\tau_\epsilon^H) := s^*(w^*; \tau_\epsilon^H) \in \mathbb{R}$ characterize a PBE.

We now turn to part (ii) of the proposition. We use the same notation as above, but we make the dependence on λ explicit. For each $\lambda > 0$, define $w^*(\lambda) = \sup\{w \in \mathbb{R} : w = f(w; \lambda)\}$, the supremum over the set of equilibrium wages, which is nonempty by part (i); by continuity of $f(\cdot; \lambda)$, $w^*(\lambda)$ itself is an equilibrium wage, and it is the maximum wage after *Conceal* across equilibria. Since a higher wage after *Conceal* gives students a higher reservation utility, the expected payoff for students of type τ_ϵ^H or τ_ϵ^L is increasing in the wage after *Conceal*, and hence the equilibrium with wage $w^*(\lambda)$ is ex ante student-optimal.

Since $s^*(w; \tau_\epsilon, \lambda)$ and $z^*(w; \tau_\epsilon, \lambda)$ are strictly increasing in λ , $f(w; \lambda)$ is also strictly increasing in λ , and hence $w^*(\lambda)$ is strictly increasing.

We have

$$\begin{aligned} z^*(w^*(\lambda); \tau_\epsilon^L, \lambda) - z^*(w^*(\lambda); \tau_\epsilon^H, \lambda) &= (w^*(\lambda) - \mu_\alpha) \left[\frac{\tau_S^L + \tau_\alpha}{\tau_S^L \nu^L} - \frac{\tau_S^H + \tau_\alpha}{\tau_S^H \nu^H} \right] \\ &\quad + \frac{\lambda \tau_\gamma}{2(\tau_\alpha + \tau_\gamma)} \left[\frac{1}{\tau_\epsilon^L \nu^L} - \frac{1}{\tau_\epsilon^H \nu^H} \right]. \end{aligned} \tag{14}$$

Both expressions in square brackets on the right hand side of (14) are positive and independent of λ . Since $w^*(\lambda)$ is increasing, the right hand side of (14) strictly increasing in λ , and it is positive for sufficiently large λ . On the other hand, since $f < \mu_\alpha$ we have $w^*(\lambda) < \mu_\alpha$, and thus it is easy to see that the right side of (14) is negative for sufficiently small $\lambda > 0$. This establishes the ‘‘if and only if’’

statement in part (ii) of the proposition.

Part (iii) of the proposition is analogous to parts (ii) and (iii) of Proposition 2, so its proof is omitted. ■

Proof of Proposition 5. Using arguments and notation from the proof of Proposition 4, establishing the existence of equilibrium reduces to finding a fixed point of the equation $f(w) = w$, where f remains well-defined; we now have $\tau_S(\tau_\epsilon^H) = \tau_\gamma$, $\nu(\tau_\epsilon^H) = \sqrt{1/\tau_\alpha + 1/\tau_\gamma}$, and

$$s^*(w; \tau_\epsilon^H) = w \left(\frac{\tau_\alpha + \tau_\gamma}{\tau_\gamma} \right) - \frac{\tau_\alpha}{\tau_\gamma} \mu_\alpha \quad (15)$$

$$z^*(w; \tau_\epsilon^H) = (w - \mu_\alpha) \frac{\tau_\gamma + \tau_\alpha}{\tau_\gamma \nu(\tau_\epsilon^H)}, \quad (16)$$

which are both finite for all $w \in \mathbb{R}$. It remains true that $f < \mu_\alpha$, so that $w > f(w)$ for all $w \geq \mu_\alpha$; hence, to establish a fixed point, it suffices to show that there exists w such that $f(w) > w$. Fix any w such that $M(w; \tau_\epsilon^L) > w$; such w must exist since $\lim_{w \rightarrow -\infty} [M(w; \tau_\epsilon) - w] > 0$ for all $\tau_\epsilon \in (0, \infty)$, as argued in the previous proof. Now as $q \rightarrow 0$, we have $f(w) \rightarrow M(w; \tau_\epsilon^L)$, so for sufficiently small q , $f(w) > w$. It follows that for sufficiently small q , a fixed point exists, and hence an equilibrium with real-valued cutoffs exists. ■

Proof of Proposition 6. As in the proof of Proposition 4, we characterize cutoff equilibria by first determining the student's best response cutoffs for an arbitrary wage w after *Conceal* and then using these cutoffs to determine the market's expectation of the student's ability given *Conceal*. From (13), the student's best response to w as a function of her risk type $\lambda \in \{\lambda_L, \lambda_H\}$ is

$$s^*(w; \lambda) = w \left(\frac{\tau_\alpha + \tau_S}{\tau_S} \right) + \frac{\lambda}{2} \frac{\tau_\gamma}{(\tau_\alpha + \tau_\gamma)\tau_\epsilon} - \frac{\tau_\alpha}{\tau_S} \mu_\alpha. \quad (17)$$

We then define $z^*(\lambda)$, $M(w; \lambda)$, $A_H(w)$, $A_L(w)$ and $f(w)$ as in the proof of Proposition 4. The existence of a fixed point $w = f(w)$ then follows by the same arguments as in that proof, and for any such fixed point w^* , there is an equilibrium with cutoffs $s^*(\lambda_L) := s^*(w^*; \lambda_L)$ and $s^*(\lambda_H) := s^*(w^*; \lambda_H)$. And in any equilibrium, since the RHS of (17) is increasing in λ , we have $s^*(\lambda_H) > s^*(\lambda_L)$. Since all other parameter values are fixed, it follows immediately that $z^*(\lambda_H) > z^*(\lambda_L)$, so students of type λ_H are more likely to conceal than students of type λ_L .

For the comparison of expected payoffs, it suffices to show that in any equilibrium of the game with uncertainty over λ , w^* satisfies $w_H^* > w^* > w_L^*$, where w_x^* in the game in which $\lambda = \lambda_x$ is common knowledge, $x \in \{L, H\}$. Let $f(w; L)$ denote the expected ability of the student given *Conceal* when $\lambda = \lambda_L$ is common knowledge and the student plays the best response cutoff $s^*(w; \lambda_L)$. Since $s^*(w; \lambda_H) > s^*(w; \lambda_L)$, we have $f(w) > f(w; L)$ for all w , and thus $w^* > w_L^*$. The proof of $w_H^* > w^*$ is symmetric. ■

Proof of Proposition 7. (i) As in the proof of Proposition 1, in any equilibrium, the student receives some wage $w_c \in \mathbb{R}$ after choosing *Conceal*. In any equilibrium, a student with signal $S = s$ prefers to disclose iff $\mathbb{E}[u(\mathbb{E}[\alpha|\gamma])|S = s] \geq u(w)$. Since the left hand side is strictly increasing in s , either (a) the equilibrium involves a cutoff $s^* \in \mathbb{R}$ as in the proposition statement, (b) all students conceal, or (c) all students disclose. We can eliminate (b) since for sufficiently high s , a student with signal $S = s$ would strictly prefer to disclose. To see this, note that from the derivation toward Proposition 1, from the student's interim perspective, $\mathbb{E}[\alpha|\gamma] = \frac{\tau_\alpha \mu_\alpha + \tau_\gamma \gamma}{\tau_\alpha + \tau_\gamma}$ can be written as $w(z) + \tilde{\epsilon}$ where $z := \frac{s - \mu_\alpha}{\nu}$ and $w(z) := \mu_\alpha + \nu z \frac{\tau_S}{\tau_\alpha + \tau_S} = \mathbb{E}[\alpha|S = s]$, and where $\tilde{\epsilon} \sim N\left(0, \frac{\tau_S(\tau_\epsilon)}{\tau_\alpha + \tau_S(\tau_\epsilon)} \frac{\tau_\gamma}{\tau_\epsilon(\tau_\alpha + \tau_\gamma)}\right)$ is independent of S . Hence, $\mathbb{E}[u(\mathbb{E}[\alpha|\gamma])|S = s]$ can be written as $\mathbb{E}[u(w(z) + \tilde{\epsilon})]$, and by the monotone convergence theorem, the limit of this expression as $s \rightarrow +\infty$ (and hence as $z \rightarrow +\infty$) is strictly greater

than $u(w_c)$.³⁵ It follows that for sufficiently high s , the student would strictly prefer to disclose. By a similar argument, we can eliminate (c), since for sufficiently low s , the student would strictly prefer to conceal. This leaves (a) as the only possibility, establishing (i).

(ii) Showing existence (and uniqueness) now reduces to showing the existence (and uniqueness) of s^* solving the indifference condition

$$\mathbb{E}[u(\mathbb{E}[\alpha|\gamma])|S = s^*] = u(\mathbb{E}[\alpha|S < s^*]). \quad (18)$$

Using the expression for $\mathbb{E}[\alpha|S < s^*]$ derived for Proposition 1, (18) becomes

$$\mathbb{E}[u(w(z^*) + \tilde{\epsilon})] = u\left(\mu_\alpha - \frac{\tau_S}{\tau_\alpha + \tau_S} \nu \frac{\phi(z^*)}{\Phi(z^*)}\right) \quad (19)$$

$$\iff \mathbb{E}[u(w(z^*) + \tilde{\epsilon})] = u\left(w(z^*) - \frac{\tau_S}{\tau_\alpha + \tau_S} \nu \left(z^* + \frac{\phi(z^*)}{\Phi(z^*)}\right)\right). \quad (20)$$

We first establish that the lhs of (19) crosses the rhs at least once from below. Note that the rhs is bounded above by $u(\mu_\alpha)$; on the other hand, by the monotone convergence theorem, the limit of the lhs as $z^* \rightarrow +\infty$ exceeds $u(\mu_\alpha)$.

Next, to establish a crossing point, we work with (20) and show that if $\liminf_{w \rightarrow -\infty} -\frac{u''(w)}{u'(w)} > 0$, then for sufficiently low z^* , the rhs exceeds the lhs. Since $z^* + \frac{\phi(z^*)}{\Phi(z^*)} \rightarrow 0$ as $z^* \rightarrow -\infty$, it suffices to show that $\liminf_{z^* \rightarrow -\infty} \pi^u(w(z^*)) > 0$, where $\pi^u(w)$ is the risk premium for u defined by $\mathbb{E}[u(w) + \tilde{\epsilon}] = u(w - \pi^u(w))$.

The main idea is to construct another C^2 function $v : \mathbb{R} \rightarrow \mathbb{R}$ that satisfies $v' > 0$ and $v'' \leq 0$ and the additional properties (1) $-\frac{v''(w)}{v'(w)}$ is nonincreasing and is not identically zero, and (2) for all $w \in \mathbb{R}$, $-\frac{u''(w)}{u'(w)} \geq -\frac{v''(w)}{v'(w)}$. From there, we argue that for any w , the risk premium $\pi^u(w)$ is bounded below by the risk premium $\pi^v(w)$ defined by $\mathbb{E}[v(w) + \tilde{\epsilon}] = v(w - \pi^v(w))$. Moreover, we then argue that $\pi^v(w) > 0$ and π^v is nonincreasing, so for arbitrary $\hat{w} \in \mathbb{R}$, $\liminf_{w \rightarrow -\infty} \pi^u(w) \geq \liminf_{w \rightarrow -\infty} \pi^v(w) \geq \pi^v(\hat{w}) > 0$, establishing the result.

To construct such v as above, let K be an arbitrary constant in $(0, \liminf_{w \rightarrow -\infty} -\frac{u''(w)}{u'(w)})$. There exists \underline{w} such that for all $w \leq \underline{w}$, $-\frac{u''(w)}{u'(w)} \geq K$. Let $w_0 < \underline{w}$ be arbitrary. Let λ be any continuous, nonincreasing function that satisfies $\lambda(w) = K$ for all $w \leq w_0$ and $\lambda(w) = 0$ for all $w \geq \underline{w}$.³⁶ Finally, define v by

$$v(w) = \begin{cases} 1 - e^{-Kw} & \text{if } w \leq w_0 \\ 1 - e^{-Kw_0} + \int_{w_0}^w K e^{-Kw_0} e^{-\int_{w_0}^z \lambda(y) dy} dz & \text{if } w > w_0. \end{cases}$$

By construction, v satisfies $v' > 0$ and $v'' \leq 0$, and as $-\frac{v''(w)}{v'(w)} = \lambda(w)$, it satisfies the additional properties (1) and (2) stated before.

By standard results, for all $w \in \mathbb{R}$, $\pi^u(w) \geq \pi^v(w)$, and as $\tilde{\epsilon}$ has full support, $\pi^v(w) > 0$. Moreover, $v(w)$ exhibits weakly decreasing absolute risk aversion, so $\pi^v(w)$ is nonincreasing.³⁷ The rest of the argument that $\liminf_{w \rightarrow -\infty} \pi^u(w) > 0$ then follows as outlined before. This establishes (ii).

³⁵The expected payoff is $\int u(w(z) + \tilde{\epsilon}) f(\tilde{\epsilon}) d\tilde{\epsilon}$, where f is the pdf of the random variable $\tilde{\epsilon}$. By the monotone convergence theorem, the limit of this integral is the integral of the pointwise limit, the pointwise limit being $\lim_{y \rightarrow \infty} u(y) > u(w_c)$.

³⁶Note that λ need not be differentiable; e.g., setting $\lambda(w) = K(\underline{w} - w)/(\underline{w} - w_0)$ for $w \in (w_0, \underline{w})$ suffices.

³⁷To see this, observe that for all $\delta > 0$, the function \tilde{v} defined by $\tilde{v}(w) = v(w - \delta)$ is a concave transformation of v , and hence for fixed w , \tilde{v} has a weakly higher risk premium; but the risk premium for \tilde{v} at w is exactly the risk premium of v at $w - \delta$.

(iii) Toward uniqueness, rewrite (20) as

$$u(w(z^*) - \pi^u(w(z^*))) = u\left(w(z^*) - \frac{\tau_S}{\tau_\alpha + \tau_S} \nu \left(z^* + \frac{\phi(z^*)}{\Phi(z^*)}\right)\right). \quad (21)$$

If $-\frac{u''(w)}{u'(w)}$ is nonincreasing and not identically zero, then $\liminf_{w \rightarrow -\infty} -\frac{u''(w)}{u'(w)} > 0$ so there exists a crossing point, and $\pi^u(w(z^*))$ is nonincreasing in x . But on the rhs, $\frac{\tau_S}{\tau_\alpha + \tau_S} \nu \left(z^* + \frac{\phi(z^*)}{\Phi(z^*)}\right)$ is strictly increasing, and hence the lhs can only cross the rhs once, from below.

(iv) By the same argument as above, the risk premia satisfy $\pi^u(w) \geq \pi^v(w)$ for all w . Now for $h \in \{u, v\}$, (20) reduces to the equation $w(z^*) - \pi^h(z^*) = w(z^*) - \frac{\tau_S}{\tau_\alpha + \tau_S} \nu \left(z^* + \frac{\phi(z^*)}{\Phi(z^*)}\right)$, where the rhs is independent of h , while for all fixed z^* the lhs is lower for u than for v . Since the lhs crosses the rhs from below, the point of intersection for u is to the right of that for v . ■

Proof of Proposition 8. The comparative static with respect to c follows from the discussion before the proposition. As for τ_ϵ , defining $C^* := \frac{c}{2}(e^*)^2$, the indifference condition (7) becomes

$$z^* - \frac{\lambda}{2} \frac{\tau_\gamma}{\nu(\tau_\alpha + \tau_\gamma)\tau_\epsilon} - \frac{C^*(\tau_\alpha + \tau_S)}{\nu\tau_S} = -\frac{\phi(z^*)}{\Phi(z^*)}, \quad (22)$$

where again $\nu^2 = 1/\tau_\alpha + 1/\tau_\epsilon + 1/\tau_\gamma$ and $\tau_S := (1/\tau_\gamma + 1/\tau_\epsilon)^{-1}$. It is straightforward to verify that the left hand side of (22) is increasing in τ_ϵ , so z^* is decreasing in τ_ϵ , and thus the probability of disclosure $1 - \Phi(z^*)$ is increasing in τ_ϵ . ■

B Appendix Tables and Figures

Table C.1: Sample restrictions and sample sizes

Restriction	# Observations (1)	# Students (2)	# Courses (3)
Initial sample	107,745	29,375	3,279
Non-missing GPA	107,723	29,369	3,279
Non-missing Canvas scores (full sample)	107,070	29,362	3,259
Canvas-registrar grades agree	75,551	28,312	2,519
At least 30 students (Analysis sample)	46,514	23,717	778

Notes: Table reports the number of observations (student-courses), number of students, and number of students, as we impose increasingly stringent sample restrictions. The initial sample consists of all undergraduates enrolled in courses at Indiana University, Bloomington, in Spring 2020, with standard final grades (A-F or P), in full-semester or second-half courses. The restriction that Canvas and registrar grades agree means that, for each class, all Canvas grades exactly match registrar grades, or disagree by one notch. We use the “analysis sample” for all analyses involving students’ signals. We use the “full sample” for our analysis of GPA notches and disclosure.

Table C.2: Determinants of course-level RMSE

	(1)	(2)	(3)	(4)	(5)
200-level	-0.13 (0.33)			0.03 (0.32)	0.19 (0.32)
300-level	-0.96 (0.29)			0.10 (0.31)	0.27 (0.31)
400-level	-1.97 (0.34)			-0.40 (0.39)	0.05 (0.38)
Average GPA		-5.02 (0.56)		-4.76 (0.66)	-4.40 (0.66)
Square root class size			0.14 (0.05)	0.06 (0.05)	0.03 (0.05)
R^2	0.04	0.11	0.01	0.11	0.19
# Classes	810	810	810	810	810
School dummies?	No	No	No	No	Yes

Notes: Table reports coefficients from a regression of course-level RMSE on the indicated controls. 100-level courses are the omitted category in columns (1), (4), and (5). The coefficients on the school dummies are in Column (2) of [Table C.3](#). The sample is the analysis sample, described in the notes to [Table 1](#). Robust standard errors in parentheses.

Table C.3: Course-level RMSE, by school

	(1)	(2)
Nursing	-4.74 (0.39)	-3.11 (0.44)
Art	-3.51 (1.04)	-3.37 (0.75)
Business	-2.03 (0.27)	-1.56 (0.28)
Education	-1.69 (1.63)	-1.28 (1.73)
Public Health	-1.69 (0.32)	-1.44 (0.30)
Policy	-1.40 (0.34)	-1.54 (0.34)
Medicine	-0.98 (2.00)	0.29 (1.96)
Media	-0.88 (0.63)	-0.81 (0.58)
Music	0.10 (0.90)	0.26 (0.89)
Informatics	0.58 (0.70)	0.71 (0.69)
International	0.78 (1.01)	0.70 (0.96)
Joint p-value for schools	< 0.001	< 0.001
R^2	0.11	0.19
# Classes	810	810
Other controls dummies?	No	Yes

Notes: Table reports coefficients from a regression of course-level RMSE on a set of school dummies, omitting the College of Arts and Sciences. The coefficients on the additional controls are reported in column (5) of [Table C.2](#). The sample is the analysis sample, described in the notes to [Table 1](#). Robust standard errors in parentheses.

Table C.4: Predictive power of signals for disclosure decision

Predictor:	May 2 Signal (1)	GPA (2)	Raw May 2 (3)	Homogeneous May 2 (4)	Final Grade (5)
R^2	0.265	0.029	0.096	0.097	0.386
% Correct	0.879	0.862	0.862	0.862	0.897
Model parameters	11	11	11	11	11
Sample size	46,514	46,514	46,514	46,514	46,514

Notes: For each column, we estimate a linear probability model of disclosure on dummy variables for each value of the indicated variable (except column (3), where we dummies for binned values for GPA grouped to grade bins (e.g. 3.3). We report the R^2 and the fraction correctly classified. We say we have correctly classified an observation if $disclose_{ic} = 1$ and the predicted probability is above 0.5, or $disclose_{ic} = 0$ and the predicted probability is less than 0.5. The May 2 signal is our main signal measure. “Raw May 2” does not adjust for drift between May 2 and final scores. Homogeneous May 2 adjusts for drifts but without class-specific coefficients. Final grade uses the Canvas final grade as the signal. The sample is the analysis sample, described in the notes to [Table 1](#).

Table C.5: Robustness of effect of uncertainty on disclosure to alternative measures of uncertainty

Measure	RMSE (1)	Pr(Switch ≥ 1) (2)	Pr(Switch ≥ 2) (3)	High RMSE (4)
Uncertainty	-0.007 (0.001)	-0.096 (0.013)	-0.158 (0.026)	-0.036 (0.005)
R^2	0.644	0.644	0.644	0.644
# Observations	37,208	37,208	37,208	37,208
# Students	14,411	14,411	14,411	14,411
# Courses	778	778	778	778
Mean uncertainty	3.913	0.436	0.143	0.553

Notes: Table reports coefficients from a linear probability model of disclosure on the indicated uncertainty measure, as well as controls for a set of dummy variables for each grade signal (A, A-, B+, etc.); dummy variables for course level (200/300/400); square root of class size; average incoming GPA of students enrolled in the class, and student fixed effects. Pr(Switch ≥ 1) is the course-specific probability the interim and final grade differ by one or more notches, and Pr(Switch ≥ 2) is the probability they differ by two or more notches. High RMSE is defined as above-median RMSE. Sample consists of full-term or second half courses, with at least 30 students in which Canvas and registrar grades always agree to within 1 notch. Robust standard errors, clustered on course, in parentheses.

Table C.6: Robustness of effect of uncertainty on disclosure to alternative signal definitions

Signal	Baseline (1)	Linear (2)	Cubic (3)	+GPA (4)	Final (5)	Round Up (6)
RMSE	-0.007 (0.001)	-0.009 (0.002)	-0.007 (0.001)	-0.007 (0.001)	-0.005 (0.001)	-0.005 (0.001)
R^2	0.644	0.603	0.619	0.647	0.716	0.717
# Observations	37,208	37,208	37,208	37,208	37,208	37,208
# Students	14,411	14,411	14,411	14,411	14,411	14,411
# Courses	778	778	778	778	778	778

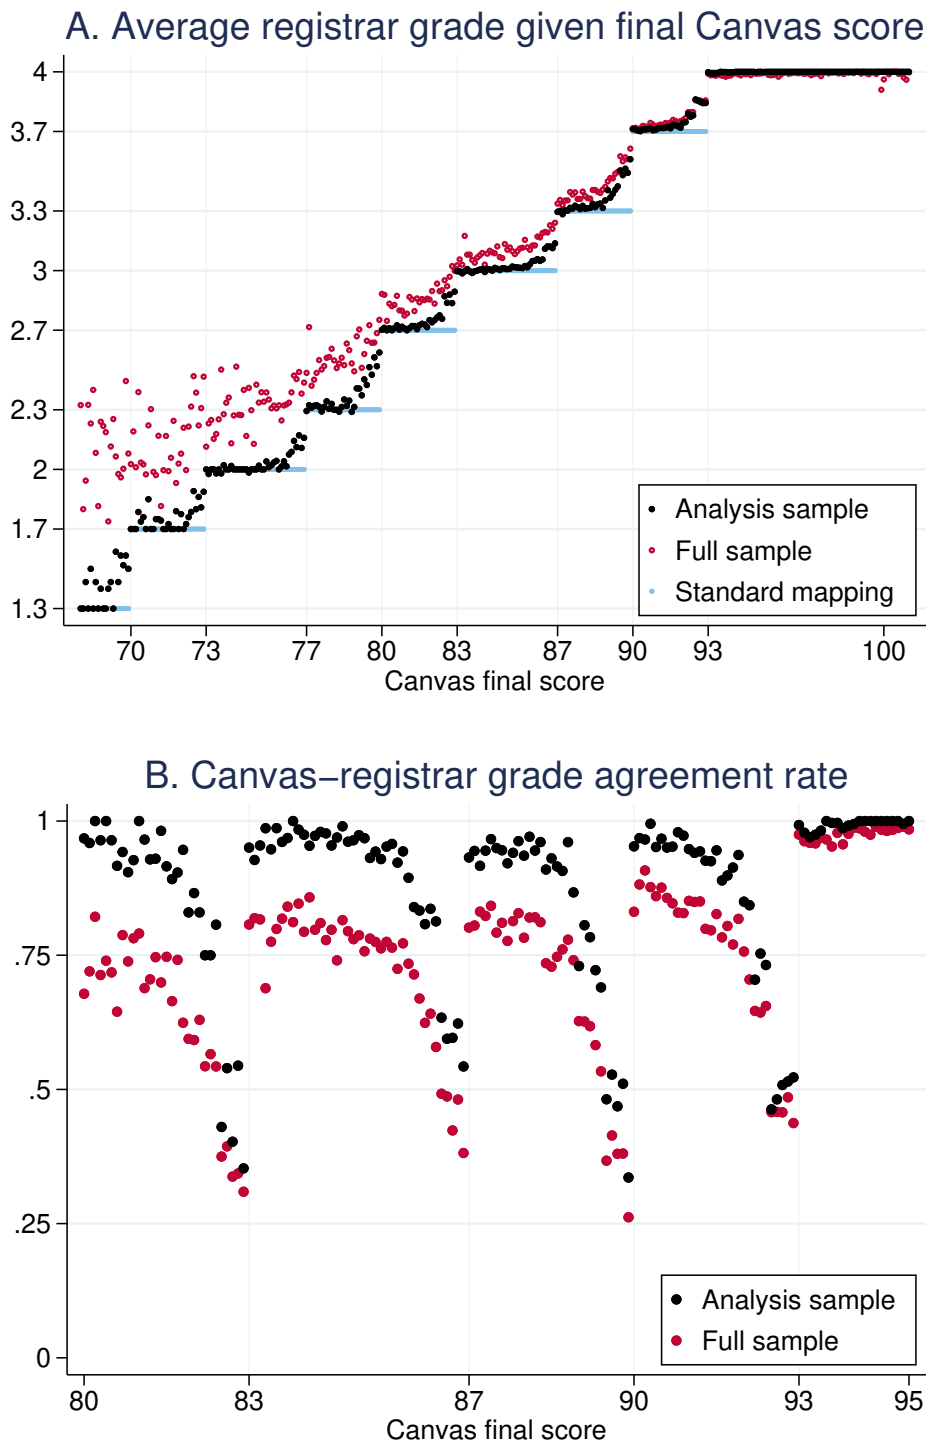
Notes: Table reports coefficients from a linear probability model of disclosure on course RMSE and the indicated signal measure, as well as dummy variables for course level (200/300/400); square root of class size; average incoming GPA of students enrolled in the class, and student fixed effects. The baseline signal measure is a set of dummy variables for predicted final grade given May 2 grade. Column (2) controls linearly for predicted final grade, and column (3) controls for a cubic in predicted final grade. Column (2) uses student GPA as well as May 2 grade to predict final grade. Column (5) controls for dummy variables for final grade rather than predicted final grade. Column (6) controls for dummy variables for final grade, with grades ending in .5 and higher rounded up. Sample consists of full-term or second half courses, with at least 30 students in which Canvas and registrar grades always agree to within 1 notch. Robust standard errors, clustered on course, in parentheses.

Table C.7: Robustness of effect of uncertainty on disclosure to alternative samples

Sample	Baseline (1)	No bunch (2)	Minority A (3)	Small (4)	Match Exact (5)	Match Round (6)
RMSE	-0.007 (0.001)	-0.007 (0.002)	-0.006 (0.002)	-0.005 (0.001)	-0.007 (0.002)	-0.006 (0.002)
R^2	0.644	0.666	0.677	0.605	0.646	0.639
# Observations	37,208	26,731	19,905	70,467	9,031	17,624
# Students	14,411	10,943	8,438	23,357	4,089	7,576
# Courses	778	634	497	2,442	330	487

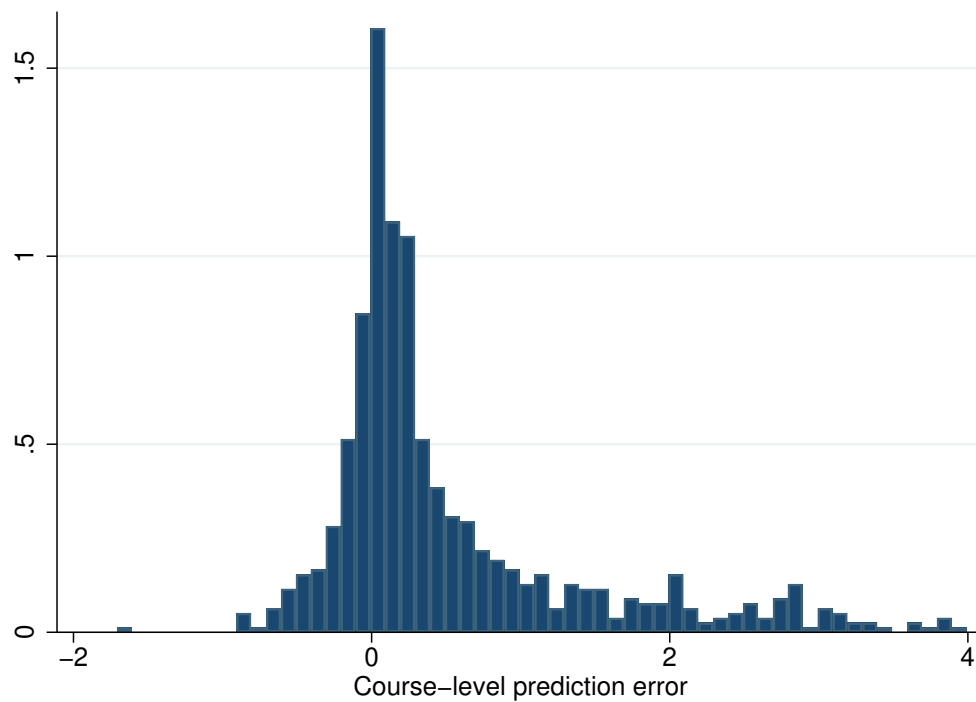
Notes: Table reports coefficients from a linear probability model of disclosure on course RMSE, as well as controls for a set of dummy variables for each grade signal (A, A-, B+, etc.); dummy variables for course level (200/300/400); square root of class size; average incoming GPA of students enrolled in the class, and student fixed effects. The baseline sample consists of full-term or second half courses, with at least 30 students in which Canvas and registrar grades always agree to within 1 notch. “No bunch” excludes classes where more than 1% of grades are exact multiples of 10. “Minority A” excludes classes where a majority of students earned an A. “Small” includes classes of all sizes. “Match exact” excludes classes where Canvas and Registrar grades disagree at all. “Match Round” is limited to classes with at least 30 students in which all Canvas and Registrar grades either match exactly, or match after rounding up Canvas grades ending in .5 and higher. Robust standard errors, clustered on course, in parentheses.

Figure C.1: Alignment between Canvas and registrar grades



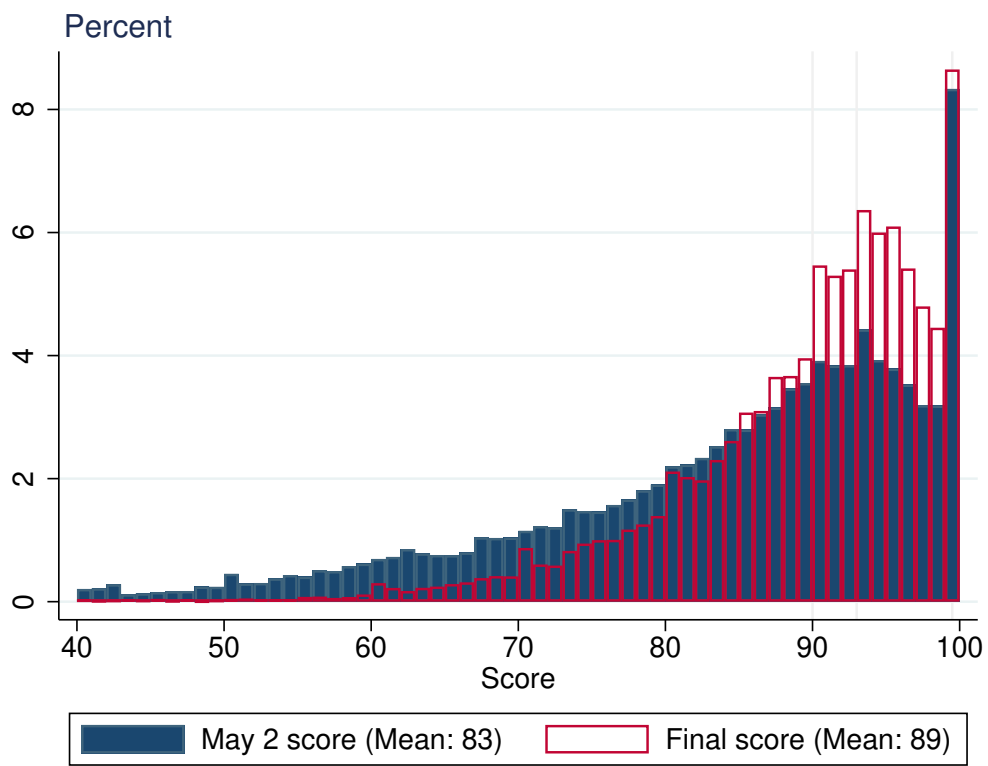
Notes: Top panel shows the average registrar grade (on a 0-4) scale for the full sample and the analysis sample, given Canvas final score, as well as the “standard mapping.” The bottom panel shows the the fraction of registrar grades that equal the standard mapping from the Canvas final score, at each 0.1 point interval. The analysis and full samples are defined in the notes to [Table 1](#).

Figure C.2: Distribution of prediction errors across courses




Notes: Figure plots the distribution of average prediction errors across courses. The prediction error is the final grade less the May 2 grade. The vertical line is at 0. If the signal were unbiased for all classes, we would expect a symmetric distribution around 0. The sample is the analysis sample as defined in the notes to [Table 1](#).

Figure C.3: Distribution of May 2 and final Canvas scores



Notes: Figure plots the distribution of final and May 2 Canvas scores in the analysis sample. The vertical lines are at 93 and 93 percent. The sample is the analysis sample as defined in the notes to [Table 1](#).

Figure C.4: Distribution of course-level RMSE



Notes: Figure plots the distribution across courses of RMSE, the root mean squared prediction error from a regression of final Canvas grade on May 2 course grade, with course-specific coefficients. The sample is the analysis sample as defined in the notes to [Table 1](#).