Coarse Grades: Informing the Public by Withholding Information

By Rick Harbaugh and Eric Rasmusen*

Certifiers of quality often report only coarse grades to the public despite having measured quality more finely, e.g., "Pass" or "Certified" instead of "73 out of 100". Why? We show that coarse grades result in more information being provided to the public because the coarseness encourages those of middling quality to apply for certification. Dropping exact grading in favor of the best coarse grading scheme reduces public uncertainty because the extra participation outweighs the coarser reporting. In some circumstances, the coarsest meaningful grading scheme, pass-fail grading, results in the most information. JEL: D82, L15.

Grades are often coarse. Rather than an exact number or rank, a grade is usually only a rough indication of quality, such as a letter grade or even just a binary pass-fail grade. Safety organizations usually certify that a product is safe with a seal of approval that does not indicate whether the product passed tests just barely or by a wide margin. Environmental organizations typically certify environmental quality with a simple "eco-label" rather than revealing the results of their more detailed evaluation. When it comes to reporting the results to the public, they throw away information.

Why this waste of information? An obvious reason is that it costs more to grade finely than coarsely. But this can't be the entire explanation, since the certifier often collects detailed in-

* Harbaugh: Kelley School of Business, Indiana University, Bloomington, Indiana, riharbau@indiana.edu. Rasmusen: Dan R. and Catherine M. Dalton Professor, Department of Business Economics and Public Policy, Kelley School of Business, Indiana University, Bloomington, Indiana, erasmuse@indiana.edu. We thank Mimi Chan, Mikkel Jakobsen, and Dan Zhao for research assistance, and Michael Baye, Chris Connell, Dan Klerman, Dmitry Lubensky, and participants in talks at the BEPP Brown Bag Lunch, Baylor University, and the IIOC and Summer Econometric Society meetings for helpful comments. We have no relevant material or financial interests that relate to the work described in this paper.

formation but refrains from reporting it.[1] When the certifier deliberately reclassifies information with a coarsening filter before reporting the results publicly, coarse grading is more expensive than exact grading, not cheaper.

This coarsening of information is a puzzle since a certifier has an incentive to provide accurate information so as to increase the value of its services to consumers and advertisers. Making coarsening even more of a puzzle, many certifiers are non-profits with the explicit goal of providing consumers with the best information. For instance, non-profits run most of the numerous eco-label schemes that provide information on products' environmental, health, and social impacts. Of 363 different schemes tracked by *Ecolabelindex.com*, 209 are controlled by non-profits, 59 by industry groups, 53 by governments, and 42 by for-profits.[2]

If a certifier really wants to provide accurate information to receivers, why make the information coarser than necessary? As discussed in the literature survey in Section V, there are several possible reasons – that the certifier aims to help the firms, not consumers; that the certifier can increase its own profits by coarseness; that coarseness can provide incentives for firms to improve their quality; that coarseness makes it easier for receivers to process the information. We will add another reason, one that applies even when the certifier's aim is to convey as much information as possible but which applies only when certification is voluntary. In situations such as certification for eco-labels, costly cooperation from firms is required; the certifier needs to get firms to participate. Just as a student would be reluctant to attend a medical school that would publicly certify him as the worst in his class, a firm would not be eager to be stamped with a seal of approval that tells the world it barely passed. (Recall the old joke, "What do you call someone who graduated from the bottom of his class in medical school? Doctor.") Hence, a certifier who wants to maximize information needs to consider how the grading scheme affects the willingness of senders to be certified at all.

Coarsening can increase the amount of information receivers get in equilibrium by inducing more participation. If the certification grade is coarse, a mediocre type is pooled with better

[1] For instance the EnergyStar label requires that a third-party measure energy usage and certify that it is below a threshold, but the label does not indicate the actual energy usage. Similarly, 95% of a product's ingredients must be organic for a product to use the label "organic", but the label does not usually indicate the exact percentage.

[2] We thank Anastasia O'Rourke for providing this information, which is for 2009.

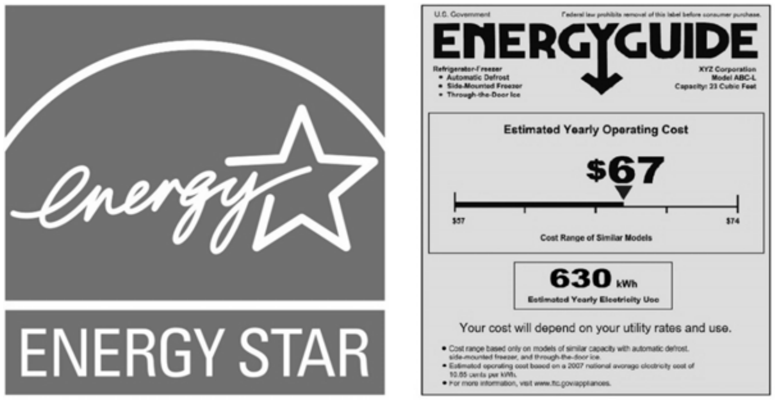

FIGURE 1. PASS-FAIL GRADE FOR A VOLUNTARY LABEL, EXACT GRADE FOR A MANDATORY LABEL.

types, so the expected quality conditional on the certification grade is higher than the true quality. Therefore the sender has more incentive to bear the costs of certification, and participation rises. We will show that at the margin the extra information from increased participation outweighs the loss from throwing away information on those that do participate. Hence, coarse grading is optimal in the sense of minimizing the mean squared error of the receiver's estimate of sender quality. Moreover, under some conditions the optimal scheme that minimizes mean squared error is maximally coarse: the firm or person being tested simply passes or fails and the exact test scores are never shown to the public.

We show that the optimal scheme is either pass-fail or what we call an "honors" scheme, in which senders who are good enough to pass remain pooled, while senders at the top are reported exactly. Schools do not publicly provide class rank information about most of their graduates, but they do publicly honor the valedictorian and other top students. Safety and environmental organizations provide product labels that certify a passing grade, and sometimes also provide public awards that highlight the best achievers. Recommendation letters work the same way: some students or employees won't even ask for a letter, some receive favorable boilerplate letters, and the best receive individuated letters with fine distinctions in commendation.

Since coarseness is used to encourage participation, the model predicts that coarseness is less likely when quality evaluation does not require the sender's cooperation. Camera companies cannot prevent consumer reviewers at *Amazon.com* or professional reviewers at *CNET* from rating their products. Without the need to encourage participation by firms, we should expect product review websites to provide fine information. Indeed, most such websites provide summary measures containing exact numeric scores, fine categorizations, or some combination thereof, and offer immediate access to detailed review information.[3]

We also expect that coarseness is less likely for mandatory labels provided by government agencies. Since they can force firms to provide information about their products, there is no need to encourage participation by clouding the truth. Using the data from *Ecolabelindex.com,* we found that of the 174 voluntary labels from OECD countries for which grading data could be found, only 5 of them provide exact grades or grades with more than a few levels. In contrast, all 5 mandatory labels provide fine or exact grades. As shown in Figure 1, the US Department of Energy's voluntary "Energy Star" label for home products only indicates that the product has met a certain standard for low energy usage, while the FTC's "EnergyGuide" label that is mandatory for large appliances provides exact information on energy usage and expected energy cost.

Our results help fill a gap in the literature. The certification literature following Lizzeri (1999) shows how a certifier trying to maximize his own profits from charging for certification can do so by reporting product information coarsely. The Bayesian persuasion literature following Kamenica & Gentzkow (2011) shows how a sender (not a certifier or receiver) can benefit from pooling types in order to "concavify" the sender's value function over receiver beliefs. In the last section of this article we will say more about these and other papers. Of the three different parties involved in certification — certifier, sender, and receiver — the existing literature assumes that the certifier acts to maximize either his own benefit or the benefit to senders, while we assume that the certifier acts to maximize the receiver's information. We find, paradoxically, that the purpose of coarsening can be to increase receiver information.

[3]Amazon, Yelp, and TripAdvisor report overall quality using star or half-star intervals, and also report exact numeric rankings. *CNET* provides numeric ratings. The consumer reviews that the ratings and rankings are based on are all linked to. In contrast, certifiers such as Underwriters Laboratory that provide pass/fail labels to participating firms typically treat the exact test results as confidential.

I. The Model

A sender (e.g., a firm selling a product) has exogenous quality q that is randomly distributed with support on $[0, 1]$ according to distribution function F with analytic density f such that $f > 0$ on $(0, 1)$. The realization of q is the sender's private information.

A certifier (e.g., an NGO) chooses a grading scheme $m(q)$ which is a function that maps quality q to a message m in a set M. We restrict the message function to rule out schemes with isolated points that are informationally equivalent to the schemes we are interested in. First, any set of types sending the same message, $\{q \in (0, 1) : m(q) = m'$ for any $m'\}$, must be a union of positive measure intervals. Second, the set of types that each send a unique message, $\{q \in (0, 1) : m(q) \neq m(q')$ for all $q' \neq q\}$, must also be a union of positive measure intervals.[4] There are no transfers. If the sender applies for certification, the certifier measures quality perfectly and reports message m based on the grading scheme. If the sender chooses not to apply for certification, the certifier reports the message $m =$ "uncertified".

The certifier chooses a grading scheme to best inform the receiver (e.g., a consumer) about the sender's quality in the sense of minimizing the expectation of a quadratic loss function $(q - E[q|m])^2$, i.e., mean squared error. This loss function could capture a preference for providing accurate information, a reputational incentive to do so, or a concern for consumer welfare.[5]

The sender incurs a fixed cost $c > 0$ in time and trouble to be certified that is independent of the sender's type.[6] The receiver updates his estimate of quality $E[q|m]$ based on the prior distribution F and the equilibrium meaning of m. Given a certification scheme, the sender chooses either payoff $E[q|m] - c$ from applying for certification or $E[q|\text{"}uncertified\text{"}]$ from not applying.[7]

[4]As suggested by a referee, these restrictions allow standard differentiation techniques and exclude, for example, schemes that assign a single message to a fat Cantor set which has positive measure but is not a union of positive measure intervals. Note that the assumption that m is a function rules out mixed-strategy grading.

[5]Quadratic loss functions, or equivalent assumptions, are widely used in cheap talk (e.g., Crawford & Sobel, 1982), signaling (e.g., Spence, 1973), disclosure (e.g., Milgrom, 1981), certification (e.g., Lizzeri, 1999), and Bayesian persuasion games (e.g., Kamenica and Gentzkow, 2011).

[6]Note that $c > 0$ prevents "unraveling" in which all senders participate. A referee suggested that, if transfers were allowed, the certifier might reimburse senders depending on their quality while maintaining a balanced budget.

[7]For simplicity we assume that there is only one certifier to apply to. Since the goal of a certifier is to maximize information to consumers, it will not engage in competition that worsens information, e.g., by providing exact grading to firms with very high

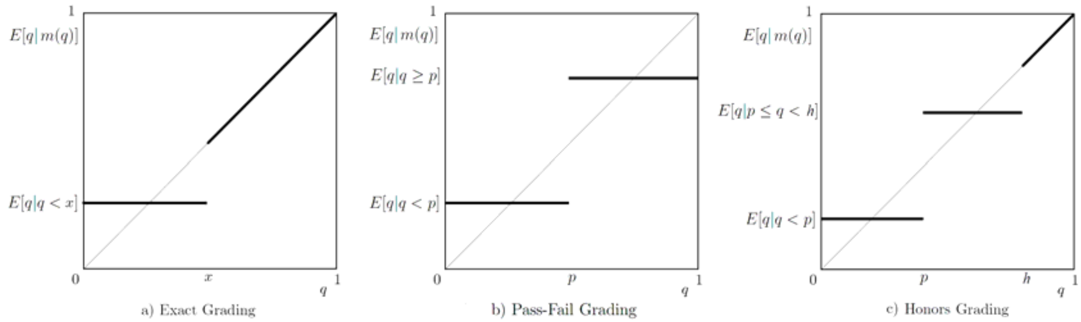

FIGURE 2. TRUE QUALITY AND CONSUMER ESTIMATED QUALITY UNDER DIFFERENT GRADING SCHEMES

Our equilibrium concept is Perfect Bayesian Equilibrium, so given any grading scheme, receiver beliefs must be consistent with sender choices and follow Bayes Rule, and sender choices must be best responses to receiver beliefs. If for a given grading scheme and cost c, there exists an equilibrium in which any positive measure of sender types applies for certification, we call that scheme *feasible*. We ignore pessimistic non-certification equilibria in which the sender never certifies because the receiver does not expect certification and off the equilibrium path views it as a sign of poor quality.

Three schemes will be of particular interest: exact, pass-fail, and honors. In Figure 2(a)'s *exact grading* scheme the product's quality is exactly revealed, with message $m = q$. Under exact grading, there will exist a quality level x such that all types $q \geq x$ have sufficient incentive to be certified, as we will explain below, but types below $q = x$ will not be certified. Since types $q \geq x$ are exactly revealed, the expected loss under exact grading consists of the loss from misestimating the quality of the uncertified senders in the quality interval $[0, x)$:

$$EL_{exact} = \int_0^x (q - E[q|q < x])^2 f(q)dq + \int_x^1 (0) f(q)dq. \quad (1)$$

If a scheme is not exact then it is *coarse*: the exact quality of at least some types is not revealed. We will later show that one of two coarse grading schemes, pass-fail or honors, will turn out to be optimal depending on the circumstances. Figure 2(b)'s *pass-fail grading* is the coarsest possible

quality to draw them away from (and destabilize) a pass-fail certifier.

meaningful grading scheme. The message is m = "uncertified" if $q \leq p$ and m = "pass" if $q > p$.[8] Assuming that p is set so all types $q > p$ have sufficient incentive to be certified, the expected loss from pass-fail grading is

$$(2) \qquad EL_{pass-fail} = \int_0^p (q - E[q|q \leq p])^2 f(q)dq + \int_p^1 (q - E[q|q > p])^2 f(q)dq.$$

The other optimal coarse grading scheme is Figure 2(c)'s *honors grading*, which sets a threshold h above which quality is revealed exactly but also divides types below h into two groups by a passing threshold p. If $q \leq p$ then m = "uncertified", if $p < q < h$ then m = "pass", and if $q \geq h$ then $m = q$. Assuming that p and h are set so that all types $q > p$ have sufficient incentive to be certified, the expected loss from honors grading is

(3)

$$EL_{honors} = \int_0^p (q - E[q|q \leq p])^2 f(q)dq + \int_p^h (q - E[q|p < q < h])^2 f(q)dq + \int_h^1 (0) f(q)dq.$$

For a scheme to be called "honors" we require strict inequalities: $p < h$ to distinguish it from exact grading, and $h < 1$ to distinguish it from pass-fail grading. In all three grading schemes, the certifier reports "uncertified" for a sender who fails to meet the certification standard, so such senders are pooled with senders who do not apply. This assumption is unimportant to the results, since in equilibrium a low-quality sender knows in advance that he would fail, and so does not incur the cost c to be certified.

II. Why Coarseness Helps

In this section we will use examples based on specific quality densities f to make three points. Figure 3(a)'s uniform density will show how coarse grading can improve on exact grading by increasing participation. Figure 3(b)'s falling triangle density will show how pass-fail grading can surpass not only exact grading but honors grading too. Figure 3(c)'s rising triangle density will show that coarse grading can be feasible when exact grading is not. These "can happen"

[8] Alternatively, the certifier could send the message "fail" if $q \leq p$, but in equilibrium low-quality firms will not apply so a "fail" grade is off the equilibrium path. Farhi, Lerner & Tirole (2013) show that revealing whether a firm tried to be certified but failed can be important when firms are uncertain of their own type, but that is not the case in our model.

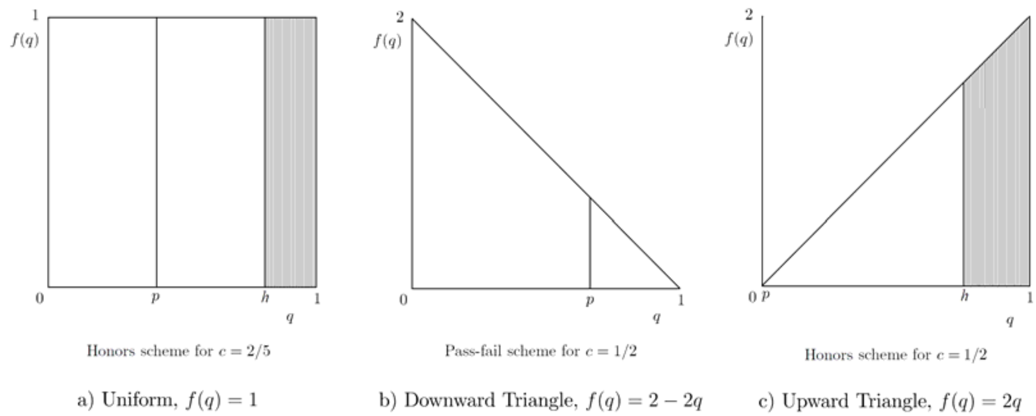

FIGURE 3. GRADING SCHEMES FOR DIFFERENT DISTRIBUTIONS OF PRODUCT QUALITY

examples will build the intuition behind Section III's general propositions.

A. *Coarse Grading Can Increase Information by Increasing Participation*

Suppose that the quality density f is uniform, as in Figure 3(a), and consider Figure 2(a)'s exact grading scheme. The sender's payoff $(q - c)$ from certification is increasing in q, so some type $q = x$ has the least incentive to be certified. Since the payoff from not being certified is $E[q|q < x] = x/2$, type $q = x$ is just indifferent between being certified and not if $x - c = x/2$, so $x = 2c$. Thus, for a quadratic loss function, the expected loss (mean squared error) is $EL_{exact} = \int_0^{2c} (q - c)^2\, dq = \frac{2}{3}c^3$ for the feasible range $c \leq 1/2$. As shown by Figure 4(a)'s "Exact" line, exact grading is perfectly informative as c approaches 0, since x also approaches 0. It is completely uninformative as c approaches $1/2$, since x approaches 1.

Now consider pass-fail grading. For the uniform density, $E[q|q \in (p, 1]] - E[q|q \in [0, p]] = (1 + p)/2 - p/2 = 1/2$, so any value of the cutoff p is feasible as long as $c \leq 1/2$. The most informative cutoff is $p = 1/2$, which from equation (2) has expected loss $EL_{pass-fail} = \int_0^{1/2} (q - 1/4)^2\, dq + \int_{1/2}^{1} (q - 3/4)^2\, dq = \frac{1}{48}$ in the feasible range. As seen from Figure 4(a)'s "P-F" line, pass-fail grading provides more information to receivers than exact grading when c is large enough that few types will be certified under exact grading. Although pass-fail grading provides only noisy information, more middling types are willing to be certified since they can pool with high types, and the extra information on these types more than compensates for the extra noise.

Pass-fail grading does better than exact grading for high c, but honors grading does even better. Honors grading cuts region $[0, x)$ in two using the passing standard p. Suppose we set $h = 2c$ instead of $x = 2c$, and set $p = c$, so the lower region is divided evenly. Types in the exact region $q \geq h$ now have more incentive to be certified, rather than look like a bad type who can't even pass. Types in the new pass region gain $E[q|q \in (p, h)] - E[q|q \in [0, p]] = (p+h)/2 - p/2 = h/2$ from passing so at $h = 2c$ this gain just covers the certification cost. Therefore all types $q \geq p = c$ will participate and $EL_{honors} = \int_0^c (q - c/2)^2 \, dq + \int_c^{2c} (q - 3c/2)^2 \, dq = \frac{1}{6}c^3$ as shown by Figure 4(a)'s "Honors" line. By allowing for more participation and continuing to provide exact information on high types, honors grading outperforms both exact and pass-fail grading.

Proposition 1 will show that this result that coarse grading is better than exact grading holds generally. Proposition 2 will show that, of all possible coarse grading schemes, either pass/fail or honors grading is the most informative.

B. Pass-Fail Grading Can Be Most Informative

Honors grading is more complex than pass-fail, but not always better. Consider Figure 3(b)'s falling triangle density $f = 2 - 2q$, which has the property that the gain from passing, $E[q|q \in (p, h)] - E[q|q \in [0, p]] = \frac{2}{3}\left(\frac{h}{2-p}\right)$, increases in p. When c is low, honors grading is best for this density. In this case, p can be set to divide the region $[0, 1]$ to minimize expected loss without the participation constraint being binding, so some additional types can be revealed exactly. From minimization of (2), the best division is at $p = \frac{3}{2} - \frac{1}{2}\sqrt{5} \approx 0.382$, which is feasible for $c < \frac{1}{3}\sqrt{5} - \frac{1}{3} \approx 0.412$, so within the range $c \in [0, .412]$ there is slack in the pass-fail participation constraint for those types that apply. Therefore types near $q = 1$ can be exactly revealed and $p \approx .382$ is still optimal and feasible, so honors grading does best. This is seen in Figure 4(b), where for low c, honors grading reduces expected loss relative to pass-fail grading.

As c rises, the participation constraint becomes binding and pass-fail becomes best. The gain from passing is increasing in the cutoff p, so p will have to be set higher to ensure participation.

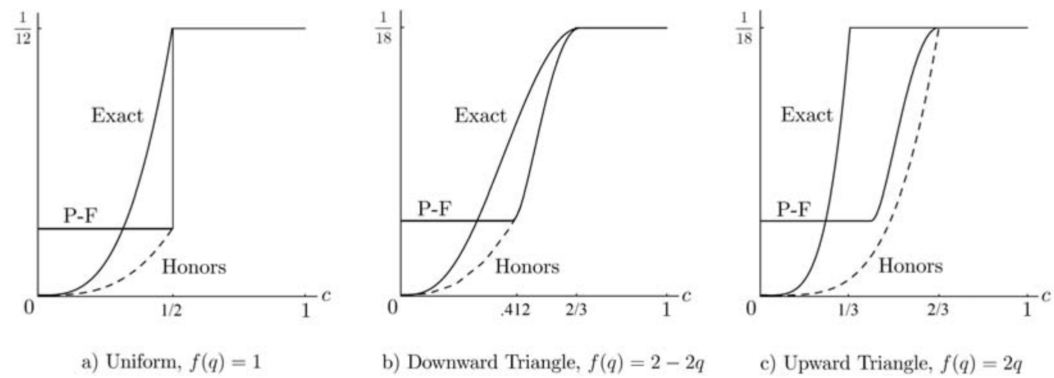

FIGURE 4. EXPECTED LOSS (MSE) VS. APPLICATION COST FOR DIFFERENT SCHEMES AND DIFFERENT DISTRIBUTIONS

If there is honors grading, so that h falls from $h = 1$ to $h < 1$, the gain from just passing falls, so p will have to be set even higher to ensure participation, to some value $p' > p$. Honors grading is no longer "for free" as with the uniform density; it comes at a tradeoff. It provides more information on types $q \in [h, 1]$, who are exactly revealed, and on types $q \in (p', h)$, who are in a smaller pooling group, but less information on uncertified types $q \in [0, p']$, who form a larger pooling group. If c is such that the uncertified group is already sufficiently large, the information loss from additional noise about the group dominates and the expected loss rises. Figure 4(b) shows this. For $c \gtrapprox .412$, honors grading is worse than pass-fail because any $h < 1$ is worse than $h = 1$, which is pass-fail grading.

This result that pass-fail grading can be most informative will be generalized and extended in Proposition 3.

C. *Coarse Grading Can Be Feasible When Exact Grading Is Not*

So far we have concentrated on how coarse grades can increase participation by more types and thereby increase informativeness. We now focus on feasibility. When can a scheme induce any participation at all? For the uniform density the three schemes induce different amounts of participation, but all of them can induce at least some participation for $c < 1/2$, as in Figure 4(a). Similarly, for the falling triangle density each scheme is feasible if $c < 2/3$, as in Figure 4(b).

However, the maximum certification cost is not always the same for all three schemes. As an example of how coarse grading can be feasible when exact grading is not, consider Figure 3(c)'s rising triangle density, $f(q) = 2q$. Under exact grading, the gap $x - E[q|q < x] = x - 2x/3 = x/3$ reaches a maximum of $1/3$ at $x = 1$, so exact grading is feasible for $c < 1/3$. Under pass-fail grading, the gap $E[q|q > p] - E[q|q \leq p] = \frac{2}{3}\left(\frac{1}{1+p}\right)$ is decreasing in p and converges to a maximum of $2/3$ as p approaches 0. Therefore, pass-fail grading is feasible for $c \leq 2/3$. Since honors grading can use an h arbitrarily close to 1, honors grading is also feasible for $c < 2/3$, as shown in Figure 4(c). For example, the honors scheme ($p = 0, h = 3/4$) shown in Figure 3(c) is feasible when $c = 1/2$, and indeed is more informative than pass-fail grading.

This result on the greater feasibility of coarse grading will be generalized and extended in Proposition 4.

III. Propositions for General Distributions of Quality

The above analysis compared the informativeness and feasibility of the exact, pass-fail, and honors grading schemes for particular distributions. We now extend our analysis to more general f and to other possible grading schemes.

Properties of means of a distribution conditional on being above or beneath a cutoff point t are central to the analysis. Section II's examples implicitly showed this and we will see it generally in this section. The crucial values are the upper mean above the cutoff, $A(t)$, the lower mean below the cutoff, $B(t)$, and the gap between the upper and lower means, $A(t) - B(t)$. Properties (i) and (ii) in Lemma 1 below are standard (see Bagnoli and Bergstrom, 2005). Properties (iii) and (iv) strengthen results from Jewitt (2004). Proofs are in the Appendix.[9]

LEMMA 1 (Properties of Upper and Lower Means): *Suppose density $f(q)$ is analytic with support on $[\underline{q}, \overline{q}]$ where $f(q) > 0$ for $q \in (\underline{q}, \overline{q})$ and define $A(t) \equiv E\left[q|q \geq t\right]$, $B(t) \equiv E\left[q|q \leq t\right]$. Then:*

(i) $A' > 0$ and $B' > 0$.

[9] We state the lemma for support on $[\underline{q}, \overline{q}]$ rather than [0, 1] to facilitate analysis of truncated distributions. Throughout the text and proofs we will use "strictly increasing" to mean the first derivative is strictly positive everywhere, which is sometimes referred to as "strongly", and similarly for strictly quasiconcave, etc.

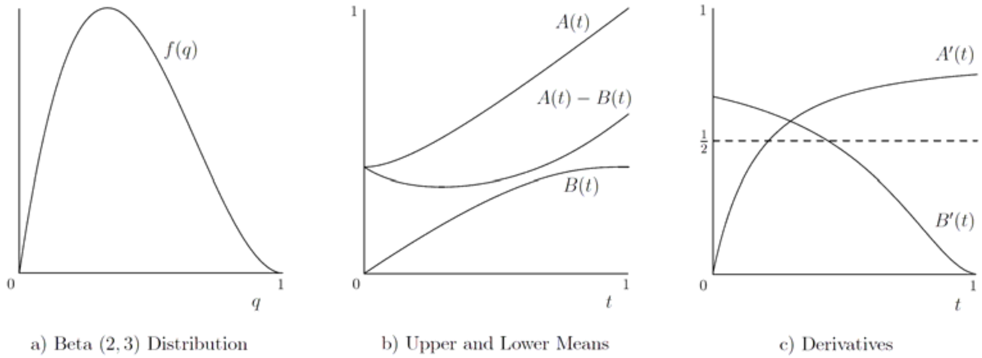

FIGURE 5. HOW THE LOWER AND UPPER MEANS AND THEIR DERIVATIVES CHANGE WITH THE CUTOFF

(ii) For strictly logconcave f, $A' < 1$ *and* $B' < 1$.
(iii) For strictly decreasing f, $A' \geq 1/2 \geq B'$ *(for strictly increasing* f, $A' \leq 1/2 \leq B'$*) with at least one inequality strict.*
(iv) For strictly quasiconcave $f(q)$*, the gap* $A-B$ *is strictly increasing iff* $f(\underline{q}) > (1/2)/(E[q]-\underline{q})$*, strictly decreasing iff* $f(\overline{q}) > (1/2)/(\overline{q}-E[q])$*, and strictly decreasing then increasing otherwise.*

As the cutoff t rises, some types who are below average in the upper region are shifted into the lower region where they are above average, so the means A and B of both the upper and lower regions rise as seen in Lemma 1(i) and Figure 5(b).[10] For exact grading we need to know more specifically how the gap $t - B$ changes and for pass-fail grading we need to know more specifically how the gap $A - B$ changes.[11] From Lemma 1(ii), if the density does not increase in slope too rapidly (i.e., if it is logconcave), then the upper and lower means do not rise too rapidly as the cutoff rises. In particular the result $B' < 1$ tells us that if there is a cutoff t such that all types above that cutoff are exactly graded ($x = t$), then raising t increases the mean of the uncertified region B at a rate slower than 1, as Figure 5(c) shows. Hence the gap $t - B$ between the marginal type who is graded exactly and the average quality of the uncertified pool

[10]E.g., “Professor Smith moved from Upstate U. to Downstate U., thus raising the quality of both.”
[11]The gap between upper and lower means is also central to binary signaling games (e.g., Benabou & Tirole 2006, 2011).

is increasing in t. Since this is the most the marginal type will pay to be certified, the cutoff for the most informative feasible exact grading scheme increases with certification costs c when f is logconcave. Most standard densities and all of our examples are logconcave (see Bagnoli and Bergstrom, 2005), so finding the most informative exact grading scheme is straightforward.

Lemma 1(iii) implies that for f increasing, $A' < B'$, and for f decreasing, $A' > B'$, which is Jewitt's result that for monotonic f the gap $A - B$ is monotonic in the opposite direction. This was the case for the triangle densities in Figure 3. A decreasing density puts relatively more mass at the lower end of each region, so a rise in t has more impact on the upper mean and the gap rises. The reverse is true for an increasing density. Monotonicity implies that the maximum gap is at either end of the support, so for a sufficiently high cost of certification the pass standard $(p = t)$ will be set either very high as in Figure 3(b) or very low as in Figure 3(c).

Lemma 1(iv) gives us exact conditions for the gap $A - B$ to be either monotonic or U-shaped in t for strictly quasiconcave (that is, unimodal) f. It implies Jewitt's result that for strictly quasiconcave densities the gap $A - B$ is strictly quasiconvex, and relaxes his conditions for $A - B$ to be monotonic by allowing for a small dip in the end of an otherwise increasing f or a small rise at the beginning of an otherwise decreasing f.[12] The distribution in Figure 5(a) has $f(0) = f(1) = 0$, so (iv) tells us that $A - B$ must be U-shaped as seen in Figure 5(b), but it also implies that a truncation sufficiently close to the mode on either side will instead ensure monotonicity of $A - B$ despite the lack of monotonicity of f.[13] Note also that $A(\underline{q}) - B(\underline{q}) = E[q] - \underline{q}$ and $A(\overline{q}) - B(\overline{q}) = \overline{q} - E[q]$, so when $A - B$ is U-shaped the maximum gap is at whichever end is further from $E[q]$.

For our first proposition we will generalize our finding that coarse grading outperforms exact grading for the uniform distribution. Recall that for the uniform distribution, turning exact into honors grading by introducing a "pass" region is always possible. We start with the exact scheme

[12] To see this suppose that f is strictly increasing on $[0, 1]$. Since F is strictly convex $\int F(q)dq$ encloses the right triangle with base on $[1 - \frac{1}{f(1)}, 1]$, hypotenuse with slope equal to $f(1)$, and height $F(1) = 1$. The area of this triangle is $\frac{1}{2f(1)}$, so $\int F(q)dq > \frac{1}{2f(1)}$, or $1 - E[q] > \frac{1}{2f(1)}$ implying the condition in (iv). This holds strictly even if $f'(1) = 0$, so extending the support to $[0, 1 + \varepsilon]$ where $f'(1 + \varepsilon) < 0$ still satisfies the condition for $\varepsilon > 0$ small enough. Similar logic holds for f decreasing except for a small initial rise.

[13] Truncation to $[.21, 1]$ implies the gap $A - B$ is increasing even though f is increasing on $[.21, 1/3]$, and truncation to $[0, .44]$ implies the gap is decreasing even though f is decreasing on $[1/3, .44]$. From the proof of Lemma 1(iv), these cutoffs follow since $A'(.21) = 1/2$ and $B'(.44) = 1/2$, as seen in Figure 5(c).

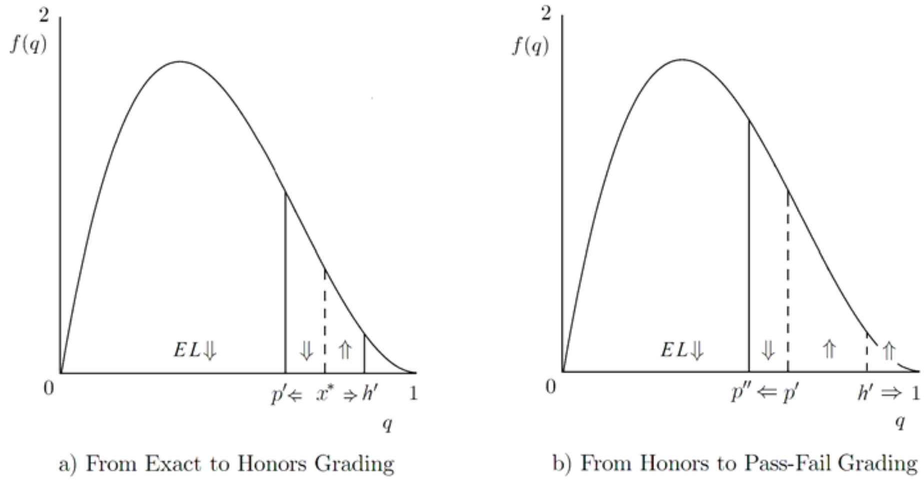

FIGURE 6. EFFECT OF COARSENESS ON EXPECTED LOSS

x^* and introduce honors grading by setting $p = p'$ slightly below $h = x^*$. Will types in (p', x^*) still participate? Yes, for uniform f, because the gain from passing relative to failing, $E[q|q \in (p, x^*)] - E[q|q \in [0, p]] = x^*/2$, is constant in p for a given x^*. For the rising triangle distribution and more generally for any rising f as shown in Lemma 1(iii), $E[q|q \in (p, x^*)] - E[q|q \in [0, p]]$ is decreasing in p, so it is always possible to introduce a pass region with $p' < x^*$ that does not affect the proportion of types who are exactly graded. More generally though, if creating the pass region does not reduce the fail region's mean enough then the participation constraint is no longer met for $h = x^*$. Hence, h will have to rise to maintain the incentive to participate. This causes a loss in information, so it is no longer clear whether there is a net benefit from introducing the pass region. This tradeoff arose with the falling triangle distribution. Lemma 1(iii) tells us that it arises for all decreasing f, and Lemma 1(iv) tells us it can arise generally with unimodal distributions.

Figure 6(a) shows a Beta (2,3) distribution where the pass cutoff is in the region of decreasing f so that one might expect the same tradeoff. Since $f(0) = f(1) = 0$, by Lemma 1(iv) the gap is increasing in p after some internal minimum, and the same tradeoff does indeed arise.[14] For

[14]The density of the Beta (α,β) distribution is logconcave (and hence quasiconcave) for all $\alpha, \beta \geq 1$ (Bagnoli and Bergstrom, 2005), is strictly increasing for all $\alpha > \beta = 1$ including the rising triangle distribution Beta (2,1), is strictly decreasing for all $\beta > \alpha = 1$ including the falling triangle distribution Beta (1,2).

cost $c = 2/5$, introducing a pass region with $p' < x^*$ requires increasing the exact cutoff from x^* to some $h' > x^*$ so as to maintain participation. This change implies types in the fail interval $[0, p']$ contribute less to expected loss than under exact grading, because the pool is smaller, with each type closer to the interval's mean. Types in (p', x^*) also contribute less, since they are moved from the larger fail interval to the new pass region. On the other hand, types in $[x^*, h')$ contribute more to expected loss since they are no longer exactly revealed.

The following proposition says that despite the loss from putting some formerly revealed senders into a pool, it is always possible to find a coarse scheme that outperforms the best exact grading scheme. For marginal reductions in p' below x^* there is a first-order effect on types in $[0, p']$, but only a second-order effect on types in (p', h') because the pass interval is so small. The result holds for Figure 6(a)'s example, and more generally. There is a tradeoff, because fewer types will be exactly revealed, but more types will have at least a pass certification.

PROPOSITION 1 (Optimality of coarse grading): *Exact grading is never optimal.*

PROOF: See the Appendix.

Moreover, of all possible grading schemes — not just the three we have highlighted — the best is either pass-fail or honors. Introducing numerous grading intervals, nonmonotonic grades, exact grading in the middle of the quality range, etc. is unnecessary.

PROPOSITION 2 (Simple schemes): *Either pass-fail or honors grading is optimal.*

PROOF: See the Appendix.

Proposition 2 says that either pass-fail or honors grading always does as well or better than more complicated schemes. The improvement over pure exact grading comes from the introduction of a pass pool that increases the number of senders certified. The certifier's fundamental problem is to make the lowest pass pool more attractive than being uncertified. If that can't be done, no higher pass pool or exact grading interval will be attractive to senders. At the same time, the ideal is to perfectly reveal the sender's type. Thus, the certifier's real tradeoff is between (a) putting high quality types into the lowest pass pool and (b) revealing them exactly. In doing this, the certifier also wants to keep the lowest pass pool connected, so that the difference between its extreme types and their mean is smaller and receivers can form a more exact estimate. This is best done with a single pass interval, with exact grading for the best types — if that does not come at too great a cost to the attractiveness of the pass pool.

We now consider which is best, pass-fail or honors grading, in particular circumstances.

PROPOSITION 3 (Pass-fail vs. honors): *(i) As the certification cost c tends to 0, honors grading is optimal, with the mass of exactly revealed types tending to 1 and the mass of pass types tending to 0. (ii) For f quasiconcave, as the certification cost c tends to the maximum feasible level, the mass of exactly revealed types tends to 0, and the mass of pass types tends to 0 if* $E[q] < 1/2$ *and to 1 if* $E[q] > 1/2$.

PROOF: See the Appendix.

The first part of Proposition 3 states that honors grading is best for c sufficiently small. In this case it is feasible to exactly grade even very low types, so it is best to have a small pass region (as required by Proposition 1) and then exactly grade better types above a low honors cutoff.

The second part states that if c is sufficiently high the best scheme has both a small pass region and a small honors region, no honors region at all, or a large pass region and a small honors region. To understand why the honors region disappears for high certification costs, look at Figure 6(b), which starts with the honors scheme (p', h') that was in Figure 6(a). Figure 6(b) then increases the pass region by dropping p' to p'', which requires increasing h to maintain feasibility. Since costs are high and we started towards the upper end of the density, as we increase the pass region, the honors region shrinks to nothing and we arrive at a pure pass-fail scheme with $h = 1$. Types in the interval $[0, p'']$ contribute less to total expected loss since the pool is tightened, so each type is closer to the conditional mean. Types in (p'', p') also contribute less since the types in this region have been moved from a larger to a smaller pool. Types in (p', h') contribute more information loss, however, because the pool these types are in has expanded, as do types in region $[h', 1]$ because they were formerly exactly revealed. Overall we have gained information on types in $[0, p']$ and lost information on types in $(p', 1]$. Numerically, this adds up to a net improvement in information for the case in the figure, so $h = 1$ is optimal and pass-fail is best.

The role of $E[q]$ in Proposition 3(ii) arises from the Lemma 1(iv) result that for quasiconcave f the gap $E[q|q \geq p] - E[q|q \leq p]$ reaches a global maximum at $p = 0$ or $p = 1$. When $E[q] < 1/2$, as in Figure 6, the $p = 1$ case holds, so for large c the optimal feasible pass-fail scheme is close to $p = 1$ and it is best to have no honors region. When $E[q] > 1/2$ the $p = 0$ case holds, so for large c feasibility requires p close to 0 and h close to or equal to 1.

So far we have focused only on the informativeness of different schemes. Sometimes exact grading is not only less informative but infeasible. A confused certifier who insisted on exact grading would find that nobody would show up to be graded! This was the case for the rising triangle distribution of Figure 4(c), where exact grading was not feasible even for many values of c for which pass-fail and honors remained feasible. Proposition 4 says that honors and pass-fail are always feasible when exact grading is, and then gives general conditions for when they are feasible even if exact grading is not.

PROPOSITION 4 (Feasibility): *(i) For any quality density f, if exact grading is feasible then so is honors grading, and if honors is feasible then so is pass-fail. (ii) Pass-fail and honors grading are both feasible for a range of grading costs so high that exact grading is not if:*
(a) $f(q)$ is strictly increasing; or
(b) $f(q)$ is strictly quasiconcave and $f(1) > (1/2)/(1 - E[q])$; or
(c) $f(q)$ is strictly logconcave and $E[q] > 1/2$.

PROOF: See the Appendix.

Proposition 4 shows the robustness of coarse grading even when the certifier might have different objectives than we have assumed. Exact grading runs a greater risk of falling apart completely because of refusal to participate.

IV. Extension: Letter Grading with Different Receiver Priors on Different Senders

We have followed the literature's standard assumption of a single sender drawn from a distribution or, equivalently, multiple senders from the same distribution. Now suppose the certifier must use the same grading scheme for multiple senders when it is common knowledge that senders have different quality distributions. For instance, consumers might know that one firm

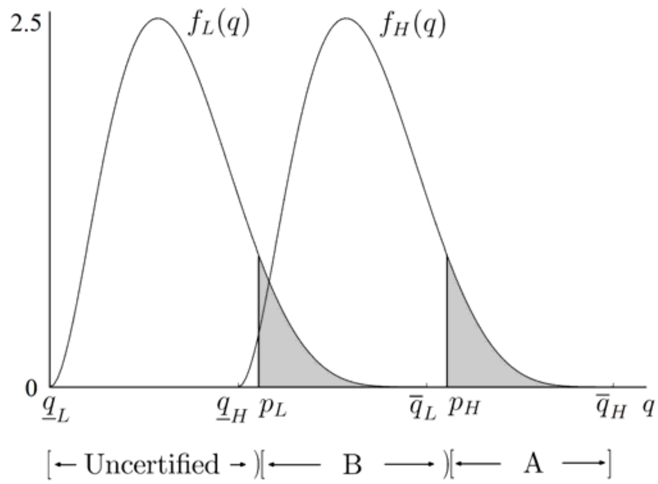

FIGURE 7. LETTER GRADING WITH DIFFERENT PRIORS

is likely to be better than another in environmental quality— they have sender-specific prior information— but only the firms know their exact quality.

We are particularly interested in when multi-tier certification in the form of "letter grades" is optimal. For example, the "LEED" certification system for building environmental impact has "Certified," "Silver", "Gold", and "Platinum" categories. For a firm with a good reputation, receiving just a "Silver" rating might not be worth the certification cost, but for a firm with a bad reputation such a rating might well be worth it. Hence, having different tiers might increase participation when consumers have different prior distributions about different firms.

Consider a setting in which sender qualities follow logconcave densities and receivers know whether a sender is drawn from a lower-range distribution with density f_L over $[\underline{q}_L, \overline{q}_L]$ and a higher-range distribution with density f_H over $[\underline{q}_H, \overline{q}_H]$. If the two distributions don't overlap ($\overline{q}_L < \underline{q}_H$), and both distributions have means closer to their lower support, the optimum for sufficiently high costs is, from Proposition 3, pass-fail for each distribution of sender. Let the optimal pass-fail cutoffs be p_L and p_H for the respective distributions. We can reframe these two pass-fail schemes as a system with four grades: "A", "B", "C", and "Uncertified". In the example, no B's would be observed since types would not want to pay c to be certified as being in the bottom region of their distribution.

Now suppose the two densities overlap moderately with $\underline{q}_L < \underline{q}_H < p_L < \overline{q}_L < p_H < \overline{q}_H$ as in Figure 7,[15] keeping the same relative values for the cutoffs $p_L - \underline{q}_L$ and $p_H - \underline{q}_H$ that were optimal when there was no overlap. Suppose the certifier assigns "Uncertified" for $q \leq p_L$, "B" for $q \in [p_L, p_H)$, and "A" for $q \geq p_H$. The senders which apply cannot do better by being uncertified, because the cutoffs were chosen in the original example to make this unprofitable and nothing has changed in the senders' incentives. The high-distribution senders in $(p_L, p_H]$ that do not apply would receive B's if they unexpectedly applied, but the cost c is too high for that to benefit them, since for any beliefs, the expected payoff is strictly less than for receiving an A and p_H has been set so that types in the A region are just indifferent to certification. Hence, given the grading scheme, it is a Perfect Bayesian Equilibrium for high-distribution senders in $(p_H, \overline{q}_H]$ to apply and get A's, low-distribution senders in $(p_L, \overline{q}_L]$ to apply and get B's, and the remaining senders to not apply. Is the grading scheme still optimal for the certifier? For sufficiently high c, p_H is arbitrarily close to $\overline{q}_H$, so $p_H > \overline{q}_L$ as in Figure 7 and there is effectively no interaction between grading for the different distributions. Hence the scheme that is individually best for each distribution remains the best for the combined case.

This shows that when receivers have information about sender quality that makes the prior distributions for each sender differ, the combination of that information with certifier grading can lead to more complicated grading schemes being optimal, including those with multiple grades.[16] Based on this example we have the following proposition.

PROPOSITION 5 (Letter grades): *If the same grading scheme must be used for senders with different quality distributions, then sometimes the most informative scheme uses multiple coarse letter grades and does not report any quality exactly.*

V. Literature Discussion and Conclusion

We have shown that a certifier who is trying to maximize information to the public should, paradoxically, coarsen his information before reporting it. Rather than simply revealing what he

[15] Figure 7 uses two Beta (3,6) distributions, one with support on [0, 1] and the other renormalized with support on $[1/2, 3/2]$. The shown cutoffs p_L and p_H are optimal for $c = 1/3$.

[16] Other models can also generate multi-tier certification. When abilities are heterogeneous, Dubey & Geanakoplos (2010) show that letter grades can maximize effort from status-concerned students by forcing them to compete for a limited number of good grades. Farhi, Lerner, and Tirole (2013) consider different pass-fail standards set by different certifiers.

has measured, the certifier will reveal only part of what he knows, and in some situations will only reveal whether a sender passes a quality threshold. The certifier faces a tradeoff between coarse grading, which attracts more senders to be certified, and fine grading, which informs the public better about the senders attracted. We show that the optimal tradeoff always involves some coarseness.

Two strands of the literature are most closely related to the situation we model. The first looks at the alternative setting in which a for-profit certification intermediary designs a scheme to maximize rent extraction from senders afraid of receiver beliefs about their quality if they lack certification. Lizzeri (1999) shows that the profit-maximizing scheme is pass-fail with all but the very worst type of sender receiving a passing grade. By certifying almost everyone, the certifier can extract a certification fee from almost everyone. Since each sender is afraid of being pooled with the minute number of bottom-quality senders who remain uncertified, the fee can be large. Our model differs in assuming that the certifier aims to maximize receiver information, as in the case of a non-profit certifier who charges just enough to cover the costs of certification or who charges nothing but requires senders to bear some cost in providing information. The certifier introduces coarseness not to gain profits, but to induce participation by senders who want to distinguish themselves from worse senders but do not want to be distinguished from even better senders.

A second strand of the literature concerns certifiers who maximize the senders' benefit by obscuring the information the public receives. Indeed, the ideal would be to perfectly fool the public. In these models the certifier can also be thought of as a sender who precommits to an information policy that maximizes his ex ante payoff before learning his type. Ostrovsky & Schwarz (2010) consider when colleges can maximize student job prospects by grading schemes which pool weaker and stronger students together. The average quality of students at Yale might be so high that they will all get jobs if grades are uniformly As, but weaker students might not get jobs if they are revealed as weak.[17] Rayo & Segal (2010) also consider the gains from pooling, and Kamenica & Gentzkow (2011) and Gentzkow & Kamenica (2014) analyze the general

[17]Pooling can also maximize a biased sender's payoff even without commitment or verifiable information (Chakraborty and Harbaugh, 2007). Note that in the current model the certifier commits to withhold some information to encourage participation despite the ex post incentive to reveal all information.

Bayesian persuasion problem of using such concavification strategies to maximize the sender's payoff. Our environment with a certification intermediary between the sender and receiver does not reduce to a standard Bayesian persuasion problem in which the sender sets an information structure. Although the certifier wants to induce participation by raising posterior estimates above a threshold, which is encouraged by pooling as in Bayesian persuasion, such pooling also reduces information to the receiver, which hurts the certifier. The certifier would always reveal his information to receivers exactly were it not for the need to encourage participation by senders.

Other important reasons for coarse grading have also been analyzed in the literature but we have abstracted from them in our model. We assume that coarse and fine grading are equally costly. Titman & Trueman (1986), Farhi, Lerner & Tirole (2011), and others consider the case where finer grading costs more. We assume that the certifier can provide verifiable information about quality, but it might be that the certifier's credibility is not assured, in which case a coarse report can be more credible than an exact report (Crawford & Sobel, 1982; Morgan & Stocken, 2003; Chakraborty & Harbaugh, 2007). If there are a large number of certifiers with different objectives, grades can be coarse because each certifier optimally chooses a different pass-fail standard (Lerner & Tirole, 2006).[18] If the audience can be overwhelmed by too much information, then even if they are rational their differing use of it may arrive at a worse result, either because they choose to acquire less (Eppler & Mengis, 2004) or because it worsens coordination problems (Chahrour, 2014).

We take quality as exogenously given, but certification can affect the incentive of senders to invest in quality. In our model informed consumers are more likely to buy higher rather than lower quality goods so there is a positive incentive effect from revealing information, but we do not model the exact incentive effect. In some situations maximizing effort itself might be the certifier's goal. Costrell (1994) considers how high to set a pass-fail standard to maximize student effort taking the pass-fail system as given. In contest environments, Moldovanu, Sela & Shi (2007) and Dubey & Geanakoplos (2010) show that coarse grades can induce more competition

[18] They assume a continuum of certifiers who each put some different, varying weight on a mix of firm profits and consumer surplus, thus leading to a continuum of different standards that firms can pick from. Since firms choose the hardest standard they can meet, exact information is revealed.

when abilities are heterogeneous. Boleslavsky & Cotton (2015) analyze competition between schools to help students in the job market, and find that schools will not provide exact information on students, which in turn induces the schools to exert more effort to increase the quality of both good and bad students.

Our approach adds to this rich literature by showing how coarseness is optimal in the environment that would seem least conducive to it — when the certifier is explicitly trying to maximize information to the public and the testing process produces a continuous score that the certifier could choose to report. Related to our approach, Rosar & Schulte (2012) address the design of tests to minimize weighted mean squared error when the quality of risk-averse agents is high or low. They find that for a risk averse agent, a pass-fail test with no false positives is often optimal because it induces agents to volunteer for the test despite the risk of ending up with a worse public image.

Our results depend on there being some cost of certification so coarseness can play an important role in encouraging participation. Studies find that certification costs can be a substantial fraction of total costs (Vitalis, 2002),[19] and that the process of applying for certification can be lengthy and compete for limited managerial resources.[20] Given these costs, the decision of whether to certify product quality is highly dependent on the effect on buyer willingness to pay. Our analysis shows how the coarseness of the grading scheme affects buyer estimates of product quality and thereby affects the incentive of marginal senders to participate.

An alternative explanation for coarse grading is that receivers have difficulty processing exact information. This explanation is at odds with the detailed information available on consumer evaluation websites as discussed in the Introduction. Nevertheless, there have been proposals for government agencies to make mandatory labels less exact so as to help consumers. The Energy Policy Act of 2005 directed the Federal Trade Commission to consider switching the mandatory EnergyGuide label to a coarse star-ranking scheme for this reason. However, after reviewing the evidence on how consumers use labels and performing its own tests, the FTC

[19]The main association of small and medium businesses in the European Union listed its primary requested revision in eco-label policy as, "An overall reduction of the costs, in particular the costs of the technical tests required in order to show the respect of the criteria." (See "UEAPME's Position on the Revision of the Eco-label Regulation," UEAPME, November, 2008), p. 4.

[20]The *2010 Global Ecolabel Monitor* found that the average time between filing the application for an ecolabel and being awarded the label was 4.3 months.

determined that consumers learned most from exact information about expected energy costs (Farrell, Pappalardo & Shelanski, 2010).[21] From the perspective of our model, this is consistent with the use of coarse labels by non-governmental organizations being driven not by consumer difficulty in understanding exact labels but by the need to encourage sender participation. For mandatory schemes where participation incentives are not a factor, shifting away from exact grades would hurt rather than help consumers.

VI. Appendix

LEMMA 1 (Properties of Upper and Lower Means): *Suppose density* $f(q)$ *is analytic with support on* $[\underline{q}, \overline{q}]$ *where* $f(q) > 0$ *for* $q \in (\underline{q}, \overline{q})$ *and define* $A(t) \equiv E\left[q|q \geq t\right]$, $B(t) \equiv E\left[q|q \leq t\right]$. *Then:*

(i) $A' > 0$ *and* $B' > 0$.

(ii) For strictly logconcave f, $A' < 1$ *and* $B' < 1$.

(iii) For strictly decreasing f, $A' \geq 1/2 \geq B'$ *(for strictly increasing* f, $A' \leq 1/2 \leq B'$*) with at least one inequality strict.*

(iv) For strictly quasiconcave $f(q)$, *the gap* $A-B$ *is strictly increasing iff* $f(\underline{q}) > (1/2)/(E[q]-\underline{q})$, *strictly decreasing iff* $f(\overline{q}) > (1/2)/(\overline{q} - E[q])$, *and strictly decreasing then increasing otherwise.*

Proof: (i) Integrating by parts,

$$E[q|q \in [a,b]] = \frac{\int_a^b f(q)q dq}{F(b) - F(a)} = \frac{bF(b) - aF(a)}{F(b) - F(a)} - \frac{\int_a^b F(q)dq}{F(b) - F(a)}. \quad (4)$$

[21] The EPA made a similar analysis of how to represent information about greenhouse gases and smog damage in the new version of its mandatory gas mileage labels, and chose to use a fine 1–10 scale. It also chose to continue reporting exact mileage and gasoline cost information rather than coarsen the information. See http://epa.gov/otaq/carlabel/labelcomparison.htm.

Applying this to $A(t)$ and $B(t)$,

$$A'(t) = \frac{d}{dt}E[q|q \in [t,\overline{q}]] = \frac{f(t)\cdot(\overline{q} - t - \int_t^{\overline{q}} F(q)\,dq)}{(1-F(t))^2}$$

$$(5) \qquad = \frac{f(t)}{1-F(t)}(A-t)$$

$$B'(t) = \frac{d}{dt}E[q|q \in [\underline{q},t]] = \frac{f(t)\int_{\underline{q}}^t F(q)dq}{F^2}$$

$$(6) \qquad = \frac{f(t)}{F(t)}(t-B).$$

Equations (5) and (6) imply that $A'(t) > 0$ for $t < \overline{q}$ and $B'(t) > 0$ for $t > \underline{q}$. As for $t = \underline{q}$, first suppose $f(\underline{q}) > 0$. We will start with $B'(t)$. By l'Hopital's rule,

$$\lim_{t\to\underline{q}} B'(t) = \lim_{t\to\underline{q}} \frac{f(t)\int_{\underline{q}}^t F(q)dq}{F(t)^2} = \lim_{t\to\underline{q}} \frac{f'(t)\int_{\underline{q}}^t F(q)dq + f(t)\cdot F(t)}{2f(t)F(t)}$$

$$= \frac{1}{2} + \lim_{t\to\underline{q}} \frac{f''(t)\int_{\underline{q}}^t F(q)dq + f'(t)F(t)}{2f'(t)F(t)+2f(t)^2}$$

$$(7) \qquad = \frac{1}{2} > 0.$$

If, instead, $f(\underline{q}) = 0$, then applying l'Hopital's rule n more times until $f^{(n)}(\underline{q}) \neq 0$, yields[22]

$$(8) \qquad B'(\underline{q}) = (n+1)/(n+2) \geq 1/2 > 0.$$

One may obtain $A'(\overline{q}) \geq 1/2 > 0$ by similar operations.

(ii) Logconcavity is inherited by integration (Prekopa, 1973), so logconcavity of f implies logconcavity of F and hence of $\int_{\underline{q}}^t F(q)dq$. Logconcavity of $\int_{\underline{q}}^t F(q)dq$ implies $f(t)\int_{\underline{q}}^t F(q)dq < F(t)^2$; see $\frac{d^2}{(dt)^2}\ln(\int_{\underline{q}}^t F(q)dq)$. From (5), $f(t)\int_{\underline{q}}^t F(q)dq < F(t)^2$ implies $B'(t) < 1$.

[22]For $f(q) > 0$, the truncated distribution converges to a uniform distribution, and for $f(q) = 0$ and $f'(q) = 0$ the truncated distribution converges to a triangle distribution. If $f^{(n)}(q) = 0$ for all n, then since f is analytic, $f(q) = 0$ in the neighborhood of $\underline{q}$, which contradicts the assumption $f > 0$ for $q \in (\underline{q}, \overline{q})$.

Similarly, logconcavity of $f(q)$ implies the reliability function $1 - F(q)$ is logconcave (see Bagnoli and Bergstrom, Theorem 3, 2005). Inheritance of logconcavity by integration therefore implies $\overline{q} - t - \int_t^{\overline{q}} F(q)dq$ is logconcave, which implies $f(t) \cdot (\overline{q} - t - \int_t^{\overline{q}} F(q)dq) < (1 - F(t))^2$, which from (5) implies $A'(t) < 1$.

(iii) Differentiating (5) and (6) and substituting,

$$
\begin{aligned}
A''(t) &= \tfrac{f'(t)}{f(t)} A' + \left(2A' - 1\right) \tfrac{f(t)}{1-F(t)} \\
B''(t) &= \tfrac{f'(t)}{f(t)} B' + \left(1 - 2B'\right) \tfrac{f(t)}{F(t)}.
\end{aligned}
\tag{9}
$$

First consider f decreasing so $f(\underline{q}) > 0$. From (7), $B'(\underline{q}) = 1/2$, and from (9), $f' < 0$ implies $B'' < 0$ evaluated at any t such that $B' = 1/2$, so B' cannot rise above $1/2$ for any t. Hence $B' \leq 1/2$ with equality only at $t = \underline{q}$. Similarly, $A'(\overline{q}) \geq 1/2$, and from (9), $f' < 0$ implies $A'' < 0$ evaluated at any t such that $A' = 1/2$, so A' cannot fall below $1/2$ for any t. So $A' \geq 1/2$ with possible equality only at $t = \overline{q}$. Applying the same logic for f strictly increasing, $A' \leq 1/2 \leq B'$, with equalities possible only at $t = \overline{q}$ and $t = \underline{q}$ respectively.

(iv) We first establish quasiconvexity. Since the gap $A - B$ is twice differentiable, it is quasiconvex if $A' = B'$ implies that $A'' \geq B''$. First consider $A' = B' > 1/2$. From (9) this implies $A'' > B''$, as required. Now consider $A' = B' \leq 1/2$. Strict monotonicity (and hence quasiconvexity) follows from Lemma 1(iii) if f is monotonic, so suppose it is not and $\widehat{q}$ is its internal mode. Strict quasiconcavity of f implies that f is strictly increasing in $[\underline{q}, \widehat{q}]$ and strictly decreasing in $[\widehat{q}, \overline{q}]$. So from Lemma 1(iii) $B' \geq 1/2$ in $[\underline{q}, \widehat{q}]$, with possible equality only at $t = \underline{q}$, and $A' \geq 1/2$ in $[\widehat{q}, \overline{q}]$ with possible equality only at $t = \overline{q}$. Hence $A' = B' < 1/2$ is not possible, and $A' = B' = 1/2$ is only possible at $t \in \{\underline{q}, \overline{q}\}$, in which case (9) implies $A'' = B''$, as required.

Given that $A - B$ is quasiconvex, it is strictly quasiconvex if the set of t such that $A' = B'$ has measure zero. There is at most one t where $A' = B' > 1/2$ since $A'' > B''$ at any such point, implying no other crossings are possible. And as shown above there are at most two points, $t \in \{\underline{q}, \overline{q}\}$, where $A' = B' = 1/2$. Hence $A - B$ is strictly quasiconvex.

By strict quasiconvexity, $A - B$ is either strictly monotonic or first strictly decreasing and then strictly increasing. Therefore for $A - B$ to be strictly increasing it is necessary and sufficient that it be strictly increasing at the lower bound. Note from (5) that $A'(\underline{q}) = f(\underline{q})(E[q] - \underline{q})$ and that $B'(\underline{q}) = 1/2$ for $f(\underline{q}) > 0$. Therefore $f(\underline{q})(E[q] - \underline{q}) > 1/2$ is equivalent to $A - B$ strictly increasing. Similarly, for $A - B$ to be strictly decreasing it is necessary and sufficient that it be decreasing at the upper bound, $A'(\overline{q}) < B'(\overline{q})$ or, $1/2 < f(\overline{q})(\overline{q} - E[q])$. If neither condition holds, then by strict quasiconvexity it must be that $A - B$ is first strictly decreasing then strictly increasing. ■

PROPOSITION 1 (Optimality of coarse grading): *Exact grading is never optimal.*

Proof: Consider a feasible exact scheme x, where all types $q \geq x$ apply for certification, and an honors scheme (p, h). If an honors scheme with $p < x$ and $h = x$ is feasible it clearly has lower expected loss than the exact scheme. Suppose this is not the case, so instead $p < x$ requires $h > x$ for feasibility. Let $p(h)$ be a continuous decreasing function on $[x, 1]$ that picks a feasible p, where $p(x) = x$ is the exact scheme. Such a function must exist in the right-neighborhood of $h = x$ since $p(x) = x$ is feasible by assumption and $E[q|q \in (p, h)] - E[q|q \in [0, p]]$ is increasing in h from Lemma 1(i).

Setting $p = p(h)$ in (3), the marginal impact on expected loss of raising h from $h = x$ to create an honors scheme with $(p(h), h)$ is

$$
\begin{aligned}
&\tfrac{dp}{dh}\left(p(h) - E[q|q \leq p(h)]\right)^2 f(p(h)) \\
&+2\left(\tfrac{d}{dp}E[q|q \leq p(h)]\tfrac{dp}{dh}\right)\int_0^{p(h)}(q - E[q|q \leq p(h)])\, f(q)dq \\
&+\left(h - E[q|q \in (p(h), h]]\right)^2 f(h) \qquad\qquad (10)\\
&-\tfrac{dp}{dh}\left(p(h) - E[q|q \in (p(h), h]]\right)^2 f(p(h)) \\
&+2\left(\tfrac{d}{dh}E[q|q \in (p(h), h]]\right)\int_{p(h)}^{h}(q - E[q|q \in (p(h), h]])\, f(q)dq.
\end{aligned}
$$

The first term of (10) is negative because $\frac{dp}{dh} < 0$ by construction, while the second term is zero since the mean $E[q|q \leq p]$ minimizes $\int_0^p (q - E[q|q \leq p])^2 f(q)dq$. The third, fourth, and fifth terms are zero evaluated at $h = x$ since $p(x) = x$ by construction. Hence, starting from the exact scheme $p = h = x$, expected loss can always be reduced by creating an honors scheme. ■

PROPOSITION 2 (Simple schemes): *Either pass-fail or honors grading is optimal.*

Proof: Proposition 1 rules out exact schemes from being optimal. Any other scheme has at least one pool, i.e., a set of types sending the same message. Denote by $\underline{p}$ the lowest cutoff for any certified pool and by $\underline{h}$ the lowest cutoff for any exact grading interval. We can rule out schemes with $\underline{p} \geq \underline{h}$ because the pooled types above $\underline{p}$ could be exactly revealed without affecting feasibility. We can also rule out schemes where there are uncertified types above $\underline{h}$ since this those types could be revealed exactly, while lowering the mean of the uncertified types, thereby improving the feasibility of all certified messages. Therefore, any potentially superior scheme different from pass-fail or honors must be one of three types: (i) there are multiple pools of certified types, or (ii) there is one pool but it is split by an interval of exact grading and/or (iii) there is one pool, but some types above $\underline{p}$ and below $\underline{h}$ are uncertified. We will show that none of these alternatives is better than pass-fail or honors grading.

(i) Suppose there are multiple certified pools with means v_i for $i = 1, ..., N$. If any pools have different means then there is slack in the feasibility constraint for all but the lowest v_i pool, and we could exactly reveal some types in those pools and maintain feasibility. So the optimal such scheme must have all v_i equal to some value $\overline{v}$ sufficiently above the mean of the certifieds. But then all the messages for the certified pools convey the same expected value. Each pooled type q contributes $(q - \overline{v})^2 f(q)$ to total loss, so one pool is as good as multiple pools.

(ii) Suppose the one pool of certified types, denoted by P, which we now refer to as the "pass types", is split so some of the pass types are above $\underline{h}$. Noting from (iii) below that there will be no uncertified types above $\underline{p}$, the lower interval of pass types is $(\underline{p}, \underline{h})$. Let $(\underline{g}, \overline{g})$ be the lowest interval of pass types above $\underline{h}$, so $\underline{g} > \underline{h}$. If $\underline{g} = \overline{g} = 1$ there is no information gain. Otherwise, we can use total differentiation to show how $\underline{h}$ must rise to keep the pool mean equal to its original mean $\overline{v}$ as we increase $\underline{g}$, and then show that this will reduce loss, so having the upper passing interval is suboptimal. Noting that there may be more than one interval of types in P above $\underline{h}$,

$$\overline{v} = \frac{\int_{\underline{p}}^{\underline{h}} q f(q)dq + \int_{\underline{g}}^{\overline{g}} q f(q)dq + \int_{q \in P \cap q > \overline{g}} q f(q)dq}{\int_{\underline{p}}^{\underline{h}} f(q)dq + \int_{\underline{g}}^{\overline{g}} f(q)dq + \int_{q \in P \cap q > \overline{g}} f(q)dq}. \tag{11}$$

If $\overline{v} > \underline{h}$ then pooled types are viewed more positively than the lowest exactly revealed type, so there is slack in the feasibility constraint for the pooled types and some can be revealed exactly to reduce loss. Thus the case of interest is $\overline{v} < \underline{h}$. Multiplying out (11) yields $\overline{v}(\int_{\underline{p}}^{\underline{h}} f(q)dq + \int_{\underline{g}}^{\overline{g}} f(q)dq + \int_{q \in P \cap q > \overline{g}} f(q)dq) - (\int_{\underline{p}}^{\underline{h}} q f(q)dq + \int_{\underline{g}}^{\overline{g}} q f(q)dq + \int_{q \in P \cap q > \overline{g}} q f(q)dq) = 0$, which when totally differentiated to raise $\underline{g}$ and have $\underline{h}$ adjust to maintain $\overline{v}$ gives

$$\overline{v} f(\underline{h})d\underline{h} - \overline{v} f(\underline{g})d\underline{g} - \underline{h} f(\underline{h})d\underline{h} + \underline{g} f(\underline{g})d\underline{g} = 0, \tag{12}$$

so

$$\frac{d\underline{h}}{d\underline{g}} = \frac{f(\underline{g})}{f(\underline{h})} \left(\frac{\underline{g} - \overline{v}}{\underline{h} - \overline{v}} \right). \tag{13}$$

Now we can differentiate the expected loss from the lower and upper intervals of the pool (the exactly revealed types in $[\underline{h}, \underline{g}]$ create zero loss) and see how it changes by substituting for $\frac{d\underline{h}}{d\underline{g}}$:

$$\begin{aligned}
&\tfrac{d}{d\underline{g}} \left(\int_{\underline{p}}^{\underline{h}(\underline{g})} (q - \overline{v})^2 f(q)dq + \int_{\underline{g}}^{\overline{g}} (q - \overline{v})^2 f(q)dq \right) \\
&= \tfrac{d\underline{h}}{d\underline{g}} \left(\underline{h} - \overline{v}\right)^2 f(\underline{h}) - (\underline{g} - \overline{v})^2 f(\underline{g}) \\
&= f(\underline{g}) \left(\tfrac{\underline{g} - \overline{v}}{\underline{h} - \overline{v}} \left(\underline{h} - \overline{v}\right)^2 - (\underline{g} - \overline{v})^2 \right) \\
&= f(\underline{g})(\underline{g} - \overline{v})(\underline{h} - \underline{g}) < 0,
\end{aligned} \tag{14}$$

where the final inequality follows from $\underline{g} > \overline{v}$ and $\underline{h} > \underline{g}$. Therefore we can increase both $\underline{g}$ and $\underline{h}$, reducing loss while preserving feasibility, and the original split pool scheme cannot have been optimal.

(iii) Suppose some types greater than $\underline{p}$ are uncertified and $(\underline{u}, \overline{u})$ is the lowest open uncertified interval above $\underline{p}$. Define μ_u as the mean of all the uncertified intervals and μ_p as the mean of all the pass types, which by feasibility implies $\mu_u < \mu_p$.

Consider the contribution to expected loss of just types in $[0, \underline{p}]$, $(\underline{p}, \underline{u})$, and $(\underline{u}, \overline{u})$:

$$\int_0^{\underline{p}} (q-\mu_u)^2 f(q)dq + \int_{\underline{p}}^{\underline{u}} (q-\mu_p)^2 f(q)dq + \int_{\underline{u}}^{\overline{u}} (q-\mu_u)^2 f(q)dq. \tag{15}$$

Let us see what happens if we increase $\underline{p}$ and $\underline{u}$ to get the same increase in probability mass, i.e., $d\underline{p}\, f(\underline{p}) = d\underline{u}\, f(\underline{u})$. This maintains feasibility since the mean of the pass pool rises, and hence since the mass of each types is kept constant, by the law of iterated expectations the mean of the uncertified pool falls. Differentiating equation (15) and substituting $d\underline{p}\, f(\underline{p}) = d\underline{u}\, f(\underline{u})$ yields

$$\begin{aligned} (16) \qquad & (\underline{p}-\mu_u)^2 f(\underline{p})d\underline{p} - (\underline{p}-\mu_p)^2 f(\underline{p})d\underline{p} + (\underline{u}-\mu_p)^2 f(\underline{u})d\underline{u} - (\underline{u}-\mu_u)^2 f(\underline{u})d\underline{u} \\ = \; & (\underline{p}-\mu_u)^2 f(\underline{p})d\underline{p} - (\underline{p}-\mu_p)^2 f(\underline{p})d\underline{p} + (\underline{u}-\mu_p)^2 f(\underline{p})d\underline{p} - (\underline{u}-\mu_u)^2 f(\underline{p})d\underline{p} \\ \propto \; & 2\left(\underline{p}-\underline{u}\right)\left(\mu_p-\mu_u\right) < 0 \end{aligned}$$

where the inequality follows since $\underline{p} < \underline{u}$ by definition and $\mu_u < \mu_p$ by feasibility. Therefore loss goes down, implying the original scheme cannot be optimal.

Thus, schemes (i), (ii), and (iii) are ruled out and the proposition is proved. ■

PROPOSITION 3 (Pass-fail vs. honors): *(i) As the certification cost c tends to 0, honors grading is optimal, with the mass of exactly revealed types tending to 1 and the mass of pass types tending to 0. (ii) For f quasiconcave, as the certification cost c tends to the maximum feasible level, the mass of exactly revealed types tends to 0 and the mass of pass types tends to 0 if $E[q] < 1/2$ and to 1 if $E[q] > 1/2$.*

Proof: (i) We know from Proposition 1's proof that the expected loss from honors grading is less than from exact grading for any c such that either scheme is feasible, as is the case for c sufficiently small. For sufficiently small c it must be that expected loss from exact grading is less than from pass-fail. As c approaches zero, the expected loss under exact grading approaches 0 if

h goes to 0, while expected loss under pass-fail is bounded from below by either the pass pool's or the fail pool's contribution to expected loss,

$$\min_p \max \left\{ \int_0^p (q - E[q|q \le p])^2 f(q)dq, \int_p^1 (q - E[q|q > p])^2 f(q)dq \right\} > 0. \tag{17}$$

Therefore, as c tends to 0 the optimal scheme is honors grading with h tending to 0. So $p(c)$ must also go to 0 since $p < h$ by definition.

(ii) Note that for any p, $E[q|q \in (p,h)]$ is increasing in h by Lemma 1(i), so as c increases to its maximum feasible level, h goes to 1. For quasiconcave f, Lemma 1(iv) tells us that, for a given h, the gap

$$E[q|q \in (p,h)] - E[q|q \le p], \tag{18}$$

is maximized at either $p = h$ or $p = 0$ depending on $E[q]$.

First, if $E[q] < 1/2$ the gap (18) is maximized at $p = h$. As established above, if c is large enough then $h = 1$ or is arbitrarily close to 1, so the mass of both exact types and pass types is arbitrarily close to 0.

Second, if $E[q] > 1/2$ the gap (18) is maximized at $p = 0$. As c tends to $E[q]$, to maintain feasibility p must tend to 0 and h must be equal to or tend 1, so the mass of pass types tends to 1. ■

PROPOSITION 4 (Feasibility): *(i) For any quality density f, if exact grading is feasible then so is honors grading, and if honors is feasible then so is pass-fail. (ii) Pass-fail and honors grading are both feasible for a range of grading costs so high that exact grading is not if:*

(a) $f(q)$ is strictly increasing; or

(b) $f(q)$ is strictly quasiconcave and $f(1) > (1/2)/(1 - E[q])$; or

(c) $f(q)$ is strictly logconcave and $E[q] > 1/2$.

Proof: Let $\overline{c}_x$, $\overline{c}_p$, and $\overline{c}_h$ represent the highest feasible certification costs for any exact, pass-fail, and honors grading scheme respectively: $\overline{c}_x \equiv \sup_{x \in [0,1)}\{c \le x - E[q|q < x]\}$, $\overline{c}_p \equiv$

$\sup_{p\in[0,1)}\{E[q|q > p] - E[q|q \le p]\}$, and $\overline{c}_h \equiv \sup_{p\in[0,h),h\in(0,1)}\{E[q|q \in (p,h)] - E[q|q \le p]\}$.

(i) We want to show $\overline{c}_p \ge \overline{c}_h \ge \overline{c}_x$. Consider any x such that exact grading is feasible and set $p = x$ and $h \in (x, 1)$ for an honors scheme. Since $E[q|q \in (p,h)] > p = x$, it must be that $E[q|q \in (p,h)] - E[q|q \le p] > x - E[q|q < x]$, so honors grading is also feasible. Similarly, for any $h < 1$, $E[q|q > p] > E[q|q \in (p,h)]$ so if honors grading is feasible so is pass-fail grading with the same p but $h = 1$.

(ii) From (i) it is sufficient to show that $\overline{c}_h > \overline{c}_x$.

(a) Start with honors. By Lemma 1(iii) the assumption $f' > 0$ implies $E[q|q \in (p,h)] - E[q|q \le p]$ is strictly decreasing in p, and by Lemma 1(i) $E[q|q \in (p,h)]$ is increasing in h. Thus to maintain feasibility as c increases we need p to approach 0 and h to approach 1. Therefore, $\overline{c}_h = E[q|q > 0] - 0 = E[q]$. Note $E[q] > 1/2$ since $f' > 0$ so $\overline{c}_h > 1/2$. As for exact grading, the assumption $f' > 0$ implies $E[q|q < x] > x/2$, so $\overline{c}_x = \sup_{x\in[0,1)}\{x - E[q|q < x]\} < \sup_{x\in[0,1)}\{x - x/2\} = 1/2$. Hence, $\overline{c}_h > 1/2 > \overline{c}_x$.

(b) Again starting with honors, by Lemma 1(iv), $f(1) > (1/2)/(1 - E[q])$ implies that $E[q|q \in (p,h)] - E[q|q \le p]$ reaches a maximum at $p = 0$ for any h. And by Lemma 1(i), $E[q|q \in (p,h)]$ is increasing in h, so again $\overline{c}_h = E[q|q > 0] - 0 = E[q]$. As for exact grading, note that $E[q|q \ge x] \ge x$, so if an exact scheme is feasible so is a pass-fail scheme with $p = x$. Hence $\overline{c}_x \le \overline{c}_p$, with strict inequality if $p = x < 1$. From Lemma 1(iv), $E[q|q \ge p] - E[q|q \le p]$ is decreasing in p. Hence $\overline{c}_p$ is maximized at $p = 0$, in which case we have the strict inequality $\overline{c}_x < \overline{c}_p$. Finally, note that for $p = 0$, $\overline{c}_p = E[q] - 0 = E[q]$. Thus, $\overline{c}_p = \overline{c}_h$. Hence, since $\overline{c}_x < \overline{c}_p$, we have $\overline{c}_h \ge E[q] > \overline{c}_x$.

(c) Honors grading with $p = 0$ and h arbitrarily close to 1 is feasible for all $c < E[q]$, so $\overline{c}_h \ge 1/2$. As for exact grading, for f logconcave, Lemma 1(ii) implies $x - E[q|q < x]$ is maximized at $x = 1$, so $\overline{c}_x = 1 - E[q]$, which is less than $1/2$ since by assumption $E[q] > 1/2$. Therefore $\overline{c}_h > \overline{c}_x$. ■

REFERENCES

BAGNOLI, MARK & TED BERGSTROM (2005): "Log-Concave Probability and Its Applications," *Economic Theory*, 26, 445–469.

BENABOU, ROLAND & JEAN TIROLE (2006): "Incentives and Prosocial Behavior," *American Economic Review*, 96, 1652–1678.

BENABOU, ROLAND & JEAN TIROLE (2011): "Laws and Norms," NBER Working Paper No. 17579.

BOLESLAVSKY, RAPHAEL & CHRISTOPHER COTTON (2015): "Grading Standards and Educational Quality," *American Economic Journal: Microeconomics*, 7, 248–279.

CHAKRABORTY, ARCHISHMAN & RICK HARBAUGH (2007): "Comparative Cheap Talk," *Journal of Economic Theory*, 132, 70–94.

CHAN, WILLIAM, LI HAO & WING SUEN (2007): "A Signaling Theory of Grade Inflation," *International Economic Review*, 48, 1065–1090.

CHAHROUR, RYAN (2014): "Public Communication and Information Acquisition," *American Economic Journal: Macroeconomics*, 6, 73–101.

COSTRELL, ROBERT M. (1994): "A Simple Model of Educational Standards," *American Economic Review*, 84, 956–971.

CRAWFORD, VINCENT P. & JOEL SOBEL (1982): "Strategic Information Transmission," *Econometrica*, 50, 1431-1451.

DUBEY, PRADEEP & JOHN GEANAKOPLOS (2010): "Grading Exams: 100,99,98... or *A*,*B*,*C*?" *Games and Economic Behavior*, 69, 72–94.

EPPLER, MARTIN J. & JEANNE MENGIS (2004): "The Concept of Information Overload: A Review of Literature from Organization Science, Accounting, Marketing, MIS, and Related Disciplines." *Information Society*, 20, 325–344.

FARHI, EMMANUEL, JOSH LERNER, & JEAN TIROLE (2013): "Fear of Rejection? Tiered Certification and Transparency," *RAND Journal of Economics*, 44, 610-631.

FARRELL, JOSEPH, JANIS K. PAPPALARDO, & HOWARD SHELANSKI (2010): "Economics at the FTC: Mergers, Dominant-Firm Conduct, and Consumer Behavior," *Review of Industrial Organization*, 37, 263–277.

GENTZKOW, MATTHEW & EMIR KAMENICA (2014): "Costly Persuasion," *American Economic Review*, 104, 457–462.

JEWITT, IAN (2004): "Notes on the Shape of Distributions," working paper, Nuffield College, Oxford.

KAMENICA, EMIR & MATTHEW GENTZKOW (2011): "Bayesian Persuasion," *American Economic Review*, 101, 2590–2615.

LERNER, JOSH & JEAN TIROLE (2006): "Forum Shopping," *American Economic Review*, 96, 1091–1113.

LIZZERI, ALESSANDRO (1999): "Information Revelation and Certification Intermediaries," *RAND Journal of Economics*, 30, 214–231.

MOLDOVANU, BENNY, ANER SELA & XIANWEN SHI (2007): "Contests for Status," *Journal of Political Economy*, 115, 338–363.

MORGAN, JOHN & PHILLIP C. STOCKEN (2003): "An Analysis of Stock Recommendations," *RAND Journal of Economics*, 34, 183–203.

OSTROVSKY, MICHAEL & MICHAEL SCHWARZ (2010): "Information Disclosure and Unraveling in Matching Markets," *American Economic Journal: Microeconomics*, 2, 34–63.

PREKOPA, A. (1973): "On Logarithmic Concave Measures and Functions," *Acta Scientarium Mathematicarum*, 34, 335–343.

RAYO, LUIS & ILYA SEGAL (2010): "Optimal Information Disclosure," *Journal of Political Economy*, 118, 949–987.

ROSAR, FRANK & ELISABETH SCHULTE (2012): "Optimal Test Design Under Imperfect Private Information and Voluntary Participation," working paper, University of Bonn.

TITMAN, SHERIDAN & BRETT TRUEMAN (1986): "Information Quality and the Valuation of New Issues," *Journal of Accounting and Economics*, 8, 159–172.

VITALIS, VANGELIS (2002) "Private Voluntary Eco-labels: Trade Distorting, Discriminatory and Environmentally Disappointing," Background paper for the round table on sustainable development, OECD, Paris, 6th December, 2002.