Coarse Grades: Informing the Public by Withholding Information

Rick Harbaugh and Eric Rasmusen

December 18, 2014

Abstract

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, is the most informative. JEL: D82, L15.

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, Dmitry Lubensky, Dan Klerman, and participants in talks at the BEPP Brown Bag Lunch, Baylor University, and the IIOC and Summer Econometric Society meetings for helpful comments.

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 information 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? We suggest that the answer often lies in certification being voluntary. In situations such as certification for eco-labels, costly cooperation from the firm is required, and the certifier needs to get

[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.

the firm to participate. Just as a student would be reluctant to attend a medical school that would publicly rank him as the worst student in his class to earn an M.D.,[3] a firm would not be eager to be stamped with a seal of approval that tells the world it barely passed. 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 information by inducing more participation. If the certification grade is coarse, a mediocre type is pooled with better types, so its expected quality conditional on the certification grade is higher than its true quality. Therefore it has more incentive to bear the costs of certification, and participation rises. We show that at the margin the extra information from increased participation outweighs the loss from throwing away information on those that do participate. Hence, in the scheme that maximizes information, the certifier always throws away information. Moreover, under plausible conditions the optimal scheme 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 differentiated exactly. Schools do not publicly provide class rank information about most of their graduates, but 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.

[3]As an old joke notes, "What do you call someone who graduated from the bottom of their class in medical school? Doctor."

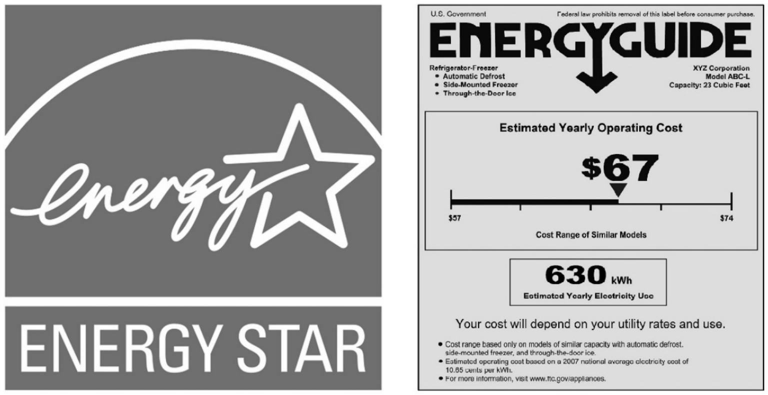

Figure 1:
Pass-fail grades for a voluntary label, exact grades for a mandatory label

Since coarseness is used to encourage participation, the model predicts that coarseness is less likely when quality evaluation does not require the cooperation of the sender. Camera companies cannot prevent consumer reviewers at *Amazon.com* or professional reviewers at *CNET* from rating their products. Thus, 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.[4]

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 of such

[4]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.

mandatory labels provide fine or exact grades.[5] 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 significant gap in the literature. Lizzeri (1999) and others show how a certifier trying to maximize his own profits from charging for certification can do so by reporting product information coarsely. Ostrovsky & Schwarz (2010), Gentzkow & Kamenica (2011) and others show how a certifier seeking to maximize perceived sender quality (and hence the payoffs of senders) can often do so by reporting information coarsely. Considering the three different parties involved in certification — certifiers, senders, and receivers — 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 information benefit to receivers. Our assumption would seem particularly unlikely to lead to coarse grading, but we find the same general result that the certifier injects some coarseness into his grading. In the last section of this article we will say more about these and other papers.

1 The Model

A sender (e.g., a firm selling a product) has exogenous quality q that is randomly distributed with full support on $[0, 1]$ according to distribution function F with an analytic density function f.[6] The realization of q is the sender's private information. A certifier (e.g., an NGO) chooses a grading scheme $m(q)$, a function that

[5] We do not analyze these differences formally. As indicated by a referee, there is an important endogeneity issue in that governments might be more likely to adopt mandatory schemes when the quality variable of interest is more easily measured.

[6] The unidimensional quality measure may be a weighted score based on multiple attributes. In some cases, such as testing for the presence of hazardous chemicals, failure on any one attribute is sufficient to award a failing score. Our model is not appropriate for such cases.

maps quality q to message m.[7] 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". An application costs the sender a fixed amount $c > 0$ in time and trouble. A receiver (e.g., a consumer) 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 the maximum of $E[q|m] - c$ from applying for certification and $E[q|\text{"}\textit{uncertified}\text{"}]$ from not applying. The certifier chooses a grading scheme to best inform the receiver in the sense of minimizing the expectation of a differentiable loss function $L(q, E[q|m])$.[8] We assume that L is convex in q with $L(q, q) = 0$ and that expected loss over $q : m(q) = m'$ is minimized at $E[q|m']$.

The certifier's loss function could capture a preference for providing accurate information, a reputational incentive to do so, or a concern for consumer welfare. For instance, in a unit demand framework, when actual quality is lower than expected quality the consumer mistakenly purchases the good, and when it is higher the consumer mistakenly does not purchase it, so more accurate information reduces these losses.[9] In Section II's examples we will use the mean-squared error loss function $L = (q - E[q|m])^2$.

[7] To reduce verbiage, we will ignore schemes that differ only on sets of zero mass, e.g., a scheme that reveals the quality of any sender who applies except that it pools types $q = .8$ and $q = .83$, or a pass-fail scheme in which types $q > .6$ pass as opposed toa scheme in which types $q \geq .6$ pass. These schemes cannot be better (or worse) than ones we discuss, since the effect a zero mass of senders.

[8] It does not matter whether there is just one certifier or many. 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 quality to draw them away from (and destabilize) a pass-fail certifier.

[9] A quadratic loss function is directly assumed in most of the strategic information transmission literature following Crawford & Sobel (1982), and behaviorally equivalent assumptions are widely used as the default case in signaling (e.g, Spence, 1973), persuasion (e.g., Milgrom, 1981), and certification games (e.g., Lizzeri, 1999).

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 degenerate pessimistic non-certification equilibria in which the receiver never expects the sender to certify its quality, and the sender never certifies because if he does, the receiver punishes unexpected certification with unfavorable off-equilibrium-path beliefs.

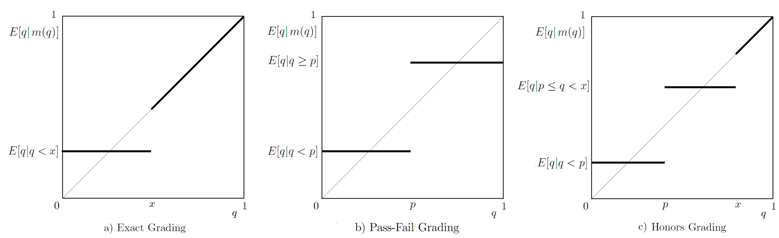

FIGURE 2:

TRUE QUALITY AND CONSUMER ESTIMATED QUALITY UNDER DIFFERENT GRADING SCHEMES

We allow for any grading scheme in which the certifier is "fair", assigning the same grade to firms of identical quality (which excludes mixing on his part). Three schemes will be of particular interest: exact, pass-fail, and honors.[10] 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

[10]In all of the schemes we assume the same certification cost c for all firm types and that the sender pays it. A referee suggested that the certifier might reimburse senders depending on their quality while maintaining a balanced budget, but we have excluded that possibility.

quality of the uncertified senders in the quality interval $[0, x)$:

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

If a scheme is not exact then it is *coarse*: the exact quality of at least some product 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 meaningful grading scheme. The message is $m =$"uncertified" if $q \leq p$ and $m =$"pass" if $q > p$.[11] Assuming that p is set so all types $q > p$ have sufficient incentive to be certified, the expected loss from pass-fail grading is

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

The other optimal coarse grading scheme is Figure 2(c)'s *honors grading*, which sets a standard x above which quality is revealed exactly but also divides types below h into two groups by a passing standard 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

$$EL_{honors} = \int_0^p L\left(q - E[q|q \leq p]\right) f(q)dq + \int_p^h L\left(q - E[q|p < q < x]\right) f(q)dq + \int_h^1 (0)\, f(q)dq. \quad (3)$$

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

[11] Alternatively, the certifier could send the message "fail", 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.

do not apply. This assumption is unimportant to the results, since in equilibrium a low-quality sender knows that in advance that he would fail, and so does not incur the cost c to be certified.

2 Why Coarseness Helps

In this section we will use examples based on specific quality densities f, the linear profit function π, and the quadratic loss (mean squared error) function $L\left(q - E[q|m]\right) = \left(q - E[q|m]\right)^2$ 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" examples will build the intuition behind Section III's propositions for broad classes of distributions and more general loss functions.

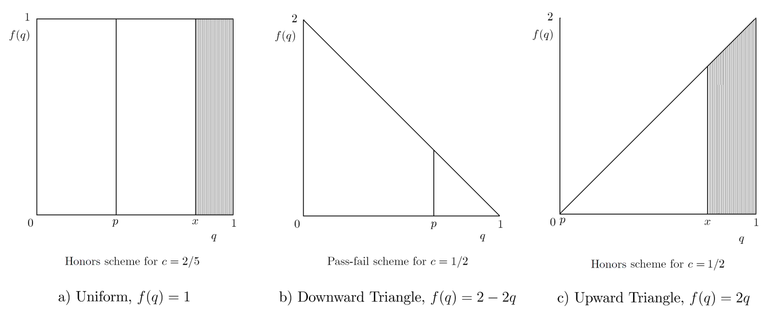

FIGURE 3:
OPTIMAL GRADING SCHEMES FOR DIFFERENT DISTRIBUTIONS OF PRODUCT QUALITY

2.1 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 under a symmetric convex loss function such as the quadratic 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 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.

This result that coarse grading — either pass/fail or honors — is better than exact grading will be shown to hold generally in Propositions 1 and 2.

2.2 Pass-Fail Grading Can Be Most Informative

Honors grading is more complex than pass-fail, but not necessarily 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. Consider pass-fail grading, so $h = 1$, and first suppose that the certification cost c is low. When c is low, p can be set to divide the region $[0, 1]$ to minimize expected loss without the participation constraint being binding. In particular, 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 $c < \frac{1}{3}\sqrt{5} - \frac{1}{3} \approx 0.412$. Within the range $c \in [0, .412]$ there is slack in the pass-fail participation constraint for those types that apply; the types near $q = 1$ can be exactly revealed and $p \approx .382$ is still optimal and feasible. Thus, as with the uniform density, the certifier can do even better by using honors grading. This is seen in Figure 4(b), where for low c, honors grading reduces expected loss relative to pass-fail grading.

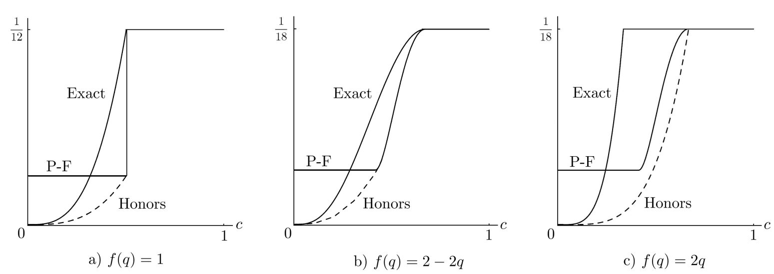

FIGURE 4:
EXPECTED LOSS (MSE) VS. APPLICATION COST FOR DIFFERENT SCHEMES AND DISTRIBUTIONS

As c rises the participation constraint becomes binding. The gain from passing is increasing in the cutoff p, so p will have to be set higher to ensure participation. If h falls from $h = 1$ to $h < 1$, to have honors grading the gain from passing must fall, 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 $c \gtrapprox .412$, the optimal honors grading scheme has h arbitrarily close to 1, which is equivalent to pass-fail grading.

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

2.3 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 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).

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 $c \leq 1/3$. 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 $c < 2/3$.

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

3 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 and the quadratic loss function. We now generalize our analysis to include any grading scheme for any quality density f and any loss function L meeting the conditions in the description of the model.

Properties of means of a distribution above or beneath a cutoff point 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 ($A(q)$), the lower mean below the cutoff ($B(q)$), and the gap between the upper and lower means. The first two Properties below are standard (see Bagnoli and Bergstrom, 2005), while the last two strengthen the monotonicity and quasiconvexity results in Jewitt (2004). The properties are proven in the Appendix. Figure 5 illustrates how the means change with q in one particular example.

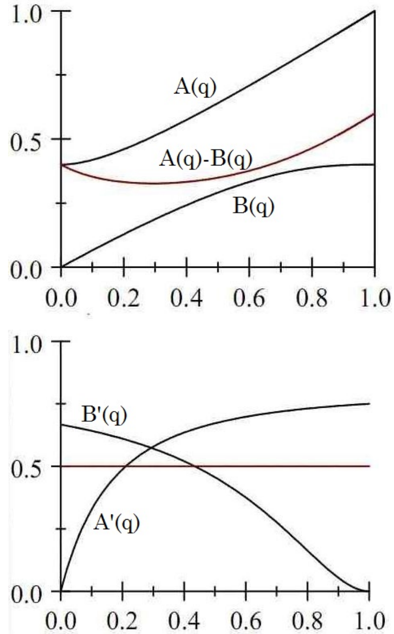

FIGURE 5:
HOW THE LOWER AND UPPER MEANS AND THEIR DERIVATIVES CHANGE AS THE CUTOFF RISES

Properties of Upper and Lower Means *Suppose density $f(q)$ is analytic*[12] *with support on $[\underline{q}, \overline{q}]$ and define $A(t) \equiv E\left[q|q \geq t\right]$, $B(t) \equiv E\left[q|q \leq t\right]$. Then*

(i) For any f, $A' > 0$ and $B' > 0$.

(ii) For strongly[13] *logconcave f, $A' < 1$ and $B' < 1$.*

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

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

[12] A function is analytic iff it is locally given by a convergent power series. All analytic functions are infinitely differentiable.

[13] Following Ginsberg (1973), we say a differentiable function $f(x)$ is *strongly* increasing if $f' > 0$ for all x. It is *strictly* increasing if $a > b$ implies $f(a) > f(b)$.

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 Property (i).[14] 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.[15] From Property (ii), if the density does not increase in slope too rapidly, 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. Hence the gap $t - B$ between the marginal type who is graded exactly and the average quality of the uncertified pool is increasing in t. Since this is the maximum 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.[16] 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.

Property (iii) implies that for f increasing $A' < B'$ and for f decreasing $A' > B'$, which is Jewitt's monotonicity result that for monotonic f the gap $A - B$ is monotonic in the opposite direction. 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, and the reverse 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). Property (iv) implies Jewitt's quasi-convexity result that for strictly quasiconcave densities (i.e., unimodal densities,

[14] As the old joke says: "Professor Smith moved from Upstate U. to Downstate U., thus raising the quality of both."

[15] The gap between upper and lower means is also central to binary signaling games (e.g., Benabou & Tirole 2006, 2011).

[16] $B' < 1$ holds more generally for F logconcave. The result $A' < 1$ for f logconcave is the standard decreasing mean residual life result from reliability theory.

which includes all logconcave densities and all strictly monotonic densities), the gap $A - B$ is either monotonic or is first decreasing and then increasing in t. It also shows that monotonicity of $A - B$ holds for unimodal f if either $f(\underline{q})$ or $f(\overline{q})$ is sufficiently large, which is a weaker condition than Jewitt's condition that f is monotonic, e.g., it allows the density to dip slightly at either end.[17] Note generally 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 quasiconvex 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 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 Property (iii), $E[q|q \in (p, x^*)] - E[q|q \in [0, p]]$ is decreasing in p so again 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, if the decrease to $p' < x^*$ reduces the pass region's mean more than the fail region's mean then the participation constraint is no longer met for $x = x^*$. Hence, x 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, Property (iii) tells us that it arises for all decreasing f, and Property (iv) tells us it can arise generally with unimodal distributions.

[17] To see that the conditions in (iv) are weaker conditions for monotonicity of $A - B$, consider any strictly increasing f with support on $[\underline{q}, \overline{q}]$ and mean $E[q]$. Take all mass below some point $y \in (\underline{q}, \overline{q})$ and add it to points $q \in [y, \overline{q})$ to create a new uniform distribution g with density $g(q) = f(\overline{q})$ for all $q \in [y, \overline{q}]$, implying $y = \overline{q} - 1/g(\overline{q})$. The mean of g is then $E_g[q] = (\overline{q} - 1/g(\overline{q}) + \overline{q})/2$ or $g(\overline{q})(\overline{q} - E_g[q]) \geq 1/2$. Since $f(\overline{q}) = g(\overline{q})$ and $E[q] < E_g[q]$, this implies $f(\overline{q})(\overline{q} - E[q]) > 1/2$. Similarly it can be shown that f decreasing implies $f(\underline{q})(E[q] - \underline{q}) > 1/2$.

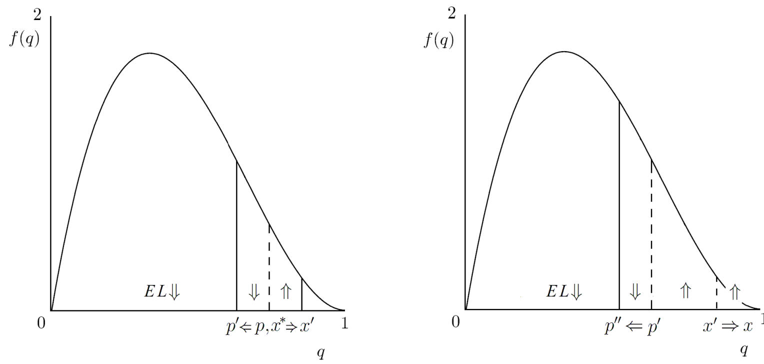

a) From Exact to Honors Grading b) From Honors to Pass-Fail Grading

FIGURE 6:
THE EFFECT OF COARSENESS ON EXPECTED LOSS

Figure 6(a) shows a Beta (2,3) distribution where the pass standard is in the region of decreasing f so that one might expect the same tradeoff. Since $f(0) = f(1) = 0$ the gap is increasing in p after some internal minimum from Property (iv), and the same tradeoff does indeed arise.[18] For costs $c = 2/5$, introducing a pass region with $p' < x^*$ requires increasing the exact cutoff from x^* to some $x' > x^*$ 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. Types in $[x^*, x')$ contribute more to expected loss since they are no longer exactly revealed.

The following Proposition will say that despite the possibility of 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. The idea is that for marginal decreases 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', x') because the pass interval is so

[18] 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), and is sufficiently "humpy" to satisfy the condition for an internal minimum from property (iv) iff $\alpha, \beta > 1$.

small. The result holds for Figure 6(a)'s example, and more generally. Now there will be 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.

Proposition 2 (Simple Schemes) *Either pass-fail grading is as good as any other scheme, or honors grading is.*

PROOF: See the Appendix.

Proposition 2 says that either pass-fail or honors grading is superior to more complicated schemes. Proposition 3 will deal with which of those two is best under different circumstances.[19] 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 and then exactly grade better types above a low honors cutoff. The second part states that for logconcave f if c is sufficiently high then either pass-fail grading is best or the optimal honors grading scheme is close to pass-fail in that an arbitrarily small share of certified types are exactly graded. With high certification costs it is only feasible to exactly grade very high types so the best scheme has a relatively large pass region and a small honors region, or instead just has no honors region at all.

Proposition 3 (Pass-fail vs. Honors) *(i) For certification cost c low enough, no grading scheme is superior to honors grading, with the pass and honors cutoffs*

[19] Recall that we require x is strictly less than one for "honors grading" so pass-fail grading is not merely a subset of honors.

p and h going to zero as c goes to zero.
(ii) For c high enough, if the quality density f is strictly logconcave the most informative grading scheme is either pass-fail, or is "almost pass fail"— honors with p going to 0 and h going to 1 as c rises

PROOF: See the Appendix.

To better understand why pass-fail grading is better for high certification costs, look at Figure 6(b). This is the same case as Figure 6(a) except that the pass region has been expanded so much that the honors region is squeezed to nothing. If we start with Figure 6(a)'s honors scheme (p', h) and increase the pass region by dropping p' to p'', we must also increase h to maintain feasibility. Since costs are high and we started towards the upper end of the density, as we increase the pass region enough, the honors region shrinks to zero 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 general result of Proposition 3(ii) uses Property (iv) that for unimodal 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 latter 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 as discussed above. And when $E[q] > 1/2$ the former case holds so for large c it is necessary for feasibility to set 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 at all would show up to be graded! This was the case in Figure 3(c) (the rising triangle density) for exact grading, even for many values of c for which pass-fail and honors remained feasible. Proposition 4 first shows that honors and pass-fail are always feasible whenever exact grading is, and then gives general conditions for when they are feasible when exact grading is not.[20]

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) $f(q)$ strictly quasiconcave and $f(1)$ is sufficiently large, 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.

4 Extension: Letter Grading with Different Consumer Priors on Different Firms

So far 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 is likely to be better than

[20]Condition (ii-a) is implied by (ii-b), but the two cases are included separately since the main intuition is from (ii-a).

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.

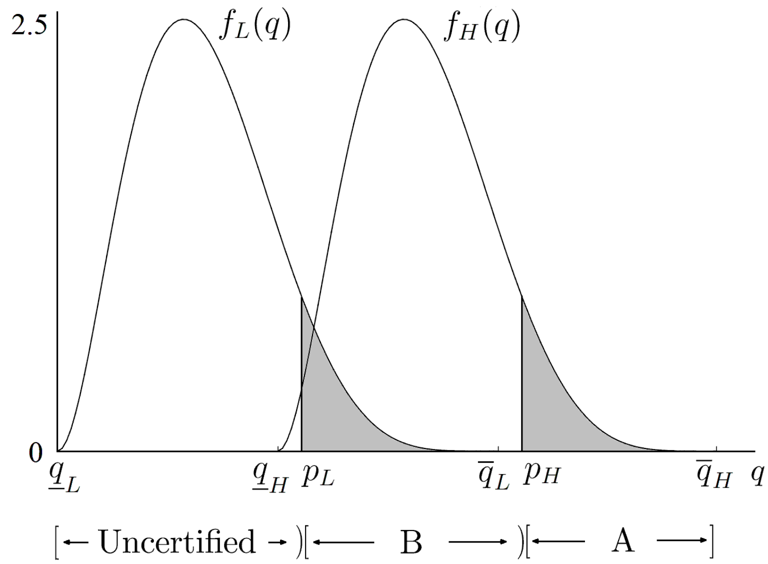

FIGURE 7:
LETTER GRADING WITH DIFFERENT PRIORS

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,[21] 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

[21]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 $c = 1/3$.

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.[22] 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.*

5 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 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 interme-

[22]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.

diary 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 be distinguished from worse senders but do not want to be distinguished from even better senders.

A second strand of the literature concerns certifiers whose goal is to maximize the benefit to the subjects being certified. That literature introduces strategic reasons to obscure the information the public receives. Indeed, the ideal would be to perfectly fool the public. Ostrovsky & Schwarz (2010) consider colleges which choose grading schemes to maximize the success of their students on the job market. Depending on the distributions of jobs and student abilities, students overall may gain if weaker students are pooled together with better students. The average quality of students at Yale might be so high that they will all get jobs if their grades are so uniformly high as to be uninformative, but weaker students might not get jobs if they are revealed as weak. In their model, students are exogenously assigned to colleges, so colleges are free to ignore the selection effect we study in the present paper. Chakraborty & Harbaugh (2007) similarly consider the gains from coarse rankings via cheap talk without precommitment to a disclosure policy. Rayo & Segal (2010) and Kamenica & Gentzkow (2011) consider the general problem of the sender committing to a disclosure policy, rather than a certification intermediary doing so as in Lizzeri (1999), and find that a coarse disclosure policy often benefits the sender. Our model differs from these because the certifier's objective is to inform the public rather than mislead it and because senders can choose not to be certified. Coarseness in these models reduces information to the public without any gains from encouraging participation.

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).[23] 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 when abilities are heterogeneous.

Our approach adds to this rich literature by showing how coarseness is optimal

[23] 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.

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),[24] and that the process of applying for certification can be lengthy and compete for limited managerial resources.[25] 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. For instance, 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 determined that consumers learned most from exact information

[24]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.

[25]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.

about expected energy costs (Farrell, Pappalardo & Shelanski, 2010).[26] 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.

6 Appendix

Following Ginsberg (1973), we say a differentiable function $f(x)$ is *strongly* increasing if $f' > 0$ for all x. It is *strictly* increasing if $a > b$ implies $f(a) > f(b)$. These are not the same, because a strictly increasing function might have $f'(c) = 0$ at some c; for example, $f(x) = x^2$ is strictly increasing on $[0, \infty)$ but $f'(0) = 0$. Similarly, a strongly concave function requires $f'' < 0$ and a strongly quasiconcave function requires either $f' > 0$, $f' < 0$, or f' first positive, then zero, then negative over its range. An analytic function $f(x)$ (a function which can be approximated by a Taylor series) is strictly increasing only if $f^n(x) > 0$ for some n'th derivative. Our results for strongly increasing, concave, and quasiconcave $f(q)$ can all be extended to "strictly" functions that are analytic, and perhaps even to transcendental functions.

Upper and Lower Means Properties *Suppose density $f(q)$ is analytic with support on* $[\underline{q}, \overline{q}]$ and define $A(t) \equiv E[q|q \geq t]$, $B(t) \equiv E[q|q \leq t]$. *Then...[The four properties are listed under the sections of the proof that prove them.]*

[26]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.

Proof:

Property (i) $A'(t) > 0$ and $B'(t) > 0$.

Integrating by parts,

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

$$= \frac{bF(b) - aF(a)}{F(b) - F(a)} - \frac{\int_a^b F(q) dq}{F(b) - F(a)}. \tag{5}$$

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

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

$$= \frac{f(t)}{1 - F(t)} (A - t) \tag{6}$$

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

$$= \frac{f(t)}{F(t)} (t - B). \tag{7}$$

Equations (6) and (7) 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,

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

$$= \frac{1}{2} + \lim_{t \to 0} \frac{f''(t) \int_{\underline{q}}^t F(q) dq + f'(t) F(t)}{2 f'(t) F(t) + 2 f(t)^2} \tag{8}$$

$$= \frac{1}{2} > 0. \tag{9}$$

If, instead, $f(\underline{q}) = 0$, then applying l'Hopital's rule n times until $f^{(n)}(\underline{q}) \neq 0$, yields[27]

$$B'(\underline{q}) = (n+2)/(n+1) \geq 1/2 > 0. \tag{10}$$

[27] The result For $f(q) > 0$, the truncated distribution converges to a uniform distribution, and

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

Property (ii) For strongly logconcave f, $A'(t) < 1$ and $B'(t) < 1$.

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 (6), $f(t)\int_{\underline{q}}^{t} F(q)dq < F(t)^2$ implies $B'(t) < 1$.

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 (6) implies $A'(t) < 1$.

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

Differentiating (6) and (7) and substituting,

$$\begin{aligned} A''(t) &= f'(t)\tfrac{A'}{f(t)} + (2A'-1)\,\tfrac{f(t)}{1-F(t)} \\ B''(t) &= f'(t)\tfrac{B'}{f(t)} + (1-2B')\,\tfrac{f(t)}{F(t)}. \end{aligned} \tag{11}$$

First consider f decreasing so $f(\underline{q}) > 0$. From (9), $B'(\underline{q}) = 1/2$, and from (11), $f' < 0$ 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 (11), $f' < 0$ 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.

Property (iv): For strongly quasiconcave $f(q)$, the gap $A(t) - B(t)$ is strictly quasiconvex and in particular is strictly increasing iff $f(\underline{q})(E[q] - \underline{q}) > 1/2$, strictly

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

decreasing iff $f(\overline{q})(\overline{q} - E[q]) > 1/2$*, and strictly decreasing then increasing otherwise.*

We first establish quasiconvexity. Since $\Delta = A - B$ is twice differentiable, Δ is quasiconvex if $A' = B'$ implies that $A'' \geq B''$. First consider $A' = B' > 1/2$. From (11) this implies $A'' > B''$, as required. Now consider $A' = B' \leq 1/2$. Strict monotonicity (and hence quasiconvexity) follows from Property (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 $[\widehat{q}, \underline{q}]$. So from Property (iii) $B' \geq 1/2$ in $[\underline{q}, \widehat{q}]$, with possible equality only at $t = \underline{q}$, and $A' \geq 1/2$ in $[\widehat{q}, \underline{q}]$ 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 (11) implies $A'' = B''$, as required.

Given that $\Delta = A - B$ 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 $\Delta = A - B$ 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 (6) 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 a 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 that picks a feasible p, where $p(x) = x$ in the limiting case where the honors scheme equals the exact scheme. Such a function must exist in the 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 Property (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

$$\begin{aligned}
&\tfrac{dp}{dh} L\left(p(h), E[q|q \le p(h)]\right) f\left(p(h)\right) \\
&- \left(\tfrac{d}{dp} E[q|q \le p(h)] \tfrac{dp}{dh}\right) \textstyle\int_0^{p(h)} L_2(q, E[q|q \le p(h)]) f(q) dq \\
&+ L\left(h, E[q|q \in (p(h), h]]\right) f(h) \\
&- \tfrac{dp}{dh} L\left(p(h), E[q|q \in (p(h), h]]\right) f(p(h)) \\
&- \left(\tfrac{d}{dh} E[q|q \in (p(h), h]]\right) \textstyle\int_{p(h)}^{h} L_2(q, E[q|q \in (p(h), h]]) f(q) dq.
\end{aligned} \tag{12}$$

The first term of (12) is negative because $\frac{dp}{dh} < 0$ by construction, while the second term is zero by the envelope theorem since the mean $E[q|q \le p]$ minimizes $\int_0^p L\left(q, E[q|q \le p]\right) f(q) dq$ by assumption. 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 grading is as good as any other scheme, or honors grading is.*

Proof: Proposition 1 rules out exact schemes from being optimal. For any other feasible scheme denote by $\underline{p}$ and $\underline{h}$ the lowest types to be certified and to be exactly revealed. We can rule out schemes with $\underline{p} \ge \underline{h}$ because the pooled types above $\underline{p}$

could be exactly revealed without affecting feasibility. For a scheme with $\underline{p} < \underline{h}$ to be different from pass-fail or honors: (i) there is more than one pool of certified but inexactly reported types, or (ii) there is one certified pool but it is split by an interval of exactly revealed types (some of its types are above $\underline{h}$) and/or (iii) there is one pass pool, but some types above $\underline{p}$ are uncertified. We will rule out these three alternatives.

(i) Suppose there are multiple pools $i = 1, ...N$. If Pool i is feasible, it must be that a Pool i type's expected valuation by buyers, v_i, is at least some value $\overline{v}$ sufficiently greater than the mean of the uncertified. If $v_i > \overline{v}$, there is slack in the feasibility constraint and we can exactly reveal some of the highest types in Pool i, reducing information loss. So the optimal scheme will have $v_i = v_j = \overline{v}$.

Since all pools have the same mean, a consumer's best estimate of a firm's type from its grade will be $\overline{v}$ regardless of grade. Each type q in Pool i makes a contribution to total loss of $L(q - \overline{v})f(q)$. Since all grades convey the same message, one nonrandomized grade would generate the same loss. But that means one pool is as good as multiple pools.

(ii) Suppose we have one pool of certified types but it is split so some of it lies above $\underline{h}$. Let $\underline{g}$ and $\overline{g}$ be the lowest and highest types in the pool above $\underline{h}$. We will use total differentiation to show how $\underline{h}$ must rise to keep $v_1 = \overline{v}$ as we increase $\underline{g}$ and then show that this will reduce loss. The mean of the pool types above the original value $\underline{h}_0$ is $\mu = \frac{\int_{\underline{h}_0}^{\underline{h}} qf(q)dq + \int_{\underline{g}}^{\overline{g}} qf(q)dq}{\int_{\underline{h}_0}^{\underline{h}} f(q)dq + \int_{\underline{g}}^{\overline{g}} f(q)dq}$. Multiplying out yields $\mu(\int_{\underline{h}_0}^{\underline{h}} f(q)dq + \int_{\underline{g}}^{\overline{g}} f(q)dq) - \int_{\underline{h}_0}^{\underline{h}} qf(q)dq - \int_{\underline{g}}^{\overline{g}} qf(q)dq = 0$, which when totally differentiated yields

$$f(\underline{h})\mu d\underline{h} - f(\underline{g})\mu d\underline{g} - \underline{h}f(\underline{h})d\underline{h} + \underline{g}f(\underline{g})d\underline{g} = 0$$

so

$$\frac{d\underline{h}}{d\underline{g}} = \frac{f(\underline{g})}{f(h)} \left(\frac{\underline{g} - \mu_+}{\underline{h} - \mu_+} \right), \tag{13}$$

Now we can differentiate the expected loss and see how it changes by substituting for $\frac{d\underline{h}}{d\underline{g}}$:

$$\begin{aligned} &\tfrac{d}{d\underline{g}}\left(\int_{\underline{h}_0}^{\underline{h}(\underline{g})} L(q)f(q)dq + \int_{\underline{g}}^{\overline{g}} L(q)f(q)dq\right) \\ &= \tfrac{d\underline{h}}{d\underline{g}} L(h_0)f(h_0) - L(\underline{g})f(\underline{g}) \\ &= f(\underline{g})\left(\tfrac{\underline{g}-\mu}{\underline{h}-\mu} L(h_0) - L(\underline{g})\right) \\ &\propto \tfrac{L(\underline{h})}{\underline{h}-\mu} - \tfrac{L(\underline{g})}{\underline{g}-\mu} < 0, \end{aligned} \tag{14}$$

where the final inequality holds by the convexity of L in q for $\mu < \underline{h}$ and by $L(\underline{h})/(\underline{h} - \mu)$ equalling zero for $\mu = \underline{h}$ (as l'Hopital's rule can show). Therefore we can reduce $\underline{g}$ and increase $\underline{h}$, reducing loss while preserving feasibility, and the split pool cannot have been optimal.

(iii) Suppose some types greater than $\underline{p}$ are uncertified and $[\underline{u}, \overline{u}]$ is the lowest uncertified interval above $\underline{p}$. Define μ_u and μ_p as the means of all the uncertified and certified intervals. Consider uncertified types in $(\mu_p, 1]$. If we exactly reveal them, μ_u falls, favoring feasibility, and loss falls too. Next, realize that we need $\mu_u \leq \underline{p}$. If that were not true, then types in the interval $[\underline{p}, \mu_u)$ could be transferred to "Uncertified" and their contribution to expected loss would fall since they are closer to μ_u than to μ_p. Since they are below μ_p, this would also maintain feasibility.

We now know any uncertified types are below μ_p]. Consider the contribution to expected loss of just types in $[0, \underline{p}]$, $[\underline{p}, \underline{u}]$, and $[\underline{u}, \overline{u}]$:

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

Let us see what happens if we reduce $\underline{u}$ and then increase $\underline{\text{p}}$ by the same probability mass. This requires $dpf(p) = duf(u)$, so $\frac{dp}{du} = \frac{f(u)}{f(p)}$. Differentiating equation (15) yields

$$L(p, \mu_u)f(p)\frac{f(u)}{f(p)} - L(p, \mu_p)f(p)\frac{f(u)}{f(p)} + L(u, \mu_p)f(u) - L(u, \mu_u)f(u) \tag{16}$$

This is proportional to $[L(p, \mu_u) - L(u, \mu_u)] + [L(u, \mu_p) - L(p, \mu_p)]$. Since p is further from μ_u than u is (because $\mu_u \leq \underline{p}$), and p is further from μ_u than u is (since $\overline{u} < \mu_p$), and since L is convex, loss has fallen.

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

Proposition 3 (Pass-fail vs. Honors) *(i) For certification cost c low enough, no grading scheme is superior to honors grading, with the pass and honors cutoffs p and h going to zero as c goes to zero.*
(ii) For c high enough, if the quality density f is strictly logconcave the most informative grading scheme is either pass-fail, or is "almost pass fail"— honors with p going to 0 and h going to 1 as c rises

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, 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,

$$\min_p \max\left\{\int_0^p L\left(q, E[q|q \leq p]\right) f(q)dq, \int_p^1 L\left(q, E[q|q > p]\right) f(q)dq\right\} > 0. \quad (17)$$

Therefore, as c approaches 0 the optimal scheme is honors grading. Now consider the cutoffs. Suppose as c goes to zero, $h(c) - p(c)$ does not go to zero. Then by the same argument as above, the expected loss is bounded away from zero, whereas the expected loss for the exact scheme goes to zero—- which contradicts our supposition, so it must be that $h(c) - p(c)$ does go to zero.
(ii) Since f is logconcave, $\frac{d}{dx}E[q|q \leq x] < 1$ for any value x by Property (ii), and hence $\frac{d}{dx}\left(x - E[q|q \leq x]\right) > 0$, so the best exact scheme x^* is increasing in c and approaches 1 as c approaches $1 - E[q]$. Since logconcave functions are unimodal, Property (iv) tells us that the gap

$$E[q|q \in (p, x)] - E[q|q \leq p] \quad (18)$$

is maximized at either $p = x^*$ or $p = 0$.

First, suppose the gap (18) is maximized at $p = x^*$, where type x^* is just willing to be exactly revealed. If c is large enough, x^* is arbitrarily close to 1. Since (18) is maximized at $p = x^*$, the lowest feasible p is also x^* Therefore, as established in that proposition, expected loss falls by introduction of an honors scheme with $h > x^*$, and $p(h) < x^*$. This implies that for c high enough that x^* is close to 1, expected loss is lower with $h = 1$ and $p = p(1) < 1$ than for the exact scheme x^*. But such a scheme is pass-fail.

Second, suppose (18) is maximized at $p = 0$. As c approaches $E[q]$, to maintain feasibility p must be equal to or approach 0 and h must be equal to or approach 1, so an arbitrarily small fraction of certified senders are exactly graded. ■

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) $f(q)$ strictly quasiconcave and $f(1)$ is sufficiently large, or
(c) $f(q)$ is strictly logconcave and $E[q] > 1/2$.

Proof: Let $\bar{c}_x, \bar{c}_p$, and $\bar{c}_h$ represent the highest feasible certification costs for any exact, pass-fail, and honors grading scheme respectively, $\bar{c}_x \equiv \sup_{x\in[0,1)}\{c \leq x - E[q|q < x]\}$, $\bar{c}_p \equiv \sup_{p\in[0,1)}\{E[q|q > p]-E[q|q \leq p]\}$, and $\bar{c}_h \equiv \sup_{p\in[0,h),h\in(0,1)}\{E[q|q \in (p, h)] - E[q|q \leq p]\}$. We want to show $\bar{c}_p \geq \bar{c}_h \geq \bar{c}_x$.
(i) Consider any x such that exact grading is feasible and set $p = x$ and $h \in (x, 1]$. Since $E[q|q \in (p, h)] > p = x$, it must be that $E[q|q \in (p, h)] - E[q|q \leq p] > x - E[q|q < x]$, 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.
(ii) From (i) it is sufficient to show that $\bar{c}_h > \bar{c}_x$.
Condition (ii-a) says $f' > 0$. Regarding $\bar{c}_h$, by property (iii) the condition implies $E[q|q \in (p, h)] - E[q|q \leq p]$ is strictly decreasing in p, and by property (i) $E[q|q \in$

$(p,h)]$ is increasing in h, so $\overline{c}_h = E[q|q > 0] - 0 = E[q]$. Note $E[q] > 1/2$ since $f' > 0$ so $\overline{c}_h > 1/2$. Regarding $\overline{c}_x$, $f' > 0$ implies $E[q|q < x] > x/2$, so $\overline{c}_x < \sup_{x\in[0,1)} \{x - x/2\} = 1/2$. Hence, $\overline{c}_h > 1/2 > \overline{c}_x$.

Condition (ii-b) says f is strictly quasiconcave and $f(1)$ is sufficiently large. Regarding $\overline{c}_h$, for $f(1) > (1/2)/(1-E[q])$, property (iv) implies that $E[q|q \in (p,h)] - E[q|q \le p]$ reaches a maximum at $p = 0$ for any h, and by property (i) $E[q|q \in (p,h)]$ is increasing in h, so again $\overline{c}_h = E[q|q > 0] - 0 = E[q]$. Regarding $\overline{c}_x$, $E[q|q \ge x] > x$ so $\overline{c}_x < \max_x\{E[q|q \ge x] - E[q|q \le x]\}$. From property (iv), $E[q|q \ge x] - E[q|q \le x]$ is maximized at $x = 1$, so $\overline{c}_x < E[q|q \ge 1] - E[q|q \le 0] = E[q]$. Therefore $\overline{c}_h > E[q] > \overline{c}_x$.

Condition (ii-c) says $f(q)$ is strictly logconcave and $E[q] > 1/2$. Regarding $\overline{c}_h$, note that an honors grading scheme with $p = 0$ and h arbitrarily close to 1 is feasible for all $c < E[q]$, so $\overline{c}_h \ge 1/2$. For f logconcave, property (ii) implies $x - E[q|q < x]$ is maximized at $x = 1$, so $\overline{c}_x = 1 - E[q] < 1/2$. Therefore $\overline{c}_h > \overline{c}_x$. ■

References

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

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

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

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

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

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

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

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

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

[10] 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.

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

[12] 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.

[13] GINSBERG, WILLIAM (1973) "Concavity and Quasiconcavity in Economics," *Journal of Economic Theory*, 6, 596-605.

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

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

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

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

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

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

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

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

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

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

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

[25] 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.