Optimal Quality Ratings and Market Outcomes

Hugo Hopenhayn* — UCLA

Maryam Saeedi† — Tepper, CMU

December 15, 2018

Abstract

Quality ratings and certification are an increasingly important source of information in many markets and trading platforms. This paper considers the impact of improvements in information quality on equilibrium market outcomes (market size, consumer surplus and profits) and the optimal design of rating mechanisms. Improved information reassigns output from lower to higher quality firms, increasing the average quality of goods and total surplus. Notwithstanding, the effect on consumer can be negative so consumers may be better off in the absence of any information on producers' qualities. In case of linear supply or Cournot competition, the optimal rating system is given by thresholds that have a very simple representation and are the solution to a standard statistical clustering problem. In a series of examples, we show that a rating system with only one or two thresholds can achieve 60%-80% of the welfare gains from full information.

Keywords: Adverse Selection, Optimal Quality Rating System, Product Differentiation. **JEL codes:** D21, D47, D60, D82, L11

*hopen@econ.ucla.edu

†msaeedi@andrews.cmu.edu

1 Introduction

The importance of certification ratings in markets where product quality is imperfectly observed has been long emphasized. In particular, this is a key consideration for the overall performance of trading platforms and online markets, which are becoming increasingly important mechanisms for trade. In these settings, market designers often provide information signals that are imperfect aggregators of trade histories and are correlated with firm quality. This is the case, for example, of quality badges given by some trading platforms that partition sellers into a small number of groups (in many cases two). This paper considers the effect of information quality on market equilibrium outcomes (e.g., market size, consumer surplus, and profits) and characterizes an optimal quality rating system.

Specifically, we consider a market where suppliers (e.g., retailers on eBay or hosts on Airbnb) differ in terms of quality, but where consumers observe (coarse) imperfect signals, which are imperfectly correlated with quality. Such is the case, for example, of quality signals in many online markets. These quality signals partition the set of sellers into a discrete set of categories, in most cases as a result of some "testing" criteria. For example, in California, restaurants are given grades A, B, C, or none based on the score obtained after a hygiene inspection is conducted. eBay's high-quality sellers are classified as Top Rated Sellers, and Airbnb awards its top quality hosts the Superhost badge. From an economic point of view, it is important to distinguish both the size of the partition (i.e., the relative shares of sellers in each group) and their informativeness (i.e., the respective conditional average qualities). While raising the bar to qualify as a Top Rated Seller on eBay might contribute to a more select group, this can come at the cost of reducing the size of the market.

When analyzing the effect of changes in information design, it is necessary to understand how they affect equilibrium market outcomes. As an example, while better information is unambiguously good for consumers when prices are held constant, we find that under some conditions the equilibrium price response can lead to a final decrease in consumer surplus. These equilibrium responses can be important in practice, too. In a recent paper, Hui et al. [2018] examine the effect of an increase in the requirements to become a badged seller on eBay. They find that this increase leads to more entry of high-quality sellers but discourages entry of sellers in the medium range of quality.

Our baseline model considers a competitive market with a large set of buyers and sellers. Firms are endowed with different levels of quality, which is the only source of product differentiation. Production technology is the same for all firms, and it is given by a continuous, strictly increasing, and strictly convex cost function and, correspondingly, a strictly increasing supply function. All buyers have the same preference for quality but are heterogeneous in their value of the outside option. In addition, buyers are risk neutral with respect to quality and price, so goods are ranked according to the difference between expected qualities and prices. Information partitions the set of sellers into different levels of expected quality. An equilibrium is a set of prices as a function of expected quality and quantities supplied by firms such that markets clear.

The information structure shared by all buyers follows the setting described in Ganuza and Penalva [2010] and Gentzkow and Kamenica [2016]. A common prior over firm qualities and a signal structure determine the distribution of expected posterior firm qualities. As an example, if we consider the case of finite rating systems as those described above, this posterior is given by a discrete distribution with point masses at the conditional mean qualities associated to each rating. This setting provides a natural ordering of the quality of information, where better information is associated to a mean-preserving spread of the distribution of expected values. At the extremes, there is no information, so all firms have an expected value equal to the common prior mean or there is full information so the posterior equals the distribution of true seller qualities.

We start with the analysis of improved information on equilibrium outcomes. In general, a higher dispersion of expected qualities results in a higher dispersion of prices and a stronger association between prices and true seller quality. As a result, output is reallocated from lower- to higher-quality firms, which has a positive effect on the average quality of goods consumed. The effect this change has on total market size and consumer surplus is ambiguous and depends on the properties of the supply function of firms. When supply is convex, the higher spread in prices results in an increase in total output, while the opposite occurs when the supply functions are concave. In the latter case, any information provided by the market designer lowers consumer surplus. This is precisely because prices respond in such a way that all gains from improved information are competed away by consumers and thus absorbed by firm profits. However, we show that better information always increases total surplus.

The second part of our paper considers the question of optimal information design, based on maximizing total surplus. The results discussed above obviously imply that better information is preferred to worse, so an unconstrained designer would prefer to release all information. In reality, rating systems are designed with a limited number of categories, leaving open the question of how these should be defined. As mentioned earlier, signal structures can be ordered by the distribution of posterior expected qualities. Given a fixed number of categories N, any information structure can be associated to a posterior distribution supported in N points, that is, a mean-preserving reduction in spread of the original distribution of qualities. The optimal information structure is the one that maximizes total surplus in this set.

We then show that the optimal information structure can be found in the set of threshold partitions, i.e., those that partition sellers into ordered intervals of quality. When $N = 2$, which can be interpreted as a certification mechanism, the quality threshold for certification has a very simple representation. In the case of linear supply, the optimal threshold is exactly the midpoint between the mean quality values of the two groups. This is a lower bound when supply is convex, and an upper bound when supply is concave. This characterization extends to the case of multiple thresholds, and it coincides with the $k-$means criteria for clustering, minimizing the expected quadratic distance between true quality and the average one in the corresponding partition interval. For the linear supply case, the optimal thresholds also maximize total profits, revenues, and the average quality of the goods supplied. Our results provide a very straightforward and easy-to-compute method for the design of rating systems.

The optimal certification threshold can be related to the properties of the distribution of qualities. For a symmetric distribution, the optimal certification threshold coincides with the mean (and median) of the distribution. When the mean exceeds the median, the optimal threshold is above the mean, and, conversely, when the median exceeds the mean, the optimal threshold is below the mean. We provide numerical computations of the optimal thresholds and the increased performance resulting from finer partitions, for several standard distribution functions. Optimal thresholds are considerably high for some distributions; for example, in the case of an exponential distribution of qualities, only 20% of sellers should be certified. We measure the performance of a partition by the percentage of the welfare gap -- between full information and no information -- that is closed. A singe threshold (two-group) partition closes from 50% to nearly 80% of the gap, depending on the underlying distribution of qualities. The gains diminish as the number of thresholds increases. Even though our results suggest that the higher the number of certification tiers, the higher the total surplus, the market designer should weigh in the cost of having a more complicated information structure against the diminishing return of having more tiers.

We apply our procedure to data from eBay. Nosko and Tadelis [2015] provides a quality measure given by the percentage time a seller got positive feedback (as opposed to negative or none). The distribution is significantly skewed to the left, and the optimal certification threshold implies that more than 65% of sellers should be certified, closing more than 63% of the utility gap. In contrast, eBay gives the badge to only 5% of sellers, closing only 13% of the utility gap.

The last part of the paper considers a series of extensions. First we consider Cournot competition with constant marginal cost, instead of the baseline case of perfect competition. Our results show that the market outcome is very similar to the competitive case with linear supply functions. In particular, total output and consumer surplus remain unchanged when the information structure is modified. Moreover, the optimal threshold is identical to the one obtained in the case of perfect competition.

As a second extension, we add entry to the baseline model. Improvements in information induce entry, as a result of the increase in profits. Entry reduces prices, increases market size and, so, consumer surplus. In the extreme case where all potential entrants are ex-ante identical, drawing their qualities after entry, all gains to producers are competed away and consumer surplus increases unambiguously.

Our last extension considers a demand system as in the standard vertical differentiation model, where agents have heterogenous preference for quality, but where firms have inelastic supply. While by construction improvements in information do not increase total quantity, they contribute to welfare by increasing the correlation between average firm quality and consumer preference for quality. The optimal threshold is defined by a slightly modified formula that weighs differences in the firm's quality gap in each interval by the respective gap in consumers' preferences. As a result, skewness in consumers' preferences for quality has similar implications to the ones observed for skewness in producers' quality.

There is a large literature, both theoretical and empirical, related to reputation mechanisms. The theoretical literature focuses either on the role of private information, unraveling, and market breakdown or on the incentives faced by certifiers. However, to the best of our knowledge, there are no papers that study the effect of changes in information on market outcomes, with the exception of a simple illustrative example in Tadelis and Zettelmeyer [2015].[1] Our extensions to seller dynamics build on Shapiro [1983], Jovanovic [1982], and Hopenhayn [1992].

The closest empirical papers to our study are Elfenbein et al. [2015], Fan et al. [2013], and Jin and Leslie [2003]. Elfenbein et al. [2015] study the value of certification badges across different markets among different types of sellers. They find that certification provides more value when the number of certified sellers is low and when markets are more competitive. Fan et al. [2013] analyze the effect of badges on Taobao.com, the leading e-commerce platform in China. They find sellers offer price discounts to move up to the next reputation level. Jin and Leslie [2003] examine the

[1] A more detailed list of these papers can be found in two surveys on certification and reputation mechanisms, Dranove and Jin [2010] and Bar-Isaac et al. [2008].

effect of an increase in product quality information to consumers on firms' choices of product quality using data on restaurant hygiene ratings. Our paper also relates to the literature that analyzes the effects of changes in eBay's feedback mechanisms on price and quality (Klein et al. [2016], Hui et al. [2016], and Nosko and Tadelis [2015]). More generally, this research also fits broadly into the theme of understanding the effect of certification and reputation in e-commerce; representative papers include Cabral and Hortacsu [2010], Wu et al. [2015], Tadelis [2016], Dellarocas et al. [2006], Bajari and Hortacsu [2004], and Chevalier and Mayzlin [2006].

Section 2 describes the model. Section 3 considers the effect of improved information on market outcomes. Section 4 derives the conditions that characterize optimal thresholds. Section 5 numerically solves for the optimal thresholds for various distribution functions and considers the application to eBay ratings. Finally, Section 6 studies the extensions to the baseline model.

2 The Model

There is a unit mass of firms with qualities z distributed according to a continuous cdf $F(z_0)$. Production technology is the same for all firms and is given by a continuous, strictly increasing, and strictly convex cost function $c(q)$ and, correspondingly, a strictly increasing supply function $q(p)$. On the demand side, there is mass M of consumers who face a discrete choice problem, with preferences

$$U(z,\theta) = z + \theta,$$

where θ is distributed according to a continuous and strictly increasing cdf $H(\theta)$, and the outside good's utility (no purchase) is normalized to zero.[2] Goods are differentiated only by quality. Given the linearity in z, we maintain a similar expression for lotteries over qualities, where z is then interpreted as the corresponding mean.

The information structure shared by consumers follows the setting described in Ganuza and Penalva [2010] and Gentzkow and Kamenica [2016]. Given a common prior $F(z_0)$ over firm qualities and a signal structure π, let $G(z)$ be the distribution of expected posterior firm quality. As an example, if we consider the case of finite rating systems, as those described above, G is a discrete distribution with point masses at the conditional mean qualities associated to each rating.

Given average quality z, equilibrium prices take the form $p(z) = p(0) + z$, where $p(0)$ corresponds to the demand price of a good of quality zero. This expression for prices guarantees that all consumers are indifferent between goods with different

[2] Alternatively, one can consider $-\theta$ to be the value of the outside good to consumers.

Figure 1: Equilibrium

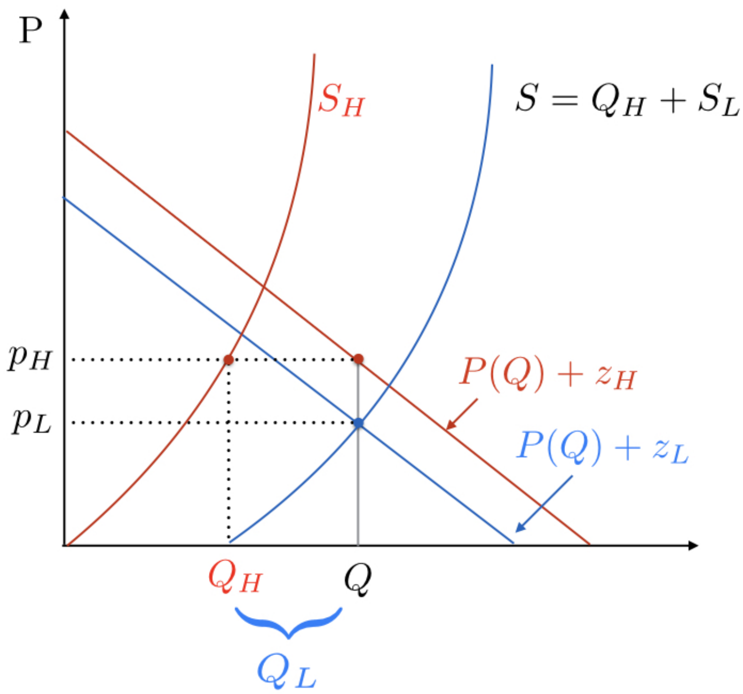

signal realizations. Each firm chooses its output after the quality signal is realized, choosing output $q(p(z))$,where z is the corresponding posterior mean, so total output $Q = \int q(p(z))\, dG(z)$. We can define a baseline inverse demand function for a good of quality zero by $P(Q) = H^{-1}(1 - Q/M)$ so that all consumers with $\theta \geq p$ purchase some good in the market. The demand system in the case of a partition into two expected quality groups is depicted in Figure 1.

Definition. An (interior) *equilibrium* for threshold given distribution of expected qualities $G(z)$ is given by prices $p(z) = P(Q) + z$,where total quantity $Q = \int q(p(z))\, dG(z)$.

Figure 1 shows graphically the derivation of an interior equilibrium for the case of a two-tier partition, where L represents the group of firms with quality below a threshold z^*, and H those above. The two curves depict the demand price for the goods in the L and H segment, respectively. Since all consumers value quality identically, the price difference $p_H - p_L = z_H - z_L$, i.e., the difference between the two respective average qualities. The first upward sloping curve is the H group's supply function: $(1 - F(z^*))\, q(p_H)$. The second one is the supply function of the L firms' segment $F(z^*)\, q(p_L)$ displaced to the right by the equilibrium quantity Q_H. The marginal consumer Q is the one that is indifferent between consuming either of these goods or none. At the equilibrium prices, Q is also the total market supply of both goods.

To prove the existence of an interior equilibrium, we make the following assumption.

Assumption 1. *There exists $\underline{\theta}$ in the support of H such that*

$$M > \int q(\underline{\theta} + z)\, dG(z).$$

This assumption rules out the possibility that all consumers make purchases in this market. While a corner equilibrium, if it exists, is also unique, we rule this out as a matter of convenience.

Proposition 1. *Under Assumption 1, there exists a unique interior equilibrium.*

Proof. Given that the cdf H is strictly increasing and continuous, the function $P(Q)$ is strictly decreasing and continuous. Define function $f(Q) = \int q(P(Q) + z)\, dG(z)$. This is strictly decreasing and continuous. It follows immediately that $f(0) > 0$. By assumption 1, $f(M) < M$ since $P(M) \leq \underline{\theta}$. Hence, there exists a unique fixed point Q^* for this mapping. □

Note that consumer surplus is determined by the area under the baseline demand function up to quantity Q,because all quality differences are arbitraged away. As a result, total consumer surplus will move together with total quantity.

3 Improved Information

This section considers the effect of improved information for general signal structures. We follow the definition of precision in Ganuza and Penalva [2010]. Given a common prior on qualities, a signal s_1 is more *precise* (*integral precise,* in Ganuza and Penalva [2010]) than signal s_2 if the induced distribution of expected qualities $G(z)$ generated by s_1 is a mean-preserving spread of the one generated by s_2. For the case of signals that partition sellers into two groups of fixed sizes with conditional means $H > L$, respectively, this coincides with the likelihood ratio ordering, i.e.,

$P(s_1 = H|v)/P(s_2 = H|v)$ is increasing in v. More generally, integral precision is weaker than the likelihood ratio and other related orderings considered in the literature (see Ganuza and Penalva [2010] for references.) We will refer to a more precise signal structure as better information.

Market Size and Consumer Surplus

We consider here the effect of precision on market size measured by the equilibrium quantity Q. For fixed quantity Q, better information implies a mean-preserving spread of prices. In the case of linear supply functions, this does not affect aggregate supply, so equilibrium quantity remains unchanged. In the case of convex supply functions, aggregate supply will exceed the original quantity Q, so aggregate output must increase. For similar reasons, in the case of concave supply functions, aggregate output decreases with the quality of information. Consumer surplus changes in the same direction as total quantity, as seen above.

Proposition 2. *Suppose signal s_1 is more precise than s_2. If in addition the supply function is convex (resp. concave), then total output $Q_1 \geq Q_2$ (resp. $Q_2 \leq Q_1$) and consumer surplus increase (resp. decrease).*

In particular, note that for the case of concave supply functions, consumers are better off with no signal provision.

Welfare Effect

Given a distribution of expected qualities G, the competitive equilibrium maximizes total surplus. Following our definitions above, a better information structure $\tilde{G}$ is a mean-preserving spread of G. Following Gentzkow and Kamenica [2016], there exists a garbling of signals that generates G from $\tilde{G}$. This means that a social planner could ignore the additional information contained in $\tilde{G}$ and reproduce the quantity-weighted distribution of average qualities corresponding to the optimal allocation under G and thus the same value. Since this is not optimal, it follows that total value increases with signal precision.

Proposition 3. *A better information system results in higher total surplus.*

In particular, the above proposition implies that average profits must rise when consumer surplus does not increase. Following Proposition 2, this will be the case when supply functions are concave or linear.

4 Optimal Signal Structure

In this section, we examine the question of optimal choice of reputation signals given the objective function, which could be total revenues or total surplus. The mechanism we consider divides sellers into a finite partition based on some criteria that are correlated with quality. As mentioned above, in the case of eBay, the partition includes two groups, the badged and unbadged. In the case of Yelp, the partition involves five stars, including the possibility of half-stars. In the case of California restaurants, the partition involves three elements, A, B, and C.

A threshold partition is a signal structure that totally orders firms into N quality intervals. Proposition 3 can be used to show the superiority of threshold partitions. Consider a partition of the set of sellers into sets $S_1, ..., S_N$. Suppose there are two sets S_k, S_{k+1} that are not totally ordered in quality. Substitute S_k and S_{k+1} by two new disjoint sets S'_k and S'_{k+1} of equal measures to the original ones that are totally ordered so that $S'_k < S'_{k+1}$. This generates a mean-preserving spread of the original distribution of means and thus higher surplus.

Corollary 1. *For any N, the optimal threshold partition maximizes total surplus (also revenues and profits) among the set of all N-group partitions.*

We now characterize the threshold partition that maximizes total surplus. Consider first the case where $N = 2$, as in the case of a certification system. This is defined by a threshold of qualities z^* and average qualities z_L and z_H for firms below and above this threshold, respectively. Total welfare is given by

$$\begin{aligned} W(z^*) &= \int_0^{Q(z^*)} P(x)\,dx + F(z^*)\,q_L z_L + (1 - F(z^*))\,q_H z_H - \\ &\quad - \left[C(q_L)\,F(z^*) + C(q_H)\,(1 - F(z^*))\right], \end{aligned} \tag{1}$$

where $Q(z^*) = F(z^*)\,q_L + (1 - F(z^*))\,q_H$. The following lemma gives the necessary conditions for an optimal z^*.

Lemma 1. *A necessary condition for an optimal threshold z^* is that*

$$(P(Q) + z^*)(q_H - q_L) = C(q_H) - C(q_L). \tag{2}$$

Proof. To totally differentiate equation (1) with respect to z^*, first note that by the envelope condition, we can ignore the effect on q_L and q_H. In particular, this implies that $\partial Q / \partial z^* = f(z^*)(q_L - q_H)$. Also note that

$$\frac{\partial z_L}{\partial z^*} = \frac{f(z^*)}{F(z^*)}(z^* - z_L), \frac{\partial z_H}{\partial z^*} = \frac{f(z^*)}{1 - F(z^*)}(z_H - z^*).$$

The result now follows by totally differentiating (1) and setting it equal to zero. □

Using the first-order condition (2), we can further characterize the optimal threshold. It is convenient first to rewrite this first-order condition as

$$(P(Q) + z^*)(q_H - q_L) = \int_{q_L}^{q_H} C'(q)\, dq.$$

In the case where $C'(q)$ is convex (concave supply function) then

$$\begin{aligned}\int_{q_L}^{q_H} C'(q)\, dq &< \frac{C'(q_L) + C'(q_H)}{2}(q_H - q_L)\\ &= \frac{p_L + p_H}{2}(q_H - q_L)\\ &= \left(P(Q) + \frac{z_L + z_H}{2}\right)(q_H - q_L),\end{aligned}$$

so it follows that $z^* < (z_L + z_H)/2$. Using a similar argument, if the supply function is convex, the opposite holds. This proves Proposition 4:

Proposition 4. *Suppose the supply function is concave (resp. convex), then z^* is lower (resp. higher) than $(z_L + z_H)/2$.*

Note that for the particular case of a linear supply function, the optimal threshold $z^* = (z_L + z_H)/2$,or equivalently,

$$z_H(z^*) - z^* = z^* - z_L(z^*). \tag{3}$$

Intuitively, at the margin the point z^* would contribute equally to the variance of each group, while the points to the right (resp. left) add less dispersion to the high-quality (resp. low-quality) group than to the other group.

While average firm quality $\bar{z}$ is independent of z^* when supply is linear, average product quality $(z_L Q_L + z_H Q_H)/Q$ is not, as the numerator changes with z^*. More precisely, average product quality,

$$\begin{aligned}z_{avg} &= z_L F(z^*) q_L + z_H (1 - F(z^*)) q_H\\ &= F(z^*) z_L (P(Q) + z_L) + (1 - F(z^*)) z_H (P(Q) + z_H)\\ &= F(z^*) z_L^2 + (1 - F(z^*)) z_H^2 + \bar{z}Q.\end{aligned}$$

Interestingly, this is also maximized at the optimal threshold:

$$2(z^* - z_L) + (z_L - z_H)(z_L + z_H) + (z_H - z^*).$$

Multiple Thresholds

We extend our previous result to the case where $z_1, z_2, ..., z_{N-1}$ is a vector of thresholds that partitions the support of F into N intervals, restricted to the case of linear supply. As seen, in this case total quantity and consumer surplus are independent of the thresholds and thus total surplus maximizes total profits. As can be shown, the thresholds that maximize seller's surplus solve

$$\max_{z_1,...,z_N} \sum_{i=0}^{N+1} \left[F(z_i) - F(z_{i-1})\right] m(z_{i-1}, z_i)^2, \tag{4}$$

where $m(z_{i-1}, z_i)$ is the conditional mean in the interval $[z_{i-1}, z_i]$, where z_0 and z_{N+1} are the lower and upper support of the distribution ($-\infty$ or $+\infty$ if unbounded), respectively. The corresponding first-order conditions are

$$z_i = \frac{m(z_{i-1}, z_i) + m(z_i, z_{i+1})}{2}.$$

This coincides with the $k-$means criteria for clustering. Uniqueness of these thresholds is guaranteed when the distribution has log-concave density (Mease and Nair [2006].) This makes estimating the optimal thresholds a trivial task, as many software programs incorporate algorithms to estimate these models. Additionally, when combining this result with the one in 4, the simple algorithm can give a market designer a good place to start given that the above thresholds will be the lower bound or upper bound, depending on the assumption for the supply function.

Imperfect Monitoring

We now relax the assumption that the certifier gets a perfect signal of quality. Let $f(z, s)$ denote the conditional pdf for quality z given signal s observed by the certifier. Let $\Phi(s)$ be the cdf of signals, and $\phi(s)$ the density. Let

$$\begin{aligned}
z_L(s^*) &= \frac{\int \int^{s^*} z f(z, s)\, d\Phi(s)\, dz}{\Phi(s^*)} \\
z_H(s^*) &= \frac{\int \int_{s^*} z f(z, s)\, d\Phi(s)\, dz}{1 - \Phi(s^*)} \\
z^*(s^*) &= \int z f(z, s^*)\, dz.
\end{aligned}$$

With linear supply, welfare is proportional to

$$W(z^*) \propto \Phi(s^*) z_L(s^*)^2 + \\ + (1-\Phi(s^*)) z_H(s^*)^2 .$$

Totally differentiating $W(z^*)$ with respect to z^* and equating to zero gives

$$z^*(s^*) = \frac{z_L(s^*) + z_H(s^*)}{2}.$$

This formula parallels the one obtained above for the case of perfect information.

5 Distribution Functions and Optimal Threshold

Equation (3) can be used to relate the threshold z^* to the properties of the distribution. In particular, note that for a symmetric distribution, the threshold $z^* = \bar{z} = z_{median}$. This follows directly from the formula noting that $F(z_{median}) = (1 - F(z_{median})) = 1/2$; so, letting $z^* = z_{median}$ and noting that $F(z^*) z_H + (1 - F(z^*)) z_L = \bar{z}$, it follows immediately from the first formula that $z^* = \bar{z}$. When the distribution $F(z)$ is asymmetric, the following result can be proved:

Proposition 5. *Suppose there is a unique point z^* satisfying (3) and $\bar{z}$ is greater (resp. smaller) than z_{median}. Then z^* is greater (resp. smaller) than $\bar{z}$.*

Proof. Let $g(z) = \frac{1}{2}(z_L(z) + z_H(z))$. The optimal threshold is a fixed point of this function. When $z \to z_{max}$ (or as $z \to \infty$ in the case of unbounded support), $g(z) \to \frac{1}{2}\bar{z} + \frac{1}{2}z < z$, and when $z \to z_{min}$(or as $z \to -\infty$ in the case of unbounded support), $g(z) \to \frac{1}{2}z_{min} + \frac{1}{2}\bar{z} > z$. Consider the case where $z_{mean} > z_{median}$, which is illustrated in Figure 2. For $z = z_{median}$, $g(z) = \bar{z} > z$. Since the function $g(z)$ is increasing and continuous, the unique fixed point z^* must be to the right of z_{median} and, as a consequence, $z^* > \bar{z}$, as illustrated in the figure. The proof for the case where $z_{mean} < z_{median}$ follows from a similar argument. □

Table 1 gives the value of z^* for a series of distributions. Except for the last case, all distributions are skewed to the right, so z^* is above the mean. The last column is instructive, as it shows the fraction of uncertified sellers. As an example, it is quite remarkable that for the Pareto distributions only a small fraction should get certified.[3] For the exponential distribution, only 20% of sellers should be certified regardless of the hazard rate.

[3] When $\alpha \leq 2$, the value of z^* is undefined, as total surplus is strictly increasing in z^* in all the support.

Figure 2: z^* when mean>median

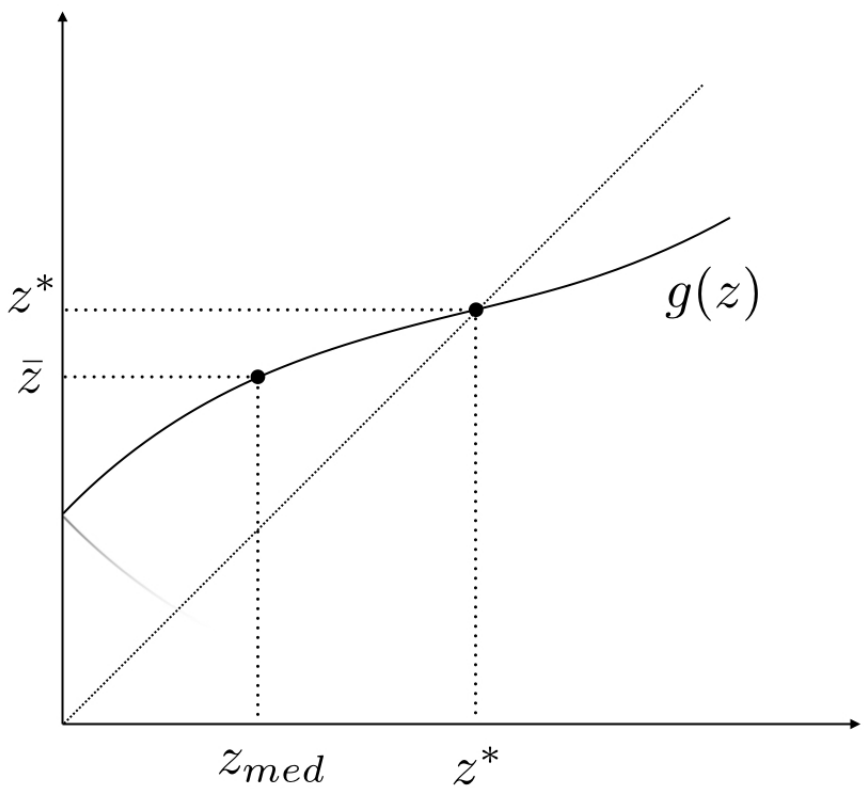

It is instructive to see how the gap between a full-information setting and no-information setting is closed as the number of groups in a partition increases. In the case of linear supply, we have shown that total output Q remains unchanged and so does consumer surplus. So, the gap in total profits equals the gap in total surplus. Profits for a firm with perceived quality z are given by $p(z)^2/2 = (P(Q)+z)^2/2$. Therefore, for any distribution G of expected quality, total profits are

$$\begin{aligned}\Pi &= \frac{1}{2}\int (P(Q)+z)^2\, dG(z) \\ &= \frac{1}{2}P(Q)^2 + P(Q)\,E(z) + \frac{1}{2}\int z^2 dG(z).\end{aligned}$$

The first two terms do not depend on G and thus on the partition. For full information, the distribution of means $G=F$, so the gap with respect to full information is $\Delta\Pi = \frac{1}{2}\left(\int z^2 dF(z) - \int z^2 dG(z)\right)$. For an n partition with thresholds $z_1, ..., z_{n-1}$ with corresponding means $\bar{z}_1, ..., \bar{z}_n$ the gap is

$$\begin{aligned}\Delta\Pi &= \frac{1}{2}\sum_{i=1}^{n}\int_{z_{i-1}}^{z_i}\left(z^2 - \bar{z}_i^2\right) dF(z) \\ &= \frac{1}{2}\sum_{i=1}^{n}\int_{z_{i-1}}^{z_i}(z-\bar{z}_i)^2\, dF(z).\end{aligned}$$

This is exactly the loss function in $k - means$ clustering, given that at the optimal thresholds as defined earlier, the expected values $\bar{z}_i$ are precisely the centroids of the corresponding interval $[z_{i-1}, z_i]$. Table 1 reports the total gaps and the percentage of the gap closed with partitions of different sizes n. As can be seen, in our simulations a one threshold (two-group) partition closes from 50% to nearly 80% of the gap, depending on the underlying distribution of qualities. The gains are diminishing as the number of thresholds increases. Even though our results suggest that the higher the number of certification tiers, the higher the total surplus, the market designer should weigh in the cost of having a more complicated information structure against the diminishing return of having more tiers.

Optimal Threshold for eBay

Nosko and Tadelis [2015] provides a quality measure given by the percentage time a seller got positive feedback (as opposed to negative or none). The distribution of this statistic across sellers is given in Figure 3 together with a density kernel estimator.[4] If we interpret this statistic as an ex-ante probability of a good (vs. a bad)

[4] The data for the histogram come directly from Table 4 in Nosko and Tadelis [2015].

Table 1: Optimal Thresholds

Distribution	Case	Mean	Median	z^*	$F(z^*)$	Share of Gap Closed			
						$n=2$	$n=3$	$n=5$	$n=10$
Pareto	$\alpha=3$	1.5	1.26	2.7	0.95	0.50	0.70	0.87	0.96
	$\alpha=4$	1.33	1.19	1.84	0.91	0.54	0.74	0.89	0.97
Exponential	$\lambda=1$	1.00	0.69	1.59	0.80	0.65	0.82	0.93	0.98
	general				0.80	same	same	same	same
$F(z)=z^{\alpha}/\alpha$	$\alpha=0.5$	0.33	0.25	0.41	0.64	0.77	0.90	0.97	0.99
	$\alpha=2$	0.67	0.71	0.62	0.38	0.72	0.87	0.95	0.99
Lognormal ($\mu=0$)	$\sigma=0.25$	1.03	1	1.09	0.64	0.63	0.81	0.92	0.98
	$\sigma=1$	1.64	1	4.23	0.93	0.55	0.75	0.89	0.97
eBay (Kernel)	band=0.2	0.65	0.67	0.61	0.34	0.63	0.80	0.92	0.98
eBay actual				0.86	0.95	0.13			

Figure 3: Distribution of % Positive Responses for eBay sellers

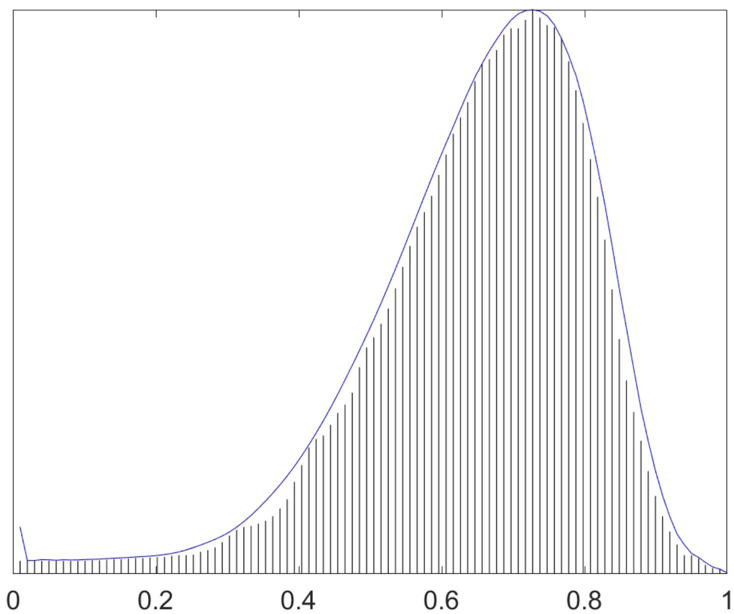

experience and the expected utility from this purchase as $(1 - P(good))\, u(bad) + P(good)u(good)$, then expected utility is an affine transformation of the probability of a good experience. Based on this interpretation, we can use this distribution to calculate the optimal thresholds and the gains from finer partitions, as done for other distributions above. Table 1 reports the results for the kernel estimate of this distribution. According to our calculations, more than 65% of sellers should be certified, closing about 63% of the utility gap. In contrast, eBay provides the badge to approximately 5% of sellers, representing a threshold of 0.86 and closing only about 13% of the utility gap.[5]

6 Extensions

In this section we consider a series of extensions. First we show the results for Cournot competition. Second, we analyze the effect of entry. Third, we consider

[5]Note that here we make very special assumptions in terms of supply and demand functions; for example, we assume a linear demand for the supply function, and if the supply function is convex the optimal z^* estimated is just a lower bound for the actual z^*.

heterogeneous preference for quality.

6.1 Cournot Competition

Here we restrict the analysis to a fixed set of firms, without considering explicitly the effect of changes in z^* on entry. There is a total of N firms (per consumer), and given signal z^*, a fraction $F(z^*)$ in the first group and $(1 - F(z^*))$ in the second. The demand structure is the same as in the competitive case considered above. Assume firms face a constant marginal cost c regardless of their type. The equilibrium conditions are

$$MR_H = P'(Q) q_h + P(Q) + z_H = c \tag{5}$$

$$MR_L = P'(Q) q_l + P(Q) + z_L = c. \tag{6}$$

Multiplying each equation by the number of firms in the respective group and adding up, we get

$$P'(Q) Q + NP(Q) + N\bar{z} = Nc,$$

where $\bar{z}$ is the mean quality for the N firms. Interestingly, this equation determines Q independently of the signal threshold z^*, as in the case of perfect competition.

While total quantity does not change, the shares of both groups do. In particular, after an increase in z^* $z_L(z^*)$ and therefore an increase in $z_H(z^*)$, the output of individual firms q_L and q_H also increases and so do p_H and p_L. This is compensated by some of the H firms becoming now L firms and lowering their output. It follows then that the output share of those firms that remain L and H increases, while that of the firms that shift category decreases.

Another implication of the invariance of total output is that, consumer surplus does not change, as in the case of linear supply with z^*. This occurs because price increases capture exactly the change in average quality in each group. It follows that optimal thresholds solve the maximization problem (4) so they are identical to those obtained above for the linear case.

6.2 Entry

In our previous analysis we did not consider explicitly the effect of changes of z^* on entry. Many of our results extend to settings where the distribution of qualities of firms is not affected by entry. We discuss here two scenarios: one where entrants are ex-ante homogenous and one when they are not.

Consider first the case where there is a mass N of entrants that are ex-ante differentiated in qualities z and fixed (or entry) costs f. Assume qualities are independent from fixed costs and are given by distribution F and Φ, respectively. For a given threshold partition z^*, we can define the aggregate supply functions S_L and S_H as follows. Let $S_H(p) = q(p) N_H(p)$, where $N_H(p) = N(1 - F(z^*)) \Phi(f(p))$, where $f(p) = \pi(p)$. This supply function combines the effect of prices on the intensive and extensive margin. Our analysis remains unchanged if we substitute $q(p)$ by $\hat{q}(p) = q(p) N\Phi(f(p))$, so total supplies are $S_L(p) = F(z^*)\hat{q}(p)$ and $S_H(p) = \hat{q}(p)(1 - F(z^*))$.[6]

For the homogenous case, assume there is a set N of potential entrants that draw their qualities independently from distribution F upon entry, after paying an entry cost f, which is distributed according to cdf $\Phi(f)$. For fixed output, improved information results in a mean-preserving spread of expected qualities and thus prices. Given that profit functions are convex in prices, this results in an increase in expected profits and a consequent increase in entry. In the case of linear supply, where in the absence of entry, total output does not change, this leads to an increase in total output and thus consumer surplus. In the case of concave supply, we have seen that total output decreases. This increases profits over and beyond what is produced by the mean-preserving spread of average qualities, thus inducing entry. This at least mitigates, if it does not undo, the drop in total output that would result in the absence of entry. Finally, note that if all potential entrants were to have the same entry cost, all surplus gains from improved information would accrue to consumers, as expected (and average) profits would remain unchanged. The above results apply in particular to the effect of introducing a certification mechanism in a market where there is none.

6.3 Heterogenous Preference for Quality

We examine briefly the determination of optimal thresholds when consumers differ in their preferences for quality. Suppose consumers' preferences are given by the utility function $u = \theta z + \theta_0 - p$ for a good of quality z. Consumers differ in their preference for quality θ and for the value they assign the inside vs. outside good θ_0, which is distributed in the population according to some joint distribution $\Psi(\theta, \theta_0)$. As earlier, firm qualities z are distributed according to cdf $F(z)$. For simplicity, we

[6]The properties of these modified supply functions will now depend both on the individual supply functions and the distribution of fixed costs. There exist assumptions on the latter that will guarantee that the modified supply functions are linear, convex, or concave when these properties hold for the original supply functions.

restrict our analysis to a partition of sellers into two groups defined by threshold z^* with qualities z_L and z_H, respectively. Given prices p_L and p_H, consumers will be split into three groups: those that do not consume and those that consume either the H or L product, with demands $D_H(p_L, p_H)$ and $D_L(p_L, p_H)$, respectively. These will be equilibrium prices provided that $D_H(p_L, p_H) = (1 - F(z^*))\, s(p_H)$ and $D_L(p_L, p_H) = F(z^*)\, s(p_L)$. As in our previous case, there is a unique equilibrium under fairly general conditions.

The optimal choice of threshold z^* satisfies the following first-order necessary condition:[7]

$$\Pi(p_H) - \Pi(p_L) \;=\; (z^* - z_L)\,\theta_L q_L + (z_H - z^*)\,\theta_H q_H, \tag{7}$$

where θ_L is the average preference for quality of the consumers that purchase the L product, and θ_H for those that purchase the H product. This formula has an intuitive explanation. The first term is the loss of profits by those firms that transition from the H to the L group, when z^* is marginally increasing. The second term measures the effect of the increase in the averages z_L and z_H as z^* is increased, valued at the quality preference of the average consumer in each group and weighted by their respective market sizes.

Vertical Differentiation with Inelastic Supply

To establish further results, we consider the canonical model of vertical differentiation where consumers differ only in their preference for quality θ and where firms supply inelastically one unit of output. Given equilibrium prices p_L and p_H, all consumers above a threshold θ^* buy an H product, while all those between $\underline{\theta}$ and θ^* buy an L product, where $\underline{\theta} z_L = p_L$ and $\theta^*(z_H - z_L) = p_H - p_L$. Substituting in equation (7) gives the condition

$$(z^* - z_L)(\theta^* - \theta_L) \;=\; (z_H - z^*)(\theta_H - \theta^*).$$

Notice that this is a modified version of equation (3), where the gaps between z^* and the respective means are weighted by the corresponding preference gaps. This highlights the role of the complementarities between average quality and preference for quality in the determination of the optimal threshold.[8] As an example, if both have uniform distributions, then when $\theta^* = z^* = 1/2$, this condition will hold:

[7] A proof is provided in the online appendix.

[8] It is interesting to note that when all consumers have the same preference for quality and supply is inelastic, welfare is independent of z^*, as the average product quality is not affected by its choice.

$$\max_{z} z_L(z)\,\theta_L(z)\,F(z) + z_H(z)\,\theta_H(z)\,(1 - F(z))\,.$$

References

Patrick Bajari and Ali Hortacsu. Economic insights from internet auctions. *Journal of Economic Literature*, 42(2):457–486, 2004. 6

Heski Bar-Isaac, Steven Tadelis, et al. Seller reputation. *Foundations and Trends® in Microeconomics*, 4(4):273–351, 2008. 5

Luis Cabral and Ali Hortacsu. The dynamics of seller reputation: Evidence from ebay*. *The Journal of Industrial Economics*, 58(1):54–78, 2010. 6

Judith A Chevalier and Dina Mayzlin. The effect of word of mouth on sales: Online book reviews. *Journal of marketing research*, 43(3):345–354, 2006. 6

Chrysanthos Dellarocas, Federico Dini, and Giancarlo Spagnolo. Designing reputation (feedback) mechanisms. 2006. 6

David Dranove and Ginger Zhe Jin. Quality disclosure and certification: Theory and practice. *Journal of Economic Literature*, 48(4):935–963, 2010. 5

Daniel W Elfenbein, Raymond Fisman, and Brian McManus. Market structure, reputation, and the value of quality certification. *American Economic Journal: Microeconomics*, 7(4):83–108, 2015. 5

Ying Fan, Jiandong Ju, and Mo Xiao. Losing to win: Reputation management of online sellers. 2013. 5

Juan-Jose Ganuza and Jose S. Penalva. Signal orderings based on dispersion and the supply of private information in auctions. *Econometrica*, 78(3):1007–1030, 2010. 3, 6, 8, 9

Matthew Gentzkow and Emir Kamenica. A rothschild-stiglitz approach to bayesian persuasion. *American Economic Review*, 106(5), 2016. 3, 6, 9

Hugo A Hopenhayn. Entry, exit, and firm dynamics in long run equilibrium. *Econometrica: Journal of the Econometric Society*, pages 1127–1150, 1992. 5

Xiang Hui, Maryam Saeedi, Zeqian Shen, and Neel Sundaresan. Reputation and regulations: evidence from ebay. *Management Science*, 62(12):3604–3616, 2016. 6

Xiang Hui, Maryam Saeedi, Giancarlo Spagnolo, and Steven Tadelis. Certification, reputation and entry: An empirical analysis. Working Paper 24916, National Bureau of Economic Research, August 2018. URL http://www.nber.org/papers/w24916. 2

Ginger Zhe Jin and Phillip Leslie. The effect of information on product quality: Evidence from restaurant hygiene grade cards. *The Quarterly Journal of Economics*, 118(2):409–451, 2003. 5

Boyan Jovanovic. Selection and the evolution of industry. *Econometrica: Journal of the Econometric Society*, pages 649–670, 1982. 5

Tobias J Klein, Christian Lambertz, and Konrad O Stahl. Market transparency, adverse selection, and moral hazard. *Journal of Political Economy*, 124(6):1677–1713, 2016. 6

David Mease and Vijayan N Nair. Unique optimal partitions of distributions and connections to hazard rates and stochastic ordering. *Statistica Sinica*, pages 1299–1312, 2006. 12

Chris Nosko and Steven Tadelis. The limits of reputation in platform markets: An empirical analysis and field experiment. Technical report, National Bureau of Economic Research, 2015. 4, 6, 15

Carl Shapiro. Premiums for high quality products as returns to reputations. *The quarterly journal of economics*, 98(4):659–679, 1983. 5

Steven Tadelis. Reputation and feedback systems in online platform markets. *Annual Review of Economics*, 8:321–340, 2016. 6

Steven Tadelis and Florian Zettelmeyer. Information disclosure as a matching mechanism: Theory and evidence from a field experiment. *The American Economic Review*, 105(2):886–905, 2015. 5

Chunhua Wu, Hai Che, Tat Y Chan, and Xianghua Lu. The economic value of online reviews. *Marketing Science*, 34(5):739–754, 2015. 6