

Quality-Ensuring Profits

31 March 2008

Eric Rasmusen

Abstract

In the reputation model of Klein & Leffler (1981) firms refrain from cutting quality or price because if they did they would forfeit future profits. Something similar can happen even in a static setting. First, if there exist some discerning consumers who can observe quality, firms wish to retain their purchases. Second, if all consumers can sometimes but not always spot flaws, firms do not want to lose the business of those who would spot them on a given visit. Third, if the law provides a penalty for fraud, but not one so high as to make fraud unprofitable, firms may prefer selling high quality at high prices to low quality at high prices plus some chance of punishment.

Dan R. and Catherine M. Dalton Professor, Department of Business Economics and Public Policy, Kelley School of Business, Indiana University. BU 438, 1309 E. 10th Street, Bloomington, Indiana, 47405-1701. Office: (812) 855-9219. Fax:812-855-3354. erasmuse@indiana.edu, <http://www.rasmusen.org>. This paper: <http://www.rasmusen.org/papers/quality-rasmusen.pdf>.

Keywords: Reputation, product quality, moral hazard, quality-guaranteeing price

I thank seminar participants at workshops at Indiana, Oxford, and Warwick Universities for their comments, and Meg Meyer in particular.

1. Introduction

Often buyers cannot rely on contracts to guarantee the quality of the goods or services they purchase. The transaction cost is too high or courts are too expensive to enforce any contracts that might be written. Yet some sellers do dependably provide high quality despite the extra profit they could generate by cheating on any one purchase. Buyers trust those sellers, and are willing to pay them a high price for their high quality. Somehow, competition does not erode the resulting profits.

Benjamin Klein & Keith Leffler (1981) give one explanation. In equilibrium a seller with a good reputation can sell at a price above marginal cost. He could produce low quality and earn a big one-time profit, but buyers would react by never buying from him again. If he produces high quality, his immediate profits will be less, but they will nonetheless be positive and he will continue to earn them indefinitely. For this to happen, the equilibrium price must be high enough to make him prefer high quality and stable profits to low quality and immediate profit. The positive profits will persist in equilibrium because buyer would correctly deduce that any seller charging a lower price would be irresistably tempted to produce low quality.

The Klein-Leffler model is based on one equilibrium of an infinitely repeated game, and repetition with a long horizon is essential to it. In a one-period version of the model, the long-term gain from stable profits would vanish and only the one-shot gain from cheating on quality would remain. The present paper will show that the idea of quality-maintaining profits is not limited to multi-period games, however. In the Klein-Leffler model, the seller's deterrent to cheating is the loss of future profits, but quality-maintaining profits can arise in a one-period model if the seller's loss from cheating is some other kind of profits. I will show this by laying out three ways that profits can arise even in a one-period setting.

First, however, it will be useful to distinguish the present setting from different ways of thinking about product quality and reputation. The main alternative is to assume that the problem is not to gain an incentive to produce high quality, but to convince other players of a quality that is already high and cannot be varied. Rogerson (1983) was an early model of this, in which high-quality firms's sales would increase over the long run. The insight of Nelson (1974) is that if a firm's quality is high, then it has more incentive

to get consumers to try its product than if its quality is low. If it spends money on advertising, that is a profitable strategy if its quality is high, but not if it is low, because the profit margin in equilibrium is not great enough to justify the advertising. Kihlstrom & Riordan (1984) and Milgrom & Roberts (1986) model the idea formally as a signalling model. Bagwell & Riordan (1991) take the same kind of signalling model of incomplete information, but make price the signal rather than advertising. Prices start high in early periods of a market, signalling high quality, but fall as information gradually becomes public. Roberson (1983) and Linnemer (2002) combine the two ideas of advertising and prices as signals into one model, in which both signals are used by firms of intermediate quality. Farrell (1986) (see also his 1980 dissertation) also takes the approach of quality that is unchanging over time, but he allows entrants to choose quality initially, and his emphasis is on how it becomes increasingly difficult to enter successfully, as the consumer surplus earned from incumbent firms makes switching increasingly unattractive.

It has also proved interesting to combine adverse selection and moral hazard, asking what happens when firms differ in type, but some firms are capable of producing either low or high quality. Diamond (1989) and Horner (2002) are two examples of this approach that yield the outcome that over time a firm's reputation and incentives for high quality can improve as it consistently shows good results to its trading partners.

Other articles move attention away from consumers and onto firms, limiting the number of firms so they interact strategically (e.g. Hertzendorf & Overgaard [2001], Fluet & Garella [2002]) or having firms compete horizontally in quality as well as vertically (Daughety & Reinganum [2006]). Kirmani & Rao(2000) survey the literature.

The product-quality literature discussed so far is about adverse selection, not moral hazard. Its starting point is quality that once fixed cannot be changed. Less attention has been given to models in which product quality is chosen anew each period. That situation, the one in Klein & Leffler (1982) is subject to the Chainstore Paradox of repeated games, that if the interaction is repeated a finite number of times, then the only subgame perfect equilibrium will have low quality in each period. Kreps, Milgrom, Roberts & Wilson (1982) show how the Chainstore Paradox can be avoided by adding a small amount of incomplete information to the game—which in the case of

product quality would be a small probability that the firm can only produce high quality. With a long enough time horizon, this can result in a high-quality equilibrium. For one example of a model constructed on this basis, see Maksimovic & Titman (1991), which shows how a firm's capital structure interacts with its product quality choice.

A number of papers have looked at moral hazard in various contexts. Some make the Informed and Uninformed Consumers assumption that I will use in this paper. In the search model of Chan and Leland (1982), consumers differ in their costs of acquiring information about price and quality. Firms can choose their level of quality. When price is known but quality information is costly to acquire, the equilibrium has a single price but two levels of quality, one for the informed and one for the uninformed consumers. That is in a search context, but several other modellers have used the informed vs. uninformed consumer assumption in a non-search context: Farrell (1980), Cooper & Ross (1984) and Tirole's *Theory of Industrial Organization* (1988, p. 107) . All three assume that some consumers are informed of quality and some are not, that there is only one period of production, and that the firm has a choice between low and high quality. In chapter 2 of Farrell's 1980 doctoral dissertation, the emphasis is on the externality that the presence of informed consumers exert on uninformed consumers. In Cooper & Ross's model, firms have U-shaped cost curves and there is free entry. If a competitive equilibrium exists, some firms sell low quality and some high, with the amount of high quality increasing with the proportion of informed consumers and the steepness of the average cost curve at subefficient scales. Tirole models a monopoly with constant returns to show that a monopolist may choose high quality even if not all consumers can observe his choice.

Wolinsky (1983) is based on something similar to the Consumer Error assumption I will use below. is a search model in which consumers visit firms and obtain noisy signals of quality. Firms are identical in costs, and choose their quality levels and prices. Consumers differ in their willingness to pay for quality, so that in equilibrium there can be a variety of price- quality pairs offered, quality rising with price. An increase in the quality of information reduces prices.

In the present article, I will go back to fundamentals, and look at moral hazard in a competitive industry where firms use identical technologies with

constant returns to scale and have a choice between high and low quality. Most of the literature has focussed on multi-period models, relying on future sales to give the firm an incentive for present high quality. I will start by assuming multiple periods, but I will follow that by letting the discount rate become infinite. The paper's message will be that the idea that a firm chooses high quality to avoid losing a fraction of its sales applies even if future sales are unimportant, and that the key result that price will exceed marginal cost even in a competitive market obtains regardless of the source of the lost sales. This can be modelled much more simply than in earlier product quality models. The ideas that differences in consumer information or errors can affect product quality are well known, but I will show how similar they are to the way that Klein-Leffler reputation can maintain product quality. In addition, the idea that even weak laws against fraud can maintain high quality in the same manner seems to be completely new.

2. The Model

We will use a formalization of Klein & Leffler (1982) similar to Rasmusen (1989) but with an exogenous measure of firms. Firms and consumers are atomistic, lying on the $[0,1]$ interval. First, the firms simultaneously choose qualities and prices. Second each consumer observes the prices, but not the qualities, and decides which one of the firms, if any, to visit. Third, after visiting a firm, each consumer decides whether to buy one unit or not based on what he observes about quality as described in Assumptions A1 and A2 below, and the government may intervene as described in Assumption A3. If the consumer does buy, the firm produces the unit and the consumer pays. Finally, the firm pays the production cost, each consumer consumes what he has bought, and all consumers discover the quality of all firms. This process is repeated an infinite number of times.

A firm pays $c = c_h$ per unit to produce high quality and $c = c_l$ to produce low quality, with $0 < c_l < c_h$.

Consumers value high quality at $v = v_h > c_h$ and low quality at $v = v_l = c_l$, so high quality is efficient.

Thus, the per-period payoff functions are, if a firm sells N units

$$\text{Payoff}(firm) = (p - c)N \tag{1}$$

and

$$Payoff f(consumer) = \begin{cases} v - p & \text{if he buys} \\ 0 & \text{if he does not buy} \end{cases} \quad (2)$$

We will assume that all payments are made at the start of periods and that the discount rate is $r > 0$.

The model departs from the standard reputation model by adding the following three assumptions:

(A1) (Consumer Error) If the quality is low, then with probability $0 \leq \alpha < 1$ the consumer observes that fact— he “spots the flaw,” but with probability $(1 - \alpha)$ he receives no information. If the quality is high, he receives no information.

(A2) (Informed and Uninformed Consumers) If a consumer is one of the fraction $0 \leq \beta < 1$ of consumers that are “discerning” he observes quality perfectly once he visit the seller.

(A3) (Weak Laws) If a firm tries to sell low quality as high, then with probability $0 \leq \gamma < 1$ independent of α and β the government interrupts the transaction and fines the seller amount F .

To simplify the strategy descriptions, we will assume that firms never follow the negative-profit strategy of $p < c_l$ (otherwise we need to encumber the consumer strategies with caveats about how they might visit low-quality firms because of give-away prices). By “equilibrium” I will mean a perfect bayesian equilibrium, where the need for out-of-equilibrium beliefs arises not because the game has incomplete information but because a player’s deviation move might change beliefs about his unobserved earlier or later moves.

Unless the expected government punishment γF is high enough for a fraudulent firm, there is no equilibrium in which quality is high and the price equals marginal cost. In such an equilibrium, $p = c_h$ and firms would earn zero profits. A firm could deviate to low quality and make positive profits, because its deviation payoff would be:

$$\pi_{firm}(low\ quality) = \theta(p - c_l) - \gamma F, \quad (3)$$

which is positive if γF is low enough, since $p = c_h$ and $\theta > 0$. We will assume that

$$\gamma F < \theta(c_h - c_l), \quad (4)$$

because otherwise we attain the first-best because the expected punishment alone makes deviation to low quality unprofitable. Given that inequality (4) is satisfied (and, as we will see later, v_h is not too low), our infinitely repeated game has an infinite number of equilibria. We will focus on two of them, which we will call the pessimistic equilibrium and the optimistic equilibrium.

Note that if we had assumed that the firm paid the production cost *before* the government detected and cancelled the fraudulent sale, and that the firm could not resell the product, then the payoff from low quality would be:

$$\pi_{firm}(low\ quality,) = \theta(p - c_l) - \gamma F - \gamma c_l. \quad (5)$$

In effect, the lost production cost would be part of the punishment, and it would allow γF to take a lower value than the bound in inequality (4) and still deter low quality.

The Pessimistic Equilibrium

The Firm: The firm chooses its quality to be low and its price to be $p = c_l$.

The Consumer: A consumer visits any of the firms with the lowest price. He buys if $p \leq v$ and he observes the quality.

He buys if $p \leq c_l$ if he does not observe the quality.

Consumer out-of-equilibrium beliefs: If $p > c_l$, quality is low.

The pessimistic equilibrium exists regardless of the parameter values. It is inefficient because both producer and consumer payoffs equal zero, whereas if the firm and consumer had traded high quality there would have been gains from trade.

The Optimistic Equilibrium

The Firm: In equilibrium, the firm chooses high quality and the price $p = p^*$, where if we define the probability of successfully completing a sale as:

$$\theta \equiv (1 - \alpha)(1 - \beta)(1 - \gamma), \quad (6)$$

then the price is:

$$p^* \equiv c_h + \frac{r\theta(c_h - c_l)}{1 + r - r\theta} - \frac{r\gamma F}{1 + r - r\theta}. \quad (7)$$

If the firm has ever deviated to low quality or to $p < p^*$ in the past, it chooses low quality and $p = c_l$.

The Consumer: The consumer never visits a firm that has produced low quality or charged $p \neq p^*$ in the past. Of the remaining firms, he visits the firm with the lowest price such that $p \geq p^*$, or no firm if all prices are less than p^* .

If he observes the quality, he buys if $p \leq v$.

If he does not observe the quality, he buys if $p \in [p^*, v_h]$.

Consumer out-of-equilibrium beliefs: If $p < p^*$, the consumer believes that the quality is low. If $p > p^*$, he believes that the quality is high.

The optimistic equilibrium requires explanation. Consumer behavior is easy to justify. In equilibrium, all firms produce high quality, so even an undiscerning consumer is safe in paying up to v_h for the product. So long as $p^* \leq v_h$ he will buy it; if $p^* > v_h$ then this equilibrium does not exist. Out of equilibrium, we are free to assign posterior beliefs for $Prob(High\ quality|deviation)$, so we assign the belief to be that a deviating firm will thenceforth produce low quality if it chose a low price. Given that belief, consumers will not visit a firm that has deviated, and such a firm has no incentive to produce high quality. This is a robust belief in the sense that the firm actually does profit more from high quality if its price is p^* or higher; it has no incentive to cheat on quality given that its price is that high.

Firms have a more delicate choice, which depends on the equilibrium price. Firms will not deviate to prices less than p^* because they would lose all their customers and earn zero payoffs. Firms could earn positive profits at $p = p^*$ either with the equilibrium high quality or by deviating to low quality, so we will have to consider both alternatives. Since quality choice will not affect how many customers a firm receives, we can look at their choices in terms of payoffs per customer.

If the price is p , then in equilibrium a firm will receive a profit of $(p - c_h)$ immediately and at the start of each future period. This is equivalent to an

undiscounted $(p - c_h)$ plus an immediate gift of a perpetuity of $(p - c_h)$ per period, so

$$\pi_{firm}(high\ quality) = (p - c_h) + \frac{p - c_h}{r}. \quad (8)$$

A firm's expected payoff per customer is a one-time payoff of $(p - c_l)$ if it gets away with fraud, which has probability θ , minus the expected government punishment, which is γF :

$$\pi_{firm}(low\ quality) = \theta(p - c_l) - \gamma F. \quad (9)$$

Thus, the firm is willing to produce high quality if

$$(p - c_h) + \frac{p - c_h}{r} = \theta(p - c_l) - \gamma F. \quad (10)$$

Solving equation (10) for p yields the value of p^* in equation (7) above.

The outcome in the optimistic equilibrium is efficient in this model, but that is only because of the assumption that each consumer buys either one or zero units. Otherwise, inefficiency would arise from the price exceeding marginal cost, because quality would be high, and consumers would buy, but less than the surplus-maximizing amount.

It is possible, as Klein and Leffer suggested, but not essential, to add some feature of rent-dissipating competition to the model, in which case the optimistic equilibrium becomes less efficient but still better than the pessimistic equilibrium. Thus, one might assume that any firm may enter the market and have an equal chance of participating in the optimistic equilibrium if it pays some fixed entry fee. This entry fee would eat up all the profits from the quality-ensuring price in equilibrium, but consumers would still earn surplus so long as $p^* < v_h$. Many of the models in the literature (e.g., Chan & Leland (1982), Wolinsky (1983), Rasmusen (1989)) use this kind of nonconvexity to dissipate rents. Rent dissipation is not essential to the model, however. It is equally valid to assume that some firms are endowed with good reputations without having to incur any entry fees, or that consumers are optimistic about the quality of incumbents and pessimistic about entrants. The essential feature of the model is that entrants cannot use price competition to secure market share, so the usual force driving even long-run profits to zero is weakened.

The Klein-Leffler Model and the Three New Assumptions

In the Klein-Leffler model, the parameters α , β , and γ all equal zero and the probability of successfully carrying out a sale is $\theta = 1$, so the quality-sustaining price in the optimistic equilibrium is, from equation (7):

$$p^* = c_h + \frac{r\theta(c_h - c_l)}{1 + r - r\theta} - \frac{r\gamma F}{1 + r - r\theta} = c_h + r(c_h - c_l). \quad (11)$$

The intuition is that if the market price is this high, the firm would make enough profit from future sales that it is unwilling to sacrifice those profits for the sake of a one-time gain from producing low-cost low-quality goods in the present period.

The Klein-Leffler model can generate high quality in a competitive market without any of the three new assumptions. Later I will show that each of the three new assumptions can independently generate high quality. First, though, let us think about all the assumptions in combination. The new assumptions to be added to the reputation model all add extra inducements for high quality. We will verify here that these new assumptions have the effects one would expect.

Proposition 1: The quality-ensuring price p^* falls in the probability of spotting a flaw α , the fraction of discerning customers β , and the government punishment's probability γ and size F . It rises in the discount rate r .

Proof: Equation (7) says that

$$p^* = c_h + \frac{r\theta(c_h - c_l)}{1 + r - r\theta} - \frac{r\gamma F}{1 + r - r\theta}. \quad (12)$$

Recalling that $\theta \equiv (1 - \alpha)(1 - \beta)(1 - \gamma)$, the derivative of p^* with respect to α is :

$$\begin{aligned} \frac{\partial p^*}{\partial \alpha} &= \frac{r \frac{\partial \theta}{\partial \alpha} (c_h - c_l)}{1 + r - r\theta} + \frac{r^2 \frac{\partial \theta}{\partial \alpha} \theta (c_h - c_l)}{(1 + r - r\theta)^2} - \frac{r^2 \frac{\partial \theta}{\partial \alpha} \gamma F}{(1 + r - r\theta)^2} \\ &= \frac{\partial \theta}{\partial \alpha} \frac{(1 + r - r\theta)r(c_h - c_l) + (r\theta)r(c_h - c_l) - r^2 \gamma F}{(1 + r - r\theta)^2} < 0, \\ &= \frac{\partial \theta}{\partial \alpha} \frac{(1 + r^2)(c_h - c_l) - r^2 \gamma F}{(1 + r - r\theta)^2} < 0. \end{aligned} \quad (13)$$

The last inequality in (13) follows because $\frac{\partial \theta}{\partial \alpha} < 0$ and we must have $\gamma F < \theta(c_h - c_l)$ as explained earlier or the punishment alone will induce firms to produce high quality.

Parameter β 's effect can be seen by substituting β for α in equation (13). As far as the equilibrium value of p^* is concerned, it does not matter whether some consumers always detect low quality or all consumers sometimes detect low quality.

Parameter γ follows the same pattern except that there is an additional term in its derivative because of its effect through the numerator of $-\frac{r\gamma F}{1+r-r\theta}$. Starting with equation (13) and adding the additional term yields:

$$\frac{\partial p^*}{\partial \gamma} = \frac{\partial \theta}{\partial \gamma} \frac{(1+r^2)(c_h - c_l) - r^2\gamma F}{(1+r-r\theta)^2} - \frac{rF}{1+r-r\theta} < 0. \quad (14)$$

Parameter F 's effect is simple:

$$\frac{\partial p^*}{\partial F} = -\frac{r}{1+r-r\theta} < 0. \quad (15)$$

Parameter r 's effect is:

$$\begin{aligned} \frac{\partial p^*}{\partial r} &= \frac{\theta(c_h - c_l) - \gamma F}{1+r-r\theta} - \frac{[1-\theta][r\theta(c_h - c_l) - r\gamma F]}{(1+r-r\theta)^2} \\ &= \frac{(1+r-r\theta)(\theta(c_h - c_l) - \gamma F) - [1-\theta][r\theta(c_h - c_l) - r\gamma F]}{(1+r-r\theta)^2} > 0. \end{aligned} \quad (16)$$

This demonstrates the comparative statics in Proposition 1. \square

Proposition 1 shows that a variety of forces help to maintain high quality. If some consumers can observe that quality is low and refuse to buy, that reduces the payoff to selling low quality, whether this be from consumers who invariably detect quality that is low or from consumers who merely have a probability of doing so. The derivatives with respect to α (and β) in the proof are complicated only by the interaction effect between the various incentives. In particular, the third term in the fraction in equation (13) has economic meaning. As the probability that the consumer is discerning rises, so does the probability that the transaction is interrupted before the government can punish it. Thus, the direct effect of the discerning consumer on making low

quality unprofitable is somewhat offset by the indirect effect of reducing the amount of government punishment that occurs, though the indirect effect cannot outweigh the direct effect.

Government detection also pushes down the profit from low quality, in two distinct ways. The effect of F is simple: fraud gets punished with some probability, even though not by enough to entirely deter low quality without the aid of a quality-inducing profits, and so it reduces the amount of quality-inducing profits needed: p^* falls. The probability of punishment, γ , also contributes to this direct effect. Second, the government prevents the fraudulent transaction from being completed.

3. The New Assumptions in a Static Setting

The model so far has been of an infinitely repeated game. What happens as the discount rate goes to infinity, converting it in effect to a static model?

We can apply L'Hospital's Rule, that $x \xrightarrow{Lim} \infty f(x)/g(x) = x \xrightarrow{Lim} \infty f'(x)/g'(x)$ when the ratio of the derivatives is finite. It is finite here if $1 - \theta > 0$. Applying the Rule to equation (7) yields:

$$r \xrightarrow{Lim} \infty p^* = c_h + \frac{\theta(c_h - c_l)}{1 - \theta} - \frac{\gamma F}{1 - \theta}. \quad (17)$$

Thus, a quality-ensuring price greater than marginal cost still exists even for the one-period model. Assumptions A1, A2, and A3, far from merely supplementing the reputation model, can each independently give rise to its feature of high quality at a price above marginal cost.

Proposition 2: Even in a one-period model, for big enough consumer reservation value v_h any one of assumptions A1, A2, and A3 yields an optimistic equilibrium in which quality is high and the equilibrium price is some p^* exceeding marginal cost.

Proof. We have seen that with infinite periods and positive r the optimistic equilibrium existed. Now the payoffs from high and low quality become, since future payments are worthless,

$$\pi_{firm}(high\ quality) = (p - c_h) + 0 \quad (18)$$

and, just as before,

$$\pi_{firm}(low\ quality) = \theta(p - c_l) - \gamma F. \quad (19)$$

Solving for p^* , these two payoffs are equal when

$$p^* = c_h + \frac{\theta(c_h - c_l)}{1 - \theta} - \frac{\gamma F}{1 - \theta}. \quad (20)$$

So long as $\theta < 1$, which it will be if $\alpha > 0$, $\beta > 0$, or $\gamma > 0$, p^* will be finite, so if v_h is large enough there exists some price that consumers are willing to pay for high quality at which firms prefer producing high quality to low. \square

Proposition 2 is somewhat surprising. It says that what is essential to a model of quality-ensuring price is not the prospect of future sales. Rather, it is that there be some loss of sales—possibly present instead of future—which results from a deviation to low quality. This loss could be from the possibility that any consumer might spot a flaw, that there is a group of discerning consumers who always spot flaws, or that the government interrupts fraudulent sales.

That the Klein-Leffler idea can be expanded to cover more than reputation is this paper’s main point. I have made it without fully exploring the equilibria of the static games, however, and without discussing their intuition in detail. The three assumptions A1, A2, and A3 each have slightly different properties. The rest of the paper will discuss in turn the three static models that they generate.

(A1) A Probability of Flaw Detection. If the quality is low, then with probability $0 \leq \alpha < 1$ the consumer observes that fact— he “spots the flaw.” If the quality is high, he receives no information.

One might think that if all consumers can spot a flaw in a product with high probability, competitive forces would lead to an equilibrium with price equal to marginal cost. Any firm, knowing it would lose most of its sales if it tried to sell flawed products, would keep its quality high even if it did not care about future periods, simply because of the potential loss in the current period.

The mistake in that reasoning is that when price equals marginal cost and marginal cost is constant, profits are zero and lost sales volume is no disincentive. Even if all consumers have probability .99 of observing low quality, a positive profit margin from the remaining fraction .01 is better than a zero profit margin from all of them. Hence there exists no equilibrium with high quality and with price equal to marginal cost.

The pessimistic equilibrium, on the other hand, can exist even in that extreme situation in which consumers have probability .99 of observing low quality. In the pessimistic equilibrium, firms produce low quality at a low price, and consumers believe that any firm which deviates to a high price will produce low quality nonetheless. A firm which deviates to high quality and high price will face two obstacles to selling its product. First, consumers will see its high price, expect low quality from it, and visit other firms. Second, even if, going outside the model, a consumer did visit the high-price high-quality firm, the fact that he would not observe a flaw would not induce him to buy at the high price. The consumer knows that he spots flaws with probability .99, but that gives him probability .01 of not spotting a flaw, and since his out-of-equilibrium belief is that the product is flawed, he can rationally retain that belief.

Competition fails to generate high quality at zero-profit prices because a firm can do better by cheating and making positive profits from however many customers fail to spot the cheat. Even if a firm loses 99% of its customers, the resulting profit is still positive. It cannot make up for a zero profit margin with any amount of volume. If the profit margin is even slightly positive, however, the calculus changes, and firms start to regret losing customers. If the profit margin is high enough, firms become willing to produce high quality.

Existence and Uniqueness of the Optimistic Equilibrium

I have already noted that the infinite-period model has multiple equilibria as an implication of the Folk Theorem of repeated games. I passed over the fact that there are other reasons for multiple equilibria that apply to either the basic Klein-Leffler model or to a single-period model generated by any of the three assumptions A1, A2, and A3. There is in addition another reason for multiple equilibria special to A2 that I will discuss later.

The single-period equilibrium is unique if and only if $p^* > v_h$, in which case only the pessimistic equilibrium survives. If flaws are so infrequently spotted that the quality-maintaining profit margin becomes too high, consumers switch to preferring low quality at a low price, even though high quality remains socially efficient.

Otherwise, both the pessimistic and optimistic equilibrium exist, but so does a continuum of other optimistic equilibria with prices in the range $p' \in [p^*, v_h]$. A strategy combination supporting the equilibrium with price p is that any firm which charges p' or higher produces high quality, any firm that charges less than p' produces low quality, consumers split equally across the firms that charge p' and otherwise do not buy, and, crucially, a consumer's out-of-equilibrium belief is that any firm charging a price less than p' chose low quality.

The equilibria with $p' > p^*$ lack plausibility. To be sure, this deviation would not be ruled out by the reasoning of the Cho-Kreps Intuitive Criterion, which says that if a deviation by one player would be profitable for him if he were of type T and the other player then believed him to be of type T , but not profitable for him if he were of type S and the other player believed him to be of type T . Here, a firm that was choosing high quality would like to deviate to p^* if that would attract consumers, but so would a firm that was choosing low quality. Whether a firm chose high or low quality it would want more customers, and so it would want customers to believe it had high quality.

Here, however, there is an even stronger reason than the Intuitive Criterion for why a firm that deviates to p^* could be expected to produce high quality: it is in the firm's interest to do so. Both high-quality and low-quality firms would benefit from higher volume, but if $p > p^*$, a firm with high quality has higher expected profit. It must not only attract customers, but also keep them once they have had a chance to look for flaws. Thus, a firm which has deviated to $p \neq p'$ where $p \geq p^*$ would also wish to choose high quality.

Formally, one way to exclude the implausible optimistic equilibria with $p' > p^*$ would be to allow firms to revise their quality choice after they make their price public. The equilibrium with $p = p^*$ is robust to the order of moves; the equilibria with $p' > p^*$ fail to survive if the firm can choose quality after price because the out-of-equilibrium belief that a firm would

choose low quality after choosing a price higher than p^* but not equal to p' would require irrational behavior by the firm.

Multiple optimistic perfect bayesian equilibria show up regardless of whether assumption A1, A2, A3, or none of the assumptions other than the Klein-Leffler assumption of multiple periods is used to generate a model with quality-inducing profits. Nonetheless, although the exact value of the equilibrium price is not pinned down unless restrictions are imposed on out-of-equilibrium beliefs, the interesting properties of the model are present in the entire continuum of equilibria: quality is high and profits are positive.

(A2) The Discerning Consumers Model. If a consumer is one of a fraction $0 \leq \beta < 1$ of consumers who are “discerning” he observes quality perfectly once he visit the seller.

Many of us who are ill-informed consumers are happy that better-informed consumers are out there giving sellers an incentive to keep quality at a reasonable level, so the underlying idea of this model fits our intuition well. A version of it can found for the monopoly context in chapter 2 of Jean Tirole’s 1988 book, *The Theory of Industrial Organization*. In his model, a monopoly seller chooses quality to be low or high. Some consumers can observe the quality before buying, while others cannot. In equilibrium, the seller will always choose high quality if there are sufficiently many informed consumers, while if there are not, he will choose high quality with some probability in the mixed strategy equilibrium that then exists. This paper’s one-period model with just assumption A2, not A1 or A3, is essentially the Tirole model transferred to a competitive market.

As we have seen with the comparative statics of parameter α for the probability of spotting a flaw and parameter β for the fraction of discerning consumers, the assumptions that all consumers have some chance of spotting flaws and that some consumers are sure to spot flaws have similar effects. As in the flaw-spotting model, there exists no equilibrium in which all the firms charge $p = c_h$ and quality is high, because in such a strategy combination a firm’s equilibrium payoff would be zero. If it deviated and chose low quality, then fraction θ of its customers would detect the low quality and turn away without buying, but it would have a positive profit margin on the remaining customers, for a positive payoff.

The pessimistic equilibrium exists in the discerning-consumer model, but it is less plausible than in the flaw-spotting model. Recall that in the flaw-spotting model there were two reasons why deviation to high price and high quality were unprofitable. The first was that the out-of-equilibrium belief was that a deviating firm had low quality, and so no consumer would visit it. The second was that even if a consumer did visit the deviating firm, he would not buy because merely not observing a flaw would not contradict his belief that the quality was low. Only the first reason applies in the discerning-consumers model. If a discerning consumer did visit a firm that had deviated to a high price, he could use direct observation to determine the quality. Hence, the equilibrium strategy of choosing to visit a firm charging a low price and low quality seems absurd compared with choosing to visit a deviating firm whose product would yield positive consumer surplus if its quality were high and zero (because it would be unbought) if it were low.

The optimistic equilibria with $p' \in [p^*, v_h]$ continue to exist, though, as in the flaw-spotting model, they depend on special out-of-equilibrium beliefs and on the assumption that a firm cannot change its quality after it has made public its price.

The discerning-consumers model differs in one important respect from the flaw-spotting model: besides the pessimistic and optimistic equilibria described already, it has an additional category of equilibrium, one which exists even if $p^* > v_h$ and the optimistic equilibrium is infeasible. The three classes of equilibria for a one-period model using assumption A2, but not A1 and A3 are listed below, where Class 1 and 3 equilibria always exist, but Class 2 equilibria require $p^* \leq v_h$.

Class 1: Pessimistic Equilibria $p = c_l$ and quality is low. Consumers are pessimistic and believe that quality is low regardless of what prices they see.

Class 2: Optimistic Equilibria. $p \in [p^*, v_h]$ and quality is high. Consumers believe that a firm that deviates to $p < p^*$ is selling low quality. The price p^* is, from our earlier calculations using L'Hospital's Rule of equation (17),

$$p^* = c_h + \frac{\theta(c_h - c_l)}{1 - \theta} - \frac{\gamma F}{1 - \theta}. \quad (21)$$

Since here $\alpha = \gamma = 0$, we can substitute $\theta = 1 - \beta$ and simplify to:

$$p^* = c_h + \frac{(1 - \beta)(c_h - c_l)}{\beta}. \quad (22)$$

Class 3: Mixed-strategy Equilibria. Firms charge \hat{p} with $\hat{p} > c_h$ and $\hat{p} \leq p^*$. They produce high quality with probability ϕ . Undiscerning consumers stay home with probability $(1 - \mu)$. They visit a random store charging \hat{p} and buy from it with probability μ , so the fraction of consumers who visit firms who are discerning is

$$d = \frac{\beta}{\beta + (1 - \beta)\mu} > \beta. \quad (23)$$

A supporting out-of-equilibrium consumer belief for any equilibrium is that any firm charging more or less than \hat{p} has low quality, in which case no consumer will visit that firm and its deviation will not yield positive profits.

In Class 3 mixed-strategy equilibria, the undiscerning consumers must be indifferent between visiting a firm and not, so:

$$\pi_{undiscerning}(buys) = \pi_{undiscerning}(stays\ home) \quad (24)$$

so:

$$\phi(v_h - p) + (1 - \phi)(v_l - p) = 0. \quad (25)$$

Then $\phi v_h - \phi p + (v_l - p) - \phi v_l + \phi p = 0$ and $\phi v_h - \phi v_l = p - v_l$. This solves to the following probability that a firm produces high quality:

$$\phi = \frac{p - v_l}{v_h - v_l} \quad (26)$$

Firms mix, so they must be indifferent between high and low quality, and the payoff per visiting consumer must be:

$$\pi_{firm}(high\ quality) = \pi_{firm}(low\ quality), \quad (27)$$

so

$$[\beta + (1 - \beta)\mu][p - c_h] = (1 - \beta)\mu(p - c_l), \quad (28)$$

so $\beta[p - c_h] + (1 - \beta)\mu[p - c_h] - (1 - \beta)\mu p + (1 - \beta)\mu c_l = 0$ and $\beta[p - c_h] - (1 - \beta)\mu c_h + (1 - \beta)\mu c_l = 0$ and $\beta[p - c_h] = \mu[(1 - \beta)c_h - (1 - \beta)c_l]$. This solves out to the probability that an undiscerning consumer buys:

$$\mu = \frac{\beta[p - c_h]}{(1 - \beta)(c_h - c_l)}, \quad (29)$$

which requires that $\beta[p - c_h] < (1 - \beta)(c_h - c_l)$, so $\beta p - \beta c_h < (1 - \beta)(c_h - c_l)$ so:

$$p \leq c_h + \frac{(1 - \beta)(c_h - c_l)}{\beta} = p^*. \quad (30)$$

The Class 3 mixed-strategy equilibrium has surprising properties, so let us take some time to explore its intuition. The mixed-strategy equilibria are discontinuously different from the optimistic equilibrium. If the price equals p^* , then firms choose high quality as a pure strategy because the loss of profit from sales to the discerning consumers would outweigh the gain from lower costs in selling to the undiscerning consumers. If the price is slightly below p^* , the firm would rather cut quality and sell only to the low-quality consumers. The mixed-strategy equilibrium works by having some of the undiscerning consumers refrain from buying, which by increasing the percentage of discerning consumers in the buying population allows the quality-ensuring price to fall. But since undiscerning consumers are earning strictly positive consumer surplus in the optimistic equilibrium, to induce them to not buy in the mixed-strategy equilibrium with only a slightly lower price requires a discontinuously greater probability of low quality. Thus, when p falls below p^* , we need a sudden jump in the probability of fraud to induce some undiscerning consumers to drop out.

Recall that \hat{p} can take any value between c_h and p^* , exclusive of those bounds. If p were to exceed p^* , the payoff from high quality would be strictly greater than from low quality for the firms, given that all the discerning consumers would buy if quality were high and not buy if it were low. One might imagine a similar category of mixed-strategy equilibria in which the number of discerning consumers was less because they mixed between buying and not buying, which would drive the quality-ensuring price above p^* in the same way as here the mixing of undiscerning consumers drives the price below p^* . That cannot happen, though, because a discerning consumer will never mix when there is some probability of high quality: he can visit a

firm and have some probability of buying high quality, which would yield a positive payoff that would be better than the zero payoff from staying home.

As p falls to c_h , the high-price firm's margin approaches zero, so the firm is more and more tempted to choose low quality. At the same time, however, the fraction of undiscerning consumers purchasing, μ , is also falling, reducing the temptation to choose low quality. That is why even a p barely above c_h can support an equilibrium.

If $p^* > v_h$, the optimistic equilibrium does not exist, but the mixed-strategy equilibria do. That is the situation where the quality-ensuring price exceeds the consumer reservation price, where the fraction of the undiscerning is so great that even at the monopoly price the firm would be willing to lose the business of the discerning in order to defraud the undiscerning. The mixed-strategy equilibria exist, however, because the fraction of undiscerning buying consumers becomes endogenous and falls to as low as necessary to support an equilibrium. Tirole (1988) shows this in the monopoly context: a monopoly can maintain at least some probability of high quality if it faces some discerning consumers, even though high quality cannot be maintained as a pure strategy and even though not all buyers can remain active.

If $p^* < v_h$, both the optimistic equilibrium and the mixed-strategy equilibria exist as perfect bayesian equilibria. Which of the two is most robust depends on the particular parameter values. The argument I made earlier against the optimistic equilibria with $p > p^*$ was that they depended critically on out-of-equilibrium beliefs and on whether the firm was able to choose quality after price or not. The high-price equilibria were fragile because both a firm and consumers would be willing to deviate to a $p = p^*$ equilibrium if they believed the other would be following it. Here, however, it is not so clear that if the expected equilibrium were some mixed-strategy equilibrium that a firm which deviated would be able to attract discerning customers.

To see why, we must think about the player's payoffs in the optimistic and the mixed-strategy equilibria. The mixed-strategy equilibria are all equally attractive to undiscerning consumers, because the undiscerning consumers, indifferent between buying and not buying, earn zero consumer surplus in all of them. The price is lower, but the undiscerning consumer buys low quality with higher probability. The pure-strategy equilibrium with p^* and high quality yields positive consumer surplus, however, so if some firm

in the mixed-strategy equilibrium were to deviate to p^* and be believed to produce high quality, it would attract all the undiscerning consumers.

Firms prefer the optimistic equilibrium for two reasons. In the Class 3 mixed-strategy equilibrium, a firm's expected payoff is equal with high and with low quality, so let us compare the high-quality payoff with $p \in (c_h, p^*)$ to the high-quality payoff with the price p^* . Clearly the higher price is better for the firm, one reason for preferring p^* . In addition the industry quantity sold (and thus the quantity per firm under our assumption of an exogenous-size interval of firms) is higher. Thus, firms prefer the optimistic equilibrium.

The discerning consumers' preference across equilibria is crucial. If they do not move to the firm charging p^* with high quality, the firm loses its incentive for quality. But the discerning consumers' face a tradeoff. They have positive expected payoffs and like the lower prices of the mixed-strategy equilibria, but the lower the price, the bigger the probability of low quality, and consumer surplus that is, ex post, zero. Thus, we must look carefully at a discerning consumer's payoff as a function of the price.

A discerning consumer's expected payoff is:

$$\pi_{discerning}(p) = \gamma(v_h - p) = \left(\frac{p - v_l}{v_h - v_l} \right) (v_h - p) = \frac{pv_h - v_l v_h + v_l p - p^2}{v_h - v_l}, \quad (31)$$

the derivative of which with respect to p is:

$$\frac{d\pi_{discerning}}{dp} = \frac{v_h + v_l - 2p}{v_h - v_l} \quad (32)$$

Equating expression (32) to zero to maximize the payoff yields the price that maximizes the discerning consumer's expected payoff:

$$\tilde{p} = \frac{v_h + v_l}{2}. \quad (33)$$

Note that \tilde{p} must exceed c_h to yield positive profits to the seller.

The discerning consumer's payoff with \tilde{p} in the mixed-strategy equilibrium

is

$$\begin{aligned}
\pi_{discerning}(\tilde{p}) &= \left(\frac{\tilde{p} - v_l}{v_h - v_l} \right) (v_h - \tilde{p}) \\
&= \left(\frac{\frac{v_h + v_l}{2} - c_l}{v_h - v_l} \right) \left(v_h - \frac{v_h + v_l}{2} \right) \\
&= \left(\frac{\frac{v_h}{2} - \frac{v_l}{2}}{v_h - v_l} \right) \left(\frac{v_h}{2} - \frac{v_l}{2} \right) \\
&= \frac{v_h - v_l}{4}.
\end{aligned} \tag{34}$$

This is to be compared with the discerning consumer's payoff in the pure-strategy equilibrium with $p = p^*$ from equation (22):

$$\begin{aligned}
\pi_{discerning}(p^*) &= v_h - p^* \\
&= v_h - c_h - \frac{(c_h - c_l)(1 - \beta)}{\beta}
\end{aligned} \tag{35}$$

Depending on the parameters, discerning consumers might prefer either the pure-strategy equilibrium with $p = p^*$ or the mixed-strategy equilibrium. If c_h is small, then $\pi_{discerning}(p^*)$ is small while $\pi_{discerning}(\tilde{p})$ is unaffected, so discerning consumers prefer the mixed-strategy equilibrium. In that case, even if a firm that deviated to $p = p^*$ were believed to produce high quality, it could not lure away the discerning consumers— and luring away only undiscerning consumers cannot support an equilibrium. On the other hand, if c_h is high, then $\pi_{discerning}(p^*) > \pi_{discerning}(\tilde{p})$ (e.g., if $v_h = 4, v_l = 0, c_h = 1, c_l = 0, \beta = .9$ then $4 - 12/4 < 4 - 1 - 4(.1)/.9$).

The discerning-consumer model is trickier than it seems at first glance, but leads us to the same essential conclusion as the flaw-spotting model: even in a one-period model, we may expect to observe high quality being produced if the price is above the marginal cost of high quality and firms fear deviation to low quality will cost them too much sales volume. The main difference between the models is that the discerning model's prediction is that even if consumer's reservation value is close to marginal cost and there is only a small probability of a consumer who observes quality, high quality is still possible in a mixed-strategy equilibrium.

(A3) The Weak Law Model: A Small Probability of Punishment by the Government. If a firm tries to sell low quality as high, then with probability $0 \leq \gamma \leq 1$ independent of α and β the government interrupts the transaction and fines the seller amount F .

Assumption A3 is different from A1 and A2 because it rules out consumer observation of quality altogether, as the Klein-Leffler model does, but adds government observation of quality. I noted earlier that if the probability of government detection of low quality is high enough and the punishment large enough then the expected punishment alone will deter fraud. Good laws can allow the market to attain the first best by heavily punishing firms that falsely advertise high quality or by requiring quality to be high. If such laws exist, but the probability of detection and the penalty are low, however, what happens? As before, let us concentrate on just one assumption, using assumption A3 but not A1 and A2 and restricting ourselves to a single period.

By our assumption of weak laws ($F < F^*$), we have ruled out an equilibrium in which firms charge $p = c_h$ and quality is high. In such a strategy combination, a firm's equilibrium payoff would be zero. If the firm deviated and chose low quality, then with probability γ it would lose the sale and incur punishment F , but by definition of F^* this would leave it with a positive expected payoff.

The pessimistic equilibrium would exist. Firms choose low quality and $p = c_l$, they advertise their low quality honestly, and consumers visit only firms with $p = c_l$. This yields payoffs of zero to all players. A firm cannot increase its payoff by raising its price because it would lose all its customers, and a consumer expects a negative payoff if he buys at any higher price. Nor can the firm increase its payoff by lying and charging a higher price, because consumers simply would not believe its claims.

In the optimistic equilibrium, firms choose high quality and $p = p^*$, where the price p^* is, from our earlier calculations for the general model,

$$p^* = c_h + \frac{\theta(c_h - c_l)}{1 - \theta} - \frac{\gamma F}{1 - \theta}. \quad (36)$$

Since here $\alpha = \beta = 0$, we can substitute $\theta = 1 - \gamma$ and simplify to

$$p^* = c_h + \frac{(1 - \gamma)(c_h - c_l)}{\gamma} - F. \quad (37)$$

A consumer randomly chooses a firm charging $p = p^*$ and buys the product. Out of equilibrium, we postulate that the consumer believes that prices below p^* imply low quality and prices above p^* imply high quality.

The intuition of the weak-law model is as follows. If $\gamma = 0$, so the government had zero probability of detecting a fraudulent transaction, then quality would of course be low. As soon as γ becomes positive, however, the firm faces a tradeoff. It can choose high quality and high production costs and sell to all consumers who visit it, or it can choose low quality and low production costs, in which case its margin will rise but its sales will fall. If the price is high enough, the lost sales are more important than the lower production cost.

In addition there is the prospect of paying the government penalty F . That is less important, however, and is not what drives the model. The government is still useful even if $F = 0$, so there is no punishment. Deterrence is still achieved with $\theta > 0$ because there is still a chance the government will interrupt the transaction and deprive the firm of its profit. To see the crucial importance of this, consider the alternative assumption A3':

(A3') If a firm tries to sell low quality as high, then with probability $0 \leq \gamma \leq 1$ independent of α and β the government fines the seller amount F . The firm is allowed to keep its profit from the transaction.

It remains true under Assumption (A3') that a large enough penalty F will deter fraud, and that $\gamma > 0$ is needed for the penalty to have any effect. If F is even slightly too small to deter fraud, however, the equilibrium moves from the first-best of no fraud and high quality to the pessimistic equilibrium; the optimistic equilibrium will not exist. In the model with our original assumption (A3), on the other hand, Proposition 1 applies, and if the government's probability of detection γ falls then p^* rises continuously. There is a discontinuous rise in fraud when p^* comes to exceed the reservation price v_h , with a sudden fall in producer profit as that happens, but consumer welfare approach zero as $p^* = v_h$ anyway, so consumers do not feel the effect of the fall in government detection discontinuously.

I have spent considerable time exploring multiple equilibria in the models based on A1 and A2. The model based on A3 behaves like the one based on A1, the flaw-spotting model: it has multiple equilibria based on the critical p lying anywhere in the continuum $[p^*, v_h]$, but it does not have the Class 3 mixed-strategy equilibria of the discerning-consumer model. As in the flaw-spotting model, only the equilibrium at $p = p^*$ is robust to out-of-equilibrium beliefs and to whether a firm can revise quality after choosing price.

One point special to the weak-law model is that it warns us to be careful about the type of law as well as its severity. Assumption A3 specifies “If a firm tries to sell low quality as high,” then it becomes subject to penalty. What if the law were somewhat different, and with probability θ , the government catches a seller who sells low quality, regardless of his claims, confiscates the profits, and imposes penalty $F \geq 0$?

If $F = 0$, nothing changes. If $F > 0$, however, the pessimistic equilibrium disappears. In the pessimistic equilibrium, $p = c_l$, yielding zero profits from sales, and if a firm also must pay F with probability θ , its net payoff will be negative. Hence, the market will dry up entirely, as firms will, if allowed, produce nothing at all.

This has an interesting policy implication— that weak measures attempting to ensure high quality may backfire and simply destroy the market. In the present model, this is no great harm, since consumer surplus is zero in the pessimistic equilibrium anyway due to the assumption that $v_l = c_l$. That was a simplifying assumption, however. More generally, $v_l > c_l$ but high quality is efficient because $v_h - c_h > v_l - c_l$. If so, the pessimistic equilibrium would yield positive consumer surplus, and punishing low quality would destroy that surplus.

4. Concluding Remarks

I have shown that the loss of a reputation and future profits is not the only way that high prices can guarantee product quality. Even in a one-period model, buyers can trust a high-price, positive-profit seller to produce high quality if a deviation to low quality would somehow cost him present profits. I have shown this in three contexts—the loss of sales to discerning buyers, the loss of sales to buyers lucky enough to spot flaws, and the loss of sales consequent to government punishment for fraud. In each setting,

sellers have an incentive to produce high quality which would be too weak to overcome the profit from cheating if price equalled marginal cost but is strong enough if prices are above marginal cost.

We should not call these three models “reputation models” since they are static and do not rely on buyer beliefs about future seller actions in the way that standard economic models of reputation do. In the everyday sense of the word, however, they are reputation models, because they are models in which buyers believe that certain sellers have more to lose than to gain by cheating and sellers earn rents to those buyer beliefs. Such beliefs might be repeated over many periods, or might just exist briefly; a reputation driven by our assumptions A1, A2, or A3 above could be either persistent or short-lived, depending on how expectations are formed. What is perhaps most important is that it is not vulnerable to the last-period problem, unlike reputations in the standard model.

On the other hand, calling these models “product quality models” may be too limiting. They are, more generally, models of moral hazard, the choice being not just high versus low quality, but honest versus deceitful action more generally. Consider, for example, an agent who might be paid an efficiency wage to induce him to choose high effort, but in a one-period setting. Applying the flaw-spotting model, the agent would choose high effort because of a probability that the principal would spot his low effort and break the relationship, which otherwise would have a wage greater than his reservation wage. Similarly, we could apply the model to a situation in which a firm is tempted to engage in opportunistic behavior with respect to its creditors, but, applying the discerning-consumer model, refrains from opportunism because it fears losing access to creditors who can observe the behavior directly. Thus, a static product-quality model may be useful in a variety of applications.

References

Bagwell, Kyle & Michael H. Riordan (1991) "High and Declining Prices Signal Product Quality," *American Economic Review*, 81(1): 224-239 (March 1991).

Chan, Yuk-Shee & Hayne Leland (1982) "Prices and Qualities in Markets with Costly Information," *Review of Economic Studies*, 49(4): 499-516 (October 1982).

Cooper, Russell & Thomas W. Ross (1984) "Prices, Product Qualities and Asymmetric Information: The Competitive Case," *Review of Economic Studies*, 51(2): 197-207 (April 1984).

Daughety, Andrew F. & Jennifer F. Reinganum (2005) "Imperfect Competition and Quality Signaling," Vanderbilt Economics Dept. and Law School working paper, <http://www.kelley.iu.edu/BEPP/documents/Reinganum.pdf> June 2005 (October 2006).

Diamond, Douglas W. (1989) "Reputation Acquisition in Debt Markets," *Journal of Political Economy*, 97(4): 828-862 (August 1989).

Farrell, Joseph (1980) "Prices as Signals of Quality," D.Phil. dissertation, Oxford University (1980).

Farrell, Joseph (1986) "Moral Hazard as an Entry Barrier," *RAND Journal of Economics*, 17(3): 440-449 (Autumn 1986).

Hertendorf, Mark N. & Per Baltzer Overgaard (2001) "Price Competition and Advertising Signals: Signaling by Competing Senders," *Journal of Economics & Management Strategy*, 10(4): 621-662 (Winter 2001).

Horner, Johannes (2002) "Reputation and Competition," *American Economic Review*, 92(3): 644-663 (June 2002).

Kihlstrom, Richard E. & Michael H. Riordan (1984) "Advertising as a Signal," *Journal of Political Economy*, 92(3): 427-450 (June 1984).

Kirmani, Amna & Akshay R. Rao (2000) "No Pain, No Gain: A Critical Review of the Literature on Signaling Unobservable Product Quality," *Journal of Marketing*, 64(2): 66-79 (April 2000).

Klein, Benjamin & Keith Leffler (1981) "The Role of Market Forces in Assuring Contractual Performance," *Journal of Political Economy*, 89(4): 615-641 (August 1981).

Maksimovic, V. & S. Titman (1991) "Financial Policy and Reputation for Product Quality," *Review of Financial Studies*, 4: 175-200.

Kreps, David M., Paul Milgrom, John Roberts & Robert Wilson (1982) "Rational Cooperation in the Finitely Repeated Prisoners' Dilemma," *Journal of Economic Theory*, 27(2): 245-252 (August 1982).

Linnemer, Laurent (2002) "Price and Advertising as Signals of Quality when Some Consumers Are Informed," *International Journal of Industrial Organization*, 20(7): 931-947 (September 2002).

Milgrom, Paul & John Roberts (1986) "Price and Advertising Signals of Product Quality," *Journal of Political Economy*, 94(4): 796-821 (August 1986) .

Nelson, Phillip (1970) "Information and Consumer Behavior," *Journal of Political Economy*, 78(2): 311-329 (March-April 1970) .

Rasmusen, Eric (1989) "A Simple Model of Product Quality with Elastic Demand," *Economics Letters*, 29(4): 281-283 (1989).

Rogerson, William P. (1983) "Reputation and Product Quality," *Bell Journal of Economics*, 14(2): 508-516 (Autumn 1983).

Shapiro, Carl (1983) "Premiums for High Quality Products as Returns to Reputations," *Quarterly Journal of Economics*, 98(4): 659-680 (Nov., 1983).

Stiglitz, Joseph E. (1987) "The Causes and Consequences of The Dependence of Quality on Price," *Journal of Economic Literature*, 25(1): 1-48 (March 1987).

Tirole, Jean (1988) *The Theory of Industrial Organization*, Cambridge: MIT Press (1988).

Wolinsky, Asher (1983) "Prices as Signals of Product Quality," *Review of Economic Studies*, 50(4): 647-658 (October 1983).