DOES IT PAY TO MAINTAIN A REPUTATION?

Quality Incentives in Financial Markets

by

Philip H. Dybvig* and Chester S. Spatt**

December 1985

* Yale University

** Carnegie-Mellon University

1. Introduction

In financial transactions, the quality of behavior often cannot be explicitly specified or guaranteed in a legal sense. For example, compliance with an oral agreement between brokers on an exchange floor is not enforceable in court. Yet behavior is often above the lowest quality. For example, floor traders do honor oral agreements, even when prices have moved against them. We attribute this phenomenon to a <u>reputation effect</u>: providers of more than the minimal passable quality level are concerned about losing their reputation for quality and consequently losing future rents.[1] In this paper, we explore how reputation plays an important role in many situations in finance. Our analysis of the impact of reputation on financial markets is based on the features of an abstract rational model of reputation featuring a monopolist supplier and competitive customers. The model has multiple equilibria: just as it is possible to have an incentive to maintain a good reputation, it is also possible to have an incentive to maintain a bad one. We obtain an explicit functional form for the best equilibrium and in the model show that the best potential quality improves with the degree of information.

Before analyzing the formal model, we start with a general discussion of reputation in financial markets. Though reputation in financial markets is generally more complicated than reputation in our formal model, a number of features of our model should apply to other markets with reputation effects. Here are some important properties of our model.

1. Reputation is important when direct legal enforcement of quality is not feasible.

2. For reputation to be effective, there must be some rents that can be threatened.

3. Reputation involves multiple equilibria. Having a good reputation can be self-reinforcing since losing the good reputation would be costly. Having a bad reputation can be self-reinforcing since trying to change the bad reputation would be too costly or even futile.

4.   The more current the available information, the higher the quality.
      Also, the smaller the interest rate the higher the quality, since
      lowering the interest rate has the same effect as speeding the
      arrival of information.

Prior to examining reputation in specific financial contexts, we will discuss

an interpretation of reputation as a solution to an agency problem.

Reputation as a Solution to an Agency Problem.

The reputation problem is closely related to the agency problems that

have received so much attention in finance, following Jensen and Meckling

[1976] in corporate finance and Ross [1973, 1974] in investments.   In an

agency context, an <u>agent</u> is providing some sort of service or good to a

<u>principal</u>.   The <u>agency problem</u> is to give the agent an incentive to do well

for the principal without destroying risk-sharing.   The traditional approach

to minimizing the agency problem uses an <u>incentive contract</u> under which the

agent's compensation depends on the realized quality of the good or service.

By contrast, with <u>reputation</u>, the agent provides good quality because

otherwise future business would be lost.   Reputation is most important when

the traditional contracting solution is infeasible or impractical, perhaps

because legal enforcement is too costly or quality cannot be legally confirmed

in court.   The presence of reputation suggests the alternative in which more

effort implies a better reputation and consequently greater rewards later.   To

bring reputation into play requires an extension of the original agency

problem to multiple periods: providing low quality now results in later loss

of business.

In the contracting approach, the agency problem can be avoided completely

(and a Pareto optimum achieved) if effort is observed directly (even if with a

lag).   Along the same lines, the agency problem is avoided if there is no

uncertainty in how effort maps into outcomes. This is not necessarily true for the reputation approach, as illustrated by our model in which the action is observable with a lag. When there is uncertainty, in general neither reputation nor contracting can completely avoid the agency problem, and the contracting and reputation solutions complement each other. Legal enforcement of contract provisions would provide incentives that would keep the agent from grossly underproviding effort. Some bad incentives would remain, since the agent knows that the principal will not find it profitable to take the agent into court for small violations. Reputation would then minimize the agent's incentive to violate the contract in small ways.

The problem of reputation is minimized in markets for goods whose characteristics are cheaply and accurately verifiable at the time of purchase. In these markets, the supplier of the goods is irrelevant, since it is easy to verify the quality of the goods. Yet there is a role for reputation even in commodity markets. For example, having a list of standard suppliers of gold bars minimizes the potential for fraud.

Investments and Security Markets

Trust is essential in executing many financial transactions. Brokers execute orders verbally on the floor of securities exchanges and customers place orders by telephone, yet commitments are rarely breached. The underlying incentive to default would seem strong, since typically one party or the other experiences an adverse price movement. In this market the quality of performance is high because of the short lag in observing quality, or equivalently, because the interest rate times the size of the time interval is small. Reputation is especially important in brokerage because the reputation of an individual broker may affect an entire exchange.

Reputation depends on a threat to future rents, whether by total exclusion from trade (as in our model) or by partial exclusion (market access is limited or on unfavorable terms). Reputation may not provide strong incentives in competitive financial markets in which firms do not earn large rents. Though one could expect that rents to sunk costs would resolve this need for economic rents (see, for example, Toaff [1979] and Diamond [1984]), even sunk costs may be saleable with the bad reputation attributed to the seller rather than to the asset. For example, major stock exchanges have a limited number of seats, and it is necessary to buy a seat (apparently a sunk cost) to do business on the exchange. However, if a member's reputation is tarnished to the point where it is no longer profitable to do business, the member can always sell the seat. So long as there is a competitive market for "business opportunities," there is some difficulty with assigning rents to sunk costs. (However, see Kreps [1984] for an interesting example in which the reputation of a business is transferred upon sale.) The sunk costs of building a customer base could work, if the customer base is not so easily transferred. To summarize the conflict, there is a distinction between the reputation of a business and the reputation of its current managers or owners.

Reputation considerations are central to a variety of financial services such as market advice letters, written personal financial plans, portfolio management, and computer programs to pick securities. In all of these instances, quality is observed with a lag. Since the amount of time it takes to discern quality of portfolio performance is lengthy in financial markets, reputation is unlikely to have much substantial basis. Inferior performers should survive for a long time, as is consistent with the empirical evidence. This suggests an important role for contracting, although it seems that

neither contracting nor reputation is particularly effective in this context. Perhaps solution lies in having more sophisticated customers who can intelligently evaluate the inputs to the investment process. For example, a reputation for prudence (e.g., not engaging in excessively risky investments) is valuable for parties with fiduciary responsibilities.

Debt Repudiation and Expropriation of Assets

The possibility of debt repudiation gives rise to another role for reputation. Since it may not be practical to seek legal redress for the repudiation of a small loan, reputation might be most important for small transactions. For example, a borrower's incentive to repay a loan may depend more on the borrower's desire to maintain a credit rating than on the ability of the lender to profitably collect via a lawsuit. Similarly, if the lender does sue a small borrower, it is most probably to maintain a reputation for being a "tough guy," rather than to make a profit net of cost.

By contrast, legal enforcement of loans to large corporations is probably sufficient to ensure compliance. Of course, there are circumstances in which even very large transactions would not be legally enforceable (thereby leading to reliance on reputation) as illustrated by the limits upon enforcing loans to foreign countries. For individuals, the laws against indentured servitude and debtor's prisons can preclude sufficient recovery even if default is cheaply verifiable and court costs are low.

Using collateral to substitute for reputation is only partially effective. When reputation is needed because legal verification is costly or impossible, giving the other party collateral simply switches the reputation problem to the opposite side of the transaction. This occurs because the party providing collateral must now trust the other party to return the

collateral. Ideally, collateral shifts the need for responsible behavior to the party with the stronger incentive to maintain a reputation. An interesting example of this occurs in the market for repurchase agreements collateralized by government securities. In practice, the amount that can be borrowed by a repurchase agreement for a given amount of collateral depends on the relative reputations of the two parties. In some cases, legal verification is possible but legal redress is impossible because of other constraints, and collateral is effective. For example, in loans to individuals who could declare bankruptcy and renege on a loan, collateral can avoid reliance on reputation.

The issue of debt repudiation is important in international finance. Indeed, Eaton and Gersovitz [1981] and Sachs [1983] model nations as acting as if they are indefinitely denied credit once they default without deriving the credit denial policies of banks. One way to describe the endogenous behavior of banks is by an explicit game in which the denial of credit to past defaulters emerges as part of the equilibrium. Useful approaches to these issues include modeling the banking sector as ex post competitive (as in the model of credit reputation and learning in Spatt [1983]) or one in which an individual bank develops a long-term relationship with the prospective borrower (e.g., Stiglitz and Weiss [1983]).

Debt repudiation is a special case of a more general phenomenon of expropriation. Reputation plays an important role in a country's choice of whether to expropriate property from other countries or foreign-owned businesses. Such expropriation can be in the form of direct nationalization, or in the form of a change in tax laws. The foreign investor must earn a competitive return taking into account the anticipated level of expropriation. If the country expropriates property it receives the rents

from the business activity but loses the benefits of future investment. The greater the weight the country places on the future, the greater the incentive to refrain from expropriating property. An important ambiguity is the extent to which reputation is associated with the country or with the particular government in power.

Corporate issues: dividend policy and managerial incentives

Some firms have reputations for paying dividends for many decades and others have reputations for frequent dividend increases. The choice of firm dividend policies is certainly influenced by the informational effects of dividend announcements. Our model suggests that dividends can signal the quality of future non-financial decisions even without learning. The multiplicity of reputation equilibria would make this an unsatisfying (even if correct) explanation of dividend policy. In any case, informational considerations should complement any reputation-based analysis of corporate dividend policy (see Bhattacharya [1979]).

The potential agency problem in the provision of managerial incentives has received considerable attention. If salaries do not depend on performance, managers have an incentive to shirk. Since reputation can provide an effective solution to multi-period agency problems, the manager's reputation could play an important role in provision of incentives. It is an open question just how important the agency problem is in this context. Holmstrom [1983] suggests that because effort is monitored directly, the main problem in managerial compensation is one of a manager's risk aversion to discovering his own type rather than underprovision of managerial effort. In any case, it is clear that managerial incentives, while perhaps related to reputation issues, also involve a number of other features specific to labor markets.

Overview of Our Formal Model and Related Literature

Our formal model abstracts from the complexity of financial markets and considers instead a simple market for a good. An infinitely lived monopolist supplier produces the good with a constant marginal cost that depends on the quality level selected. On the demand side, there is a different set of competitive customers in each time period, and all buyers have a common reservation price equal to the anticipated quality. Customers live one period, and in a given time period they differ from each other only in the length of the time lag after which they learn the past qualities offered by the supplier. All customers know the current price and all past prices.

We study Nash equilibria in our model, and we restrict ourselves to pure strategies in this exposition, since our investigation of mixed strategies did not yield anything interesting. (It can be shown that our equilibria are in fact perfect equilibria. The natural definition of perfect equilibrium for our model is the concept of sequential equilibrium of Kreps and Wilson [1982a]. Since our theorems are true whether we use Nash equilibrium or perfect equilibrium, we use the simpler Nash equilibrium notion for exposition.)

Since our basic model has many equilibria, we focus upon the equilibrium which is most profitable for the firm. This assumption is consistent with the intuition that the firm is a dominant player in the model, and it allows us to analyze systematically what happens when the model's parameters vary. In this most profitable equilibrium, the supplier offers the same quality in all time periods. We have a simple explicit formula for this equilibrium quality level, which depends on the discount rate, how fast demand is growing, the cost function, and consumers' information lags.[2]

The equilibrium quality level (in this most profitable equilibrium) is the perfect information quality level when that can be achieved in a Nash equilibrium. Otherwise, the quality level will make the firm indifferent between maintaining that quality level forever, on the one hand, and cheating by producing at the lowest quality level and selling to customers until they find out (after which they don't buy), on the other hand. Because only the perfect information quality level is efficient, we can use our explicit solution to characterize when the reputation effect is sufficiently strong to prevent the imperfect information from yielding an inefficient outcome.

Reputation does not require learning. For example, in our model there are no unobservable characteristics for agents to learn. It may seem that the multiplicity of equilibria in our model is an artifact of the infinite horizon, since infinite horizon models typically have many equilibria. However, many equilibria arise in finite horizon models with exogenous uncertainty, but in a slightly disguised form. In these models, the solution depends upon the particular small degree of uncertainty introduced into the payoff functions. To the extent that this uncertainty is unobservable in practice, the finite horizon models in a sense have multiple equilibria. In fact, "the" folk theorem from game theory asserts that the set of equilibria obtained under small perturbations of a finite horizon model is approximately the same as the set of equilibria in an infinite horizon model. This discussion explains the apparent difference between our model and those of Kreps and Wilson [1982b], Milgrom and Roberts [1982] and Kreps, Milgrom, Roberts and Wilson [1982]. In a sense, these papers implicitly include partial proofs of the folk theorem in a reputation context (besides establishing that an infinite horizon is not required for reputation effects in a perfect equilibrium).

Several other papers have discussed the reputation problem in the context of product quality. (See, for example, Allen [1984], Heal [1976], Klein and Leffler [1981], Rogerson [1983], Shapiro [1982], and Toaff [1979].) Our results differ: we obtain multiple equilibria (Allen [1984] also discusses multiple equilibria) and we show that product quality improves as consumers become better informed. One reason that we get multiple equilibria when others don't is because several of the other papers assume that reputation is important (by putting lagged quality in the demand function), rather than deriving it. Just as writing down a demand for money function doesn't explain money demand, putting lagged quality in the product demand function doesn't explain reputation. Some of the papers have a single equilibrium because they implicitly assume that there is only one. The presence of multiple equilibria is important since it means otherwise identical individuals can have different reputations.

Abstract theoretical models of reputation outside the product quality context include Kreps and Wilson [1982b], and Milgrom and Roberts [1982] on the chainstore paradox, Kreps, Milgrom, Roberts, and Wilson [1982] on the repeated prisoner's dilemma, and Radner [1981] on the repeated principal-agent problem. All of these except Radner's introduce uncertainty to prevent the "folding-back" that can preclude successful reputation in a finite horizon model. Radner [1981] uses a concept of approximate equilibrium to accomplish this. In contrast, by working with an infinite horizon we can model reputation without having any exogenous uncertainty in our model or the necessity of working with mixed strategies. Therefore, we can keep the model simple enough to solve explicitly in spite of our general lag structure of consumer information. This allows us to show how product quality depends on the level of customer awareness of past quality.

In Section 2 of the paper we set up our model and derive the most profitable equilibrium path. In Section 3 we discuss the comparative statics properties of the solution and we present concluding remarks in Section 4.

2. The Model and its Solution

At each time period t, the firm faces a different group of buyers. Each buyer has the same payoff function:

$$u_t(d_t, p_t, q_t) = d_t(q_t - p_t),$$

where $d_t \in [0,1]$ is the quantity demanded, p_t is the price, and q_t is the quality measured in the same units as the customers' common reservation price. At t, there is a continuum of agents, numbering g^t in all, where g is a positive constant (which allows us to study exponentially rising or declining demand).

The sequence of events in period t is indicated by the tree shown in Figure 1. First, the firm chooses and announces publicly the price it is charging, and chooses what quality to produce. Second, agents choose their demands. In the diagram, we have drawn the information sets (dashed lines around nodes which cannot be distinguished by the agent making the choice) to represent a special case in which all buyers know past quality levels with a one-period lag, but not the current quality level. In our model, we allow for the more general case in which customers learn qualities with differing positive lags: a buyer with a short lag has better (more current) information about quality than an agent with a long lag. We will always assume that all buyers know today's price and previous prices. Drawing the game tree for the

general case is hopelessly complicated: for only 2 types of buyers and 2 periods there are 256 terminal nodes.[3] However, we can describe the information structure explicitly by saying that a buyer with k-period lag can distinguish two nodes for that buyer at t (i.e., the nodes are in different information sets) if and only if those nodes correspond to different prices $p_{t'}$ for $t' \leq t$ or different qualities $q_{t'}$ for $t' \leq t-k$. We assume that the firm has full information so the firm knows all past prices, qualities, and demand by various types. (If firms cannot observe past demands by types or even past aggregate demand, our solution is still an equilibrium, because the firm's strategy in our solution is independent of past demands.) We do not permit the monopolist to distinguish the lag types of customers at the time of purchase.

The fraction of all buyers having an information lag of k periods in any time period is n_k. We will let $n_\infty = 1 - \sum_{k=1}^{\infty} n_k$ be the fraction of all buyers having an infinite information lag. (At a specific date t, those agents for which $k > t$ also have no past quality information, since, at t, their lags would allow them to know only quality levels from before the initial period 0.) Note that our assumption that buyers cannot observe the demands of other agents ensures that buyers can learn the information of other agents only through the firm's choice of price and quality, if at all. Also, recall that the agents of each lag type are different agents in every period. This assures that our equilibrium will not rely on any "strategic" behavior (e.g., threats) by the buyers: since the buyers live only one period, they have no stake in what happens subsequently and will always behave myopically based on their expectations.

Firm profits are given by the present value of revenues less costs:

$$\pi = \sum_{t=0}^{\infty} (d_t p_t - d_t c(q_t)) \rho^t g^t,$$

where d_t is total demand (per buyer) at t, $c(q)$ is the unit production cost of producing each unit of quality q, q_t is the actual quality at t, and ρ is a discount factor $(\rho = 1/(1+r)$ where r is the interest rate). We assume $\rho g < 1$ (to rule out infinite profit) and we also make the following assumptions about the cost function.

$c(0) = 0$

$c'(0) = 0, \quad c'(\infty) = \infty$

$c'(q) > 0 \qquad\qquad\qquad\qquad$ all $q > 0$

$c''(q) > 0 \qquad\qquad\qquad\qquad$ all $q > 0$

$AC(q) \equiv c(q)/q$

$MC(q) \equiv c'(q)$

These are ordinary sorts of convexity and Inada conditions.[4,5]

Now that we have specified buyers' preferences, the firm's profit function, and the form of the game tree, we are ready to write down notation for strategies so we can proceed to define pure strategy Nash equilibrium. The firm's strategy is a choice of price and quality for each t. The choice may depend on all previous prices, qualities, and demands by lag k, since the firm has perfect information. The firm's strategy is therefore represented by functions

$$q_t^S(Q_{t-1}, P_{t-1}, D_{t-1}) \quad \text{and}$$

$$p_t^S(Q_{t-1}, P_{t-1}, D_{t-1}),$$

for all t, where Q_t represents quality levels at all times up to and including t, P_t prices at all times and up to and including t, and D_t demands at all times up to and including t. Customers of type k at t know all current and past prices, know quality with a k-period lag, and know nothing of past demands, so that their strategy is written as

$$d_t^{sk}(Q_{t-k}, P_t).$$

For each specification of strategies for all players, the <u>realization</u> is the outcome which will arise if all agents use those strategies. It is defined recursively by

$$\hat{q}_t(<q_t^S>, <p_t^S>, <d_t^{sk}>) = q_t^S(\hat{Q}_{t-1}, \hat{P}_{t-1}, \hat{D}_{t-1})$$

$$\hat{p}_t(<q_t^S>, <p_t^S>, <d_t^{sk}>) = p_t^S(\hat{Q}_{t-1}, \hat{P}_{t-1}, \hat{D}_{t-1})$$

$$\hat{d}_t^k(<q_t^S>, <p_t^S>, <d_t^{sk}>) = d_t^{sk}(\hat{Q}_{t-k}, \hat{P}_t).$$

Note that the recursion "gets started" at t = 0 since $\hat{Q}_{-1}$, $\hat{P}_{-1}$, and $\hat{D}_{-1}$ are empty strings, making q_0^S and p_0^S constants.

Now we are ready to define pure strategy Nash equilibrium in our model:

Definition: A list of strategies $(<q_t^S>, <p_t^S>, \text{and} <d_t^{sk}>)$ is called a pure strategy Nash equilibrium if

(a) $\quad (\forall <\tilde{q}_t^S>, <\tilde{p}_t^S>) \sum_{t=0}^{\infty} (\hat{p}_t - c(\hat{q}_t))\hat{d}_t \rho^t g^t \geq \sum_{t=0}^{\infty} (\tilde{p}_t - c(\tilde{q}_t))\tilde{d}_t \rho^t g^t$

where ^'s indicate the realization given strategies $(\langle q_t^s \rangle, \langle p_t^s \rangle, \langle d_t^{sk} \rangle)$, and ~'s indicate the realization given $(\langle \tilde{q}_t^s \rangle, \langle \tilde{p}_t^s \rangle, \langle d_t^{sk} \rangle)$,

and   (b) $(\forall \tilde{t}, \tilde{k}, \langle \tilde{d}_t^{sk} \rangle) \hat{d}_{\tilde{t}}^{\tilde{k}} (\hat{q}_{\tilde{t}} - \hat{p}_{\tilde{t}}) \geq \tilde{d}_{\tilde{t}}^{\tilde{k}} (\tilde{q}_{\tilde{t}} - \tilde{p}_{\tilde{t}})$

where   ^'s   again   indicate   the   realization   given   strategies $(\langle q_t^s \rangle, \langle p_t^s \rangle, \langle d_t^{sk} \rangle)$,   but   ~'s   now   indicate   the   realization   given $(\langle q_t^s \rangle, \langle p_t^s \rangle, \langle \tilde{d}_t^{sk} \rangle)$ where $\tilde{d}_t^{sk} \equiv d_t^{sk}$ whenever $t \neq \tilde{t}$ or $k \neq \tilde{k}$.

As we have set up our model so far, it has many equilibria. Theorem 1 gives a class of equilibria having constant quality and price over time.

Theorem 1: Let $\bar{q}$ be any quality level with $0 \leq \bar{q} \leq AC^{-1}(\sum_{k=1}^{\infty} n_k (\rho g)^k)$. (By our assumptions on the cost function, this means $0 \leq c(\bar{q}) \leq \bar{q} \sum_{k=1}^{\infty} n_k (\rho g)^k$.) Also, let $q_t^s(\cdot) = \bar{q}$ whenever all price and quantity arguments are $\bar{q}$ and let $q_t^s(\cdot) = 0$ otherwise, let $p_t^s(\cdot) = \bar{q}$ always, and let $d_t^{sk}(\cdot) = 1$ when all arguments are $\bar{q}$ and let $d_t^{sk}(\cdot) = 0$ otherwise. Then $(\langle q_t^s \rangle, \langle p_t^s \rangle, \langle d_t^{sk} \rangle)$ is a Nash equilibrium for which the realizations are all $\hat{q}_t = \bar{q}$, $\hat{p}_t = \bar{p}$, and $\hat{d}_t^{sk} = 1$.

Proof: The realizations are clearly as stated. Condition (b) of the definition of Nash equilibrium is satisfied since any buyer with lag $\tilde{k}$ at $\tilde{t}$ cannot change the price or quality she faces by changing her strategy, and is indifferent about how much to demand given that $\hat{p}_{\tilde{t}} = \hat{q}_{\tilde{t}} = \bar{q}$. To see that condition (a) is satisfied, first note that if the firm ever deviates, it may as well offer zero quality then and subsequently, since once any buyer has

learned about this deviation, under the strategy the buyer will have already "decided" never to buy again. Also, if the firm deviates now it may as well offer the anticipated price, since otherwise no demand will occur now or subsequently -- with zero quality there are no costs so the lost sales can only eliminate revenue and therefore reduce profits. We therefore need only consider deviations which consist of reducing quality to zero, but leaving price at $\bar{q}$. Such a deviation is unprofitable if

$$\sum_{t=\tau}^{\infty} c(\bar{q})(\rho g)^t \leq \sum_{t=\tau+1}^{\infty} \bar{q}(\sum_{k=1}^{t-\tau} n_k)(\rho g)^t,$$

since the left hand side is the cost savings for now and the future, and the right hand side is the loss in future revenues. Reversing the order of summation,[6] this can be rewritten as:

$$\frac{c(\bar{q})}{q} \sum_{t=\tau}^{\infty} (\rho g)^t \leq \sum_{k=1}^{\infty} n_k \sum_{t=\tau+k}^{\infty} (\rho g)^t.$$

Using the definition of AC(q), summing the geometric series, and cancelling $(\rho g)^\tau/(1-\rho g)$ on both sides, we get equivalently

$$AC(\bar{q}) \leq \sum_{k=1}^{\infty} n_k(\rho g)^k$$

or

$$\bar{q} \leq AC^{-1} (\sum_{k=1}^{\infty} n_k(\rho g)^k)$$

which is what we have assumed. This verifies condition (a), so both conditions of the Nash equilibrium definition are satisfied.   Q.E.D.

Theorem 1 gives an interesting and simple class of Nash equilibria for our model. In fact, there are many other equilibria including equilibria in which the firm and all customers have surplus. Theorem 2 shows that one of the equilibria described in Theorem 1 is distinguished by the fact that no other Nash equilibrium is more profitable for the firm. Therefore, that equilibrium can be interpreted as an outcome of a two-stage game in which the second stage is the game we have described, and the first stage is the choice by the firm of a Nash equilibrium for the second stage. This procedure is consistent with the intuition that the monopolist is a dominant player, and corresponds to the relatively efficient equilibrium that places all weight on the firm. We will interpret the equilibrium arising from this two-stage game as our solution, and in Section 3 we use it to study the impact of consumer information on the equilibrium quality level.

Theorem 2: No Nash equilibrium has more profits than the equilibrium described in Theorem 1 with

$$\bar{q} = \min(AC^{-1}(\sum_{k=1}^{\infty} n_k(\rho g)^k), MC^{-1}(1)).$$

Proof: For $\bar{q} = MC^{-1}(1)$, the quality level produced is the efficient level which maximizes the customer's surplus net of cost. Since the firm extracts all of the buyers' surplus, this is the most profitable Nash equilibrium, since customers can avoid having negative surplus net of price paid by not purchasing.

For $\bar{q} = AC^{-1}(\sum_{k=1}^{\infty} n_k(\rho g)^k)$ the argument is much more difficult. For exposition, our proof covers the case $n_1 = 1$ and $n_k = 0$ for $k \neq 1$ (the general result is proved in an earlier version of the paper). For this case,

we need to show that when $AC^{-1}(\rho g) < MC^{-1}(1)$, no equilibrium with realizations $\tilde{q}_t$, $\tilde{p}_t$ gives more profits than the equilibrium described in the statement of the theorem.

Any Nash equilibrium realizations $(\tilde{q}_t, \tilde{p}_t)$ must satisfy the condition that the gain from cheating in quality at each time $\bar{t}$ (equal to production cost at $\bar{t}$) must be less than the cost of cheating at $\bar{t}$ (which is no larger than future anticipated profits, since the firm can always set quality to zero in all future periods in order to earn zero future profit):

$$
(1) \qquad c(\tilde{q}_{\bar{t}}) \, \tilde{d}_{\bar{t}} g^{\bar{t}} \rho^{\bar{t}} \le \sum_{t+\bar{t}+1}^{\infty} (\tilde{p}_t - c(\tilde{q}_t)) \tilde{d}_t g^t \rho^t
$$

$$
\le \sum_{t=\bar{t}+1}^{\infty} (\tilde{q}_t - c(\tilde{q}_t)) \tilde{d}_t g^t \rho^t
$$

where the second inequality exploits the fact that $\tilde{d}_t = 0$ whenever $\tilde{q}_t < \tilde{p}_t$ as a result of condition (b) in the definition of equilibrium.

Since the proposed solution has positive profits, we can restrict ourselves to alternative strategies having positive profits. Positive profits implies that $\sup_t (\tilde{q}_t - c(\tilde{q}_t)) \tilde{d}_t$ is positive, and since for each t the value is bounded by $MC^{-1}(1) - c(MC^{-1}(1))$, the sup exists. Therefore, for each $\gamma > 1$, there is some τ such that

$$
\gamma(\tilde{q}_\tau - c(\tilde{q}_\tau)) \tilde{d}_\tau > \sup_t (\tilde{q}_t - c(\tilde{q}_t)) \tilde{d}_t
$$

(In particular, if the sup is achieved, the τ for which the sup is achieved will work.) Using Equation 1 with this choice of τ we have

$$c(\tilde{q}_\tau)\tilde{d}_\tau g^\tau \rho^\tau \leq \sum_{t=\tau+1}^{\infty} (\tilde{q}_t - c(\tilde{q}_t))\tilde{d}_t g^t \rho^t$$

$$< \sum_{t=\tau+1}^{\infty} \gamma(\tilde{q}_\tau - c(\tilde{q}_\tau))\tilde{d}_\tau g^t \rho^t$$

$$= \gamma(\tilde{q}_\tau - c(\tilde{q}_\tau))\tilde{d}_\tau g^{\tau+1}\rho^{\tau+1}/(1-\rho g)$$

which implies that

$$\frac{c(\tilde{q}_\tau)}{\tilde{q}_\tau - c(\tilde{q}_\tau)} < \gamma\frac{\rho g}{1-\rho g} .$$

Using the definition of $AC(\cdot)$ and solving for $AC(\tilde{q}_\tau)$, we have:

$$AC(\tilde{q}_\tau) < \gamma\rho g(1+(\gamma-1)\rho g)^{-1} .$$

Since $AC^{-1}(\cdot)$ is increasing, we have:

$$(2) \qquad \tilde{q}_\tau < AC^{-1}(\gamma\rho g(1+(\gamma-1)\rho g)^{-1}) .$$

Now we can check firm profits:

$$\pi = \sum_{t=0}^{\infty} (\tilde{p}_t - c(\tilde{q}_t))\tilde{d}_t g^t \rho^t$$

$$\leq \sum_{t=0}^{\infty} (\tilde{q}_t - c(\tilde{q}_t))\tilde{d}_t g^t \rho^t$$

$$< \sum_{t=0}^{\infty} \gamma(\tilde{q}_\tau - c(\tilde{q}_\tau))\tilde{d}_\tau g^t \rho^t$$

$$= \gamma(\tilde{q}_\tau - c(\tilde{q}_\tau))\tilde{d}_\tau /(1-\rho g)$$

$$< \gamma(AC^{-1}(\gamma\rho g(1+(\gamma-1)\rho g)^{-1})-c(AC^{-1}(\gamma\rho g(1+(\gamma-1)\rho g)^{-1})))/(1-\rho g)$$

where, by Equation 2, the last inequality is valid for γ close enough to 1, since $AC^{-1}(\rho g) < MC^{-1}(1)$ and $q-c(q)$ is increasing for $q < MC^{-1}(1)$, and since $d_\tau \leq 1$. Since the final right hand side does not depend on the choice of $\gamma > 1$ and the corresponding choice of τ, we have that:

$$\Pi \leq (AC^{-1}(\rho g)-c(AC^{-1}(\rho g)))/(1-\rho g),$$

which is the profit level associated with the equilibrium defined in the statement of this theorem so that no other solution yields higher profit.

Q.E.D.

As noted in the theorem, $MC^{-1}(1)$ is the perfect information quality level. This also corresponds to the price since the firm extracts the entire surplus given our form of demand. Theorem 2 tells us that the perfect information quality (i.e., no incentive to shirk in our agency interpretation) is the solution if achievable. This is the equilibrium level that would result if all customers were informed of current quality and is efficient.

In contrast, if all customers were completely uninformed ($n_\infty = 1$ and $n_k = 0$ for $k < \infty$), the minimum quality level would be the solution. In this solution, the firm picks the entire quality and price sequence and each customer picks his demand, taking the strategies of other agents as given. Given any demand sequence, the firm will specify $q_t \equiv 0$. But since $q_t \equiv 0$, $d_t \equiv 0$ (unless $p_t = 0$). Therefore, in equilibrium the firm

offers the minimum quality level (a prediction inconsistent with the data in many, but not all, contexts) and earns zero profit. This is a lemons solution.

The threat of switching to the zero quality level imposes discipline on the firm. This threat makes the equilibrium seem very stark in that small deviations induce substantial punishment. This is not essential to our model. Equilibria can be obtained with continuous strategies that support the realizations in our profit maximal Nash equilibrium solution. In the case n_1 = 1 one continuous strategy is $q_0^S = \bar{q}$,

$$q_{t+1}^S(\tilde{q}_0,\ldots,\tilde{q}_t) = R^{-1}\left[R(q_t^S) - \frac{c(q_t^S)-c(\min(\tilde{q}_t,q_t^S))}{\rho g}\right],$$

and $p_j^S(\cdot) = q_j^S(\cdot)$ for $j = 0,1,\ldots$, where ~'s indicate realizations and $R(x) = \frac{x-c(x)}{1-\rho g}$ is the present value of the residual profit from following an anticipated strategy of offering $p = q = x$ with $d = 1$ forever.

3. The Determinants of Quality

Theorem 2 provides an explicit formula given equilibrium quality in terms of the cost function c, the discount factor ρ, the growth rate g of demand, and the distribution of customer information lags $(n_k$'s), specifically

$$\bar{q} = \min\left[MC^{-1}(1),\ AC^{-1}\left(\sum_{k=1}^{\infty} n_k(\rho g)^k\right)\right].$$

We can use this formula to study the impact of the customers' information lags and other parameters on quality.

First, we consider the effect on quality of improving customers' information. To do so, we construct the distribution function $F(k) = \sum_{k'=1}^{k} n_{k'}$, of buyer lags, which gives the number of customers with a lag of less than or equal to k periods. An improvement of customer information means that we are decreasing lags for some buyers, and corresponds to increasing the value of the distribution function at some or all values without decreasing it anywhere. (Readers familiar with stochastic dominance theory will realize that this is analogous to first-order stochastic dominance of random variables.) The following simple derivation allows us to write $\bar{q}$ in terms of F:

$$\sum_{k=1}^{\infty} n_k (\rho g)^k = \sum_{k=1}^{\infty} n_k \sum_{k'=k}^{\infty} (1-\rho g)(\rho g)^{k'}$$

$$= \sum_{k'=1}^{\infty} (1-\rho g)(\rho g)^{k'} \sum_{k=1}^{k'} n_k$$

$$= \sum_{k'=1}^{\infty} (1-\rho g)(\rho g)^{k'} F(k').$$

Therefore, equilibrium quality can be written as:

$$= \min[(MC^{-1}(1),\ AC^{-1}(\sum_{k=1}^{\infty} (1-\rho g)(\rho g)^k F(k))].$$

Since each coefficient of F(k) is nonnegative, an improvement of customer information (implying an increase in F(k) for some k and no decrease in any k) yields increased quality, unless quality was already at the perfect information level, in which case quality remains at the level. This is an analytic derivation of the intuition that better customer information results

in more discipline so that the equilibrium quality level is larger. If consumers are better informed, then demand dries up more quickly provided the firm produces an inferior product. Therefore, it is possible for the firm to credibly promise a higher level of quality.

Since AC^{-1} is increasing, the equilibrium quality is an increasing function of ρg, if not everywhere, then at least up to a critical value, above which quality is constant at the full information level. This critical value is the unique solution of the equation

$$MC^{-1}(1) = AC^{-1}(\sum_{k=1}^{\infty} n_k (\rho g)^k),$$

provided this equation has a solution for $\rho g < 1$, as it will for $\sum_{k=1}^{\infty} n_k \geq AC(MC^{-1}(1))$. Otherwise, the perfect information quality is not achieved for any $\rho g < 1$. This will occur, for example, if too many agents are completely uninformed, i.e., if $n_{\infty} > 1-AC(MC^{-1}(1))$. On the other hand, so long as some customer has some information, i.e., $n_{\infty} < 1$, then $\bar{q} > 0$. (This result relies crucially on the Inada condition at 0.)

To illustrate this we consider the example $c(q) = \alpha q^{\beta}$, $\beta > 1$, so that $MC(q) = \alpha \beta q^{\beta-1}$ and $AC(q) = \alpha q^{\beta-1}$. Then

$$\bar{q} = \min[(\alpha\beta)^{\frac{-1}{\beta-1}}, \ (\frac{1}{\alpha} \sum_{k=1}^{\infty} n_k (\rho g)^k)^{\frac{1}{\beta-1}}].$$

For $\beta = 2$ (quadratic cost), the quality takes a simpler form:

$$\bar{q} = \min[\tfrac{1}{2}, \ \sum_{k=1}^{\infty} n_k (\rho g)^k]/\alpha.$$

Further specializing to $n_1 = 1$ (and therefore $n_k = 0$ for $k \neq 1$), we have:

$$\bar{q} = \min[\tfrac{1}{2}, \rho g]/\alpha,$$

i.e., $\bar{q}$ is linear in ρg, then constant.

The increase in quality as ρg increases has a simple intuition. A large growth rate and a large discount factor (small interest rate) both make the future more important to the firm, increasing the punishment for losing a reputation, since both make the future loss from cheating large. This increased discipline means that the firm can more credibly offer higher quality than if the discount factor and the growth rate were smaller, and the discipline less stringent. The parameters ρ and g affect the equilibrium only through ρg. It is therefore clear that the equilibrium quality is unaffected by changes in inflation, provided the real interest rate and real growth rate of demand are unaffected. Therefore, the interest rate comparative static does not apply to increases in the nominal interest rate which are offset by inflation in the value of demand and costs (through the Fisher equation), at least in the absence of taxes.

As may be expected, $\bar{q}$ falls when the cost of providing quality rises. Formally, suppose that $c^*(q) > c(q)$ and $c^{*\prime}(q) > c'(q)$ for all q. Then $MC^*(q) > MC(q)$ and $AC^*(q) > AC(q)$ for all q. Since AC, MC, AC*, and MC* are increasing, this implies that $MC^{*-1}(m) < MC^{-1}(m)$ and $AC^{*-1}(m) < AC^{-1}(m)$ for all m. Therefore, by the formula for $\bar{q}$, the higher costs (under c*) imply lower equilibrium quality.

4. Conclusion

Many phenomena are based upon the incentive to maintain a reputation. In our analysis of the quality decision of a firm when the quality is not immediately observable, we derive endogenously the incentive for the firm to maintain a reputation for quality. The firm cheats unless it anticipates losing too much in the future as a consequence.

Many limitations remain with our analysis. It would be interesting to model customers who live several periods so that customers can act strategically (as in a bilateral bargaining problem). The treatment of uncertainty would be interesting. If the firm selects a cost level that determines (random) quality up to a noise term, then several cases arise. The firm might observe quality before setting price (selling a burnt pizza at a discount), observe quality after price is determined (but before future quality), or never observe quality. If the firm selects quality and the cost factor is random, then the firm might select quality before or after the current cost factor is known. In the former case our model trivially extends in that our cost function can be reinterpreted as an expected cost function, but in other cases the extensions require further analysis. Uncertainty considerations seems basic to the analysis of how the incentive to maintain a reputation changes over the business cycle with expectations of future states (i.e., cost or demand).[7] Reputation issues should also be analyzed in other market structures. The study of reputational issues in many of the substantive contexts discussed in the introduction would be useful.

Despite the limitations of our analysis we hope that this approach can be applied to understanding an important array of issues.

Footnotes

[1]Of course, there are "non-economic" reasons, such as ethics, pride, and altruism, but it is unlikely that these reasons explain the whole phenomenon.

[2]Further justifying the use of this equilibrium as a benchmark is the fact that this equilibrium maximizes supplier plus customer surplus, at least for the identical customers case. (We believe this result is general but we have no proof.)

[3]In Figure 1 there are 16 outcomes from one period, as in the first of two periods. In the second of two periods there is a tree (like Figure 1) from each of the 16 outcomes from the first period. Since each such tree has 16 terminal nodes, there are 16 x 16 or 256 terminal nodes in all. And this is assuming only 2 choices per node!

[4]The assumption that $c(0) = 0$ can be relaxed by interpreting all costs and prices as the costs and prices in excess of the cost of producing one unit of the good having zero quality. The assumption that $c'(0) = 0$ and $c'(\infty) = \infty$ are Inada conditions, without which we would have to worry about corner solutions. The assumption $c'(q) > 0$ for $q > 0$ asserts that it is costly to produce higher quality, and the assumption $c''(q) > 0$ for $q > 0$ assures first order conditions are sufficient.

[5]Another interpretation of this model is as an infinite horizon agency problem facing a long-lived agent who acts on behalf of a sequence of principals. We assume the production technology is non-stochastic and for

ease of exposition we normalize the outcome to the action choice. The agent maximizes his discounted intertemporal utility, where the individual period utility is just his fee less the disutility of the action choice (which is increasing and convex, i.e., the marginal disutility is positive and increasing). Increasing marginal disutility of the action choice is common in the types of agency problems that arise in financial market as illustrated by the diminishing marginal utility of expropriated capital, increasing marginal managerial aversion to effort. The utility of the principal (who lives one period) is just the outcome less the fee. The inelastic demand specification for the short-lived party is analogous to competition among principals (in equilibrium the long-lived agent receives all the rents). The current performance of an agent determines his future assignments in this interpretation. For ease of exposition we interpret the model as describing the product quality choices of a long-lived monopolist firm.

[6]Reversing the order of summation is the discrete analog of reversing the order of integration. The formula

$$\sum_{t=\tau+1}^{\infty} \sum_{k=1}^{t-\tau} \alpha_{tk} = \sum_{k=1}^{\infty} \sum_{t=\tau+k}^{\infty} \alpha_{tk}$$

is valid whenever the α_{tk}'s are positive and the sums exist.

[7]For example, the Wall Street Journal reported that several firms did not fulfill entry level job offers during the 1980 recession, attributing the failure to maintain a reputation to the economic downturn.

References

Allen, F., "Reputation and Product Quality," Rand Journal of Economics 15 (Autumn 1984), 311-327.

Bhattacharya, S., "Imperfect Information, Dividend Policy, and 'The Bird in the Hand' Fallacy," Bell Journal of Economics 10 (1979), 259-270.

Diamond, D., "Reputation Acquisition in Debt Markets," Presented at the American Finance Association Meetings, December 1984.

Dybvig, P. and C. Spatt, "Does It Pay to Maintain a Reputation? Consumer Information and Product Quality," Princeton University Financial Research Center, August 1980.

Eaton, J. and M. Gersovitz, "Debt with Potential Repudiation: Theoretical and Empirical Analysis," Review of Economic Studies 48 (1981), 289-309.

Heal, G., "Do Bad Products Drive Out Good?" Quarterly Journal of Economics 90 (1976), 499-502.

Holmstrom, B., "Managerial Incentive Problems--A Dynamic Perspective," in Essays in Economics and Management in Honour of Lars Wahlbeck (1983), pp. 209-230.

Jensen, M. and W. Meckling, "Theory of the Firm: Managerial Behavior, Agency Costs and Capital," Journal of Financial Economics 3 (1976), 305-360.

Klein, B. and K. Leffler, "The Role of Market Forces in Assuring Contractual Performance," Journal of Political Economy 89 (1981), 615-641.

Kreps, D., "Corporate Culture and Economic Theory," unpublished manuscript, Stanford University, 1984.

Kreps, D., P. Milgrom, J. Roberts, and R. Wilson, "Rational Cooperation in the Finitely Repeated Prisoners' Dilemma," Journal of Economic Theory 27 (August 1982), 245-252.

Kreps, D. and R. Wilson, "Sequential Equilibria," Econometrica 50 (July 1982), 863-894. [1982a]

Kreps, D. and R. Wilson, "Reputation and Imperfect Information," Journal of Economic Theory 27 (August 1982), 253-279. [1982b]

Milgrom, P. and J. Roberts, "Predation, Reputation, and Entry Deterrence," Journal of Economic Theory 27 (August 1982), 280-312.

Radner, R., "Monitoring Cooperative Agreements in a Repeated Principal-Agent Relationship," Econometrica 49 (1981), 1127-1148.

Rogerson, W., "Reputation and Product Quality," Bell Journal of Economics 14 (Autumn 1983), 508-516.

Ross, S., "The Economic Theory of Agency: The Principal's Problem," _American Economic Review_ 63 (May 1973), 134-139.

Ross, S., "The Economic Theory of Agency and the Principle of Similarity," in M. Balch, D. McFadden and S. Wu, eds., _Essays on Economic Behavior under Uncertainty_, Amsterdam: North-Holland Publishing Co., 1974.

Sachs, J., "Theoretical Issues in International Borrowing," NBER Working Paper Series No. 1189, August 1983.

Shapiro, C., "Consumer Information, Product Quality, and Seller Reputation," _Bell Journal of Economics_ 13 (1982), 20-35.

Spatt, C., "Credit Reputation Equilibrium and the Theory of Credit Markets," GSIA Working Paper No. 40-81-82, Carnegie-Mellon University, revised May 1983.

Stiglitz, J. and A. Weiss, "Incentive Effects of Terminations: Applications to the Credit and Labor Markets," _American Economic Review_ 73 (Dec. 1983), 912-927.

Toaff, E., "Models of Product Diversity in Large Markets with Imperfectly Informed Buyers," Ph.D. dissertation, Carneige-Mellon University (1979).