Product Quality, Reputation and Turnover*

Rafael Rob† and Tadashi Sekiguchi‡

February 15, 2001

Abstract

We consider a repeated duopoly game where each firm chooses its investment in quality, and realized quality is a noisy indicator of the firm's investment. We derive reputation equilibria, whereby consumers 'discipline' a firm by switching to its rival in case its realized quality is too low. The model predicts that firms with good reputation charge a higher price, sell a bigger quantity and have a higher stock-market capitalization. Every so often, the market is subjected to turnover, whereby the high-quality / good reputation firm loses market share, lowers its price and its capitalization suffers, while its rival gains market share, raises its price and enjoys increased capitalization. We examine properties of reputation equilibria. In particular, we show that the equilibrium is not efficient or nearly efficient even as the discount factor goes to 1.

*Part of this research was conducted while Sekiguchi was visiting the University of Pennsylvania. Sekiguchi thanks the University of Pennsylvania for its kind hospitality. We are also grateful to Michihiro Kandori and Hitoshi Matsushima, as well as seminar audiences at Kyusyu University and the University of Tokyo, for their helpful comments and discussions.

†Department of Economics, University of Pennsylvania, 3718 Locust Walk, Philadelphia, PA. 19104-6297

‡Faculty of Economics, Kobe University, 2-1 Rokkodai-cho, Nada-ku, Kobe 657-8501, Japan

1 Introduction

When product quality is not apparent at the point of sale a consumer runs the risk of buying a low-quality product at a high-quality price. As a result, and as Akerlof (1970) has pointed out some 30 years ago, the market outcome may be one where low-quality products are being traded, although buyers and sellers derive a higher surplus by trading higher-quality products. Numerous market and non-market mechanisms have been suggested to cope with this problem. In this paper we analyze one such mechanism, which is operative when consumers purchase the product repeatedly and are able to observe its realized quality (once they "bring it home" and use it), although they are still not able to observe the firm's investment in quality. In this set of circumstances consumers are able to reward a firm whose realized quality is high by raising their future demand for the firm's product. Conversely, consumers are able to penalize a firm whose realized quality is low by lowering their future demand for the firm's product. Therefore, if high investment in quality is more likely to result in high realized quality the firm has an incentive to invest in quality because of the threat of consumers' disciplinary actions.

Rather than describe this as "disciplinary actions" we can describe it as a reputational mechanism. Suppose a firm is known, or reputed, to have produced high-quality products in the past, and suppose that that makes consumers believe the firm will also produce high-quality products in the future. Then having the reputation of being a high-quality provider is valuable because consumers expect the delivery of high-quality products in the future and are, hence, willing to pay a premium for the firm's product. Therefore, investing in product quality may be viewed as an attempt to sustain one's reputation and the flow of extra profit that comes with it. Conversely, not investing in quality may result in poor realized quality, tarnished reputation and the denial of future profit.

There is ample empirical evidence to suggest that consumers discipline firms in this fashion and that firms react to this by trying to sustain their reputation. This evidence is most transparent, probably, in the case of catastrophic losses. A recent example is the Ford / Bridgestone / Firestone debacle, which occupied the media during the summer of 2000; see, for example, the Chicago Tribute 8 / 24 / 2000. As reported in that article Bridgestone / Firestone decided to fire 25% of its workforce in one plant and suspend production for more than a month in another plant, and these cuts might have been even bigger if not for the tire-replacement program which kept production temporarily high. A similar event took place some 20 years ago when Tylenol capsules were poisoned at the retail level in the summer of 1982. As documented by Mitchell (1989), Johnson and Johnson decided to cut the price of Tylenol by 2.50 dollars (by distributing a coupon), and had lost significant market share. Likewise, following an airplane crash, consumers switch to competitors, which affect adversely the market share and profits of the airline whose plane has crashed; see Chalk (1986). While such events receive a lot of publicity, this phenomenon is certainly not limited to catastrophic losses. For instance, Ippolito (1992) persuasively documents the reaction of investors to performance in the mutual-fund industry. Not surprisingly, he finds that poor relative performance results in investors shifting their assets into other funds.

The following empirical regularities are common in the above (and other) cases. When a firm is hit by an adverse event (car accidents, drug poisoning, etc.) the demand for its product is dramatically reduced. As a result, the firm loses market share, it has to decrease its product price and its stock-market capitalization suffers. While the exact combination of these regularities, and where consumers end up diverting their expenditures varies, naturally, from one industry (or case) to another, these regularities can be found quite commonly in the aftermath of many adverse events.

In this paper we construct and analyze a model based on these empirical regularities. More specifically, the model is based on the premise that consumers react to an adverse event by lowering their demands or, more precisely, by expecting the firm to deliver a poor-quality product from this point onwards, which, in turn, translates into low demand. Given this low demand the firm has indeed little incentive to invest in product quality, which, in turn, fulfills consumers' expectations. Conversely, when realized quality is high, consumers' expectations and demands are high and the firm invests more in product quality, which, again, fulfills consumers' expectations. The aim of our model is to embed this idea into an equilibrium framework and show the set of circumstances under which such *reputation equilibria* exist and to study their properties. In particular the model highlights the type of dynamic that this mechanism gives rise to; namely, it shows that the time path of an industry is subject to turbulences, whereby market shares, prices, and stock-market capitalizations fluctuate over time.

The first paper to suggest consumers-switching-alliances as a disciplinary device is the one by Klein and Leffler (1981). A subsequent paper by Shapiro (1983) significantly extended this idea by developing a competitive model, where different firms specialize in different quality products, and where entry and exit are possible. Both papers as well as subsequent contributions consider, however, only the "deterministic case," i.e., where the link between investment in quality and realized quality is deterministic. By contrast, we consider the stochastic case, where higher investment leads to stochastically higher quality but may, nonetheless, result in low quality (for reasons beyond the firm's control.) In that case, and unlike the deterministic case, a firm cannot maintain indefinitely a reputation for being a high-quality provider. To the contrary, sooner or later, the quality of a firm's product falls short of consumers' tolerance level, consumers' expectations and demands for the product drop, and so does the firm's product quality. In this way market shares, prices and stock-market capitalizations fluctuate over time - the market is subjected to turnover - which is consistent with the empirical literature mentioned earlier.

In greater detail our analysis delineates a whole set of equilibria that can be supported via such mechanism. These equilibria are differentiated based on firms' investment in quality and consumers' tolerance level - the level of realized quality below which a firm "falls from grace," and a punishment is triggered. Not surprisingly given the self-fulfilling nature of equilibrium there are multiple, Pareto-ranked equilibria.

The variables which differentiate the equilibria - consumers' tolerance level and firms' investment in quality - are not directly observable. Nonetheless, the values of these variables determine the values of two other variables, which *are* observable, and which can be used to study the relationship between the various equilibria. Indeed, a firm's investment in quality determines its period profit, an observable. And a firm's investment in quality along with consumers' tolerance level determine the probability of turnover, also an observable (known in the empirical literature as the "hazard rate.")

When we compare reputational equilibria, for instance, by looking at the same product traded at geographically separated markets, the following relationship between profitability and turnover is predicted by our model. If the discount factor is high, then, locally, high-profit firms are subjected to less turnover, i.e., there is a negative relationship between profits and turnover. This accords with the intuition that profitable and patient firms have a lot to lose by tarnishing their reputation and, hence, invest enough in product quality to make the likelihood of turnover low.

However, the qualification to high discount factors and to local changes is indispensable. Without this qualification, one could get a *positive* relationship between profits and turnover. This hinges on whether consumers' tolerance level and firms' investment in quality are complements or substitutes, i.e., whether an increase in consumers' tolerance level leads to an increase or a decrease in firms' investments. Both reactions are economically plausible: As consumers' become more tolerant firms might either be "encouraged" to increase their investment in quality, or they might "abuse" consumers' tolerance and decrease their investment in quality. Which reaction prevails depends on the discount factor and on consumers' initial tolerance level. And, which reaction prevails determines whether we have a negative or a positive correlation between profits and turnover. In fact, in one of our examples we show that the correlation between profits and turnover is "hump-shaped."

The rest of the paper is organized as follows. In the next section we set up the model. In section 3 we characterize reputation equilibria. In section 4 we state conditions under which a reputation equilibrium exists. In section 5 we study the comparative static properties of the equilibrium. And, in section 6, we study the efficiency of reputation equilibria.

2 The Model

We consider a model of duopoly, denoting the firms by $i = 1, 2$. There is a continuum of consumers whose measure is 1. The model is formulated as an infinitely-repeated game in discrete time. We start with a description of the stage game.

At the beginning of each period each firm i chooses an action x_i, interpreted as "investment in quality." We restrict $x_i \in X \equiv [\underline{x}, \infty)$ and assume $\underline{x} \geq 0$. When a firm chooses x_i it spends $c(x_i)$. At the same time, firm i also chooses the price of its product, $\widehat{p}_i \geq 0$. If a fraction $\gamma \in [0, 1]$ of consumers buy from firm i, firm i's stage game payoff is $\gamma\widehat{p}_i - c(x_i)$.[1]

Once x_i is chosen and assuming the firm sells a positive quantity, $\gamma > 0$, the quality of its product q_i is realized. q_i is a random variable with support $q_i \in [\underline{q}, \infty)$, where $\underline{q}$ is exogenously given and can be $-\infty$. q_i is the same for all consumers.[2] x_i is privately chosen and observed, whereas q_i is publicly observed. The distribution of q_i depends on x_i and we denote the conditional density by $f(q_i \mid x_i)$ with a corresponding cumulative distribution function $F(q_i \mid x_i)$. Throughout the paper, we make the following assumptions.

Assumption 0.

1. c and F are twice continuously differentiable.

2. c is strictly increasing and strictly convex, with $c(\underline{x}) = c'(\underline{x}) = 0$ and $c(\infty) = c'(\infty) = \infty$.

3. $\int qf(q \mid x) = x$, and f satisfies stochastic dominance: $F_2(q \mid x) < 0$ for all $q > \underline{q}$, where F_2 is the derivative of F with respect to its second variable, x.

The following examples satisfy Assumption 0.

Example 1 The cost is $c(x) = (x - \underline{x})^\alpha$, where $\alpha > 1$, and the quality is normally distributed with variance 1 and mean value x, $f(q \mid x) = \frac{1}{\sqrt{2\pi}} e^{-(q-x)^2}$. In this example, we have $\underline{q} = -\infty$.

[1] Since the total cost of production is independent of the number of consumers who buy it, γ, this formulation assumes that x_i is a fixed cost. A more general assumption is that the cost depends both on the investment and on the fraction of consumers who buy from the firm. In this case, we would denote the cost by $c(x_i, \gamma)$, and we would assume that $c(x_i, \gamma)$ satisfies Assumption 0 for any fixed γ and that it is linear in γ for any fixed x_i. This more general approach yields the same qualitative results, so we here limit attention to the special case of $c(x_i) = c(x_i, \gamma)$, for the sake of simplicity.

[2] In some real world applications, for example airlines, not all consumers buy the product in the same period. In that case the assumption is that realized quality (airplane crashed or did not crash) becomes *known* to all consumers, whether they bought the product or not. Formally, the assumption is that realized quality becomes common knowledge.

Example 2 The same cost as in Example 1, but quality is exponentially distributed, $f(q \mid x) = \frac{1}{x} e^{-\frac{q}{x}}$, $F(q \mid x) = 1 - e^{-\frac{q}{x}}$. In this example $\underline{q} = 0$.

Consumers move after firms have chosen their x_i and $\widehat{p}_i$, and have unit demands in each period. Namely, given the firms' price vector $(\widehat{p}_1, \widehat{p}_2)$, and given consumers' beliefs regarding the equilibrium x_i, consumers decide whether to buy or not, and in case they decide to buy, from which firm. A consumer's gross benefit when she buys one unit, whose realized quality is q, is q. If the consumer paid price p, her net benefit is $q - p$. Consumers are risk neutral with respect to quality shocks.[3] Thus, if a consumer buys from a firm whose choice of action is (x, p), her expected benefit is $x - p$ by Assumption 0. If a consumer buys nothing she gets zero benefit.

We consider a repeated game in which the stage game described above is played infinitely often. In this repeated game, we have two types of public information: The price and realized quality of the firms, and the fractions of consumers who buy from each firm in each period (since the consumers are anonymous, firms know the measure of consumers who bought from them, but not *who* bought from them). Those two types of information constitute the *public history* of the game. The private history of the game is as follows. For each firm the sequence of actions, x_i, it chose, and, for each consumer, the sequence of firms she bought from. A player's strategy in this repeated game specifies behavior as a function of both the public and her own private history. In later sections, though, we restrict attention to *public strategy equilibria*, in which each player's action depends solely on the public history.

The objective of the firms is to maximize average discounted profits, where the discount factor is $\delta \in (0, 1)$. Each consumer, on the other hand, maximizes her payoff *in the current period* given a particular history. There are several ways to justify this assumption of myopia. First, consumers may be literally myopic, i.e., their discount factor equals zero. Second, which is a subcase of the first, consumers may live just one period and each consumer is replaced by a successor and successors learn the public history of the game before making their choices. Third, and most relevant to the case of non-atomic consumers, if all consumers behave myopically, then, a single consumer has no effect on the public history and, hence, on the game's continuation; thus, the best a single consumer can do is maximize her current period payoff.

This completes the description of the repeated game. We proceed now to solve it, using sequential equilibrium as the solution concept.

[3] In other words, we assume that the von Neumann-Morgenstern utility of the consumer paying p for the good whose quality is q is $q - p$.

3 Reputation Equilibrium

Now that we have boiled down the situation into a repeated game, we should expect multiple equilibria, i.e., a Folk-Theorem type result.[4] This holds true even if we confine attention to strategies that depend on public-information only.[5] However, in this paper we focus on a particular subset of the set of equilibria, which we call *reputation equilibria*. These equilibria are defined more precisely immediately below.

The reason for restricting attention to particular equilibria is three-fold. First, as discussed in the introduction, there is empirical support for firms' attempts to establish and maintain reputation for product quality, and a reputational equilibrium is the theoretical counterpart of such behavior. Thus, by studying reputational equilibria we capture and understand the logic of a frequently observed phenomenon. Second, some of the Folk-Theorem equilibria are supported by strategies that are very complicated for players to carry out, whereas reputational equilibria are simple. Thus, reputational equilibria may be more realistic descriptors of actual behavior than other repeated-game equilibria. Third, some repeated game equilibria achieve high payoffs when the discount factor is close to one (see Fudenberg, Levine and Maskin (1994), which is a normative reason for favoring them. However, we analyze the game under a *fixed* discount factor. So these results are not necessarily valid, and it is less than clear that higher payoffs (than the ones we exhibit) can be achieved by means of relatively simple strategies.

Let us first describe the strategies that players use in the equilibria we study, calling them *reputation strategies*. In a reputation strategy, only one firm makes sales in each period, and all consumers buy from that firm. The reason only one firm makes positive sales is that consumers believe it invests more in product quality.[6] We call the lone seller (in some period) the *high quality firm* (of that period), or, in short, the HQ firm. The firm that sells nothing is called the *low quality firm*, or, in short, the LQ firm. In each period, the HQ firm chooses its investment in quality, as well as an associated price, and we denote this pair by (x_H, p_H). Similarly, the LQ firm chooses (x_L, p_L) in each period.

Next we turn to the rule that specifies who is the HQ firm in each period. The rule is specified inductively. First, the HQ firm of period 1 is firm 1.[7]

[4] This is true even in the presence of imperfect monitoring; see Fudenberg, Levine and Maskin (1994). Although that paper considers discrete action sets, we can discretize our action sets, and obtain a similar game to which their result is applicable.

[5] See Abreu, Pearce and Stacchetti (1986, 90) for studies of the structure of the public strategy equilibrium set.

[6] In fact the equilibrium of a static Bertrand game in which one firm is believed to produce a higher-quality product is such that the high-quality firm supplies the whole market while the low-quality firm sells nothing. Therefore, reputation equilibria extend this feature from the static game to the repeated game.

[7] That, obviously, is an arbitrary specification and a 'twin' equilibrium exists whereby firm 2 is the HQ firm of period 1.

Second, let a period $t > 1$ and the HQ firm of period $t-1$ be given. If HQ sold a positive quantity[8] in period $t-1$ and if the corresponding quality was above some threshold, denoted by $\overline{q}$, the firm continues to be the HQ firm of period t. Otherwise, it becomes the LQ firm of period t, and the firm that was the LQ firm of period $t-1$ becomes the HQ firm of period t. $\overline{q}$ is endogenous and has yet to be determined. We refer to $\overline{q}$ as *consumers' tolerance level.*

Let us express this idea in formal terms. A reputation strategy profile is a 5-tuple $(x_H, x_L, p_H, p_L, \overline{q})$. This 5-tuple along with the rule stated in the previous paragraph determines which firm is the HQ firm at any given history (whether deviation had occurred or not.) Thus, at any information set of a consumer, she knows who is the HQ firm and the prices chosen by the two firms, $(\widehat{p}_H, \widehat{p}_L)$. Using this data, a consumer's reputation strategy is as follows. At any information set, buy from the HQ firm if $x_H - \widehat{p}_H \geq \max\{0, x_L - \widehat{p}_L\}$; otherwise, buy from the LQ firm if $x_L - \widehat{p}_L > 0$, and buy from neither firm otherwise. A firm's reputation strategy is as follows. Whenever it is the HQ firm, choose (x_H, p_H); otherwise, choose (x_L, p_L).

A reputation equilibrium is a reputation strategy profile that is a sequential equilibrium. Given that a reputation strategy profile is characterized by five parameters, we have, in principle, a large number of equilibrium candidates. The following result, however, reduces considerably the set of these candidates.

Proposition 1 *In any reputation equilibrium the following must hold*
(i) $x_L = \underline{x}$.
(ii) $x_H \geq c(x_H) + \underline{x}$*, where equality holds if and only if* $x_H = \underline{x}$.
(iii) $p_H = x_H - \underline{x}$.
(iv) $p_L = 0$.

Proof. Fix a reputation equilibrium characterized by $(x_H, x_L, p_H, p_L, \overline{q})$ and fix a period. Since the LQ firm cannot sell anything in this period unless the HQ firm or consumers deviate, it chooses the action that minimizes its cost. Hence (i) follows.

To prove (ii), suppose $\underline{x} < x_H \leq c(x_H) + \underline{x}$. If $p_H \leq c(x_H)$, then HQ's overall payoff is nonpositive. If HQ chooses $(\underline{x}, p_H)$ instead of (x_H, p_H), consumers cannot observe this deviation so they buy from HQ anyway. Therefore, since $x_H > \underline{x}$, this deviation gives HQ larger current-period profit. Moreover, choosing $\underline{x}$ maximizes the probability of turnover, which is a good thing for HQ because $p_H \leq c(x_H)$. Thus playing $(\underline{x}, p_H)$ and then conforming to the reputation strategy is a profitable deviation. Thus $p_H \leq c(x_H)$ is not consistent with equilibrium.

So suppose $p_H > c(x_H)$. Because $x_H \leq c(x_H) + \underline{x}$, there exists an $\varepsilon > 0$ such that $0 \leq x_H - p_H < \underline{x} - \varepsilon$. The first inequality follows because consumers

[8] Although the HQ sells positive quantity *in equilibrium* the rule is specified independently of whether firms are playing their equilibrium strategies or not. In particular, if the HQ firm deviates it may sell nothing.

must be willing to buy from HQ in equilibrium. Suppose LQ chooses $(\underline{x}, \varepsilon)$ in period 1. Since (i) holds, consumers buy from LQ (by the choice of ε), and therefore LQ can earn a positive profit in this period. Moreover, this deviation automatically causes a turnover, because HQ would sell nothing. Hence for LQ, the deviation not only produces a larger profit in this period but also maximizes the probability of turnover, which is a good thing because $p_H > c(x_H)$. Thus the deviation is profitable, and therefore we must have (ii) in equilibrium.

To prove (iii), suppose first $p_H < x_H - \underline{x}$. This implies the existence of an $\varepsilon > 0$ such that $x_H - (p_H + \varepsilon) > \max\{0, \underline{x} - p_L\}$, and therefore HQ can still make sales in the next period if it chooses $(x_H, p_H + \varepsilon)$. Since this deviation has no effect on the probability of turnover, it is profitable. Thus $p_H < x_H - \underline{x}$ cannot hold in equilibrium. If $p_H > x_H - \underline{x}$, then there exists an $\varepsilon > 0$ such that $\underline{x} - \varepsilon > x_H - p_H \geq 0$. Thus if LQ chooses $(\underline{x}, \varepsilon)$ in the current period, consumers would buy from it. Hence the deviation raises both the current-period profit and the probability of turnover, which makes it a profitable deviation. Therefore, $p_H > x_H - \underline{x}$ is also inconsistent with equilibrium, and we have (iii).

Finally, to prove (iv), suppose $p_L > 0$. Then there exists an $\varepsilon > 0$ such that $\underline{x} - p_L < x_H - (p_H + \varepsilon)$, because of (iii). Therefore, HQ can still make sales in the current period when it chooses $(x_H, p_H + \varepsilon)$, which raises the current period profit without affecting the probability of turnover. Thus $p_L > 0$ cannot be true in equilibrium. ■

Proposition 1 tells us that a reputation equilibrium is characterized by two parameters only, $(x_H, \overline{q})$.

Let us fix a reputation strategy profile characterized by $(x_H, \overline{q})$, which satisfies (i)-(iv) of Proposition 1. Let v_H and v_L be the overall (average) payoffs of HQ and LQ, respectively, given the discount factor δ. Since turnover occurs with probability $F(\overline{q}|x_H)$, we have the following recursion equations.

$$v_H = (1-\delta)[\widehat{x}_H - c(x_H)] + \delta\{v_H[1 - F(\overline{q}|x_H)] + v_L F(\overline{q}|x_H)\}, \tag{1}$$

where $\widehat{x}_H = x_H - \underline{x}$, and

$$v_L = \delta\{v_L[1 - F(\overline{q}|x_H)] + v_H F(\overline{q}|x_H)\}. \tag{2}$$

Solving (1) and (2), we obtain

$$v_H = \frac{1 - \delta[1 - F(\overline{q}|x_H)]}{1 - \delta[1 - 2F(\overline{q}|x_H)]}[\widehat{x}_H - c(x_H)] \tag{3}$$

and

$$v_L = \frac{\delta F(\overline{q}|x_H)}{1 - \delta[1 - 2F(\overline{q}|x_H)]}[\widehat{x}_H - c(x_H)]. \tag{4}$$

Given that LQ follows the reputation strategy, let $v(x; \delta, \overline{q}, x_H)$ be the payoff of HQ when it chooses (x, p_H) in the current period and then follows the reputation strategy from the next period onwards. HQ is still the lone seller given consumers' belief, but turnover now occurs with a different probability, $F(\overline{q}|x)$. Hence we obtain

$$v(x;\delta,\overline{q},x_H) = (1-\delta)[\widehat{x}_H - c(x)] + \delta\{v_H[1-F(\overline{q}|x)] + v_L F(\overline{q}|x)\}. \quad (5)$$

Using (3) and (4), (5) is rewritten as

$$v(x;\delta,\overline{q},x_H) = (1-\delta)[\widehat{x}_H - c(x)] + H(x;\delta,\overline{q},x_H)[\widehat{x}_H - c(x_H)], \quad (6)$$

where

$$H(x;\delta,\overline{q},x_H) \equiv \delta\frac{(1-\delta)[1-F(\overline{q}|x)] + \delta F(\overline{q}|x_H)}{1-\delta[1-2F(\overline{q}|x_H)]}. \quad (7)$$

Hence, in order for the above reputation strategy profile to be an equilibrium, it must be a maximizer of the RHS of (6); namely,

$$v(x_H;\delta,\overline{q},x_H) \geq v(x;\delta,\overline{q},x_H) \text{ for all } x \in X. \quad (8)$$

The above argument shows that (8) is a necessary condition for a reputation equilibrium. The following result demonstrates that it is also a sufficient condition.

Proposition 2 *Fix a reputation strategy profile characterized by* $(x_H,\overline{q})$*, satisfying (i)-(iv) of Proposition 1. Then the profile is a sequential equilibrium if and only if (8) holds.*

Proof. We have already proven the "only-if" part. Thus, let us suppose (8) holds. Let consumers' belief at any period t be that HQ has chosen x_H and LQ has chosen x_L in that period, given any actual choice of prices, $(\widehat{p}_H,\widehat{p}_L)$. This belief is consistent, and the consumers' purchase decision specified by the reputation strategy is sequentially rational given this belief.

Let us then turn to the incentives of firms. Note that at the beginning of any period, the HQ firm (the LQ firm, respectively) of that period has a continuation payoff of v_H (v_L), given its strategy. Thus (8) implies that any one-shot deviation with respect to x of the HQ firm would be unprofitable. Furthermore, we have seen in Proposition 1 that deviating with respect to price is not profitable, either, which establishes sequential rationality of the HQ firm's behavior. The LQ firm makes no sales irrespective of its choice of actions, as long as all other players follow the given strategy profile. Hence its behavior is also sequentially rational, which completes the proof. ■

Thus (8) is the equilibrium condition for our solution concept. With an abuse of terminology, we say that $(x_H,\overline{q})$ is a reputation equilibrium under δ if (8) is satisfied.[9]

[9] Although $x_H,\overline{q}$ are jointly determined in equilibrium and, thus, have the same "status" of

4 Existence

In this section we prove the existence of reputation equilibria. To start with, let us note that a trivial (or degenerate or Akerlof-type) reputation equilibrium always exists for *any* fixed tolerance level $\overline{q}$. Namely, $(\underline{x}, \overline{q})$. In this equilibrium, consumers' expectation is so low that firms never bother to invest in quality (above the minimum, $\underline{x}$) irrespective of consumers' tolerance level.

Obviously, we are interested in "better" equilibria, i.e., equilibria which have higher social surplus associated with them. And the whole point of introducing repeat purchase and consumers' disciplinary actions is to ensure the existence of such equilibria. With this idea in mind we seek now equilibria with $x_H > \underline{x}$. To prove the existence of such equilibria, we need an additional assumption and a definition.

Assumption 1 $c''(\underline{x}) = 0$.

Assumption 1 is satisfied in Examples 1 and 2 if $\alpha > 2$.

Next, we define

$$Q = \{\overline{q} > \underline{q} : -\frac{c''(x)}{c'(x)} + \frac{F_{22}(\overline{q}|x)}{F_2(\overline{q}|x)} < 0, \quad \forall x > \underline{x}\},$$

where $F_2(\overline{q}|x)$, $F_{22}(\overline{q}|x)$ are the first and second derivatives of F with respect to its second variable, x.

The significance of the set Q is that the function $v(x; \delta, \overline{q}, x_H)$ is "hump-shaped" in x for any fixed δ and x_H, if $\overline{q} \in Q$. More precisely, we have:

Lemma 3 *Fix $\overline{q} \in Q$, $\delta \in (0,1)$ and x_H. Then the function $v(\cdot; \delta, \overline{q}, x_H)$ attains a global maximum on X, and has no other local maxima. It has at most one local minimum, which occurs at $x = \underline{x}$.*

Proof. Fix $\overline{q} \in Q$ and $\delta \in (0,1)$. For any fixed x_H, (6) and (7) imply that $v(x; \delta, \overline{q}, x_H)$ is a continuous function of x. Moreover, since $H(x; \delta, \overline{q}, x_H)$ is bounded and since $c(x) \to +\infty$ as $x \to \infty$, $v(x; \delta, \overline{q}, x_H) \to -\infty$ as $x \to \infty$. Thus, $v(x; \delta, \overline{q}, x_H)$ attains a maximum on X.

To show there is no other local maximum, suppose we have two local maxima, x_1 and x_2, where $x_1 > x_2$. Let x_3 be an element of $[x_2, x_1]$ that minimizes $v(x; \delta, \overline{q}, x_H)$ on $[x_2, x_1]$. Let us confine attention to $x_3 \in (x_2, x_1)$, for otherwise, if $x_3 \in \{x_1, x_2\}$, $v(x; \delta, \overline{q}, x_H)$ is constant on $[x_2, x_1]$. Since $x_3 > x_2 \geq \underline{x}$, both

being self-enforcing, there is a fundamental difference between x_H and $\overline{q}$. x_H is the behavior of a single player, the HQ firm, so it is conceivable that this player will eventually converge on it (via, say, a process of trial and error.) On other hand, $\overline{q}$ is the behavior of a continuum of players so it is far from obvious how they would converge on the same behavioral rule. We interpret $\overline{q}$ as a "convention" or a "norm," although our model is silent on how such norm may spontaneously arise.

the first-order and the second-order conditions must be satisfied at x_3. Hence, we obtain

$$-(1-\delta)c'(x_3) + H_1(x_3;\delta,\overline{q},x_H)[\widehat{x}_H - c(x_H)] = 0 \tag{9}$$

and

$$-(1-\delta)c''(x_3) + H_{11}(x_3;\delta,\overline{q},x_H)[\widehat{x}_H - c(x_H)] \geq 0. \tag{10}$$

Since $x_3 > \underline{x}$, it follows that $c'(x_3) > 0$, which in turn implies $H_1(x_3;\delta,\overline{q},x_H) \neq 0$ by (9). Thus, combining (9) and (10), we have

$$-\frac{c''(x_3)}{c'(x_3)} + \frac{H_{11}(x_3;\delta,\overline{q},x_H)}{H_1(x_3;\delta,\overline{q},x_H)} \geq 0. \tag{11}$$

However, since (7) implies

$$\frac{H_{11}(x_3;\delta,\overline{q},x_H)}{H_1(x_3;\delta,\overline{q},x_H)} = \frac{F_{22}(\overline{q}|x_3)}{F_2(\overline{q}|x_3)}, \tag{12}$$

combining (11) and (12) yields a contradiction to the fact that $\overline{q} \in Q$. ■

Therefore, $\overline{q} \in Q$ implies that $v(\cdot;\delta,\overline{q},x_H)$ has a unique maximizer *and* no local minimizers other than $\underline{x}$. Thus, there is a one-to-one correspondence between the set of equilibria with $x_H > \underline{x}$, and the set of solutions to $v_1(x_H;\delta,\overline{q},x_H) = 0$. Or, in other words, the "first-order approach" à la Rogerson (1985) is valid.[10]

Note that the Lemma 3 does not rely on Assumption 1. On the other hand, our next result (the proof of existence) does.

Proposition 4 *Assume Assumption 1 holds. Then, for any $\overline{q} \in Q$ and $\delta \in (0,1)$,*

(i) an $x_H > \underline{x}$ exists so that $(x_H, \overline{q})$ is a reputation equilibrium under δ.

(ii) The set of all x_H's such that $(x_H, \overline{q})$ is a reputation equilibrium is compact.

Proof. Fix $\overline{q} \in Q$ and $\delta \in (0,1)$. By the foregoing Lemma, there exists a unique maximum to $v(x;\delta,\overline{q},x_H)$ on X. Let us denote it by $R(x_H;\delta,\overline{q})$. In view of Proposition 2, $(x_H, \overline{q})$ is a reputation equilibrium under δ if and only if x_H is a fixed point of $R(\cdot;\delta,\overline{q})$, i.e., $x_H = R(x_H;\delta,\overline{q})$.

We first show that $R(x_H;\delta,\overline{q})$ is a continuous function of x_H. To see this, fix x_H and let $\{\xi_n\}_{n=1}^{\infty}$ be a sequence of real numbers converging to zero. Then for any x and any n, we have

$$v(R(x_H+\xi_n;\delta,\overline{q});\delta,\overline{q},x_H+\xi_n) \geq v(x;\delta,\overline{q},x_H+\xi_n). \tag{13}$$

[10] Hump-shapedness is slightly weaker than strict quasi-concavity and, a fortiori, strict concavity. Thus, if $F(\overline{q}|x)$ is convex in x (as is assumed by Rogerson (1985)), $\overline{q} \in Q$ follows. Indeed, in our examples together with Assumption 1, $\overline{q} \in Q$ is equivalent to convexity of $F(\overline{q}|x)$ in x.

Let $\overline{R}$ be a limit point of the sequence $\{R(x_H+\xi_n;\delta,\overline{q})\}_{n=1}^{\infty}$. Then taking a limit in (13) yields

$$v(\overline{R};\delta,\overline{q},x_H)\geq v(x;\delta,\overline{q},x_H)$$

for any x. Therefore, by Lemma 3, $\overline{R}=R(x_H;\delta,\overline{q})$. Since the choice of $\{\xi_n\}_{n=1}^{\infty}$ is arbitrary, R is continuous in x_H.

If x_H is so large that $x_H-c(x_H)<\underline{x}$, then it is easily seen that $R(x_H;\delta,\overline{q})=\underline{x}$. Next, let us define

$$h(x)=-(1-\delta)c'(x)+H_1(x;\delta,\overline{q},x)[x-\underline{x}-c(x)].$$

It is easily seen that $h(\underline{x})=0$. By Assumption 1, we have $h'(\underline{x})>0$. Hence there exists an $x_1>\underline{x}$ such that $R(x_1;\delta,\overline{q})>x_1$. Since R is continuous, there exists an $x_H>x_1$ so that $x_H=R(x_H;\delta,\overline{q})$. Hence (i) is proved.

Since R is continuous, the set of fixed points is closed. The above argument shows that the set is also bounded, which proves (ii). ■

Let us examine what the set Q looks like in our examples. In Example 1, a simple computation shows that $F_2(\overline{q}|x)=-f(\overline{q}|x)$, and therefore

$$\frac{F_{22}(\overline{q}|x)}{F_2(\overline{q}|x)}=\frac{\partial\ln F_2(\overline{q}|x)}{\partial x}=2(\overline{q}-x). \tag{14}$$

Since c is strictly convex, (14) implies that any $\overline{q}\leq\underline{x}$ belongs to Q. In Example 2, $F_2(\overline{q}|x)=\frac{-\overline{q}}{x^2}e^{-\frac{\overline{q}}{x}}$ and $F_{22}(\overline{q}|x)=\frac{-\overline{q}}{x^3}e^{-\frac{\overline{q}}{x}}(\frac{\overline{q}}{x}-2)$. If we additionally assume that $\underline{x}>0$, F_{22}/F_2 is well-defined whenever $\overline{q}\neq 0$. In this case, $\frac{F_{22}(\overline{q}|x)}{F_2(\overline{q}|x)}=\frac{1}{x}(\frac{\overline{q}}{x}-2)$, which is non-positive if and only if $\frac{\overline{q}}{x}\leq 2$. Therefore, $(0,2\underline{x}]\subseteq Q$. Thus, the set Q is non-empty in these examples so Proposition 4 holds non-vacuously. Note also that Proposition 4 applies to any $\delta\in(0,1)$, as long as $\overline{q}\in Q$. Thus the existence of a non-trivial equilibrium is guaranteed independent of firms' rate of time preference. On the other hand, the set of x_H's that can be supported in equilibrium depends on δ, as the comparative static result of the next section shows.

The approach of Proposition 4 is to fix $\overline{q}$ first, and ask whether an x_H exists so that along with $\overline{q}$ it forms a reputation equilibrium. While our analysis does not show existence for an *arbitrary* $\overline{q}$, $\overline{q}\in Q$ is a sufficient condition for the existence of such an x_H and, hence, for the existence of a reputation equilibrium. It is interesting to note that the set Q in our examples forms an interval (or a half line), which includes all sufficiently small $\overline{q}$. This implies two things. First, there exists a continuum of tolerance levels that are consistent with equilibrium. Second, reputation equilibria exist only when consumers are relatively tolerant.

Since the equilibrium x_H in the above proof is generated as a fixed-point of a continuous function and since a continuous function may have more than one fixed point, the nontrivial reputation equilibrium corresponding to some $\overline{q}\in Q$

need not be unique.[11] However, uniqueness can be shown in some situations. To see this, recall that a reputation equilibrium $(x, \overline{q})$, where $\overline{q}$ and $x > \underline{x}$, must satisfy the first-order condition (9). Rewriting it yields

$$c'(x) = \frac{-\delta F_2(\overline{q}|x)}{1-\delta[1-2F(\overline{q}|x)]}[x - c(x) - \underline{x}]. \tag{15}$$

By Lemma 3, (15) is also a sufficient condition for $(x, \overline{q})$ to be a reputation equilibrium.

Let us specialize (15) to the case of $c(x) = (x - \underline{x})^3$. If we limit attention to $x > \underline{x}$, (15) is rewritten as

$$3 = -\frac{\delta F_2(\overline{q}|x)}{1-\delta[1-2F(\overline{q}|x)]} \cdot \frac{1-(x-\underline{x})^2}{x-\underline{x}}. \tag{16}$$

Fix an $x > \underline{x}$ such that (16) holds. Then the derivative of the RHS of (16) at x is

$$3\frac{F_{22}(\overline{q}|x)}{F_2(\overline{q}|x)} - [5 - \frac{1}{(x-\underline{x})^2}]\frac{\delta F_2(\overline{q}|x)}{1-\delta[1-2F(\overline{q}|x)]} < 3\frac{F_{22}(\overline{q}|x)}{F_2(\overline{q}|x)} - 4\frac{\delta F_2(\overline{q}|x)}{1-\delta[1-2F(\overline{q}|x)]}, \tag{17}$$

where the inequality follows from Proposition 1(ii). If F is given as in either Example 1 or 2, $\frac{F_{22}(\overline{q}|x)}{F_2(\overline{q}|x)} < 0$ because $\overline{q} \in Q$. Since $F_2(\overline{q}|x)$ is bounded in these examples, the RHS of (17) is negative for all sufficiently small δ. Therefore, the RHS of (16) is decreasing in x whenever x is a solution to (16). This proves uniqueness of the solution to (16) and, hence, uniqueness of reputation equilibrium given $\overline{q} \in Q$.

The analysis shows that the multiplicity of reputation equilibrium for a fixed $\overline{q}$ is possible, for some parameterization, only for large δ. This is parallel to Folk-Theorem-type results, where a large δ is associated with multiple equilibria; see Fudenberg, Levine and Makin (1994). However, the equilibrium set in the Folk Theorem is usually a continuum, whereas in Examples 1 and 2, or, more generally, whenever c and F are smooth, the set of equilibria, which is the set of solutions to (15), need not be a continuum. Hence, we do not have nearly as many equilibria as the Folk Theorem would suggest. One should bear in mind, though, that we still have a continuum of $\overline{q}$ that are consistent with equilibrium, and that we focus on a *subset* of the repeated game equilibria (those we call reputation equilibria.)

5 Other Properties of Reputation Equilibria

The compactness result of Proposition 3(ii) ensures that, for given $\overline{q}$ and δ, the maximum of the set of equilibrium investment levels is attained; let us denote

[11] Recall that "nontrivial equilibrium" means $x_H > \underline{x}$. Also, uniqueness is under a *given* $\overline{q} \in Q$. As usual, different $\overline{q} \in Q$ give rise to different equilibria.

it by $\overline{x}(\delta,\overline{q})$. Note that $\overline{x}(\delta,\overline{q})$ is not necessarily continuous in its arguments. $\overline{x}(\delta,\overline{q})$, in a sense, reflects consumers' ability to discipline the firms by the threat of turnover, because this is the maximum investment in quality that can be sustained in equilibrium.

In this section, we examine the local behavior of the function $\overline{x}(\delta,\overline{q})$. We start with examination of the effect of the firms' rate of time preference on the equilibrium investment in quality. The following statement confirms the intuition that firms who care more about the future will try harder to keep a good reputation and, therefore, will invest more in quality.

Proposition 5 *Fix a $\overline{q}\in Q$. Then $\overline{x}(\delta,\overline{q})$ is increasing in δ.*

Proof. We first show that $R(x_H;\delta,\overline{q})$ is increasing in δ for (arbitrarily) fixed x_H and $\overline{q}\in Q$. To that end we show that for any $\delta\in(0,1)$ and $\delta'>\delta$,

$$R(x_H;\delta',\overline{q})\geq R(x_H;\delta,\overline{q}). \tag{18}$$

Let us write $y=R(x_H;\delta,\overline{q})$. Then (18) immediately follows if $y=\underline{x}$, so let us assume otherwise. Then the first-order condition implies

$$-(1-\delta)c'(y)+H_1(y;\delta,\overline{q},x_H)[\hat{x}_H-c(x_H)]=0. \tag{19}$$

By (7), it is easy to verify that $\frac{H_1(y;\delta,\overline{q},x_H)}{1-\delta}$ is increasing in δ for any x and x_H. Therefore (19) implies

$$-(1-\delta')c'(y)+H_1(y;\delta',\overline{q},x_H)[\hat{x}_H-c(x_H)]\geq 0 \tag{20}$$

since $\delta'>\delta$. Since $v(x;\delta',\overline{q},x_H)\to-\infty$ as $x\to\infty$, (20) implies the existence of a $y'\geq y$ which achieves a local maximum of $v(\cdot;\delta',\overline{q},x_H)$. And, by Lemma 3, y' is also the global maximum. Hence $R(x_H;\delta',\overline{q})\geq y=R(x_H;\delta,\overline{q})$.

In particular,

$$R(\overline{x}(\delta,\overline{q});\delta',\overline{q})\geq R(\overline{x}(\delta,\overline{q});\delta,\overline{q})=\overline{x}(\delta,\overline{q}),$$

where the last equality follows from the definition of $\overline{x}(\delta,\overline{q})$. Since $R(x_H;\delta',\overline{q})=\underline{x}$ for all sufficiently large x_H, there exists an $x_H\geq\overline{x}(\delta,\overline{q})$ such that $R(x_H;\delta',\overline{q})=x_H$. Hence $\overline{x}(\delta',\overline{q})\geq\overline{x}(\delta,\overline{q})$. ■

The monotonicity result obtained here holds only for the equilibrium corresponding to the largest investment level. Indeed, we have proven our result by showing that R shifts upwards when δ increases and, therefore, that the maximum (and the minimum) fixed point of R increases. This is analogous to monotonicity results reported by Milgrom and Roberts (1990) for supermodular games. On the other hand, if we fix δ and $\overline{q}$, if R crosses the 45^0 line at some point in $(\underline{x},\overline{x}(\delta,\overline{q}))$ and if R has slope >1 at that point, the equilibrium investment *decreases* in δ. Thus, monotonicity is only with respect to the maximal equilibrium.

Next, we consider the behavior of $\overline{x}(\delta, \overline{q})$ with respect to $\overline{q}$. Since $\overline{q}$ is endogenous to the model, this should *not* be construed as a comparative static exercise. Rather, what we are doing is comparing different equilibria. For example, the same product may be traded in geographically separated markets, which, somehow, have settled on different equilibrium tolerance levels. To make the comparison between equilibria possible, we invoke the following additional assumption. This assumption concerns the stochastic link between x and q only, and not the cost function or the discount factor.

Assumption 2 For any $q \in Q$ and $x > \underline{x}$,

$$\frac{f_2(q|x)}{f(q|x)} > \frac{F_2(q|x)}{F(q|x)}. \tag{21}$$

Assumption 2 is satisfied for our two examples.

Proposition 6 *Assume Assumption 2 holds, and fix a $\overline{q} \in Q$. Then there exists an $\varepsilon > 0$ and a $\underline{\delta}$ so that for any $\delta > \underline{\delta}$, $\overline{x}(\delta, q)$ is decreasing in q on $Q_\varepsilon \equiv Q \cap [\overline{q} - \varepsilon, \overline{q} + \varepsilon]$.*

Proof. Fix a $\overline{q} \in Q$. Let $\overline{x} \in X$ be such that $\overline{x} = c(\overline{x}) + \underline{x}$. Assumption 0(ii) ensures that $\overline{x}$ uniquely exists. By Proposition 1(ii), $\overline{x}(\delta, q) \leq \overline{x}$ always holds. Choose an $\varepsilon > 0$ sufficiently small that $f_2(q|x)F(q|x) - f(q|x)F_2(q|x)$ is bounded away from zero on $Q_\varepsilon \times [\underline{x}, \overline{x}]$. We can do this because, by (21), $f_2(q|x)F(q|x) - f(q|x)F_2(q|x)$ is continuous and positive on the compact set $Q_\varepsilon \times [\underline{x}, \overline{x}]$.

Also since $f_2(q|x)$ is continuous and, hence, bounded on $Q_\varepsilon \times [\underline{x}, \overline{x}]$, there exists a $\underline{\delta}$ such that $\delta > \underline{\delta}$ implies

$$(1-\delta) f_2(q|x) + 2\delta [f_2(q|x)F(q|x) - f(q|x)F_2(q|x)] > 0 \tag{22}$$

for any $q \in [\overline{q} - \varepsilon, \overline{q} + \varepsilon]$ and any $x \in [\underline{x}, \overline{x}]$.

Consider some $\delta > \underline{\delta}$, $q^1 \in Q_\varepsilon$ and $q^2 \in Q_\varepsilon$, where $q^1 > q^2$. Let us write $\overline{x}^i = \overline{x}(q^i, \delta)$. We show that $\overline{x}^2 \geq \overline{x}^1$. Since $\overline{x}^1 > \underline{x}$ by Proposition 4(i), the first-order condition implies

$$-(1-\delta)c'(\overline{x}^1) + H_1(\overline{x}^1; \delta, q^1, \overline{x}^1)[\overline{x}^1 - c(\overline{x}^1) - \underline{x}] = 0. \tag{23}$$

Proposition 1(ii) implies $\overline{x}^1 - c(\overline{x}^1) - \underline{x} \geq 0$. Since (22) implies that $H_1(\overline{x}^1; \delta, q, \overline{x}^1)$ is decreasing in q on Q_ε (recall that $\overline{x}^1 \leq \overline{x}$), (23) implies

$$-(1-\delta)c'(\overline{x}^1) + H_1(\overline{x}^1; \delta, q^2, \overline{x}^1)[\overline{x}^1 - c(\overline{x}^1) - \underline{x}] \geq 0.$$

Thus the same argument used in the proof of Proposition 5 demonstrates that $R(\overline{x}^1; \delta, q^2) \geq \overline{x}^1$. Therefore there exists a $y \geq \overline{x}^1$ so that $R(y; \delta, q^2) = y$, which proves $\overline{x}^2 \geq \overline{x}^1$. ■

The role of Assumption 2 is as follows. Assumption 2 implies that for a fixed x, $\frac{F_2(q|x)}{F(q|x)}$ is increasing in q. Hence, the marginal return to investment in product quality (i.e., the added probability of remaining the HQ firm as the investment in product quality is marginally increased) is decreasing in q. Therefore, additional investment in product quality becomes less desirable the more stringent the tolerance level is (the bigger is q.) In this situation, setting a more stringent tolerance level simply discourages firms from investing in quality. So Assumption 2 is critical to the result. However, our two examples, which are fairly natural, do satisfy it.

The second assumption that Proposition 6 relies upon is patience, which implies that firms care enough about future profits and are, hence, reluctant to lose them. This assumption is only natural since the essence of the disciplinary mechanism is to deny future profit, so this must indeed be a sufficiently strong deterrent to make firms want to sustain their investment in quality.

From Proposition 6 we can infer how the probability of turnover varies as we go across equilibria and how it relates to the period profit of the HQ firm. If we are on the region where Proposition 6 holds, then an equilibrium with a smaller $\overline{q}$ has a greater x and, therefore, a smaller probability of turnover. Thus, on this region, the equilibrium investment is negatively related to the probability of turnover. Furthermore, as we show in the next section, the period profit of the HQ firm is locally increasing in x. Thus, our model predicts negative correlation between profits and the probability of losing customers. As stated in the introduction, this confirms the intuition that high-profit and patient firms are reluctant to lose their profits and, hence, will invest enough in quality to make the probability of losing customers small.

However, Proposition 6 only establishes *local* monotonicity. Namely, it confines attention to small changes in the tolerance level, $\overline{q}$, and to patient firms. We now show, by means of a counter example, that, these qualifications cannot be removed.

Let us fix some $\overline{q} \in Q$ and x_H. Then, specializing (15) to Example 2 with $\alpha = 3$ and $\underline{x} = 1$ yields

$$\frac{\delta z e^{-z}}{1 + \delta - 2\delta e^{-z}} = 3\frac{x(x-1)}{2-x}, \tag{24}$$

where $z = \frac{\overline{q}}{x}$. The LHS of (24), as a function of z, is increasing for all sufficiently small z, and is decreasing for all sufficiently large z. Thus the effect of $\overline{q}$ on x depends on whether the corresponding z is small or large, which is equivalent to whether $\overline{q}$ itself is small or large. Indeed, the local monotonicity of Proposition 6 corresponds to the case of a large $\overline{q}$, i.e., where the LHS of (24) is decreasing. If, on the other hand, we start from a sufficiently small $\overline{q}$, then the LHS of (24) is increasing in $\overline{q}$. Therefore, if we increase $\overline{q}$, x must *increase* in order to restore the equality. Thus, $\overline{q}$ and x are positively related in this region.

More importantly, if (24) is to hold x must increase less than proportionately when $\overline{q}$ increases , i.e., z must increase. Hence if we consider a case like Example 2 in which the equilibrium turnover probability (recall that this is $F(\overline{q}|x)$) is positively related to z, there is a region where the investment and the probability of turnover are *positively* correlated. And, likewise, profits and probability of turnover are positively correlated in this region.

6 Efficiency

So far we have pre-specified a $\overline{q}$ and asked whether there exists an equilibrium x_H under this $\overline{q}$. Conversely, we can pre-specify an x_H and ask whether a $\overline{q}$ exists so that $(x_H, \overline{q})$ is a reputation equilibrium. In particular, can the optimal x or an x that is nearly optimal be an equilibrium under some $\overline{q}$? The optimum in our environment is the x that maximizes the difference between the expected benefit to consumers, x, and the cost $c(x)$. This optimum, call it x^o, is characterized by $c'(x^o) = 1$.

This section proves a negative result, stating that x^o (and any greater x) cannot be implemented. The result suggests limits to consumers' ability to discipline the firms. We show this for the multiplicative case where there exists a random variable z, with mean 1, so that $q = zx$. We assume z has a twice continuously differentiable c.d.f., $G(z)$, and a p.d.f., $g(z)$. Then, $F(q|x) = G(\frac{q}{x})$ and $F_2(q|x) = -\frac{q}{x^2}g(\frac{q}{x})$. Note that Example 2 is a special instance of the multiplicative case once we set $G(z) = 1 - e^{-z}$.

Let us specialize (15) to the multiplicative case and characterize an equilibrium via:

$$\frac{\delta \frac{\overline{q}}{x} g(\frac{\overline{q}}{x})}{1 - \delta[1 - 2G(\frac{\overline{q}}{x})]} = \frac{xc'(x)}{x - c(x) - \underline{x}}.$$

Or,

$$\frac{\delta z g(z)}{1 - \delta[1 - 2G(z)]} = \frac{xc'(x)}{x - c(x) - \underline{x}}, \tag{25}$$

where $z = \overline{q}x$.

Proposition 7 *Assume $g(z) > zg'(z)$ and $\lim_{z \to 0} zg(z) = 0$. Then, there exists an $\widetilde{x}$, $\widetilde{x} < x^o$, so that no x, $x > \widetilde{x}$, is implementable as an equilibrium.*

Proof. Let us first show that the LHS of (25) is smaller than 1 for any z and any δ. Note that $\frac{\delta z g(z)}{1-\delta[1-2G(z)]} < 1$ if and only if $\delta z g(z) < 1 - \delta[1 - 2G(z)]$. This inequality holds at $z = 0$. And, if we take derivatives on both sides, $\delta[g(z) + zg'(z)] < 2\delta g(z) \Leftrightarrow zg'(z) < g(z)$, which is what we assume.

Second, at the social optimum $c'(x) = 1$ and $\frac{x}{x-c(x)-\underline{x}} > 1$. Therefore the equality in (25) cannot hold when we set $x = x^o$. By continuity it cannot hold if $x > \widetilde{x}$ for some $\widetilde{x} < x^o$. ■

Since $g'(z) < 0$ in Example 2, the premise underlying Proposition 7 clearly holds and, therefore, we have shown that x^o cannot be implemented in the example. Note that Proposition 7 holds independently of the value of δ. Hence, the inefficiency of equilibria persists even as we let $\delta \to 1$.

Given Proposition 7, any equilibrium investment level is suboptimal. Thus, given the convexity of c, a bigger equilibrium x means bigger social surplus and bigger HQ's period profit (which equals the social surplus.) Thus, as argued in Section 5, equilibria with higher x are equilibria with higher firm profits.

References

[1] Akerlof, George The Market for Lemons: Qualitative Uncertainty and the Market Mechanism, *Quarterly Journal of Economics* 84 (1970), 488-500.

[2] Abreu, Dilip, David Pearce and Ennio Stacchetti, Optimal Cartel Equilibria with Imperfect Monitoring, *Journal of Economic Theory* 39 (1986), 251-269.

[3] Abreu, Dilip, David Pearce and Ennio Stacchetti, Toward a Theory of Discounted Repeated Games with Imperfect Monitoring, *Econometrica* 58 (1990), 1041-1063.

[4] Chalk, Andrew, Market Forces and Aircraft Safety: The Case of the DC-10 *Economic Inquiry* 24 (1986), 43-60.

[5] Fudenberg, Drew, David Levine and Eric Maskin, The Folk Theorem with imperfect public information, *Econometrica* 62 (1994), 997-1039.

[6] Ippolito, Richard A., Consumer Reaction to Measures of Poor Quality - Evidence from the Mutual Fund Industry, *Journal of Law and Economics* 35 (1992), 45-70.

[7] Klein, Benjamin and Keith B. Leffler, The Role of Market Forces in Assuring Contractual Performance, *Journal of Political Economy* 89 (1981), 615-641.

[8] Milgrom, Paul and John Roberts, Rationalizability, Learning, and Equilibrium in Games with Strategic Complementarities, *Econometrica* 58 (1990), 1255-1277.

[9] Mitchell, Mark, L. The Impact of External Parties on Brand-Name Capital: The 1982 Tylenol Poisoning and Subsequent Cases, *Economic Inquiry* 27 (1989), 601-618.

[10] Mitchell, Mark, L. and Michael T. Maloney, Crisis in the Cockpit? The Role of Market Forces in Promoting Airline Safety, *Journal of Law and Economics* 32 (1989), 329-355.

[11] Rogerson, William P., The First-Order Approach to Principal-Agent Problems, *Econometrica* 53 (1985), 1357-1367.

[12] Shapiro, Carl, Premiums for High-Quality Products as Returns to Reputation, *Quarterly Journal of Economics* 98 (1983), 659-680.