Vertical Competition and Collusion in Information Acquisition*

Andrea Canidio† and Thomas Gall‡

March 16, 2016

Abstract

Competing sellers of heterogeneous products may generate public information about their goods, e.g., in the form of transparent quality tests, as a way to increase vertical product differentiation. We show that sellers' aggregate benefit from transparent quality testing exceeds its social benefit. Therefore sellers overinvest in quality testing, particularly if they can compensate each other, e.g., if quality is tested by an industry association. If product quality is endogenous, however, sellers' quality choices are distorted by a desire to vertically differentiate. In this case, quality testing can be socially beneficial because it reduces harmful quality degradation.

JEL classification: D21, D62, D83, L13, L15, L49

Keywords: Quality Testing, Quality Certification, Market Power, Collusion, Vertical Competition.

---

*We are grateful to Patrick Legros, Yossi Spiegel and Timothy Van Zandt for their valuable comments and suggestions. The usual caveat applies.

†Economics and Political Science Area, INSEAD, Boulevard de Constance, 77300 Fontainebleau, France; email: andrea.canidio@insead.edu

‡Department of Economics, University of Southampton, Southampton SO17 1BJ, United Kingdom; email: T.Gall@soton.ac.uk.

1 Introduction

Sellers often submit their products to experts and organizations for an objective and publicly observable evaluation. This public assessment can take the form of product classifications by industry bodies (e.g., for Bordeaux wines the wine classification of 1855, its more recently updated offshoot Cru Bourgeois, and similar systems for wines from Burgundy, Champagne, Douro, and other regions), reviews by customers or experts in the media (e.g., Consumer Reports, a non-profit organization that reviews consumer goods), quality tests and certification by professional agencies (e.g., rating agencies for financial products, TÜV for industrial goods). Such mechanisms generate public signals correlated with product quality, either in absolute terms or relative to the quality of competing products.

Generating these public signals has ambiguous effects on welfare. On the one hand, more precise information about product quality will improve the allocative efficiency of the market by facilitating an efficient assignment of heterogeneous goods to heterogeneous consumers. However, this information effect is countered by a competition effect: acquiring information through publicly observable signals will increase the expected distance between the qualities of different products, thus softening vertical competition and increasing the sellers' market rents.

In this paper, we use a textbook duopoly model of vertical competition to study private and social incentives to acquire information by means of costly quality testing. An immediate result is that the private benefit of quality testing exceeds its social benefit. The reason is that quality testing increases the expected distance between products, and therefore sellers' respective profits. Hence, each seller has an incentive to increase the level of quality testing beyond the social optimum, resulting in overinvestment in signals that are informative about product quality.

More intriguingly, we show that *each* seller benefits from an investment in better quality testing made by either seller, because any new piece of information increases the expected distance between quality levels, regardless of whether this information is relative to one, the other, or both products. As a consequence, the ability of sellers to coordinate their investment in learning is detrimental to social welfare. That is, even though there is perfect price competition in the product market, sellers may soften competition by colluding in generating publicly available information about products quality. This could be done, for instance, by classification systems or competitions organized by industry bodies.[1]

As a corollary, agents on one side of the market will not value public learning about attributes of agents on the other side of the market—especially coordinated public learning—as this will be accompanied by a loss of market rents. This appears consistent with the fact that in entry-level labor markets public signals about applicants' qualities have become less informative because of grade inflation (see, for instance, www.gradeinflation.com). In light of our model, this trend does not appear surprising: employers would have no incentive to demand more accurate information, while borrowing constraints and collective action problems may prevent applicants from generating more accurate signals.

Our results may reverse, however, when quality is endogenous and can be affected by sellers before they decide on quality testing. Without quality testing, sellers can increase the expected quality distance in the product market via their quality choices, which can lead to quality degradation. This problem is mitigated by the possibility of generating public information, because learning provides an alternative means to generate quality dispersion in the market, resulting in less quality degradation. Hence, when quality degradation is a concern, the sellers' private benefit of learning may fall short of its social benefit, so that encouraging sellers to cooperate in quality testing may be desirable from an aggregate welfare point of view.

The two papers that are closest to ours are Bouton and Kirchsteiger (2015)

[1] For example, the wine classifications mentioned above are usually sponsored by associations of vintners in the area.

and Bergemann and Välimäki (2000). Bouton and Kirchsteiger (2015) consider a model of vertical competition with uncertain quality levels, and show that the provision of reliable rankings can reduce consumers' welfare. We extend their findings by showing that, in general, any type of information (not only rankings) increases sellers' market power and decreases consumers' welfare.[2] Furthermore, we are interested here in sellers' incentives to generate information (which could take the form of establishing a ranking), both when they act individually and when they can coordinate. Bergemann and Välimäki (2000) consider a dynamic model of vertical competition in which sellers can generate information about their products. In their model information is acquired exclusively through repeated purchases. They also find that information acquisition increases sellers' market power and may reduce social welfare. Here we are interested in both individual and joint incentives to generate information before the product is brought to market. Moreover, we allow sellers to affect expected quality, which may make information acquisition socially valuable.

Within the auction literature, a number of authors raised points that are related to ours. Lewis and Sappington (1994) first argued that higher information dispersion among buyers leads to higher information rents. Closer to our model, Ganuza and Penalva (2010) analyze a problem in which one side of the market (the seller) generates information related to the *other* side of the market (the buyers' evaluations). They argue that the amount of information generated by the seller will fall short of the social optimum, because information increases the dispersion in buyers' evaluations and information rents. Here we focus on a market environment where one side can generate information about itself, and argue that information acquisition will be inefficiently high, especially if this side can act in a coordinated way.

A related literature has studied information production and information disclosure in the context of horizontal competition. In these model, sellers

[2]However, Bouton and Kirchsteiger (2015) analyze three channels through which rankings can affect consumers' welfare—rationing, consumption externalities and market power—whereas we focus exclusively on the latter channel.

may want to reveal information to the market whenever it increases market segmentation (see Johnson and Myatt, 2006, who call this effect a "demand rotation"). In particular, Anderson and Renault (2000, 2009) show that information acquisition may decrease consumers' welfare. Levin, Peck, and Ye (2009) consider a Hotelling model and argue that when sellers operate as a cartel they may disclose more information than under a duopoly. These results are clearly related to ours. However, in contrast to a model of horizontal product differentiation, information acquisition in a model of vertical competition will affect welfare directly by increasing the highest expected quality level in the market. Also, when the quality choice is endogenous, increasing vertical product differentiation by quality degradation is unambiguously harmful for welfare.

2 A model of vertical competition

To convey our argument in the simplest manner possible we employ a canonical, textbook model of vertical competition (as presented, for example, in Tirole, 1988, Chapter 7). The market consists of two sellers and a mass 1 of buyers. Call the two sellers 1 and 2, each producing a good of quality $\tilde{s}_i \in [\underline{s}, \overline{s}]$ for $i \in \{1, 2\}$. A buyer's utility is given by

$$U = \begin{cases} \theta \tilde{s}_i - p_i & \text{if good } s_i \text{ is purchased} \\ 0, & \text{if no purchase} \end{cases}$$

where p_i is the price of the good produced and θ is an i.i.d. taste parameter distributed uniformly over $[\underline{\theta}, \overline{\theta}]$ with $\overline{\theta} = \underline{\theta} + 1$.

As for the sellers, the marginal cost of production is zero, so that each seller's profit is given by price times quantity sold.

Information and Learning

We depart from the canonical model by assuming that the quality levels $\tilde{s}_1$ and $\tilde{s}_2$ are unknown to both buyers and sellers, who have common ex-ante beliefs $q_1 = E[\tilde{s}_1]$ and $q_2 = E[\tilde{s}_2]$. Quality levels can, however, be learned (at least partially) through information acquisition. For example, the technical specifications of each product may be perfectly known by all market participants, which generates symmetric information and the common priors q_i. From the consumers' point of view, however, the true product quality $\tilde{s}_i$ that determines the consumption utility generated by purchasing the product, may depend on subtle details unknown to the seller. Learning about a product's quality $\tilde{s}_i$ requires *quality testing*, which takes the form of generating an informative, noisy signal σ_i of $\tilde{s}_i$ at a cost k.[3] This is best understood as submitting the product to a public quality review process, for instance quality assessment and certification by industry or professional bodies, product testing by independent experts, or official quality evaluation by government institutions. Information generated by quality testing is public: all market participants receive the signal and update their belief about quality.[4]

Formally, we express the market belief over the quality level of seller i using a c.d.f. $F_i(s) : [\underline{s}, \overline{s}] \rightarrow [0, 1]$, representing the probability that the "true" quality $\tilde{s}$ is below some level s. Given $F_i(s)$, we can compute the expected quality level of the good sold by seller $i \in \{1, 2\}$ as

$$q_i = \int_{\underline{s}}^{\overline{s}} s \; dF_i(s).$$

Without loss of generality, let seller 1 be the quality leader ex ante, i.e., $q_1 > q_2$. Let $F_i(s|\hat{\sigma})$ denote the posterior belief distribution and define the posterior

[3]Note that all derivations and results presented in this paper extend to situations in which the signal drawn by a seller is also informative about the other seller's quality.

[4]We abstract away from the issue of disclosure; see Section 3 for a discussion regarding the effect of asymmetric information on our results.

expected quality as

$$q_i^{\hat{\sigma}} = \int_{\underline{s}}^{\overline{s}} s \; dF_i(s|\hat{\sigma}),$$

where $\sigma = (\sigma_1, \sigma_2)$ is the vector of signals, and $\hat{\sigma}$ is a specific realization of the vector of signals. We adopt the convention that $\sigma_i = \emptyset$ if seller i did not submit the product to quality testing, so that $q_i = q_i^{(\emptyset,\emptyset)}$. Note, also, that by iterating expectations $E\left[q_i^{\hat{\sigma}}\right] = q_i$, meaning that before the signal is drawn the expected posterior quality in case of a future draw is equal to the prior expected quality.

Timing

To summarize, the timing of the game is as follows.

1. $\tilde{s}_1$, $\tilde{s}_2$, $F_i(s)$, $F_2(s)$, are exogenously determined.[5]
2. Sellers decide whether to submit their good for quality testing at cost k, yielding a signal σ_i.
3. Realizations of signals are publicly revealed.
4. Sellers simultaneously announce prices. Consumers decide if and from whom to buy and consume. Payoffs are realized.

Covered Market

Finally, following the textbook model, we impose a parametric assumption that guarantees that for every expected quality levels $q_i^{\hat{\sigma}}$ each consumer buys from at least one seller, i.e., that the market is "covered", which simplifies the sellers' pricing problems.

Assumption 1. *There is sufficient heterogeneity in tastes across consumers: $\overline{\theta} \geq 2\underline{\theta}$, and the valuation of the least quality-sensitive consumer is sufficiently large: $\underline{\theta} \geq \overline{\theta}\frac{\overline{s}-\underline{s}}{2\overline{s}-\underline{s}}$.*

---

[5] See Section 2.4 for the possibility that sellers can affect expected quality q_i.

Note that, because $\overline{\theta} = \underline{\theta} + 1$, these two conditions are equivalent to $1 - \frac{\underline{s}}{\overline{s}} \leq \underline{\theta} \leq 1$. As will be shown below, under the above assumption the solution to each seller's pricing problem is always interior, and both sellers' profit functions are simple piecewise linear functions.[6]

2.1 Benchmark: Constrained efficient allocation

We start by deriving the constrained efficient allocation: the production, consumption and testing choice of a social planner who has the same ex-ante beliefs as buyers and sellers. Recall that the cost of production is zero, so that the social planner maximizes welfare by allocating the highest-quality good to all buyers. Hence, for given posterior quality levels $q_1^{\hat{\sigma}}$ and $q_2^{\hat{\sigma}}$, aggregate welfare is

$$W(q_1^{\hat{\sigma}}, q_2^{\hat{\sigma}}) = \max\{q_1^{\hat{\sigma}}, q_2^{\hat{\sigma}}\} \left(\frac{\overline{\theta} + \underline{\theta}}{2}\right)$$

Consider now the social planner's optimal choice of quality testing. Noting that $W(q_1^{\hat{\sigma}}, q_2^{\hat{\sigma}})$ is convex in both arguments and applying Jensen's inequality reveals that $E\left[\max\{q_1^{\hat{\sigma}}, q_2^{\hat{\sigma}}\}|\hat{\sigma}\right] \geq \max\{q_1, q_2\}$. The social benefit of learning in the constrained efficient allocation is given by

$$\begin{aligned}
E\big[W(q_1^{\hat{\sigma}}, q_2^{\hat{\sigma}})\big] - W(q_1, q_2) &= \left(\frac{\overline{\theta} + \underline{\theta}}{2}\right)\big(E\big[\max\{q_1^{\hat{\sigma}}, q_2^{\hat{\sigma}}\}|\sigma\big] - q_1\big) = \left(\frac{\overline{\theta} + \underline{\theta}}{2}\right)\big(E\big[\max\{q_1^{\hat{\sigma}}, q_2^{\hat{\sigma}}\}|\sigma\big] - E[q_1^{\hat{\sigma}}]\big) \\
&= \left(\frac{\overline{\theta} + \underline{\theta}}{2}\right)\big(E[q_1^{\hat{\sigma}}|q_1^{\hat{\sigma}} \geq q_2^{\hat{\sigma}}]\mathrm{pr}\{q_1^{\hat{\sigma}} \geq q_2^{\hat{\sigma}}\} + E[q_2^{\hat{\sigma}}|q_2^{\hat{\sigma}} \geq q_1^{\hat{\sigma}}]\mathrm{pr}\{q_2^{\hat{\sigma}} \geq q_1^{\hat{\sigma}}\} \\
&\quad - E[q_1^{\hat{\sigma}}|q_1^{\hat{\sigma}} \geq q_2^{\hat{\sigma}}]\mathrm{pr}\{q_1^{\hat{\sigma}} \geq q_2^{\hat{\sigma}}\} - E[q_1^{\hat{\sigma}}|q_2^{\hat{\sigma}} \geq q_1^{\hat{\sigma}}]\mathrm{pr}\{q_2^{\hat{\sigma}} \geq q_1^{\hat{\sigma}}\}\big) \\
&= \left(\frac{\overline{\theta} + \underline{\theta}}{2}\right) \mathrm{pr}\{q_2^{\hat{\sigma}} \geq q_1^{\hat{\sigma}}\} E[q_2^{\hat{\sigma}} - q_1^{\hat{\sigma}}|q_2^{\hat{\sigma}} \geq q_1^{\hat{\sigma}}] \\
&\equiv \left(\frac{\overline{\theta} + \underline{\theta}}{2}\right) \Delta^{\sigma}(q_1, q_2),
\end{aligned}$$

[6] In Section 3 we argue that our main results hold qualitatively in any model of vertical competition, including those in which Assumption 1 may be violated.

where

$$\Delta^{\sigma}(q_1, q_2) \equiv \text{pr}\{q_2^{\hat{\sigma}} \geq q_1^{\hat{\sigma}}\} E[q_2^{\hat{\sigma}} - q_1^{\hat{\sigma}} | q_2^{\hat{\sigma}} \geq q_1^{\hat{\sigma}}],$$

is the *expected quality gain*, i.e., the average gain in quality conditional on the event that the quality ranking changes, weighted by the probability that the quality ranking changes.

That is, if there is a realization of the signal vector σ, such that the conditional expectation of seller 2's quality is greater than that of seller 1, then quality testing strictly increases the expected aggregate welfare from the consumption allocation (i.e., not taking into account the cost of quality testing). The social benefit of learning depends on aggregate valuations and is linear in the expected quality gain $\Delta^{\sigma}(q_1, q_2)$, i.e., quality testing generates a social benefit if, and only if, the identity of the quality leader can potentially change. Another observation is noteworthy at this point: characteristics of the quality distribution determine the expected quality gain and thus the social benefit of learning. For instance, the closer the priors q_1 and q_2, the higher the probability of a ranking reversal and the greater the average gain in quality should a reversal occur.

The socially optimal investment in quality testing then solves

$$\max_{\hat{\sigma} \in \{\emptyset; \sigma_1\} \times \{\emptyset; \sigma_2\}} E\big[W(q_1^{\hat{\sigma}}, q_2^{\hat{\sigma}})\big] - \begin{cases} 2k \text{ if } \hat{\sigma} = (\sigma_1, \sigma_2) \\ k \text{ if } \hat{\sigma} \in \{(\emptyset, \sigma_2), (\sigma_1, \emptyset)\} \\ 0 \text{ if } \hat{\sigma} = (\emptyset, \emptyset) \end{cases} .$$

Intuitively, social welfare is a linear function of the highest quality in the market, so that learning is beneficial if a reversal of the quality ranking is a possible outcome and if the cost of obtaining a signal is not too high. Whether it is optimal to learn about only the quality leader, only the quality follower, or both will depend on the distributions of qualities, priors and signals.

2.2 Competitive Market Equilibrium

Turn now to the outcome in a competitive market, where sellers compete for buyers, given their expected product quality conditional on any signal obtained through quality testing. Invoking Assumption 1 the market will be covered and all buyers buy from one of the sellers. For given $q_1^{\hat{\sigma}}$ and $q_2^{\hat{\sigma}}$, (following Tirole, 1988, page 296) in equilibrium sellers set prices equal to:

$$p_i = \begin{cases} \frac{2\bar{\theta}-\underline{\theta}}{3}|q_i^{\hat{\sigma}} - q_{-i}^{\hat{\sigma}}| & \text{if } q_i^{\hat{\sigma}} > q_{-i}^{\hat{\sigma}} \\ \frac{\bar{\theta}-2\underline{\theta}}{3}|q_i^{\hat{\sigma}} - q_{-i}^{\hat{\sigma}}| & \text{if } q_i^{\hat{\sigma}} \leq q_{-i}^{\hat{\sigma}} \end{cases}$$

All buyers with $\theta \geq \frac{\bar{\theta}+\underline{\theta}}{3}$ buy from the quality leader and all other buyers buy from the quality follower. As a consequence, profits are given by

$$\pi_i = \begin{cases} \frac{(2\bar{\theta}-\underline{\theta})^2}{9}|q_i^{\hat{\sigma}} - q_{-i}^{\hat{\sigma}}| & \text{if } q_i^{\hat{\sigma}} > q_{-i}^{\hat{\sigma}} \\ \frac{(\bar{\theta}-2\underline{\theta})^2}{9}|q_i^{\hat{\sigma}} - q_{-i}^{\hat{\sigma}}| & \text{if } q_i^{\hat{\sigma}} \leq q_{-i}^{\hat{\sigma}} \end{cases} \tag{1}$$

and are shown in Figure 1.

Note that both profit functions π_1 and π_2 are increasing in the quality distance $|q_1^{\hat{\sigma}} - q_2^{\hat{\sigma}}|$. Intuitively, the distance between the two quality levels is a measure of vertical differentiation and thus sellers' market power. A large distance implies that each seller can extract higher rents from its respective consumers. A distance of zero implies that sellers compete à la Bertrand, set their prices equal to their marginal cost, and obtain zero profits.

This observation becomes relevant when analyzing the incentive to invest in learning through quality testing, because distance is a convex function of the difference in quality $q_i^{\hat{\sigma}} - q_{-i}^{\hat{\sigma}}$. The increase in expected distance due to

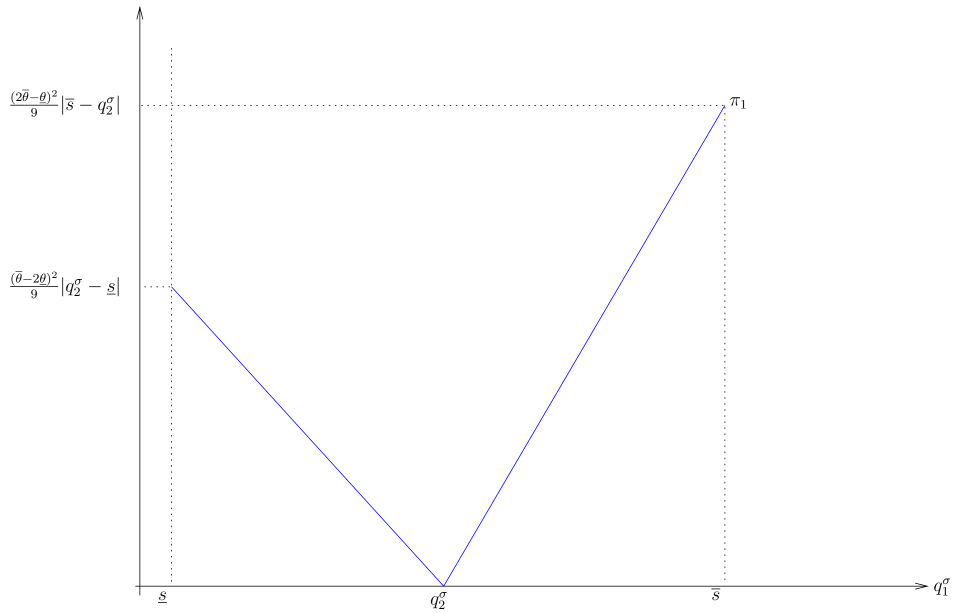

Figure 1: Seller 1's profits as a function of own expected quality.

quality testing can be computed as:

$$
\begin{aligned}
&E\left[|q_1^{\hat{\sigma}} - q_2^{\hat{\sigma}}|\right] - |q_1 - q_2| = E\left[|q_1^{\hat{\sigma}} - q_2^{\hat{\sigma}}|\right] - E\left[q_1^{\hat{\sigma}} - q_2^{\hat{\sigma}}\right] \\
=&\mathrm{pr}\{q_1^{\hat{\sigma}} \geq q_2^{\hat{\sigma}}\}E\left[q_1^{\hat{\sigma}} - q_2^{\hat{\sigma}} \ |q_1^{\hat{\sigma}} \geq q_2^{\hat{\sigma}}\right] + \mathrm{pr}\{q_2^{\hat{\sigma}} \geq q_1^{\hat{\sigma}}\}E\left[q_2^{\hat{\sigma}} - q_1^{\hat{\sigma}} \ |q_2^{\hat{\sigma}} \geq q_1^{\hat{\sigma}}\right] \\
&- \mathrm{pr}\{q_1^{\hat{\sigma}} \geq q_2^{\hat{\sigma}}\}E\left[q_1^{\hat{\sigma}} - q_2^{\hat{\sigma}} \ |q_1^{\hat{\sigma}} \geq q_2^{\hat{\sigma}}\right] - \mathrm{pr}\{q_2^{\hat{\sigma}} \geq q_1^{\hat{\sigma}}\}E\left[q_2^{\hat{\sigma}} - q_1^{\hat{\sigma}} \ |q_2^{\hat{\sigma}} \geq q_1^{\hat{\sigma}}\right] \quad = 2 \cdot \Delta^\sigma(q_1, q_2).
\end{aligned}
$$

Therefore, the expected change in distance is proportional to the expected quality gain $\Delta^\sigma(q_1, q_2)$, and is strictly positive if the ranking between sellers may change for some realizations of σ. Hence, quality testing by any seller increases both sellers' expected market profits. More precisely, the benefit from submitting one's own product to quality testing while the other player does not is given by (the details are in the appendix, proof of Proposition 1):

$$
\pi_i^{(\sigma_i,\emptyset)} - \pi_i^{(\emptyset,\emptyset)} = \left(\frac{(\overline{\theta} - 2\underline{\theta})^2 + (2\overline{\theta} - \underline{\theta})^2}{9}\right)\Delta^{(\sigma_i,\emptyset)}(q_1, q_2) \text{ for } i = 1, 2. \qquad (2)
$$

This expression is positive whenever $\Delta^{\sigma}(q_1, q_2)$ is positive, i.e., whenever for some realizations of the signals the ranking between the two sellers may change. Similarly, seller i's benefit from investing in quality testing when the other player $(-i)$ invests as well is given by:

$$\pi_1^{(\sigma_i,\sigma_{-i})} - \pi_1^{(\emptyset,\sigma_{-i})} = \left(\frac{(\overline{\theta} - 2\underline{\theta})^2 + (2\overline{\theta} - \underline{\theta})^2}{9}\right) \left(\Delta^{(\sigma_i,\sigma_{-i})}(q_1, q_2) - \Delta^{(\emptyset,\sigma_{-i})}(q_1, q_2)\right). \quad (3)$$

This expression is positive whenever drawing two signals rather than one either increases the probability of a change in ranking, or increases the distance between quality levels conditional on a change in the quality ranking, or both.

Interestingly the benefits of quality testing are symmetric: they are identical for both the quality leader and follower. The reason is that the profit functions of both leader and follower are linear in the expected quality difference (see Figure 1). Also, for both sellers the expected posterior qualities are equal to the qualities before testing. Therefore each seller's expected payoff gain is exactly zero if the quality ranking remains the same, but strictly positive for signal realizations that reverse the quality ranking. The expected payoff gain in case of a ranking change is the same for both sellers, because both sellers switch payoff functions at the inflection point in Figure 1, each gaining the difference between the new payoff function and the old, which are the same for both sellers as they switch payoffs.

We can also compute the seller i's benefit from seller $-i$'s investment in quality testing, which is

$$\pi_i^{(\emptyset,\sigma_{-i})} - \pi_1^{(\emptyset,\emptyset)} = \left(\frac{(\overline{\theta} - 2\underline{\theta})^2 + (2\overline{\theta} - \underline{\theta})^2}{9}\right) \Delta^{(\emptyset,\sigma_{-i})}(q_1, q_2). \quad (4)$$

Hence, quality testing by one seller imposes a positive externality on the other seller, so that sellers' joint benefit from quality testing is larger than the sum of the two individual benefits. Therefore there are cost parameters k, for which no seller invests in learning in equilibrium but, when sellers can jointly decide on the level of quality testing and share its cost, quality testing by one seller may occur. Similarly, for some level of k only one seller learns in equilibrium, and both sellers learn when they can collude. The following proposition summarizes these observations.

Proposition 1. *A seller's benefit from obtaining signals $\hat{\sigma}$ through quality testing is the same for quality leader and follower. The benefit is positive if, and only if, the expected quality gain $\Delta^{\hat{\sigma}}(q_1, q_2)$ is positive, and the benefit increases proportionally in the expected quality gain.*
If sellers can collude and jointly decide on their investment in quality testing, they will invest (weakly) more than if they decide individually.

Proof. In appendix. □

2.3 Welfare

Deriving a normative statement requires us to compare the private benefit of quality testing (as derived in Section 2.2) with its social benefit. The latter is given by the change in aggregate utility caused by quality testing. Given prior beliefs $q_1^{\hat{\sigma}}$ and $q_2^{\hat{\sigma}}$ and using the results above, aggregate utility from consumption (excluding possible cost of quality testing) is given by:

$$\begin{aligned} S(q_1^{\hat{\sigma}}, q_2^{\hat{\sigma}}) &= \max\{q_1^{\hat{\sigma}}, q_2^{\hat{\sigma}}\} \left(\frac{4\overline{\theta} + \underline{\theta}}{6}\right) \left(\overline{\theta} - \frac{\overline{\theta} + \underline{\theta}}{3}\right) + \min\{q_1^{\hat{\sigma}}, q_2^{\hat{\sigma}}\} \left(\frac{\overline{\theta} - 2\underline{\theta}}{6}\right) \left(\frac{\overline{\theta} + \underline{\theta}}{3} - \underline{\theta}\right) \\ &= \max\{q_1^{\hat{\sigma}}, q_2^{\hat{\sigma}}\} \left(\frac{2\overline{\theta}^2 + (\overline{\theta} - \underline{\theta})^2}{6}\right) - |q_1^{\hat{\sigma}} - q_2^{\hat{\sigma}}| \frac{(\overline{\theta} - 2\underline{\theta})^2}{18}. \end{aligned} \quad (5)$$

This expression highlights the two competing effects that learning has on aggregate welfare. On the one hand, quality testing generates an increase in the expected highest quality. On the other hand, quality testing has an adverse competition effect by increasing the expected quality distance as shown in Section 2.2 above, which increases the second term and thus depresses social welfare. This second effect stems from the vertical competition: quality testing expands the expected quality distance between the two sellers' qualities, increasing sellers' market power and rents. This competition effect is absent from the constrained optimum in Section 2.1, where welfare increases proportionally to $\max\{q_1^{\hat{\sigma}}, q_2^{\hat{\sigma}}\}$.

The expected change in aggregate welfare in the competitive equilibrium outcome due to quality testing and generating signals σ can be computed as:

$$E[S(q_1^{\hat{\sigma}}, q_2^{\hat{\sigma}})] - S(q_1, q_2) = \left(\frac{2\overline{\theta}^2 + (\overline{\theta} - \underline{\theta})^2}{6} - \frac{(\overline{\theta} - 2\underline{\theta})^2}{9} \right) \Delta^\sigma(q_1, q_2). \qquad (6)$$

It is quickly verified that this expression is strictly positive whenever the expected quality gain $\Delta^\sigma(q_1, q_2)$ is positive, i.e., if learning induces a quality rank reversal with positive probability. Therefore, some degree of learning about qualities is indeed beneficial from the social perspective whenever learning may result in a change of the quality leader.

Due to the competition effect of learning the social benefits from learning will differ from the private benefits, however. Comparing the private benefit of information acquisition by a given seller (Lemma 1) to the social benefit yields

$$\frac{(\overline{\theta} - 2\underline{\theta})^2 + (2\overline{\theta} - \underline{\theta})^2}{9} = \frac{2\underline{\theta}^2 + 2\underline{\theta} + 5}{2} > \frac{4\underline{\theta}^2 + 16\underline{\theta} + 7}{18} = \frac{2\overline{\theta}^2 + (\overline{\theta} - \underline{\theta})^2}{6} - \frac{(\overline{\theta} - 2\underline{\theta})^2}{9},$$

where we used the fact $\overline{\theta} = \underline{\theta} + 1$. This observation allows us to state the following:

Proposition 2 (Overinvestment in Quality Testing)**.** *In a competitive market each seller's private benefit of quality testing is strictly greater than the social benefit of quality testing. Therefore, sellers will invest (weakly) too much in learning about qualities in a competitive market relative to the social optimum.*

That is, each seller has greater benefit from investing in quality testing than the economy as a whole. The reason for this is that the social benefit comes solely from a more precise expectation of qualities. This will trigger an adjustment of prices and quantities sold and have a neutral effect on each seller's profit in expectation. However, each seller benefits additionally from the expected increase in the distance of expected quality. Since each seller's private benefit from investing exceeds the social benefits, the buyers' aggregate utility must decrease.

Corollary 1. *Quality testing always decreases consumers' welfare.*

In the previous section we established that quality testing by one seller has a positive externality on the other seller. It follows that if sellers coordinate and jointly decide on an investment in quality testing the over-investment problem becomes more severe, because the joint benefit of investing exceeds that of each individual seller.

Corollary 2 (Collusion and Overinvestment)**.** *Coordination in quality testing is welfare decreasing.*

2.4 Endogenous quality choice

One aspect of the canonical model of vertical competition that we have so far ignored is that sellers have a strong incentive to degrade their products to increase the distance between their products' qualities (see, for example, Tirole, 1988). Quality degradation is observed in some instances,[7] but it is much less prevalent than what the standard theory of vertical competition predicts. In this section we argue that the expectation of future quality testing may eliminate the incentives to degrade quality.

Denote by $q_i^0 \in [\underline{s}, \overline{s}]$ seller i's initial quality level with the convention that $q_1^0 > q_2^0$. Before the market opens, both sellers simultaneously can decrease their quality at zero cost to any $q_i \in [\underline{q}, q_i^0]$ with $\underline{s} \leq \underline{q} \leq q_i^0$ for $i = 1, 2$. Actions are publicly observable. As both sellers' profits increase in the distance

[7] Returning to our example of wine-making, sellers may adopt inferior cultivation techniques or extend their vineyard to less desirable territory and then blend their grapes. Similarly, several producers of electronic devices are known to intentionally reduce the performance and functionality of their products, the case of IBM printers being the most well-known example.

between their expected quality levels, absent quality testing the possibility of quality degradation results in maximum differentiation. Indeed, in the pure strategy Nash equilibrium of this game one seller will choose to maintain the default quality q_i^0 and the other to degrade the quality as much as possible to $q_{-i} = \underline{q}$.

To account for quality testing, suppose that after sellers decide on degrading their quality to $q_i \in [\underline{q}, q^0]$ they can decide whether to submit their product to quality testing for a cost k. Introducing quality testing into the framework with quality degradation may change the equilibrium, because quality testing provides an alternative means to increase the quality distance between sellers. However, in contrast to degradation, quality testing allows for upwards revisions of the expected quality as well as downwards revisions, increasing the expected highest quality and thus aggregate surplus. Also, both means of vertical differentiation are interdependent, because the benefit from learning through quality testing decreases in the difference in expected quality. That is, if one seller engages in degradation, the benefit of quality testing decreases and thus quality testing and degradation are imperfect substitutes for each seller.

The following statement shows that, in line with this intuition, quality testing will occur in a subgame perfect Nash equilibrium when the profits earned by the quality follower are sufficiently low and the expected gain in quality through learning sufficiently high given its cost k.

Proposition 3. *Suppose that $k > 0$ is sufficiently small:*

$$\frac{(2\overline{\theta} - \underline{\theta})^2 + (\overline{\theta} - 2\underline{\theta})^2}{9} \Delta^\sigma(q_1^0, q_2^0) > k \tag{7}$$

for either $\sigma = (\sigma_1, \emptyset)$ or $\sigma = (\emptyset, \sigma_2)$. If $\overline{\theta}$ is sufficiently close to $2\underline{\theta}$, then there will be quality testing in equilibrium and both sellers will choose not to degrade their quality (i.e., $q_i = q_i^0$ for $i = 1, 2$).

Proof. In appendix. □

If sellers anticipate quality testing (by either seller), and assuming that $\overline{\theta}$ is close enough to $2\underline{\theta}$, then both sellers' payoff increases in own quality q_i and there is no incentive to degrade quality anymore. If instead sellers anticipate that there will be no quality testing (which may be the case, e.g., if the distance $q_1^0 - \underline{q}$ is sufficiently large), then both sellers' payoffs decrease in the lower quality and thus full degradation is optimal as above. When $\overline{\theta}$ is close enough to $2\underline{\theta}$, however, the quality follower's expected payoff from quality testing without degradation is higher than from full degradation without quality testing. Note also that condition (7) requires that there is a benefit of quality testing (i.e., learning leads to reversal of quality ranks with positive probability), and that this benefit exceeds the cost of quality testing for at least one seller given the default quality levels. This will ensure that at least one seller will submit their product to quality testing and neither player chooses to degrade.

Relative to the canonical model, quality testing leads to lower dispersion in ex-ante quality levels. However, the sellers' market power is ultimately determined by the dispersion in *posterior* quality levels (i.e., quality levels after the signal is drawn). Whereas without quality testing the distance in quality levels is at least $q_2^0 - \underline{q}$, with quality testing the expected distance is $2\Delta^\sigma(q_1^0, q_2^0)$ and depends on the precision of the signals and the amount of quality testing in equilibrium (i.e., whether both sellers test their goods or only one).

Turning to social welfare, we compare an outcome with degradation (by the quality follower, which is the Nash equilibrium of the degradation game that maximizes aggregate payoff) but no quality testing to one with quality testing, but no degradation. Letting $|\sigma|$ denote the cardinality of the signal configuration (i.e., whether learning is about one or both sellers) and using (5) and (6), the outcome with quality testing is socially preferable if:

$$q_1^0\left(\frac{2\overline{\theta}^2+(\overline{\theta}-\underline{\theta})^2}{6}\right)-(q_1^0-\underline{q})\frac{(\overline{\theta}-2\underline{\theta})^2}{18} < q_1^0\left(\frac{2\overline{\theta}^2+(\overline{\theta}-\underline{\theta})^2}{6}\right)-(q_1^0-q_2^0)\frac{(\overline{\theta}-2\underline{\theta})^2}{18} + \left(\frac{2\overline{\theta}^2+(\overline{\theta}-\underline{\theta})^2}{6}-\frac{(\overline{\theta}-2\underline{\theta})^2}{9}\right)\Delta^\sigma(q_1^0,q_2^0)-|\sigma|k.$$

That is, if

$$|\sigma|k < (q_2^0-\underline{q})\frac{(\overline{\theta}-2\underline{\theta})^2}{18}+\left(\frac{2\overline{\theta}^2+(\overline{\theta}-\underline{\theta})^2}{6}-\frac{(\overline{\theta}-2\underline{\theta})^2}{9}\right)\Delta^\sigma(q_1^0,q_2^0). \quad (8)$$

The second term on the RHS is in fact the social benefit of learning derived above without degradation, which will fall short of the private benefit given in condition (7). The first term on the RHS reflects the social benefit of avoiding quality degradation, which may indeed even render the social benefit of quality testing greater than the private one. One immediate consequence is that, in the presence of quality degradation, the maximum cost such that quality testing by one seller is socially desirable is higher than in the case without quality degradation.

A slightly different question is whether a social planner should allow for quality testing, and then rely on the Nash equilibrium outcome to bring about the socially desirable learning. For the case of obtaining one signal for one of the sellers this will happen when conditions (7) and (8) coincide. This is a knife-edge condition depending on the signal technology and initial parameters q_i^0 as well as the minimum quality $\underline{q}$. Allowing for quality testing and observing whether this action is chosen in equilibrium may, however, offer valuable information on the social optimum.

The following proposition sums up these observations.

Proposition 4. *For intermediate values of k, quality testing by a single seller decreases aggregate welfare when quality is exogenous, but increases welfare when quality is endogenous. For low (high) values of k quality testing increases* *(decreases) aggregate welfare both with endogenous and exogenous quality levels.*

If quality is endogenous and $\overline{\theta}$ is sufficiently close to $2\underline{\theta}$, then observing no quality testing in a Nash equilibrium implies that quality testing is socially not desirable either.

Proof. The first part follows directly from condition (8). For the second part note that if $\overline{\theta}$ is sufficiently close to $2\underline{\theta}$, then condition (8) implies condition (7). Hence, if there is no quality testing in equilibrium of the quality degradation and testing game condition (7) fails, and thus also condition (8) will fail. □

3 Conclusions

The results derived so far rely on strong assumptions. In a more general setup, consumers may be distributed non uniformly, Assumption 1 may be violated, consumers may demand a good only if its quality is sufficiently high, or they may be risk averse and be willing to pay more for less uncertainty on quality. The general mechanism driving our results is, however, quite robust.

Suppose that all consumers have the same preference ranking over the goods (i.e., assume a model of vertical competition). This ranking could equal the quality level, or a function of it, for example, if consumers are risk averse. For simplicity, let the marginal cost of production be constant, equal across sellers, and independent of the sellers' quality level. If both sellers are offering the same product quality (i.e., $q_1 = q_2$), in equilibrium they both charge their marginal cost and make zero profits. Hence, sellers will have a strong incentive to increase the distance between their quality levels.

Therefore, our results extend to a large class of models of vertical competition: when quality levels are close, each firm benefits from investment in information as long as there is some ex-post quality distribution such that at least one of them makes positive profits. The reason is that information increases the expected distance in quality, regardless of the specific model parameters. Hence, sellers will acquire information if the investment cost is not too high, and have an incentive to collude in information acquisition. As for the social welfare, the basic trade off discussed above is still present. If the two quality levels are sufficiently close, quality testing may reveal what product is most valuable to the consumers, but will also increase sellers' market power.

Finally, sellers can increase the distance between their quality levels by degrading quality or by acquiring information. Acquiring information allows for a potential reversal in the quality ranking, and may be preferred over quality degradation when the follower's profits are sufficiently low. Because the chance to overtake the market leader increases in the follower's quality level, anticipating later quality testing sellers have less incentive to degrade their quality and information acquisition may be socially beneficial. This basic intuition is, of course, considerably more general than our model.

Our setup also has abstracted from asymmetric information on the side of the sellers. A straightforward extension could entail each seller having private information about their quality, i.e., their expectation of signal generated by quality testing equals their own prior, but differs from the public prior belief about quality. If public prior beliefs about qualities are sufficiently close, our results remain unchanged: quality testing will increase expected distance in qualities, more so than in the common information change, and thus the incentives for overinvestment and collusion in quality testing will remain.

Finally, the intuition developed in this paper does not depend on "business stealing effects" and extends to assignment problems. In the appendix, we assume that there are two types of consumers, and that the market share of each seller is kept constant — which effectively transforms our vertical competition problem into a matching problem. We show that our results go through unchanged. The reason is that, also in this case, sellers' collective market power increases with the distance between competitors. The only difference with the model discussed in the body of the text is that, in an assignment problem, increasing distance between competitors always increases the leader's profits but may decrease the follower's profits. However, as long as there is some probability that by acquiring information the follower overtakes the leader, quality testing by either the follower or the leader is valuable both to the follower and to the leader.

References

Anderson, S. P. and R. Renault (2000). Consumer information and firm pricing: negative externalities from improved information. *International economic review 41*(3), 721–742.

Anderson, S. P. and R. Renault (2009). Comparative advertising: disclosing horizontal match information. *The RAND Journal of Economics 40*(3), 558–581.

Bergemann, D. and J. Välimäki (2000). Experimentation in markets. *The Review of Economic Studies 67*(2), 213–234.

Bouton, L. and G. Kirchsteiger (2015). Good rankings are bad-why reliable rankings can hurt consumers.

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

Johnson, J. P. and D. P. Myatt (2006). On the simple economics of advertising, marketing, and product design. *The American Economic Review*, 756–784.

Levin, D., J. Peck, and L. Ye (2009). Quality disclosure and competition. *The Journal of Industrial Economics 57*(1), 167–196.

Lewis, T. R. and D. E. M. Sappington (1994, May). Supplying Information to Facilitate Price Discrimination. *International Economic Review 35*(2), 309–27.

Tirole, J. (1988). *The theory of industrial organization.* MIT press.

A Appendix: Assignment problem

In this section we assume that each seller's supply of goods is limited to 1/2. Hence, we preclude any change on the size of the market served by sellers, and we isolate a single channel through which information affects sellers profits: the sorting of consumers. We also assume that a measure 1/2 of buyers have low valuation $\theta = \underline{\theta} > 0$, and measure 1/2 have high valuation $\theta = \overline{\theta} > \underline{\theta}$. We impose no further restrictions on $\overline{\theta}$ and $\underline{\theta}$, which implies that Assumption 1 may be violated. The structure of the problem is otherwise identical to the one presented in the body of the text.

Since quality and taste are complements in the buyers' payoffs, high valuation buyers will always be able to outbid low valuation buyers for the higher quality product in the market. Because capacity is fixed, market prices will make low valuation buyers indifferent between the low-quality good and no consumption, and high valuation buyers indifferent between the two products. This observation implies the following statement:

Lemma 1 (Market Equilibrium)**.** *For given q_1^σ and q_2^σ, equilibrium profits are:*

$$\pi_i(q_i^\sigma, q_{-i}^\sigma) = \begin{cases} \frac{1}{2}\left(q_i^\sigma\overline{\theta} - q_{-i}^\sigma(\overline{\theta} - \underline{\theta})\right) & \textit{if } q_i^\sigma > q_{-i}^\sigma \\ \frac{1}{2}\ \underline{\theta} q_i^\sigma & \textit{otherwise}\ . \end{cases}$$

Proof. Omitted. □

Public learning will change the distribution of qualities in the market. This may alter both the quality ranking and buyers' equilibrium outside option (i.e., the payoff from consuming the next best alternative). For instance, if learning reveals that $q_i = 0$, seller $-i$ will serve high instead of low valuation buyers, and high valuation buyers' outside option (purchasing from i) decreases to 0. Hence, also here information acquisition affects the distance between the qualities in the market, and the market power of high-quality sellers. The following statement formalizes this intuition:

Lemma 2. *Suppose seller* 2 *is not testing its product. seller* 1*'s net benefit from quality testing is*

$$\left(\frac{\overline{\theta}-\underline{\theta}}{2}\right)\Delta^{(\sigma_1,\emptyset)}(q_1,q_2). \tag{9}$$

Suppose seller 1 *is not testing its product, and seller* 2 *switches from no quality testing to quality testing. The change in seller* 1*'s profits is*

$$\left(\frac{\overline{\theta}-\underline{\theta}}{2}\right)\Delta^{(\emptyset,\sigma_2)}(q_1,q_2). \tag{10}$$

Suppose seller 2 *is testing its product. seller* 1*'s net benefit from quality testing is*

$$\left(\frac{\overline{\theta}-\underline{\theta}}{2}\right)\left(\Delta^{(\sigma_1,\sigma_2)}(q_1,q_2)-\Delta^{(\emptyset,\sigma_2)}(q_1,q_2)\right). \tag{11}$$

Proof. Omitted because identical to the proof of Lemma 1 - modulo the parameters of the profit function. □

Hence, also here, by acquiring information seller i imposes a positive externality on seller $-i$. Intuitively, when i acquires information, i's product may turn our better than expected or worse than expected. In the first case seller $-i$ will be hurt, in the second case seller $-i$ will benefit. Crucially, if $-i$ benefits, the size of the gain depends linearly on the ex-post quality of seller i's product. If $-i$ is hurt, its profits will be $\underline{\theta}q_{-i}$ which is independent on the ex post quality of seller i's product. In other words, when i acquires information, $-i$ faces an upside that depends on the ex-post realization of q_i^σ, and a downside that is bounded. Because of this positive externality, also here the following corollary holds.

Corollary 3 (Collusion)**.** *If sellers can collude and jointly decide on their level of quality testing, they will test for quality (weakly) more than in equilibrium.*

Finally, we can compare the sellers' private benefits from investing in learning with the social benefits. Social welfare is equal to $\max\{q_i^\sigma, q_{-i}^\sigma\}\frac{\overline{\theta}}{2} + \min\{q_i^\sigma, q_{-i}^\sigma\}\frac{\underline{\theta}}{2}$, and the change in social welfare when seller 1 tests for quality but seller 2 does not is:

$$\Delta^{(\sigma_1,\emptyset)}(q_1,q_2)\left(\frac{\overline{\theta}-\underline{\theta}}{2}\right),$$

which is equal to the private benefit of quality testing.[8] The change in social welfare when seller 1 starts testing for quality when seller 2 is already testing for quality is:

$$\left(\frac{\overline{\theta}-\underline{\theta}}{2}\right)\left(\Delta^{(\sigma_1,\sigma_2)}(q_1,q_2)-\Delta^{(\emptyset,\sigma_2)}(q_1,q_2)\right),$$

which is also equal to the private value of quality testing. The proposition follows immediately.

Proposition 5 (Investment Equilibrium)**.** *For any investment cost k the Nash equilibrium level of information acquisition maximizes aggregate surplus. seller's collusion in quality testing is welfare decreasing.*

Proof. In the text. □

Contrary to the case presented in the body of the text, here the individual incentive to acquire information is efficient. However, similarly to the the case presented in the body of the text, also here collusion leads more information acquisition than socially optimal, and is therefore welfare decreasing. This also implies that quality testing has either no effect on consumers' welfare (when it is efficient) or decreases consumers' welfare (when it is above its efficient level due to collusion).

[8]The expression follows from noting that social welfare can be written as $\max\{q_i^\sigma, q_{-i}^\sigma\}\frac{\overline{\theta}+\underline{\theta}}{2} - |q_i - q_{-i}|\underline{\theta}$, and then using the expressions for $E\left[\max\{q_i^\sigma, q_{-i}^\sigma\}|\sigma\right] - \max\{q_i, q_{-i}\}$ and $E\left[|q_i^\sigma - q_{-i}^\sigma||\sigma\right] - |q_i - q_{-i}|$ derived in the body of the paper.

B Appendix: Mathematical derivations

Proof of Proposition 1

Computing the difference in sellers' payoffs if neither, one, or both of them invest in quality testing is relatively straightforward. If there is no quality testing the expected market profits for seller 1 can be rewritten as

$$\pi_1 = \left(\frac{(2\overline{\theta} - \underline{\theta})^2}{9}\right)(q_1 - q_2),$$

and similarly, for seller 2:

$$\pi_2 = \left(\frac{(\overline{\theta} - 2\underline{\theta})^2}{9}\right)(q_2 - q_1).$$

Suppose now that some quality testing has occurred, resulting in signals $\sigma \neq (\emptyset, \emptyset)$. The expected market payoff of seller $i \in \{1; 2\}$ depends now on whether i is the quality leader or follower after taking into account the realization $\hat{\sigma}$ of the quality signal σ. That is, seller i's expected market payoff conditional on quality testing is now the leader's payoff π_1, if i turns out to be the quality leader, and the follower's payoff π_2, if i turns out to be the follower:

$$\pi_i^\sigma = \text{pr}\{q_i^{\hat{\sigma}} \geq q_{-i}^{\hat{\sigma}}\}\left(\frac{(2\overline{\theta} - \underline{\theta})^2}{9}\right)E\left[q_i^{\hat{\sigma}} - q_{-i}^{\hat{\sigma}}\,|q_i^{\hat{\sigma}} \geq q_{-i}^{\hat{\sigma}}\right] + \text{pr}\{q_{-i}^{\hat{\sigma}} \geq q_i^{\hat{\sigma}}\}\left(\frac{(\overline{\theta} - 2\underline{\theta})^2}{9}\right)E\left[q_{-i}^{\hat{\sigma}} - q_i^{\hat{\sigma}}\,|q_{-i}^{\hat{\sigma}} \geq q_i^{\hat{\sigma}}\right].$$

Note that, when no quality testing occurs, market payoffs for the quality leader can be rewritten as

$$\pi_1 = \left(\frac{(2\overline{\theta} - \underline{\theta})^2}{9}\right)(q_1 - q_2) = \left(\frac{(2\overline{\theta} - \underline{\theta})^2}{9}\right)E[q_1^{\hat{\sigma}} - q_2^{\hat{\sigma}}],$$

for any signal configuration $\sigma \neq (\emptyset, \emptyset)$. The identity follows from the law of iterated expectations, which implies $E[q_i^{\hat{\sigma}}] = q_i$, because $E[q_i^{\hat{\sigma}}]$ is the conditional expectation of q_i given a specific realization of the signal σ. That is, before a signal is drawn, the expected value of q_i conditional on the realization of the signal(s) is equal to the initial, unconditional expectation q_i. Decomposing the expected quality difference $E[q_1^{\hat{\sigma}} - q_2^{\hat{\sigma}}]$ into the two events $q_1^{\hat{\sigma}} \geq q_2^{\hat{\sigma}}$ and $q_1^{\hat{\sigma}} \leq q_2^{\hat{\sigma}}$ yields:

$$\pi_1 = \left(\frac{(2\overline{\theta} - \underline{\theta})^2}{9}\right) \left\{ \text{pr}\{q_1^{\hat{\sigma}} \geq q_2^{\hat{\sigma}}\} E[q_1^{\hat{\sigma}} - q_2^{\hat{\sigma}} | q_1^{\hat{\sigma}} \geq q_2^{\hat{\sigma}}] + \text{pr}\{q_2^{\hat{\sigma}} \geq q_1^{\hat{\sigma}}\} E[q_1^{\hat{\sigma}} - q_2^{\hat{\sigma}} | q_2^{\hat{\sigma}} \geq q_1^{\hat{\sigma}}] \right\}.$$

Identical steps lead to a similar expression for the quality follower's payoff:

$$\pi_2 = \left(\frac{(\overline{\theta} - 2\underline{\theta})^2}{9}\right) \left\{ \text{pr}\{q_1^{\hat{\sigma}} \geq q_2^{\hat{\sigma}}\} E[q_1^{\hat{\sigma}} - q_2^{\hat{\sigma}} | q_1^{\hat{\sigma}} \geq q_2^{\hat{\sigma}}] + \text{pr}\{q_2^{\hat{\sigma}} \geq q_1^{\hat{\sigma}}\} E[q_1^{\hat{\sigma}} - q_2^{\hat{\sigma}} | q_2^{\hat{\sigma}} \geq q_1^{\hat{\sigma}}] \right\}.$$

Finally, simple algebra reveals that for $i = 1, 2$ and $-i \neq i$:

$$\pi_i^{\sigma} - \pi_i = \left(\frac{(\overline{\theta} - 2\underline{\theta})^2 + (2\overline{\theta} - \underline{\theta})^2}{9}\right) \text{pr}\{q_2^{\hat{\sigma}} \geq q_1^{\hat{\sigma}}\} E\left[q_2^{\hat{\sigma}} - q_1^{\hat{\sigma}} \; | q_2^{\hat{\sigma}} \geq q_1^{\hat{\sigma}}\right].$$

Using these expressions we can compare configurations $\sigma = (\sigma_1, \sigma_2)$ and $\sigma' = (\sigma_1, \emptyset), (\emptyset, \sigma_2)$:

$$\begin{aligned} \pi_i^{\sigma} - \pi_i^{\sigma'} = & \left(\frac{(\overline{\theta} - 2\underline{\theta})^2 + (2\overline{\theta} - \underline{\theta})^2}{9}\right) \\ & \times \left(\text{pr}\{q_2^{\hat{\sigma}} \geq q_1^{\hat{\sigma}}\} E\left[q_2^{\hat{\sigma}} - q_1^{\hat{\sigma}} \; | q_2^{\hat{\sigma}} \geq q_1^{\hat{\sigma}}\right] - \text{pr}\{q_2^{\hat{\sigma}'} \geq q_1^{\hat{\sigma}'}\} E\left[q_2^{\hat{\sigma}'} - q_1^{\hat{\sigma}'} \; | q_2^{\hat{\sigma}'} \geq q_1^{\hat{\sigma}'}\right] \right) \end{aligned}$$

The expressions (2), (3) and (4) above follow immediately and imply the proposition.

Proof of Proposition 3

To solve the game we proceed backwards in time as usual. Given quality choices q_1 and q_2 the benefits of quality testing outweigh the cost if $\pi_i^{(\sigma_i, \emptyset)} - k >$

$\pi_i^{(\emptyset,\emptyset)}$, that is:

$$\left(\frac{(\overline{\theta}-2\underline{\theta})^2+(2\overline{\theta}-\underline{\theta})^2}{9}\right)\Delta^{(\sigma_i,\emptyset)}(q_1,q_2)>k \text{ for } i=1,2.$$

Denote the quality distance by $d=|q_1-q_2|$. Since $\Delta^{(\sigma_i,\emptyset)}(q_1,q_2)$ continuously decrease in d there is d^* such that

$$\left(\frac{(\overline{\theta}-2\underline{\theta})^2+(2\overline{\theta}-\underline{\theta})^2}{9}\right)\Delta^{(\sigma_i,\emptyset)}(q_1,q_1-d^*)=k \text{ for } i=1,2.$$

Hence, there will be quality testing by at least seller in any pure strategy equilibrium for all (q_1,q_2) with $d<d^*$, and no quality testing for all (q_1,q_2) with $d>d^*$. Condition (7) in the proposition implies that $q_1^0-q_2^0<d^*$. Note that $\Delta^{(\sigma_i,\emptyset)}(q_1,q_2)>0$ implies that the signal σ induces a reversal of the quality ranking with positive probability. Anticipating quality testing in the second stage expected payoffs given q_1 and q_2 and the signals σ are for $i=1,2$:

$$\pi_i^\sigma=\text{pr}\{q_i^{\hat{\sigma}}\geq q_{-i}^{\hat{\sigma}}\}\left(\frac{(2\overline{\theta}-\underline{\theta})^2}{9}\right)E\left[q_i^{\hat{\sigma}}-q_{-i}^{\hat{\sigma}}\,|q_i^{\hat{\sigma}}\geq q_{-i}^{\hat{\sigma}}\right]+\text{pr}\{q_{-i}^{\hat{\sigma}}\geq q_i^{\hat{\sigma}}\}\left(\frac{(\overline{\theta}-2\underline{\theta})^2}{9}\right)E\left[q_{-i}^{\hat{\sigma}}-q_i^{\hat{\sigma}}\,|q_{-i}^{\hat{\sigma}}\geq q_i^{\hat{\sigma}}\right].$$

First, note that $\pi_i^\sigma>0$ for both players. Condition (7) implies that the signal σ induces a reversal of the quality ranking with positive probability, i.e. $\text{pr}\{q_i^{\hat{\sigma}}\geq q_{-i}^{\hat{\sigma}}\}\in(0,1)$. Then a decrease in q_i, which is equal to $E[q_i^{\hat{\sigma}}]$, will decrease the first term and increase the second. Therefore, if $\overline{\theta}$ is sufficiently close to $2\underline{\theta}$ the first effect outweighs the second and each player's payoff strictly increases in q_i. Hence, both sellers' payoffs increase in own quality if there is quality testing in the second stage.

If there is no quality testing anticipated, both sellers' payoffs increase in the higher quality and decrease in the lower quality. Hence, given the other player's action q_{-i} for any degradation choice in the first stage that results in $|q_1-q_2|>d^*$ it is optimal to choose $q_i=\underline{q}$, and for any degradation choice in the first stage that results in $|q_1-q_2|<d^*$ it is optimal to choose $q_i=q_i^0$. Therefore $q_i=q_i^0$ for $i=1,2$ and at least one seller investing in quality testing

is a subgame perfect Nash equilibrium of the quality degradation and testing game. To see that it is indeed the only one when $\overline{\theta}$ is sufficiently close to $2\underline{\theta}$, suppose that $|q_1 - q_2| > d^*$ in equilibrium. Then the lower quality choice must be $q_i = \underline{q}$, otherwise there would be a profitable deviation degrading some more, and the other seller will not degrade. This pins down equilibrium payoffs $(q^0_{-i} - \underline{q})(\overline{\theta} - 2\underline{\theta})/9$ for the seller who has degraded. By not degrading and choosing $q_i = q^0_o$ this seller would get at least (if $q^0_i < q^0_{-i}$):

$$\pi^\sigma_2 - k = \mathrm{pr}\{q^{\hat{\sigma}}_2 \geq q^{\hat{\sigma}}_1\}\left(\frac{(2\overline{\theta} - \underline{\theta})^2}{9}\right)E\left[q^{\hat{\sigma}}_2 - q^{\hat{\sigma}}_1\,|q^{\hat{\sigma}}_2 \geq q^{\hat{\sigma}}_1\right] + \mathrm{pr}\{q^{\hat{\sigma}}_1 \geq q^{\hat{\sigma}}_2\}\left(\frac{(\overline{\theta} - 2\underline{\theta})^2}{9}\right)E\left[q^{\hat{\sigma}}_1 - q^{\hat{\sigma}}_2\,|q^{\hat{\sigma}}_1 \geq q^{\hat{\sigma}}_2\right] - k$$

Under Condition (7) there exists $\overline{\theta}$ sufficiently close to $2\underline{\theta}$ such that $\pi^\sigma_2 - k$ is strictly positive and greater than $(q^0_1 - q^0_2)(\overline{\theta} - 2\underline{\theta})/9$.