

Penn Institute for Economic Research
Department of Economics
University of Pennsylvania
3718 Locust Walk
Philadelphia, PA 19103-6297
pier@ssc.upenn.edu
<http://www.econ.upenn.edu/pier>

PIER Working Paper 98-006

“Your Reputation is Who You’re Not, Not Who You’d Like To Be”

by

George J. Mailath and Larry Samuelson

Your Reputation Is Who You're Not, Not Who You'd Like To Be*

George J. Mailath
Department of Economics
3718 Locust Walk
University of Pennsylvania
Philadelphia, PA 19104-6297
gmailath@econ.sas.upenn.edu

Larry Samuelson
Department of Economics
1180 Observatory Drive
University of Wisconsin
Madison, WI 53706-1320
LarrySam@ssc.wisc.edu

August 11, 1998

*First draft August 1996. This paper has been presented at many universities and conferences, and we have received helpful comments from many people, too numerous to list here. Financial support from the National Science Foundation is gratefully acknowledged. This version was written while Mailath was visiting the Australian National University, whom he thanks.

Your Reputation Is Who You're Not, Not Who You'd Like To Be

by George J. Mailath and Larry Samuelson

Abstract

We construct a model in which a firm's reputation must be built gradually, is managed, and dissipates gradually unless appropriately maintained. Consumers purchase an experience good from a firm whose unobserved effort affects the probability distribution of consumer utilities. Consumers observe private, noisy signals (consumer utilities) of the behavior of the firm, yielding a game of *imperfect private monitoring*. The standard approach to reputations introduces some "good" or "Stackelberg" firms into the model, with consumers ignorant of the type of the firm they face and with ordinary firms acquiring their reputations by masquerading as Stackelberg firms. In contrast, the key ingredient of our reputation model is the continual possibility that the ordinary or "competent" firm might be replaced by a "bad" or "inept" firm who never chooses the Stackelberg action. Competent firms then acquire their reputations by convincing consumers that they are *not* inept. Building a reputation is an exercise in separating oneself from inept firms *who one is not*, rather than pooling with Stackelberg firms *who one would like to be*. We investigate how a firm manages such a reputation, showing, among other features, that a competent firm may not always choose the most efficient effort level to distinguish itself from an inept one.

Journal of Economic Literature Classification number: C70.

Keywords: Reputation, incomplete information, Stackelberg types, commitment types.

Your Reputation Is Who You're Not, Not Who You'd Like To Be

by George J. Mailath and Larry Samuelson

1. Introduction

A firm's reputation, often described by terms such as customer loyalty, market trust, and name recognition, is an important asset. Kreps [16], for example, suggests that a key reason for the existence of firms is to serve as a repository for reputations.

This paper examines the economics of building and maintaining a reputation. While we motivate the analysis as a model of a firm, our results apply more generally.

We begin from the viewpoint that a reputation is an asset. Like any other capital stock, a reputation requires investment to create and maintain. We are accordingly interested in reputation models that capture three features:

First, reputations are built gradually. The typical experience for new firms and products is that the consumer goodwill or trust that lies at the heart of a reputation can only be built slowly, by compiling a record of continued high-quality performance. Only after the firm has a demonstrated record of high quality does it enjoy the fruits of its reputation.

Second, reputations dissipate gradually. A firm that no longer invests in its physical capital stock typically experiences no immediate decline in productivity, with its capital stock depreciating as it is consumed. Similarly, a firm that no longer maintains its reputation can initially rest on its laurels, with its reputation only gradually losing its value as consumers adjust to the new performance level.

Finally, reputations can be managed. A firm's physical-capital investment profile is seldom uniform over its lifetime. Relatively high investment levels early in a firm's lifetime are often followed by lower investment levels when mature. Similarly, we expect firms to manage their reputations, with initial periods of high investment in reputation building possibly being followed by subsequent periods in which the reputation is sustained with lower investment levels. As a result, it may often be the case that "number two" tries harder.¹

Reputations are valuable because they help firms solve moral hazard problems. We model the moral hazard as arising out of an effort choice on the part of the

¹The analysis closest in spirit to ours is Gale and Rosenthal [11], who examine a model in which a firm gradually acquires a reputation and then depletes it upon learning that an exogenously generated shock will soon force the firm out of the market. Papers which adopt a quite different approach to the need to build relationships from small beginnings include Datta [6], Diamond [7], and Watson [22], [23].

firm, which affects the probability distribution over the utilities consumers derive from consuming an experience good. If the firm could credibly commit to high (rather than low) effort, perhaps by writing an appropriate contract, then the firm would be profitably rewarded through a higher willingness to pay on the part of consumers. Unfortunately, consumers cannot observe the firm's effort level, and cannot observe their idiosyncratic, unverifiable utility realization until after they have consumed the good. It is then impossible for the firm to credibly commit to high effort, dooming the firm to an inefficient, low-effort equilibrium, unless reputation considerations come to its rescue.

We are interested in cases in which firms face two important obstacles in attempting to surmount this moral hazard problem. First, consumers observe *noisy* signals of the behavior of the firm. High effort makes a good utility outcome more likely, but does not ensure a good outcome, just as low effort does not condemn the consumer to a bad outcome. Even after the fact, consumers cannot be sure whether the firm has chosen high or low effort, making the game one of *imperfect* monitoring. This precludes the use of simple trigger strategies, familiar from the theory of repeated games with perfect monitoring, as a commitment device. Second, different consumers receive different realizations of the signals, making the game one of imperfect *private* monitoring. This rules out the coordination in consumer beliefs and firm behavior that allows trigger strategies to serve as a commitment device in games of imperfect public monitoring.²

The standard approach to reputations, beginning with Kreps, Milgrom, Roberts, and Wilson [14], Kreps and Wilson [15], and Milgrom and Roberts [20], introduces the possibility that the firm is a "good" or "Stackelberg" type into a game of perfect monitoring. Consumers know that such firms exist, but are ignorant of the type of the firm they face. Ordinary firms then acquire their reputations by masquerading as Stackelberg firms. In our setting, a Stackelberg firm always chooses high effort.³

Fudenberg and Levine [9] extend the standard reputation model by assuming that consumers can observe only noisy signals of the firm's action, so that the game has imperfect public monitoring. They show that, if consumers assign

²The trigger strategies differ in the observations that trigger the "punishments." In games of perfect monitoring, deviations trigger the punishments, and so are not observed in equilibrium. In contrast, for games of imperfect public monitoring, "bad" signals trigger punishments, and since these signals arise with positive probability in equilibrium, punishments are also triggered in equilibrium. Section 6 returns to these issues with a more detailed discussion.

³More generally, the Stackelberg action is the action to which the firm would credibly commit if it could. Equivalently, it is the firm's optimal action if its choice is observed by the consumers (so the firm behaves as a "Stackelberg leader"). This type is also sometimes referred to as a commitment type.

positive prior probability to the firm's being a Stackelberg type and if the firm is sufficiently patient, then an ordinary firm achieves an average discounted payoff close to its commitment payoff (i.e., its payoff if it could credibly commit to the Stackelberg action). In order to obtain this payoff, the ordinary firm spends long periods of time choosing high effort with high probability. The simplest standard reputation model has only two types, ordinary and Stackelberg. In that setting, there is no reputation building. Rather than being built gradually, reputations spring to life, in the sense that a *de novo* firm begins with consumers immediately assigning high probability to the firm's choosing high effort, and the probability assigned to the firm's being a Stackelberg type then steadily declines.⁴ Equilibria that exhibit reputation building can be constructed by expanding the model to include three types, a Stackelberg type, an ordinary type, and an "inept" type, where the inept type always chooses the myopically optimal action. In the initial periods, consumers then put substantial probability on the firm's being inept, and an ordinary firm builds its reputation as this probability falls. We also note that while Fudenberg and Levine [9] restrict attention to games of imperfect *public* monitoring, their analysis applies without change to the case of imperfect private monitoring, as long as the game is sufficiently simple that the Stackelberg action is a constant, as it is in our case.

The standard reputation model can thus be extended to capture many of the features of a reputation that we consider important, but it relies crucially on consumers believing in the possibility of a Stackelberg type. This may be reasonable in some circumstances, but is questionable in others. In our model, for example, the Stackelberg type must choose high effort in every period. Such behavior is typically justified by assuming that the Stackelberg action is a strictly dominant action.⁵ But in our case, constant high effort will be a dominant strategy in the stage game only if the Stackelberg type has quite different payoffs from the ordinary firm.

An inept type, in contrast, may require much less stringent assumptions. In our model, an inept type need differ from ordinary types only in having higher

⁴In any equilibrium, the Stackelberg type chooses high effort with probability one, while the ordinary firm mixes between high and low effort. If the ordinary firm were to choose high effort with probability one in some period, then consumers would not adjust posteriors in response to the signal in that period, and so the ordinary firm would optimally choose low effort, disrupting the equilibrium.

⁵Fudenberg and Levine [9] are concerned with providing bounds on equilibrium payoffs, and so it is important for them that the Stackelberg type necessarily choose the Stackelberg action, which requires that the action be strictly dominant in the *repeated* game. On the other hand, if one is only concerned with the existence of an equilibrium in which the Stackelberg type always chooses the Stackelberg action, it is enough that the action be dominant for that type in the *stage* game.

costs of high effort. Perhaps pessimistically, we think that inept types are more plausible than Stackelberg types in many circumstances. It is thus important to investigate whether reputation building can occur in the absence of Stackelberg types.

In this paper, we examine reputation equilibria in markets containing only ordinary and inept firms. Our ordinary firms acquire their reputations in the course of convincing consumers that they are *not* inept. Building a reputation is an exercise in separating oneself from inept firms rather than pooling with Stackelberg firms. In this sense, reputation in our model involves establishing *who one is not*, rather than *who one would like to be*.

Ordinary types (henceforth competent firms) distinguish themselves from inept firms in our model by choosing high effort. They do so because a firm's value is increased when consumers think it is competent, as long as consumers also believe that competent firms choose high effort. However, if the firm's type is determined once and for all in the initial period, then consumers will eventually become virtually convinced that the firm is competent, with an additional signal prompting only a minuscule revision in their beliefs. As a result, the incentive to choose high effort collapses and the equilibrium unravels, ensuring that there is no pure-strategy high effort equilibrium.⁶ The key ingredient of our reputation model is the *continual* possibility that a competent firm might be replaced by an inept firm. This constant possibility that the type of the firm has changed, and hence the continued desire of the competent firm to separate itself from bad firms, leads to pure-strategy equilibria in which competent firms always choose high effort, gradually building and then maintaining a reputation for competency and high effort.⁷

Once a reputation can be supported by the desire to separate, the question arises as to how the firm manages its reputation. If there is more than one effort level that could distinguish a competent firm from an inept one, which one will the competent firm choose? Will it always choose the same one? To address such questions, we expand the model so that the competent firm has more than one effort level available that inept firms cannot choose. Higher effort levels lead to higher probabilities of good outcomes, but at higher cost. The informativeness of consumer outcomes depends on consumers' beliefs about firm

⁶A related feature was first described, in the context of a signal jamming model, by Holmstrom [12], who also described the role of changing characteristics in removing it. Section 3 discusses the important respects in which our analysis differs from Holmstrom's [12].

⁷The restriction to pure strategies is not as troubling here as it is in the standard reputation model, where there are no pure strategy equilibria (see footnote 4). Mixed strategy equilibria for games with imperfect private monitoring are not well understood. We return to this in Section 3.

behavior. If consumers believe competent firms, like inept firms, choose low effort, then consumer outcomes are viewed as being completely uninformative. On the other hand, the higher the effort level consumers believe the competent firm has chosen, the more consumers' posteriors will adjust in response to a good or bad outcome. As a consequence, there can be equilibria in which the competent firm builds a reputation, but does so by choosing an inefficiently low level of effort. Choosing the higher, efficient level is not a profitable deviation because consumers are updating according to a rule which rewards good outcomes less than would be the case if the consumers expected the efficient effort level to be taken. There can also be equilibria in which an inefficiently high level of effort is chosen, supported by an updating rule which punishes bad outcomes with large downward revision in consumer posteriors. In the course of examining these possibilities, Section 5 shows that there is a sense in which inefficiently-built reputations are more likely to involve too little rather than too much effort. Finally, we show that equilibrium may require the competent firm to choose different effort levels for different configurations of consumer posteriors.

Is a reputation the only means by which the firm can commit to high effort? Why doesn't the firm simply offer a guarantee, compensating any consumer who receives a bad outcome so as to bring their utility up to that of a good outcome and hence providing the firm with the appropriate incentives to exert high effort?

There are three related obstacles to such guarantees. First, high effort still leaves some probability of a bad utility outcome. An effective guarantee of a good outcome may then be prohibitively expensive, while a guarantee of a high *probability* of a good outcome may be impossible to enforce. Second, in order to provide an outcome equivalent to receiving a good outcome with certainty, a guarantee will typically have to involve much more than simply replacing any product delivering low utility, since the process of experiencing a bad outcome and claiming compensation is itself costly. A restaurant's offer to prepare another meal cannot fully compensate for having received a bad dinner. A surgeon's offer to repeat a procedure cannot compensate for having done the first one incorrectly. Finally, because consumer utility outcomes are not verifiable, a guarantee provides consumers receiving good as well as bad outcomes with an incentive to claim compensation.

Committing to good outcomes, even with the help of devices such as guarantees, is thus likely to be too costly. The firm will optimally choose to have some chance of a bad outcome remain. We view the uncertainty over the utility provided by the good in our model as that uncertainty which remains after the firm has taken whatever cost-effective measures it can, including guarantees, to ensure high quality, and it is this residual uncertainty that gives rise to the value

of a reputation.⁸

There are many other devices besides guarantees that might be used as signals that the firm has exerted high effort, including prices. For example, firms might use prices as a device to signal their effort levels by “burning money.” The possibility of signaling is important, but we focus on reputation formation issues in this paper by assuming that the only choice the firm makes is whether to exert high or low effort. Notice, in addition, that any model in which such signaling is possible always has a pooling equilibrium that duplicates the equilibrium of the model we examine.

Section 2 describes the interaction between the firm and its consumers. Section 3 illustrates the importance of constant adverse selection in allowing the firm to overcome its moral hazard problem. Section 4 presents an equilibrium in which a reputation is built and maintained, with the competent firm constantly striving to produce high quality in order to assure consumers that it has not been replaced by an inept firm. Section 5 expands the model to accommodate a richer set of actions for the firm, examining circumstances under which the firm may use inefficient means to build a reputation and circumstances under which a firm may work very hard to build up a mediocre reputation, but then coast on the strength of the resulting high reputation. Section 6 describes a generalization of the model and relates our analysis to the literature.

Our view of reputations as assets directs attention to the market for reputations. When can reputations be purchased? Which reputations will competent firms find most valuable, and which reputations will inept firms find most valuable? We turn to these issues in a companion paper ([19]).

2. The Market

Consumers. Consider a continuum of consumers who repeatedly purchase an experience good or service from a single firm. Time is discrete and consumers are infinitely lived.⁹ The experience good generates two possible utility levels, which we normalize to 0 and 1. We describe a utility of 1 as a good outcome, denoted by g , and a utility of 0 as a bad outcome, denoted by b .

In each period, the firm exerts effort e , which determines the probability of a good outcome, ρ_e . There are two possible efforts, high (H) and low (L), with the

⁸Klein and Leffler [13] stress the importance of reputational considerations in enforcing commitments to high quality when contractual devices are prohibitively expensive.

⁹Equivalently, since consumers will be behaving myopically, each consumer lives one period and has one offspring. The critical assumption is then that an offspring inherits the history of her parent, but not of the other consumers.

probabilities satisfying

$$0 < \rho_L < \rho_H < 1.$$

Section 5 extends the analysis to three effort levels, while Section 6 describes how our results extend to a more general technology linking effort levels and utilities.

The consumer population is of unit measure. Each consumer purchases one unit of the good in each period. In each period t , each individual consumer receives a good outcome with probability ρ_{e_t} , where e_t is the firm's effort in period t . Moreover, in each period and for each group of consumers having experienced a common history of utility realizations, a proportion ρ_{e_t} of this group receives the good outcome. A deterministic sequence of effort choices thus yields a deterministic sequence of distributions of consumers' histories. In describing the model, we proceed as if in each period each consumer receives an independent draw from the distribution of utility outcomes, the probability of a good outcome in period t is ρ_{e_t} , and the continuum of consumers ensures that there is no uncertainty about the aggregate outcome.¹⁰

After purchasing the good, each consumer observes her realized utility. This provides a noisy signal of the firm's effort, and consumers cannot observe the firm's effort levels directly. The aggregate distribution of the utilities received by consumers in any period is perfectly informative about the firm's effort choice in that period: consumers need only observe the fraction of good outcomes in order to infer effort. However, a consumer observes neither the aggregate distribution nor the outcome of any other consumer, and hence can only draw inferences about the firm's behavior from her own realized utility.

Since there is a continuum of consumers, no consumer's action can affect

¹⁰There are well-known technical complications in modeling a continuum of independent random variables (see Al-Najjar [2]). In our case, independence serves only an expositional role. The critical feature of the model is that, in period t , if the firm is expected to exert effort level e_t , then each consumer assigns a probability of ρ_{e_t} to the event that she receives a good outcome, and believes that a fraction ρ_{e_t} of consumers will receive the good outcome. This can be achieved rigorously as follows: Let the space of consumers be $[0, 1]$. Fix the first-period effort choice e . Let u be the realization of a uniform random variable on $[0, 1]$. Then, consumers in $([u - 1, u + \rho_e - 1] \cup [u, u + \rho_e]) \cap [0, 1]$ receive the good outcome and consumers in the complement the bad outcome. The subpopulation who received a good outcome in the first period can be viewed in a natural way as a population distributed on I_g , an interval of length ρ_e , while the subpopulation who received a bad outcome in the first period can be viewed as a population distributed on I_b , an interval of length $1 - \rho_e$. Second-period outcomes are determined from second period effort and the realizations of a uniformly distributed random variable on I_g and one distributed on I_b as above. These two random variables are independent of each other and of the realization of u . This construction has an obvious recursion, and yields the desired pattern of utility realizations. Notice that this construction requires a countable number of independent random variables (one for each possible finite history of utility realizations).

the future play of the game, and hence consumers optimize myopically. Period t revenues of the firm then depend entirely on consumers' expectations of the firm's period- t effort choice. Let F be a distribution function, with $F(x)$ identifying the proportion of consumers who expect the firm to exert high effort with probability less than or equal to x , and let \mathcal{F} be the set of distributions. Revenue is given by $p(F)$, where $p : \mathcal{F} \rightarrow [0, 1]$ is a function representing the market interaction by which prices are formed. We assume p is increasing, so that higher expectations of high quality lead to higher revenue:

$$F' \succ F \Rightarrow p(F') > p(F),$$

where \succ is first-order stochastic dominance. We also assume that $p(F_n) \rightarrow p(F)$ for all sequences $\{F_n\}$ converging weakly to F . Finally, if a sequence of distributions $\{F_n\}$ first order stochastically dominates F and is bounded away from it, in the sense that $\int_0^1 x dF_n - \int_0^1 x dF = \int_0^1 (F(x) - F_n(x)) dx$ is bounded away from zero, then $p(F_n)$ is bounded away from $p(F)$. An obvious example of a market interaction with these properties is perfect price discrimination, where the firm charges each consumer her reservation price.

If every consumer has the same belief about effort choice, we suppress F and write $p(x)$, where $x \in [0, 1]$ is the common probability assigned to the event that the firm exerts high effort. If consumers have homogeneous beliefs and competition among consumers bids the price of the good up to its expected utility, or, equivalently, if the firm can perfectly price discriminate, then revenue is given by

$$p(x) = \rho_H x + \rho_L (1 - x).$$

The Firm. There are two types of firm, *competent* and *inept*. Low effort is costless for both types. A competent firm can choose high effort at a cost of $c > 0$. An inept firm can only choose low effort, or equivalently, a choice of high effort (at a cost of c) generates the same distribution of consumer outcomes as low effort. Both types of firm maximize the discounted sum of expected profits, with discount factor β .

We assume that the increase in the firm's revenue from exerting high rather than low effort would more than compensate for the cost of effort *if* consumers knew the effort choice, i.e.,

$$p(1) - p(0) > c. \tag{2.1}$$

With perfect price discrimination, this assumption reduces to

$$\rho_H - \rho_L > c.$$

This ensures that the efficient action is to exert high effort, and so if the firm could commit to an effort level, it would choose high effort.¹¹

A firm's type, competent or inept, is determined by a variety of factors, including the composition and skills of the work force, technology and capital stock, access to appropriate materials, management style and organization, and workplace culture. For concreteness, we speak of a firm's type as being determined by the characteristics of its owner, and speak interchangeably of the type of owner or type of firm.

Before play begins, Nature determines the original type of the firm, choosing competent with probability ϕ_0 and inept with probability $1 - \phi_0$. The firm learns its type, but consumers do not. Moreover, in each subsequent period, there is a probability λ that the current owner of a firm is replaced by a new owner.¹² We think of changes in personal circumstances, such as reaching retirement age, as occasionally causing current owners to leave the market (or, more generally, we think of exogenous variations in any of the factors that affect a firm's type). Regardless of the type of owner who left, the new owner's type is competent with exogenously fixed probability $\theta \in (0, 1)$. Consumers do not observe the type of a replacing firm, and cannot observe whether a replacement has occurred. For example, the ownership of a restaurant might change without changing the restaurant's name and without consumers being aware of the change.¹³

Intuitively, replacements are designed to capture the fact that consumers, faced with unexpected outcomes, are always willing to consider the possibility that "things have changed." In practice, we expect consumers who have received a bad outcome after an extremely long string of good outcomes to resort to "something about the firm has changed," rather than bad luck from an unchanged firm, as their primary explanation, and this is just what happens once replacements are

¹¹In general, since consumers may be retaining some surplus, efficiency may require the firm to exert high effort even if $p(1) - p(0) \leq c$. If consumers are paying their reservation utilities, low effort is efficient if $p(1) - p(0) \leq c$.

¹²This makes the effective discount factor is $(1 - \lambda)\beta$.

¹³In other cases where reputations are thought to be important, ownership changes can clearly be observed. For example, ownership changes of medical practices typically cannot be concealed. We regard this as a fundamentally different case and would not characterize the purchase of such a practice as the purchase of a reputation. In yet other cases, new owners may have the option of announcing their presence. For example, restaurants often boldly display banners announcing new ownership, or change their names. The analysis of such cases is complicated by the observation that existing owners may also announce that they are new, or change names. As a result, observed ownership changes may signal information that unobserved changes cannot, but only if something endows the observed change with "credibility." New owners may then be compelled to engage in costly remodelling, even if the new establishment does not differ significantly from the old. We consider these issues in [19].

added to the model.

Sections 3 and 4 explain why the possibility of changing type is crucial to our results. We have described the change in the firm's type as involving a change of ownership, but with only superficial changes, we could model the change of type as a change in any of the other factors that influence the firm's ability to produce better consumer outcomes by taking high effort. For example, we could think of a single owner who faces random shocks to his cost of production, with high costs corresponding to ineptness and low costs to competence.

Strategies and Equilibrium. At the beginning of period t , each consumer i is characterized by her posterior probability that the firm is competent, denoted ϕ_t^i , and her posterior probability that the firm will exert high effort, denoted v_t^i . If the firm is competent, it makes its unobserved effort choice. The output is then produced, and the firm receives revenues that depend on the distributions of ϕ_t^i and v_t^i , but not on the firm's type or action in that period. Consumers observe their own realized utilities and update beliefs about the type of firm. Finally, with probability λ , the owner is replaced.

For consumer i , a period- t history is a t -tuple of realized utility outcomes, $h_t^i \in \{g, b\}^t \equiv K_t^c$, describing the utilities the consumer has received in periods 0 through $t - 1$ (note that the space of possible histories is the same for all consumers and that the initial period is period zero). The set of all consumer histories is $\mathcal{K}^c = \cup_{t \geq 0} K_t^c$, where $K_0^c = \emptyset$. A *belief function* for consumer i is a function $v^i : \mathcal{K}^c \rightarrow [0, 1]$, where $v^i(h_t^i)$ is the probability consumer i assigns to the firm exerting high effort in period t , given history h_t^i . Since every history of utility realizations has positive probability under any sequence of effort choices of the firm, it is necessarily the case that, as long as consumers use Bayes' rule and consumers start with a common prior, two consumers observing the same sequence of utility realizations have the same beliefs about the firm's behavior. In particular, $v^i(h_t^c) = v^j(h_t^c)$ for all $h_t^c \in \mathcal{K}^c$ and all $i, j \in [0, 1]$. We impose this requirement without further comment, and describe consumers' beliefs by a single function $v : \mathcal{K}^c \rightarrow [0, 1]$. Given a sequence of realized effort choices by the firm, h_t^f , there is an induced probability measure on K_t^c , denoted $\mu_t(\cdot | h_t^f)$. Then, given v and h_t^f , $F_{v, h_t^f}(x) = \mu_t(\{h_t^c : v(h_t^c) \leq x\} | h_t^f)$ and the revenue in period t after the history h_t^f is given by $p(F_{v, h_t^f})$.

We take a period- t history h_t^f for the firm to be the t -tuple of realized effort choices, $h_t^f \in \{L, H\}^t \equiv K_t^f$, describing the choices made in periods 0 through $t - 1$. Defining $K_0^f = \emptyset$, the set of all possible firm histories is $\mathcal{K}^f = \cup_{t \geq 0} K_t^f$. We consider only pure strategies. A pure strategy for a competent firm is a

strategy, $\tau : \mathcal{K}^f \rightarrow \{L, H\}$, giving the effort choice after observing history h_t^f . The collection of all pure strategies for the firm is denoted \mathcal{S}^f . An inept firm always exerts low effort. After a history h_t^f , the firm's continuation strategy $\tau|_{h_t^f} : \mathcal{K}^f \rightarrow \{L, H\}$ is given by $\tau|_{h_t^f}(h_{t'}^f) = \tau(h_t^f h_{t'}^f)$, where $h_t^f h_{t'}^f$ is the concatenation of h_t^f and $h_{t'}^f$.

If $\lambda > 0$, then with probability one, there will be an infinite number of replacement events, infinitely many of which will introduce new, competent owners into the game. Our description of histories ignores these replacement events. By restricting attention to firm histories in \mathcal{K}^f , we are requiring that a new competent owner, entering after the effort history h_t^f , behave in the same way as an existing competent owner after the same history. Notice, however, that a competent firm replacing an inept firm need not behave the same as would a continuing competent firm who chose high effort in the previous period. While our restriction may rule out some equilibria, any equilibrium under this assumption will again be an equilibrium without it.¹⁴ We will sometimes refer to a strategy $\tau \in \mathcal{S}^f$ as the competent owner's (or competent type's) strategy, although it describes the behavior of all new competent owners as well.

The strategies (τ, ν) will be an equilibrium if $\tau(h_t^f)$ is maximizing for competent owners after every effort history $h_t^f \in \mathcal{K}^f$, and consumers' beliefs about effort choice, ν , are correct. While the precise specification of these equilibrium conditions is conceptually straightforward, it involves tedious notation. The difficulty is that a mixed strategy, or a pure strategy in which the firm sometimes takes high effort and there are replacements that might be either competent or inept, gives rise to a random sequence of effort levels.¹⁵ Since the firm's strategy may call for different effort choices after different effort histories, a consumer must then use her outcome history to form a posterior over the effort histories that the firm has observed, yielding an intricate updating process. In particular, the posterior probability that a consumer assigns to the firm being competent is *not* a sufficient statistic for her history of outcomes.

The next two sections examine two important cases in which the equilibrium specification and analysis are greatly simplified by that fact that a consumer's

¹⁴If we were not restricting attention to pure strategies, then this restriction of histories would sacrifice no generality. In particular, if competent types of different generations behave differently, then consumers will treat this the same as they would a mixed strategy of a single competent type. The equivalence is an implication of Kuhn's theorem on the realization equivalence of mixed and behavior strategies.

¹⁵Even if $\tau(h_t^f) = H$ for all $h_t^f \in \mathcal{K}^f$ and $\lambda \in (0, 1)$ and $\theta \in (0, 1)$, the induced effort path will switch between high and low effort whenever a replacement changes the type of the firm's owner.

posterior belief concerning the firm's competency *is* a sufficient statistic for her outcome history. The first is the case of no replacements ($\lambda = 0$, discussed in the next section), where any pure strategy τ induces two possible deterministic sequences of effort choices, one corresponding to a competent firm and occurring with ex ante probability ϕ_0 and one corresponding to an inept firm and occurring with probability $1 - \phi_0$. The second, analyzed in Section 4, is a "high effort" equilibrium in a model with replacements, in which competent firms always choose high effort. While the pure strategy does not induce a deterministic sequence of effort choices in the presence of replacements, a consumer's posterior is still a sufficient statistic for her outcome history, since the competent firm takes the same equilibrium action after any realized effort-level history.

3. Getting Ahead

In this section we illustrate the role of replacements in leading to equilibria with reputations by examining the case in which $\lambda = 0$, so that replacements do not occur.

Suppose first that $\phi_0 = 1$, in which case the original and only owner of the firm is known to be competent. Then it is immediate that the only *pure-strategy* equilibrium calls for the firm to always exert low effort. A consumer who receives a bad outcome when the firm is supposed to exert high effort will attribute this bad outcome to an event in which the firm exerted high effort but the consumer received an unlucky draw from a distribution which gives a good outcome with probability $\rho_H < 1$. But bad outcomes will then prompt no punishments from consumers, creating an irresistible incentive for the firm to exert low effort, destroying the equilibrium.

The logic of the complete-information case also holds with incomplete information ($\phi_0 < 1$), in the absence of replacements, though the argument is more involved. Suppose the competent firm is following the pure strategy τ , and recall that an effort choice for the firm maps deterministically into a distribution of utility realizations for the consumers. Then τ induces the effort path, $h_\infty^f(\tau) \equiv \{e_t^\tau\}_{t=0}^\infty$, for the competent firm through the recursion: $e_0^\tau = \tau(\emptyset)$ and $e_t^\tau = \tau(h_{t-1}^f(\tau))$ for $t \geq 1$. Let $\varphi(\phi|x)$ denote the posterior probability that the firm is competent, given a prior probability of ϕ that the firm is competent and a single realized utility outcome of $x \in \{g, b\}$:

$$\varphi(\phi | g) = \frac{(e_t^\tau \rho_H + (1 - e_t^\tau) \rho_L) \phi}{(e_t^\tau \rho_H + (1 - e_t^\tau) \rho_L) \phi + \rho_L (1 - \phi)} \quad (3.1)$$

and

$$\varphi(\phi | b) = \frac{(e_t^\tau(1 - \rho_H) + (1 - e_t^\tau)(1 - \rho_L))\phi}{(e_t^\tau(1 - \rho_H) + (1 - e_t^\tau)(1 - \rho_L))\phi_t + (1 - \rho_L)(1 - \phi)}. \quad (3.2)$$

Extending the notation in an obvious manner, we write $\varphi(\phi | h_t^c)$ for the update from a prior ϕ after the history h_t^c . If a consumer believes the competent firm is following the pure strategy τ , his belief about the firm's effort choice after observing h_t^c is $\varphi(\phi_0 | h_t^c)e_t^\tau$.¹⁶

Definition 1. A pure strategy sequential equilibrium for the no-replacement game ($\lambda = 0$) is a pair (τ, v) satisfying

1. $\tau|_{h_t^f}$ is maximizing for the competent firm, given v , for all $h_t^f \in \mathcal{K}^f$, and
2. $v(h_t^c) = \varphi(\phi_0 | h_t^c)e_t^\tau$ for all $h_t^c \in \mathcal{K}^c$.

Proposition 1. If there are no replacements ($\lambda = 0$), there is a unique pure-strategy sequential equilibrium, and in this equilibrium the competent type exerts low effort in every period.

Proof. See appendix. ■

In the case of complete information, the firm cannot sustain high effort as a pure-strategy equilibrium choice because a consumer receiving a bad outcome attributes this outcome to an unlucky draw from the stochastic technology rather than shirking on the part of the firm, and hence inflicts no punishment on the firm. The possibility of an inept firm provides the firm with an incentive to exert high effort, since a consumer who receives a bad outcome now “punishes” the firm by increasing the probability with which the consumer thinks the firm is inept.

The difficulty is that a firm who builds a reputation does too good a job of building the reputation. Eventually, almost all consumers become almost certain that the firm is competent, in the sense that the posterior probability attached to a competent firm gets arbitrarily close to one for an arbitrarily large subset of consumers. The incentive to exert high effort arises only out of the desire to affect consumers' beliefs about the firm. As the posterior probability of a competent firm approaches unity, the effect of a good or bad outcome on this belief becomes smaller and smaller. At some point, the current outcome will have such a small

¹⁶Since every consumer history occurs with positive probability, all beliefs are given by Bayes rule, obviating the need to specify beliefs at unreached information sets or verify consistency when discussing sequential equilibria.

effect on the current belief that the cost c of high effort overwhelms the very small difference in beliefs caused by a good rather than a bad outcome, and the competent firm will find it optimal to revert to low effort. Consumers and the firm can foresee that this will happen, however, causing the equilibrium to unravel, so that the only pure-strategy equilibrium calls for only low effort to be exerted.

This effect was first described by Holmstrom [12] in the context of a signal-jamming model of managerial employment. The wage of the manager in Holmstrom's model is higher if the market posterior over the manager's type is higher, even if the manager chooses no effort. In contrast, the revenue of a firm in our model is higher for higher posteriors only if consumers also believe that the competent firm is choosing high effort. Since the market directly values managerial talent, Holmstrom's manager always has an incentive to increase effort, causing the market to overestimate his talent. In contrast with our Proposition 1, an equilibrium then exists in which the manager chooses effort levels that are higher than the myopic optimum. In agreement with spirit of our Proposition, however, this overexertion disappears over time, as the market's posterior concerning the manager's type approaches one. Holmstrom [12] uses the same device as we do, namely changing types, to obtain a sustained reputation effect.

Holmstrom [12] is concerned with the case of imperfect public monitoring, while our interest in imperfect *private* monitoring significantly complicates matters. More importantly, neither the market *nor the manager* knows the talent of the manager in Holmstrom's model, so that there is no incomplete information. Because the manager does not know his own talent, his effort cannot depend upon his talent, and the manager's evaluation of the profitability of effort reflects only market beliefs. As a result, the manager's expectation is that his efforts will not affect the average market beliefs about his talent. In contrast, our competent firms are more optimistic than the market, condition their actions on this information, and correctly believe that, by investing in high effort, their reputations will on average improve. We view this as an important aspect of reputation building, and this possibility provides much of the motivation for the analysis in Section 5.

Proposition 1 immediately extends to mixed strategies in which the competent firm's period- t randomization does not depend upon his private history.¹⁷ The implications of allowing more general mixed strategies, in which period- t effort choices depend upon earlier effort-choice realizations, are unclear. Sekiguchi [21] and Bhaskar and van Damme [4], discussed in Section 6, provide examples of

¹⁷Fudenberg, Levine, and Maskin [10] similarly observe that the construction of Abreu, Pearce, and Stacchetti [1], originally given for pure strategies in games with imperfect public monitoring, also holds for mixed strategies in which the mixtures depend only on the public signal.

games of imperfect monitoring in which pure strategy and mixed strategy equilibria have quite different properties. We have not been able to rule out mixed strategy equilibria but have not been able to construct a mixed equilibrium in which the firm sometimes attaches positive probability to high effort.¹⁸

4. Keeping Your Place

We now let $\lambda > 0$ and show that replacements allow the firm to maintain a reputation. Since $\lambda > 0$ and $\theta \in (0, 1)$, consumers can never become too confident about the type of firm they are facing. No matter what history a consumer has observed, she assigns a probability of at least $\lambda\theta$ and no more than $1 - \lambda(1 - \theta)$ to the firm's being competent. These bounds capture the intuition that a consumer, receiving a bad outcome after a very long string of good outcomes, will entertain the explanation that "something about the firm may have changed" as well as the possibility that "it's obviously a good firm and I've been unlucky."

We examine a pure-strategy equilibrium in which the competent firm always chooses high effort. The behavior of the competent firm is independent of the effort history in this case, and so a consumer's posterior probability that the firm is competent is a sufficient statistic for her history of outcomes.

We now let $\varphi(\phi|x)$ denote the posterior probability that the firm is competent, after the consumer has received a single realized utility outcome, $x \in \{g, b\}$, given a prior probability of ϕ that the firm is competent and that the competent firm chooses high effort:

$$\varphi(\phi | g) = (1 - \lambda) \frac{\rho_H \phi}{\rho_H \phi + \rho_L (1 - \phi)} + \lambda \theta \quad (4.1)$$

and

$$\varphi(\phi | b) = (1 - \lambda) \frac{(1 - \rho_H) \phi}{(1 - \rho_H) \phi + (1 - \rho_L)(1 - \phi)} + \lambda \theta. \quad (4.2)$$

These differ from (3.1)–(3.2) by incorporating the information that a competent firm always chooses high effort and incorporating the possibility of replacement.

¹⁸To support high effort, current high effort must lead to higher expected future prices than does low effort. With pure strategies, such a link cannot arise, as consumer expectations evolve deterministically and independently of the firm's actions. With mixed strategies, the possibility arises that the firm mixes between high and low effort in period t , and then is more likely to choose high effort if period $t + 1$ if the current realization is high effort. Consumers would then use their outcomes to infer the likelihood that the current choice was high effort and hence that future choices are likely to be high effort. As a result, current high effort would cause consumers to have a higher expectation of high effort in the next period, and hence lead to a higher price in the next period, raising the possibility of a mixed equilibrium.

Definition 2. A profile (τ, v) is a high effort sequential equilibrium in the game with replacements if

1. $\tau(h_t^f) = H$ is maximizing for the competent firm, given v , for all $h_t^f \in \mathcal{K}^f$,
and
2. $v(h_t^c) = \varphi(\phi_0 \mid h_t^c)$ for all $h_t^c \in \mathcal{K}^c$.

Proposition 2. Suppose $\lambda \in (0, 1)$, let ϕ' solve $\varphi(\phi \mid b) = \phi$, and let ϕ'' solve $\varphi(\phi \mid g) = \phi$. If $\phi_0 \in [\phi', \phi'']$, then there exists $\bar{c} > 0$ such that a high effort equilibrium exists for all $0 \leq c < \bar{c}$.

The proof, contained in the appendix, is complicated, because we have placed little structure on the revenue function. A more direct proof of this result is possible if we assume that the revenue function has the additional structure used in the next section (see Proposition 3).

Propositions 1 and 2 combine to provide the seemingly paradoxical result that it can be good news for the firm to have consumers constantly fearing that the firm might “go bad.” The purpose of a reputation is to convince consumers that the firm is competent and will produce high quality. As we saw in the previous section, the problem with maintaining a reputation in the absence of replacements is that the firm essentially succeeds in convincing consumers it is competent. If replacements continually introduce the possibility that the firm has turned bad, then the firm can never do “too good” a job of convincing consumers it is competent. Formally, there is an upper bound, short of unity, on the posterior ϕ_t^i .¹⁹ The incentive to exert high effort in order to convince consumers that the firm is still competent always remains, opening the door to an equilibrium in which the competent firm exerts high effort.

An equilibrium in which the competent firm exerts high effort exists as long as the cost of high effort is sufficiently low. This restriction on costs is expected. We can always ensure that high effort will not be undertaken by making its cost prohibitively high. The smaller is the probability of a replacement, and the smaller the probability that a replacement is inept, the smaller must the cost of high effort be in order to support an equilibrium in which the firm chooses high effort. To see why, notice that a necessary equilibrium condition is that the difference between the value of high effort and the value of low effort exceed the cost of high effort. These value functions approach each other as the ϕ_t^i approach one, because the values diverge only through the effect of current outcomes on future

¹⁹There is also a lower bound, and the assumption $\phi_0 \in [\phi', \phi'']$ ensures that consumers' prior probabilities that the firm is competent lie between these bounds.




Figure 5.1: The different effort levels. As drawn, E and M are equally efficient.

expectations, and current outcomes have very little effect on future expectations when consumers are currently quite sure of the firm's type (i.e., current posteriors are very close to one). The smaller is the probability of an inept replacement, the closer can the posterior expectation of a competent firm approach unity, and hence the smaller must be the cost in order to support high effort.

5. Catching Up

We now give the competent firm the ability to manage its reputation by considering three possible effort levels. We retain the choices of high effort (H) and low effort (L), where high effort is efficient and the inept firm necessarily exerts low effort. We add an additional effort level that gives a higher probability of a good utility outcome, at higher cost, than does low effort, but is inefficient.

The inefficient effort level may be either too high or too low. We illustrate these possibilities in Figure 5.1. Each of the line segments identifies a combination of the probability of a good utility outcome and a cost level that preserves the

difference $\rho - c$. Every probability/cost combination on a lower line segment is more efficient than every combination on a higher segment. The points H and L correspond to the high-effort and low-effort combinations considered in previous sections. Effort M is an effort level that would allow a reputation to be built, because it gives a higher probability of a good outcome than does L , but is inefficiently low. We refer to this as *medium effort* because it gives a smaller probability of a good outcome than does high effort. Effort E is an inefficient effort level that is *excessive*, in that it gives a higher probability of a good outcome than high effort.

If the competent firm builds a reputation by choosing an inefficient effort level, is this more likely to involve an inefficiently high or inefficiently low effort level? Are there cases in which a firm chooses a relatively high effort level to build a reputation, and then coasts on this reputation by choosing a lower effort level?

Let e' denote the third, inefficient, effort level. The action space for the competent firm is now $\{L, H, e'\}$ and a strategy for the competent type is $\tau : \mathcal{K}^f \rightarrow \{L, H, e'\}$, where $\mathcal{K}^f \equiv \cup_{t \geq 0} \{L, H, e'\}^t$. A belief function of the consumer is now a function $v : \mathcal{K}^c \rightarrow \Delta\{L, H, e'\}$, which assigns a probability distribution over possible effort choices to every consumer history (the space of which is unchanged, since consumer outcomes are unchanged).

Revenue is a function of the distribution of consumer beliefs over effort choices, i.e., $p : \Delta(\Delta(\{L, H, e'\})) \rightarrow \mathbb{R}_+$. Let Λ be a probability measure over the possible beliefs of consumers over the effort choices, so that $\Lambda \in \Delta(\Delta(\{L, H, e'\}))$. We impose the following additional structure on the revenue function: there exists a continuous function $\rho : \Delta(\{L, H, e'\}) \rightarrow \mathbb{R}_+$ such that

$$p(\Lambda) = \int \rho d\Lambda, \quad (5.1)$$

where ρ is increasing in the expected probability of a good outcome. One possibility for ρ is that $\rho(\alpha)$ is the maximum price that a consumer would be willing to pay, given a belief that the firm's distribution over effort levels is $\alpha \in \Delta(\{L, H, e'\})$, in which case $\rho(\alpha) = \alpha_L \rho_L + \alpha_H \rho_H + \alpha_{e'} \rho_{e'}$. Condition (5.1) asserts that, however individual prices are formed, revenue is a weighted average of these individual prices.

The remainder of this subsection identifies conditions for equilibria in which the competent firm always chooses the same effort level $e \in \{L, H, e'\}$. The definition of equilibrium for this constant-effort case is the obvious modification of Definition 2. Because the competent firm's effort is constant, a consumer's posterior probability that the firm is competent is again a sufficient statistic for her history of outcomes.

We first collect some information concerning beliefs and revenues. Let G_t be the distribution in period t of consumer posteriors that the firm is competent. The replacement process ensures that the probability assigned to the firm being competent must be at least $\lambda\theta$. We denote by \mathcal{G} the collection of distributions that assign a probability to the firm of being competent of at least $\lambda\theta$. If all consumers expect all competent firms to be choosing e , then

$$G_{t+1}(\phi) = \rho_e G_t(\varphi^{-1}(\phi | g)) + (1 - \rho_e) G_t(\varphi^{-1}(\phi | b)), \quad (5.2)$$

where $\varphi(\phi | x)$ is the Bayes' update under effort level e .

Let $V(\cdot | e) : \mathcal{G} \rightarrow \mathbb{R}_+$ denote the discounted expected value of the revenue generated given that consumers expect all competent firms to be choosing e , as a function of the distribution of posteriors. Then,

$$V(G_t | e) = \sum_{\tau=0}^{\infty} (\beta(1 - \lambda))^\tau p(\Lambda_{t+\tau}^e),$$

where $\Lambda_{t+\tau}^e$ is the probability measure on $\Delta(\{L, H, e'\})$ implied by competent firms only choosing e and the distribution of consumer beliefs being $G_{t+\tau}$ (i.e., $\Lambda_{t+\tau}^e(\{\alpha : \alpha_e \leq x, \alpha_L = 1 - \alpha_e\}) = G_{t+\tau}(x)$). Using (5.1), we have

$$V(G_t | e) = \sum_{\tau=0}^{\infty} (\beta(1 - \lambda))^\tau \int_0^1 \rho(\phi \delta_e + (1 - \phi) \delta_L) dG_{t+\tau}(\phi), \quad (5.3)$$

where δ_x is the degenerate measure that assigns probability one to x .

We establish that $V(\cdot | e)$ is a positive linear functional on \mathcal{G} . To do this, suppose G_{t+1} and \hat{G}_{t+1} are the updated distributions, according to (5.2), from G_t and \hat{G}_t , respectively. Then if the period t posterior distribution is $\gamma G_t + (1 - \gamma) \hat{G}_t$, the period $t + 1$ distribution is

$$\begin{aligned} & \rho_e \left\{ \gamma G_t(\varphi^{-1}(\phi | g)) + (1 - \gamma) \hat{G}_t(\varphi^{-1}(\phi | g)) \right\} \\ & + (1 - \rho_e) \left\{ \gamma G_t(\varphi^{-1}(\phi | b)) + (1 - \gamma) \hat{G}_t(\varphi^{-1}(\phi | b)) \right\} \\ = & \gamma \left\{ \rho_e G_t(\varphi^{-1}(\phi | g)) + (1 - \rho_e) G_t(\varphi^{-1}(\phi | b)) \right\} \\ & + (1 - \gamma) \left\{ \rho_e \hat{G}_t(\varphi^{-1}(\phi | g)) + (1 - \rho_e) \hat{G}_t(\varphi^{-1}(\phi | b)) \right\} \\ = & \gamma G_{t+1}(\phi) + (1 - \gamma) \hat{G}_{t+1}(\phi). \end{aligned}$$

The linearity of $V(\cdot | e)$ then follows from (5.3).

Writing $V_e(\phi)$ for $V(\delta_\phi | e)$, we thus have

$$V(G_t | e) = \int V_e(\phi) dG_t(\phi).$$

It is useful to rewrite this as

$$V(G_t | e) = p(\Lambda_t^e) + \beta(1 - \lambda) \int \{\rho_e V_e(\varphi(\phi | g)) + (1 - \rho_e) V_e(\varphi(\phi | b))\} dG_t(\phi).$$

We now turn to equilibrium conditions. Consider the competent firm, after some effort history $h_t^f \in \mathcal{K}^f$ (and hence consumer distribution G_t) considering an effort level $\hat{e} \neq e$ for one period. The expected value of the revenue stream from this effort is

$$V(G_t | \hat{e}; e) = p(\Lambda_t^e) + \beta(1 - \lambda) \int \{\rho_{\hat{e}} V_e(\varphi(\phi | g)) + (1 - \rho_{\hat{e}}) V_e(\varphi(\phi | b))\} dG_t(\phi).$$

It will then be an equilibrium for the competent firm to always take effort e if, for all distributions G ,

$$V(G | e) \geq V(G | \hat{e}; e)$$

for all $\hat{e} \in \{L, H, e'\}$. This is equivalent to

$$(\rho_e - \rho_{\hat{e}})\beta(1 - \lambda) \int \{V_e(\varphi(\phi | g)) - V_e(\varphi(\phi | b))\} dG(\phi) \geq c_e - c_{\hat{e}}. \quad (5.4)$$

5.1. Inefficiently Low Effort

We first consider the implications of allowing an inefficient effort level that is lower than high effort, denoted by M and referred to as medium effort. The cost to the competent firm of effort $e \in \{L, M, H\}$ is c_e . We assume

$$\rho_H > \rho_M > \rho_L, \quad c_H > c_M > c_L = 0, \quad (5.5)$$

and

$$\rho_H - c_H > \rho_M - c_M > \rho_L. \quad (5.6)$$

The competent firm would find both high effort and medium effort better than low effort, given that it could commit to such effort levels and consumer expectations were formed accordingly. However, medium effort would be inferior to high effort.

By making medium effort both quite costly and quite inefficient, it is easy to construct examples in which medium effort can never be taken in equilibrium. The more interesting possibility is that there exist markets with equilibria in which competent firms always exert medium effort. Why would a firm use medium

effort to build a reputation when high effort is more efficient? Given a candidate equilibrium in which medium effort is always exerted, there are two obstacles to the profitability of high effort. The first involves timing: high effort requires an immediate expenditure in return for favorable adjustments in beliefs that are not realized until the future. The second involves consumer expectations: if the candidate equilibrium involves medium effort, then high effort increases the likelihood that consumers attach to the event that the firm is a competent firm *exerting medium effort*. There are cases in which high effort would be undertaken if it could quickly convince the consumers that the firm was competent and would exert future high effort, but the combination of the delay in consumer reactions and the continued equilibrium expectation of medium effort can make high effort unprofitable, leading to an equilibrium in which the firm always exerts medium effort.

Proposition 3. Let (5.5)–(5.6) hold. Then, for any $\lambda > 0$,

(3.1) There exist markets with equilibria in which the competent firm always chooses high effort, but no equilibria in which the firm always chooses medium effort. This is the case when medium effort is relatively efficient, but the cost of medium effort is relatively high.

(3.2) There exist markets with equilibria in which the competent firm always chooses medium effort, but no equilibria in which the firm always chooses high effort (though there would be an equilibrium in which the firm always chooses high effort if medium effort were not feasible). This is the case when medium effort is relatively efficient, but the cost of medium effort is relatively low, and when the firm's discount factor is relatively low.

(3.3) There exist markets with pure-strategy equilibria in which the competent firm always chooses high effort, as well as pure-strategy equilibria in which the competent firm always chooses medium effort. This is the case when medium effort is relatively inefficient and relatively inexpensive.

Result (3.3) highlights the role of consumer expectations. In this case, the firm would find it advantageous to take high effort if doing so allowed it to establish a reputation for choosing high effort, but will not do so if the payoff is a reputation for taking medium effort.

Proof. (3.1): From (5.4), in a candidate high effort equilibrium, medium effort is not a profitable deviation if, for all feasible distributions G ,²⁰

$$\beta(1 - \lambda) \int \{V_H(\varphi(\phi | g)) - V_H(\varphi(\phi | b))\} dG(\phi) \geq \frac{(c_H - c_M)}{(\rho_H - \rho_M)}, \quad (5.7)$$

²⁰ A feasible distribution is one that can be generated by perpetual medium effort choices from an initial degenerate distribution placing all mass on ϕ_0 .

and low effort is not a profitable deviation if, for all feasible distributions G ,

$$\beta(1 - \lambda) \int \{V_H(\varphi(\phi | g)) - V_H(\varphi(\phi | b))\} dG(\phi) \geq \frac{c_H}{(\rho_H - \rho_L)}. \quad (5.8)$$

On the other hand, medium effort is destabilized as an equilibrium by high effort if for at least one feasible distribution G_t ,

$$\beta(1 - \lambda) \int \{V_M(\varphi(\phi | g)) - V_M(\varphi(\phi | b))\} dG_t(\phi) > \frac{(c_H - c_M)}{(\rho_H - \rho_M)}. \quad (5.9)$$

Note that $V_H(\phi)$ is independent of c_M and ρ_M . Fixing all parameters other than c_H , c_M and ρ_M , $V_H(\varphi(\phi | g)) - V_H(\varphi(\phi | b))$ is bounded away from zero for $\phi \in [\lambda\theta, 1 - \lambda(1 - \theta)]$, and so (5.8) is satisfied strictly for c_H sufficiently small. Now, any choice of (c_M, ρ_M) so that $(c_H - c_M)(\rho_H - \rho_L) = c_H(\rho_H - \rho_M)$ and $\rho_M < \rho_H$ ensures that (5.7) is also satisfied strictly. Note that such a choice of (c_M, ρ_M) must lie on the line segment LH in Figure 5.2, and is necessarily inefficient. Moreover, since $V_M(\phi)$ is a function of ρ_M but not ρ_H , we have $V_M \rightarrow V_H$ as $\rho_M \rightarrow \rho_H$. Hence, the previous choice of (c_M, ρ_M) can be made with ρ_M sufficiently close to ρ_H that (5.9) also holds.

(3.2): From (5.4), in a candidate medium effort equilibrium, high effort is not a profitable deviation if, for all feasible distributions G ,

$$\beta(1 - \lambda) \int \{V_M(\varphi(\phi | g)) - V_M(\varphi(\phi | b))\} dG(\phi) \leq \frac{(c_H - c_M)}{(\rho_H - \rho_M)}, \quad (5.10)$$

and low effort is not a profitable deviation if, for all feasible distributions G ,

$$\beta(1 - \lambda) \int \{V_M(\varphi(\phi | g)) - V_M(\varphi(\phi | b))\} dG(\phi) \geq \frac{c_M}{(\rho_M - \rho_L)}. \quad (5.11)$$

On the other hand, high effort is destabilized as an equilibrium by medium effort if for at least one feasible distribution G_t ,

$$\beta(1 - \lambda) \int \{V_H(\varphi(\phi | g)) - V_H(\varphi(\phi | b))\} dG_t(\phi) < \frac{(c_H - c_M)}{(\rho_H - \rho_M)}. \quad (5.12)$$

Fix all of the parameters except β , c_H , c_M and ρ_M . By choosing β small, the left sides of (5.10) and (5.12) can be guaranteed to be smaller than 1, for all possible ρ_M . By choosing c_H to be sufficiently small, (5.8) will be satisfied, so that high effort would be an equilibrium if medium were not available. Finally, we can then choose (c_M, ρ_M) so that $(c_H - c_M)/(\rho_H - \rho_M)$ is arbitrarily close to 1. This corresponds to making the slope of the segment MH in Figure 5.2 close




Figure 5.2: The slopes of the edges of the triangle are $c_H/(\rho_H - \rho_L)$, $c_M/(\rho_M - \rho_L)$, and $(c_H - c_M)/(\rho_H - \rho_M)$.

to 1, and ensures that (5.10) and (5.12) hold. We are then free to choose c_M sufficiently small, while preserving $\rho_M - c_M$, so that (5.11) and (5.5)–(5.6) hold.

(3.3) Fix values of the parameters, other than ρ_M and c_M , such that the relevant inequalities in (5.5)–(5.6) hold and such that there exists an equilibrium, in the model *without* medium effort, in which the good firm always takes high effort and in which

$$\beta(1 - \lambda) \int \{V_H(\varphi(\phi | g)) - V_H(\varphi(\phi | b))\} dG(\phi) \geq \frac{c_H}{(\rho_H - \rho_L)} + \eta,$$

for some $\eta > 0$ and for all $G \in \mathcal{G}$. Then let $\rho_M = \rho_L + \varepsilon$ and $c_M = 0$. Then, for sufficiently small ε , we have

$$\frac{c_H - c_M}{\rho_H - \rho_M} = \frac{c_H}{\rho_H - \rho_L - \varepsilon} < \frac{c_H}{\rho_H - \rho_L} + \eta,$$

ensuring that high effort remains an equilibrium when medium effort is also added to the model. In addition, setting ε sufficiently small ensures that

$$\beta(1 - \lambda) \int \{V_M(\varphi(\phi | g)) - V_M(\varphi(\phi | b))\} dG(\phi) < \frac{c_H}{(\rho_H - \rho_M)},$$

for all $G \in \mathcal{G}$, since the left hand side converges to zero as ε converges to zero. This is a reflection of the difficulty that if the equilibrium calls for medium effort to be exerted, then choosing high effort only enhances one's reputation for choosing medium and not high effort. Setting such an ε , we note that we still have

$$\beta(1 - \lambda) \int \{V_M(\varphi(\phi | g)) - V_M(\varphi(\phi | b))\} dG(\phi) \geq \frac{c_M}{(\rho_M - \rho_L)} = 0,$$

for all $G \in \mathcal{G}$. These inequalities ensure that it is an equilibrium for the competent firm to always take medium effort. Because each inequality in the proof is strict, we can preserve the inequalities while increasing c_M slightly to ensure that (5.5)–(5.6) are satisfied. ■

The existence of the multiple equilibria in Proposition 3.3, one in which medium effort is always taken and one in which high effort is always taken, is established with the help of a pair of arguments, each resembling the proof of Proposition 2 (establishing the existence of a high-effort equilibrium in the two-effort case). By making ρ_M and c_M small and hence close to ρ_L and $c_L (= 0)$, the demonstration that a competent firm in a high-effort equilibrium will not deviate to medium effort parallels the demonstration that he will not deviate to low effort. But by letting c_M be especially small relative to ρ_M , the argument

that a firm in a medium-effort equilibrium will not deviate to low effort parallels the proof that a high-effort equilibrium exists in the two-effort model.

Propositions 3.1 and 3.2 address a different issue. In each case, we exploit the fact that medium effort can be relatively efficient to construct examples in which only one of a medium-effort and high-effort equilibrium exists. Which of these two equilibria exists hinges on the ratio $c_M/(\rho_M - \rho_L)$, with only a high-effort equilibrium existing when this ratio is high and only a low-effort equilibrium existing when it is low. These results serve as a reminder that the costs of effort are borne immediately while the benefits are delayed. An inefficient, medium effort level can then be preferred to an efficient, high effort level if the former has the advantage of involving especially small immediate expenditures (i.e., $c_M/(\rho_M - \rho_L)$ is small), but is unlikely to be preferred without this advantage.

5.2. Inefficiently High Effort

The previous proposition's focus on the cost-benefit ratio $c_M/(\rho_M - \rho_L)$ directs our attention to a fundamental disadvantage of excessive effort levels: the cost of excessive effort is bounded above the cost of high effort. Unlike medium effort, the immediate costs of excessive effort cannot be arbitrarily small.

Suppose that the three effort levels available to the firm are $\{L, H, E\}$, where E denotes excessive effort, with effort $e \in \{L, H, E\}$ yielding a good outcome with probability ρ_e and yielding a bad outcome with complementary probability, at cost c_e . We assume

$$\rho_E > \rho_H > \rho_L \quad c_E > c_H > c_L = 0, \quad (5.13)$$

and

$$\rho_H - c_H > \rho_E - c_E > \rho_L - c_L. \quad (5.14)$$

The competent firm thus finds both excessive and high effort better than low effort, given that it could commit to such effort levels and consumer expectations were formed accordingly. However, high effort is more efficient than excessive effort.

The disadvantage of excessive effort is that c_E must be relatively large. Notice in particular that the ratio $(c_E - c_H)/(\rho_E - \rho_H)$ must exceed one, placing a lower bound on the ratio of the marginal cost of excessive (rather than high) effort to the marginal benefit of excessive effort. This lower bound implies that for any specification of parameters satisfying (5.13)–(5.14), there will be cases in which an excessive-effort equilibrium does not exist because high effort is more attractive:

Proposition 4. Fix $\rho_E, \rho_H, \rho_L, c_E, c_H$, and c_L consistent with (5.13)–(5.14). Then the set of values of λ for which an equilibrium exists in which the competent firm always takes excessive effort is a subset of the corresponding set for the model in which high effort is unavailable, and is a strict subset whenever the latter set is nonempty. In particular, the smallest value of λ for which an excessive-effort equilibrium exists, if one does, is strictly higher when high effort is available than when it is not available.

We prove this result by showing that, for small values of λ , an excessive-effort equilibrium is destabilized by the temptation to exert high (but not low) effort. This result contrasts with the case of medium effort, where we can find examples in which the cost of medium effort is relatively small and in which the option of exerting high effort is suboptimal, for *all* values of λ , so that the presence or absence of high effort is irrelevant for determining which values of λ are consistent with the existence of a medium-effort equilibrium.

Proof. If only low and excessive efforts are available, then an equilibrium in which the competent firm always takes excessive effort requires, for all feasible G_t ,

$$\beta(1 - \lambda) \int \{V_E(\varphi(\phi | g)) - V_E(\varphi(\phi | b))\} dG_t(\phi) \geq \frac{c_E}{(\rho_E - \rho_L)}, \quad (5.15)$$

where $c_E/(\rho_E - \rho_L) < 1$. If high effort is also available, then an equilibrium in which the competent firm always takes excessive effort also requires

$$\beta(1 - \lambda) \int \{V_E(\varphi(\phi | g)) - V_E(\varphi(\phi | b))\} dG_t(\phi) \geq \frac{(c_E - c_H)}{(\rho_E - \rho_H)} > 1 \quad (5.16)$$

for all feasible G_t . Now, G_t converges weakly to $\delta_{\phi''}$, where ϕ'' solves $\phi'' = \varphi(\phi'' | g)$. Moreover, $\phi'' \rightarrow 1$ as $\lambda \rightarrow 0$. Since $\varphi(\phi | g) - \varphi(\phi | b)$ is small for ϕ close to one, by choosing λ sufficiently small, (5.16) is eventually violated. The result then follows from noting that (5.16) is a strictly more demanding condition than (5.15). ■

The comparison between inefficiently low and inefficiently high effort levels is instructive. Fix ρ_H, ρ_L, c_H and c_L . Then for every positive value of λ , no matter how small, we can find an equilibrium in which the competent firm always exerts medium effort. However, for sufficiently small λ , there is no equilibrium in which excessive effort is always exhibited. In this sense, if the firm is to inefficiently build a reputation, the inefficiency is more likely to involve too little effort. The

key is that an equilibrium in which the firm always exerts a given effort level can be broken by the temptation to choose a lower effort level if the value functions following the two effort levels are too close together. By giving medium effort a sufficient cost advantage over high effort, we can always ensure that the value functions are not too close together, no matter how large is λ . By definition, however, excessive effort must always exhibit a cost disadvantage over high effort. As a result, small values of λ will create cases in which high effort is preferred to excessive effort, ensuring there is no equilibrium in which the competent firm always exerts excessive effort.

5.3. Spending a Reputation

In this section, we discuss the possibility that the firm may sometimes, but not always, choose excessive effort. The competent firm may have an incentive to take excessive effort when many of the consumers' posteriors are low (such as after replacing a long-lived inept firm), since excessive effort increases the firm's reputation with a larger fraction of the consumers than does high effort. The competent firm may then eventually revert to high effort, "enjoying" the fruits of this reputation building, when many of the consumers are more optimistic about the firm.

We first verify the intuition that the firm may have an incentive to take excessive effort when many of the consumers' posteriors are low.

Proposition 5. *Suppose $\rho(\alpha) = \alpha_L \rho_L + \alpha_H \rho_H + \alpha_E \rho_E$. There exist markets (preserving (5.13)–(5.14)) for which excessive effort is a profitable deviation from the always-high-effort profile, whenever enough consumers' posteriors are low.*

Proof. We first show that excessive effort is a profitable deviation in the initial period, if the common initial belief, ϕ_0 , is not too large. Excessive effort is a profitable deviation in the initial period if

$$\beta(1 - \lambda)V_H(\varphi(\phi_0 | g)) - V_H(\varphi(\phi_0 | b)) \geq \frac{c_E - c_H}{\rho_E - \rho_H} > 1. \quad (5.17)$$

For fixed ρ_H and ρ_E , we can choose $c_E - c_H$ arbitrarily close to $\rho_E - \rho_H$ while still preserving (5.13)–(5.14). Thus, by appropriate parameter choices, (5.17) is implied by

$$\beta(1 - \lambda)(V_H(\varphi(\phi_0 | g)) - V_H(\varphi(\phi_0 | b))) > 1. \quad (5.18)$$

Suppose temporarily that $\lambda = 0$ and $\rho_L = 1 - \rho_H$. Note that $\rho_H = 1 - \rho_L$ implies $\varphi(\phi | g | b) = \phi$. Let $\Phi \equiv \{\phi_\tau\}_{\tau=-\infty}^{\infty}$ denote the set of beliefs defined by

$$\varphi(\phi_\tau | g) = \phi_{\tau+1}$$

and

$$\varphi(\phi_\tau | b) = \phi_{\tau-1},$$

for every integer τ (ϕ_0 having already been fixed). Of course $\phi_t > \phi_{\tau-1}$. Let $h(t, \phi_\tau | \phi)$ be the probability that a consumer period- t posterior is given by ϕ_τ , given a period-zero belief of $\phi \in \Phi$. Since $h(t, \phi_\tau | \varphi(\phi | b)) = h(t, \phi_{\tau+2} | \varphi(\phi | g))$,

$$V_H(\varphi(\phi_0 | g)) = V_H(\phi_1) = \sum_{t=0}^{\infty} \beta^t \sum_{\tau} h(t, \phi_\tau | \phi_1) (\phi_\tau \rho_H + (1 - \phi_\tau) \rho_L)$$

and

$$V_H(\varphi(\phi_0 | b)) = V_H(\phi_{-1}) = \sum_{t=0}^{\infty} \beta^t \sum_{\tau} h(t, \phi_\tau | \phi_1) (\phi_{\tau-2} \rho_H + (1 - \phi_{\tau-2}) \rho_L),$$

and hence

$$V_H(\phi_1) - V_H(\phi_{-1}) = \sum_{t=0}^{\infty} \beta^t \sum_{\tau} h(t, \phi_\tau | \phi_1) (\rho_H - \rho_L) (\phi_\tau - \phi_{\tau-2}).$$

Fix $\varepsilon > 0$. Lemma 1 (in the appendix) demonstrates that the sum

$$\sum_{t=0}^T \sum_{\tau} h(t, \phi_\tau | \phi_1) (\phi_\tau - \phi_{\tau-2})$$

is larger than $2(1 - \phi_0) - \varepsilon$ for sufficiently large T . By taking β close to 1 (and maintaining $\lambda = 0$), we have

$$\beta (V_H(\varphi(\phi_0 | g)) - V_H(\varphi(\phi_0 | b))) > 2(\rho_H - \rho_L)(1 - \phi_0) - \varepsilon/2.$$

We can then ensure that (5.18) holds, with strict inequality, by taking ρ_H sufficiently large (which implies ρ_L small), and ϕ_0 not too large.

We now relax the restrictions $\lambda = 0$ and $\rho_H = 1 - \rho_L$. All the relevant functions are continuous, so that for λ close to 0 and ρ_H close to $1 - \rho_L$, (5.18) still holds (although it is no longer the case that $\varphi(\varphi(\phi | g) | b) = \phi$).

Consider now parameter values such that (5.18) holds for $\phi_0 = \phi'$, where ϕ' solves $\varphi(\phi | b) = \phi$. Since the distribution of consumer posteriors weakly converges to $\delta_{\phi'}$ in the length of an inept firm's life, and the competent firm's value function, $V(\cdot | e) : \mathcal{G} \rightarrow \mathfrak{R}$, is continuous, a competent firm that replaces a sufficiently long-lived inept firm will also find it profitable to deviate to excessive effort. ■

By taking λ sufficiently small (as was done in the proof of Proposition 5), we can preserve the result of Proposition 5 while also ensuring that there is no equilibrium in which the competent firm *always* exerts excessive effort. The problem with sustaining permanent excessive effort is the familiar one that when a competent firm is long-lived, eventually most consumers are quite confident about the type of owner, and the competent firm now has no incentive to choose excessive effort.

We thus have a market in which, if the competent firm is to build a reputation, it must do so by sometimes exerting high effort and sometimes exerting excessive effort. Our intuition is that competent firms will exert excessive effort whenever replacing a long-lived inept firm, and hence facing the task of quickly building a relatively low reputation, with excessive effort giving way to high effort as more typical competent-firm reputations are attained.

We have not constructed such equilibria because of the significant difficulties that arise in calculating best replies when the firm's actions are not constant. In particular, suppose that the set of firm histories is partitioned into two sets, \mathcal{K}_H^f and \mathcal{K}_E^f , and that firm behavior is $\tau(h_t^f) = H$ if $h_t^f \in \mathcal{K}_H^f$ and $\tau(h_t^f) = E$ if $h_t^f \in \mathcal{K}_E^f$. In order to calculate her beliefs over the firm's effort choice in some period t , consumer i must now not only assign a probability to the firm being competent in this period (as before), but also to the firm's history being in the set \mathcal{K}_H^f . In addition, consumer i 's updating of the probability assigned to the firm being competent will also depend on the probability she assigns to the firm's history being in the set \mathcal{K}_H^f . The payoff to the competent firm of different effort levels then depends on the evolution of consumer beliefs over the firm's type and the set \mathcal{K}_H^f . Finally, note that simply appealing to standard existence theorems does not establish even the existence of such an equilibrium, since the trivial behavior in which the competent firm always chooses low effort is always an equilibrium.

6. A Model with Random Quality and The Imperfect Monitoring Literature

This section explores the relationship of our analysis to the literature on repeated games of imperfect monitoring. To do so, it is helpful to sketch a straightforward generalization of our model. Suppose that the firm's effort choice now *randomly* determines the quality of the good, which in turn randomly determines consumer outcomes. The good can be either high quality, which we refer to as a success (s), or low quality, which we refer to as a failure (f). An effort level of e yields a success with probability ψ_e and a failure with probability $1 - \psi_e$, with $\psi_E >$

$\psi_H > \psi_M > \psi_L$. The quality of the good is observed by the firm, but not by the consumers. A success provides a good outcome with probability ζ_s and a bad outcome with probability $1 - \zeta_s$, while a failure provides a good outcome with probability ζ_f and a bad outcome with probability $1 - \zeta_f$, where $1 \geq \zeta_s > \zeta_f \geq 0$.

Our earlier analysis corresponds to setting $\psi_H = 1$ and $\psi_L = 0$. However, extending Propositions 1 and 2 to the general case in which $0 < \psi_L < \psi_H < 1$ requires only some more elaborate notation in the proof. Moreover, Propositions 3-5 hold, without any modification, if we interpret ρ_e as the probability with which effort level e gives a good utility outcome, so that $\rho_e = \psi_e \zeta_s + (1 - \psi_e) \zeta_f$.

The interpretation of our formal analysis becomes richer in the context of this more elaborate model. The evolution of consumer histories, and hence the evolution of consumer posteriors facing a single competent firm, is now stochastic rather than deterministic. The behavior described in Section 5.3, for example, is now one in which a single firm may *alternate* between periods of high effort, designed to capitalize on a healthy reputation, and periods of excessive effort, in which the firm rebuilds a reputation depleted by a series of unfortunate failures.

Now consider the literature on repeated games and reputations. The simplest special case is obtained by setting $\phi_0 = 1$ and $\lambda = 0$, so that there is one firm, who is competent, and $\zeta_s = 1$, $\zeta_f = 0$, $\psi_H = 1$, and $\psi_L = 0$, so that all consumers observe the actions of the firm. The result is a *repeated game with perfect monitoring*. Even though the consumers behave as short-run players, this game has many equilibria (Fudenberg, Kreps, and Maskin [8]). High effort can be supported (if the firm is sufficiently patient) by an equilibrium profile in which the firm begins with high effort and chooses high effort as long as it has previously chosen high effort, while choosing low effort if it has ever chosen low effort in the past; and in which consumers expect high effort as long as high effort has been chosen in the past, and expect low effort if low effort has ever been chosen in the past. It is then convenient to describe the firm as maintaining a reputation for high effort, with any deviation instantly destroying that reputation. In this setting, a reputation consists of beliefs about the future *actions* of the firm. For example, we might speak of firms as maintaining a reputation for not cutting prices, or a monetary authority maintaining a reputation for not inflating, or countries maintaining reputations for setting low tariffs.

Consumers may be unable to perfectly observe the firm's actions. We capture such a case by setting $0 < \psi_L < \psi_H < 1$, but maintaining the parameter restrictions $\phi_0 = 1$, $\lambda = 0$, $\zeta_s = 1$, and $\zeta_f = 0$. This yields a game of *imperfect public monitoring* (Abreu, Pearce and Stacchetti [1]). The critical feature in this model is that both consumers and the firm observe the noisy signal of firm behavior, and hence can condition behavior on that signal. It is always an equilibrium in

this model for the firm to always choose low effort, and for consumers to expect low effort in every period. If the firm is sufficiently patient and ψ_H is not too small, then there also exists an equilibrium in which the firm initially takes high effort, with a low quality outcome triggering a switch to the low effort equilibrium.²¹ The key to these equilibria, as for games with perfect monitoring, is that the equilibrium behavior of firm and the beliefs of consumers can be perfectly coordinated.

Our analysis is motivated by the observation that the beliefs that appear in repeated games with perfect or imperfect public monitoring are notoriously volatile. The reputations in these models typically spring instantly to life rather than being gradually built, and are often instantly (though perhaps temporarily) destroyed upon a firm's first miscue. This volatility arises because consumer beliefs are concerned with the future actions of the firm, which can change abruptly with different realizations of the public outcome. The resulting equilibria require an implausible degree of coordination between the firm's behavior and consumer beliefs, especially when there is a continuum of consumers.

The natural way to preclude this coordination is to assume that each consumer receives a private realization of the outcome, so that $0 < \zeta_f < \zeta_s < 1$. In this situation, the firm cannot predict the realization received by any particular consumer. Moreover, each consumer knows that the firm does not know the signal she observed. This is now a game with *imperfect private monitoring*, and coordination between firm behavior and consumer beliefs is not easily achieved. It is clear from the literature that even small amounts of imperfection in a private monitoring model can pose significant barriers to sustaining a reputation, as well as significantly complicating the analysis. Compte [5], for example, shows that small amounts of idiosyncrasy can destroy the ability of trigger strategies to sustain cooperation in the repeated prisoners' dilemma. Sekiguchi [21] shows that if one allows a broader class of strategies, then equilibria exist featuring a positive probability of cooperation, but only in mixed strategies. Even in simple two-stage games with independent private monitoring, cooperation is impossible in pure strategies, while it can be sustained with positive probability in mixed strategies if publicly observable randomizations are also available (Bhaskar and van Damme [4]). We avoid much of this complication by exploiting the presence of replacements to support pure-strategy equilibria with nontrivial reputations.

The difficulty in coordinating consumer reactions and firm behavior could be removed if consumers could publicly reveal their utility realizations, as in

²¹Other equilibria, in which bad consumer outcome triggers a temporary reversion to low effort, followed by subsequent high effort, can achieve payoffs close to the efficient level. See Fudenberg, Levine, and Maskin [10].

Aoyagi [3], since the difficulty is that consumers have private rather than different realizations. It is then no surprise that many markets contain quality-rating services, often sponsored by the firms in the market, though this possibility is excluded from our model.²² However, because even small amounts of imperfection in a private monitoring technology can have profound effects, the effectiveness of such rating services may require that every consumer receives precisely the same observation of the rating, a possibility that seems unlikely in light of the common existence of multiple, imperfect rating services.

Mailath and Morris [18] prove a Nash-reversion type folk theorem in games with imperfect private monitoring when the signals exhibit a form of strong correlation. Suppose the stage game has a strict pure strategy Nash equilibrium. Then for any action profile whose payoffs \bar{u} Pareto dominate the Nash equilibrium, there is a mixed equilibrium of the repeated game whose average discounted expected payoff is close to \bar{u} if players are sufficiently patient and signals are sufficiently correlated. The expanded interpretation of our current model, with effort generating either successful or failing products which in turn stochastically yield consumer outcomes, introduces some correlation into consumers' signals. However, the underlying signaling technology differs significantly from that of Mailath and Morris. In order to make the models comparable, the probability of a good outcome in the current model would need to depend on both the effort choice of the firm *and* the beliefs of the consumer. In particular, Mailath and Morris [18] make crucial use of the fact that the distribution of signals depends on the entire action profile, not just one player's action.

We do not know if our game of imperfect private monitoring, with $\phi_0 = 1$ and $\lambda = 0$ (and hence neither incomplete information nor replacements), has a mixed strategy equilibrium in which the firm chooses high effort. In particular, we do not know if the firm can choose high effort infinitely often with probability bounded away from one. Even if such equilibria are possible, we have doubts concerning both their plausibility and their suitability for capturing the important features of reputation building. The introduction of types ($\phi_0 \neq 1$ and $\lambda > 0$) allows reputations to arise and to have a plausible interpretation.

7. Conclusion

Pooling models of reputations are well established. Why do we need a new model of reputation based on separation? In many circumstances, we think the type of

²²Levine and Martinelli [17] examine a model in which information about utility realizations is made public, finding that such information can have unexpected effects in the presence of complementary information.

beliefs needed for pooling models of reputation are sufficiently implausible as to motivate a search for alternative models. This belief is reinforced by our finding that the separating model yields properties, such as gradual reputation building, that match common intuition as to how reputations are managed.

These observations suggest the next direction for this research. It is common to speak of reputations as being bought and sold, though current reputation models provide few clues as to what commodity is traded in such cases. In [19], we take the first steps in embedding our model in a market where we can examine the conditions under which reputations can be traded.

8. Appendix

Proof of Proposition 1: Suppose (τ, v) is an equilibrium and τ is a pure strategy. Suppose that τ calls for the competent firm to sometimes exert high effort in equilibrium. Since there are no replacements, the firm's history evolves deterministically, and hence τ determines the periods in which the firm exerts high effort.

It is immediate that τ must call for high effort infinitely often. If this were not the case, there is a final period t^* in which the firm exerts high effort. The revenue in every subsequent period would then be $p(0)$, regardless of the outcome in period t^* , ensuring that high effort is suboptimal in period t^* .

Hence, let $T_n(H) \equiv \{t \geq n : e_t^\tau = H\}$ be those periods larger than n in which high effort is exerted (recall that e_t^τ is the effort choice of the competent type in period t under τ). Similarly, let $T_n(L) \equiv \{t \geq n : e_t^\tau = L\} = \{n, n+1, \dots\} \setminus T_n(H)$. Then, for $t \in T_n(L)$, all consumers expect the firm to choose low effort with probability one, i.e.,

$$p(F_{v, h_{t-1}^f(\tau)}) = p(0).$$

For any period $t \in T_n(H)$, we have, for any history h_t^c ,

$$v(h_t^c) = \varphi(\phi_0 \mid h_t^c),$$

and so

$$p(F_{v, h_{t-1}^f(\tau)}) = p(F_t),$$

where F_t is the probability distribution function of posterior beliefs $\varphi(\phi_0 \mid h_t^c)$ of consumers in period t . Note that if $t \in T_n(L)$, then $F_{t+1} = F_t$. Moreover, since $T_n(H) \neq \emptyset$ for all n , conditional on the firm being competent, $F_t(x) \rightarrow 0$ for all $x < 1$ as $t \rightarrow \infty$. Thus, for all $\varepsilon > 0$, there exists $t(\varepsilon)$ such that for all $t \geq t(\varepsilon)$, $F_t(1-\varepsilon) < \varepsilon$. That is, at least a fraction $1-\varepsilon$ of consumers have observed a private

history h_t^c that yields an update $\varphi(\phi_0 | h_t^c) > 1 - \varepsilon$. Observe that for all $\eta > 0$ and $m \in \mathcal{N}$, there exists $\varepsilon(\eta, m) > 0$ such that $\phi > 1 - \varepsilon(\eta, m) \Rightarrow \varphi(\phi | b^{(m)}) > 1 - \eta$, where $b^{(m)}$ is the history of m consecutive bad realizations, which implies that for any m period history h_m^c , $\varphi(\phi | h_m^c) > 1 - \eta$. Let F^η be the distribution function given by $F^\eta(x) = \eta$ for $x < 1 - \eta$ and $F^\eta(x) = 1$ for $x \geq 1 - \eta$. Then, for all $t' \geq t(\varepsilon(\eta, m))$ and all $t \in \{t', \dots, t' + m\}$, F_t first order stochastically dominates F^η . Thus, the continuation payoff from deviating in a period $t' \in T_{t(\varepsilon)}(H)$ is at least

$$p(F^\eta) + \sum_{\substack{t \in T_{t'}(H), \\ t'+1 \leq t \leq t'+m}} \beta^{t-t'} (p(F^\eta) - c) + \sum_{\substack{t \in T_{t'}(L), \\ t'+1 \leq t \leq t'+m}} \beta^{t-t'} p(0) + \frac{\beta^{m+1}}{(1-\beta)} (p(0) - c),$$

for any $m > 0$. By choosing m large and η small ($p(F^\eta) \rightarrow p(1)$ as $\eta \rightarrow 0$), this lower bound can be made arbitrarily close to

$$p(1) + \sum_{\substack{t \in T_{t'}(H), \\ t'+1 \leq t}} \beta^{t-t'} (p(1) - c) + \sum_{\substack{t \in T_{t'}(L), \\ t'+1 \leq t}} \beta^{t-t'} p(0).$$

Since the continuation payoff from following τ is no more than

$$\sum_{t \in T_{t'}(H)} \beta^{t-t'} (p(1) - c) + \sum_{t \in T_{t'}(L)} \beta^{t-t'} p(0),$$

the good firm has a profitable deviation. Thus, there is no equilibrium in pure strategies with the competent firm ever choosing high effort. ■

Proof of Proposition 2: Suppose the competent firm always chooses high effort, so that $e_t^T = H$ for all t . By hypothesis, $\varphi(\phi | x) \in [\lambda\theta, 1 - \lambda + \lambda\theta]$ for all $\phi \in [0, 1]$ and $x \in \{g, b\}$. Moreover, $\lambda\theta < \phi' < \phi'' < 1 - \lambda + \lambda\theta$ and $\varphi(\phi | x) \in [\phi', \phi'']$ for all $\phi \in [\phi', \phi'']$ and $x \in \{b, g\}$. For $x \in \{b, g\}$, $\varphi^{-1}(\phi | x)$ denotes the inverse image of ϕ under $\varphi(\cdot | x)$; for $\phi < \min_{\phi} \varphi(\phi | x)$, set $\varphi^{-1}(\phi | x) = 0$ and for $\phi > \max_{\phi} \varphi(\phi | x)$, set $\varphi^{-1}(\phi | x) = 1$. If $\phi_0 \in [\phi', \phi'']$, the distribution of consumer posteriors over the type of the firm in period t , denoted G , has support in $[\phi', \phi'']$. Since $\varphi^{-1}(\phi | b) - \varphi^{-1}(\phi | g) > 0$ for all $\phi \in (\lambda\theta, 1 - \lambda + \lambda\theta)$, there is a constant $z > 0$ such that $\varphi^{-1}(\phi | b) - \varphi^{-1}(\phi | g) \geq z$ for all $\phi \in [\phi', \phi'']$. Let G_e denote the distribution over posteriors that the firm is competent in period $t+1$ that results from a choice of effort e , given the distribution of consumer posteriors, G . Then, $G_e(\phi) = \rho_e G(\varphi^{-1}(\phi | g)) + (1 - \rho_e) G(\varphi^{-1}(\phi | b))$, so that

$$G_L(\phi) - G_H(\phi) = (\rho_H - \rho_L) (G(\varphi^{-1}(\phi | b)) - G(\varphi^{-1}(\phi | g))) \geq 0.$$

Choose ε satisfying $0 < \varepsilon < \min\{z, \phi' - \varphi^{-1}(\phi' | g), 1 - \phi''\}$. Then, since $\varepsilon < z$,

$$\int_0^1 G(\varphi^{-1}(\phi | b)) - G(\varphi^{-1}(\phi | g)) d\phi \geq \int_{\phi'}^{\phi''} G(\varphi^{-1}(\phi | g) + \varepsilon) - G(\varphi^{-1}(\phi | g)) d\phi.$$

Let M be the largest integer k for which $\varphi^{-1}(\phi' | g) + k\varepsilon \leq \phi''$; by construction $\varphi^{-1}(\phi' | g) + (M+1)\varepsilon < 1$. Construct an increasing sequence $\{\phi'_k\}_{k=0}^{M+1}$ by setting $\phi'_k = \varphi(\varphi^{-1}(\phi' | g) + k\varepsilon | g)$; note that $\phi'_0 = \phi'$. For $k = 0, \dots, M$, if $f_k : [\phi'_0, \phi'_1] \rightarrow [\phi'_k, \phi'_{k+1}]$ is the function $f_k(\phi) = \varphi(\varphi^{-1}(\phi | g) + k\varepsilon | g)$, then $f_k(\phi'_0) = \phi'_k$ and $f_k(\phi'_1) = \phi'_{k+1}$; and since f_k is continuous, it is onto. Moreover, $f_{k+1}(\phi) = \varphi(\varphi^{-1}(f_k(\phi) | g) + \varepsilon | g)$, and since $\varphi(\cdot | g)$ is concave, $f'_{k+1}(\phi) \leq f'_k(\phi)$. Then for $k = 0, \dots, M$,

$$\begin{aligned} \int_{\phi'_k}^{\phi'_{k+1}} G(\varphi^{-1}(\phi | g) + \varepsilon) d\phi &= \int_{\phi'_0}^{\phi'_1} G(\varphi^{-1}(f_k(\phi) | g) + \varepsilon) f'_k(\phi) d\phi \\ &= \int_{\phi'_0}^{\phi'_1} G(\varphi^{-1}(f_{k+1}(\phi) | g)) f'_k(\phi) d\phi \\ &\geq \int_{\phi'_0}^{\phi'_1} G(\varphi^{-1}(f_{k+1}(\phi) | g)) f'_{k+1}(\phi) d\phi. \end{aligned}$$

Then (since $\phi'_0 = \phi'$, $\phi'_{M+1} > \phi''$, and G 's support is a subset of $[\phi', \phi'']$)

$$\begin{aligned} &\int_{\phi'}^{\phi''} G(\varphi^{-1}(\phi | g) + \varepsilon) - G(\varphi^{-1}(\phi | g)) d\phi \\ &= \sum_{k=0}^M \int_{\phi'_k}^{\phi'_{k+1}} G(\varphi^{-1}(\phi | g) + \varepsilon) - G(\varphi^{-1}(\phi | g)) d\phi \\ &\geq \int_{\phi'_0}^{\phi'_1} \sum_{k=0}^{M-1} G(\varphi^{-1}(f_{k+1}(\phi) | g)) f'_{k+1}(\phi) - \sum_{k=0}^M G(\varphi^{-1}(f_k(\phi) | g)) f'_k(\phi) d\phi \\ &\quad + \int_{\phi'_M}^{\phi'_{M+1}} G(\varphi^{-1}(\phi | g) + \varepsilon) d\phi \\ &= \int_{\phi'_M}^{\phi'_{M+1}} G(\varphi^{-1}(\phi | g) + \varepsilon) d\phi - \int_{\phi'_0}^{\phi'_1} G(\varphi^{-1}(\phi | g)) d\phi \\ &\geq G(\varphi^{-1}(\phi'_M | g) + \varepsilon)(\phi'_{M+1} - \phi'_M) - G(\varphi^{-1}(\phi'_1 | g))(\phi'_1 - \phi'_0) \\ &\geq \phi'_{M+1} - \phi'_M, \end{aligned}$$

where the last inequality is implied by $\varphi^{-1}(\phi'_1 | g) \leq \phi'$ and $\varphi^{-1}(\phi'_M | g) + \varepsilon \geq \phi''$. Thus, $\int_0^1 G_L(\phi) - G_H(\phi) d\phi$ is bounded away from zero (the bound depends only on ε , and not on G or t), and (since the competent firm is always expected to be choosing high effort) so is $\int_0^1 F_L(\phi) - F_H(\phi) d\phi$, where F_e is the distribution in period $t+1$ of consumer expectations of the probability that the firm exerts high effort. This in turn implies that there is a constant Δ (independent of t) such that period $t+1$ revenues after high effort in period t exceed those after low effort in period t by at least Δ . Thus, a sufficient condition for an equilibrium with high effort is that the discounted value of this difference exceeds the cost, $c < \beta(1 - \lambda)\Delta$. ■

Lemma 1. Fix $\varepsilon > 0$. For T sufficiently large,

$$\sum_{t=0}^T \sum_{\tau} h(t, \phi_{\tau} | \phi_1) (\phi_{\tau} - \phi_{\tau-2}) \geq 2(1 - \phi_0) - \varepsilon.$$

Proof. Represent the stochastic process determining the evolution of consumer posteriors as follows. Denote by $z \in \mathbf{Z}^T$ (where \mathbf{Z} is the set of integers) the realization of an outcome path obtained by setting $z(0) = 1$, and determining the value of $z(t)$ recursively according to the rule,

$$z(t+1) = \begin{cases} z(t) + 1, & \text{with probability } \rho_H, \text{ and} \\ z(t) - 1, & \text{with probability } 1 - \rho_H. \end{cases} \quad (0.1)$$

Denote by $Z(T)$ the set of possible realizations of this process. Each $z \in Z(T)$ corresponds to the T -period sequence of realized posteriors $\{\phi_{z(t)}\}_{t=1}^T$. The probability distribution on $Z(T)$ implied by (0.1) is denoted H_T .

Since

$$h(t, \phi_{\tau} | \phi_1) = \sum_{\{z: z(t)=\tau\}} H_T(z)$$

for $t \leq T$, we have

$$\sum_{t=0}^T \sum_{\tau} h(t, \phi_{\tau} | \phi_1) (\phi_{\tau} - \phi_{\tau-2}) = \sum_{z \in Z(T)} \left(\sum_{t=0}^T (\phi_{z(t)} - \phi_{z(t)-2}) \right) H_T(z).$$

Fix $\eta > 0$, and note that if $\phi_{z(T)-1} > 1 - \eta$, then $\phi_{z(T-1)}, \phi_{z(T)} > 1 - \eta$. For outcome paths z such that $\phi_{z(T)-1} > 1 - \eta$, we have (where τ^* is the smallest index τ such that ϕ_{τ} is larger than $1 - \eta$)

$$\sum_{t=0}^T (\phi_{z(t)} - \phi_{z(t)-2}) > 2(\phi_{\tau^*} - \phi_0) \geq 2(1 - \eta - \phi_0).$$

Since consumer posteriors are weakly converging to 1, there exists T' such that

$$\sum_{\{z: \phi_{z(T)-1} > 1-\eta\}} H_T(z) \geq 1 - \eta$$

for all $T > T'$. Then

$$\sum_{z \in Z(T)} \left(\sum_{t=0}^T (\phi_{z(t)} - \phi_{z(t)-2}) \right) H_T(z) \geq 2(1 - \eta)(1 - \eta - \phi_0),$$

and the right hand side is larger than $2(1 - \phi_0) - \varepsilon$ if $\eta < \varepsilon/4$. ■

References

- [1] Dilip Abreu, David Pearce, and Ennio Stacchetti. Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica*, 58:1041–1063, 1990.
- [2] Nabil Ibraheem Al-Najjar. Decomposition and characterization of risk with a continuum of random variables. *Econometrica*, 63:1195–1224, 1995.
- [3] Masaki Aoyagi. A collusion mechanism in dynamic bertrand oligopoly under private observation of correlated demand shocks. Mimeo, University of Pittsburgh, 1997.
- [4] V. Bhaskar and Eric van Damme. Moral hazard and private monitoring. Mimeo, University of St. Andrews and CentER Tilburg, 1997.
- [5] Olivier Compte. On failing to cooperate when monitoring is private. Mimeo, C.E.R.A.S.-E.N.P.C. and CEPRAMAP, 1996.
- [6] Saikat Datta. Building trust. STICERD Discussion Paper 96/305, London School of Economics, 1996.
- [7] Douglas W. Diamond. Reputation acquisition in debt markets. *Journal of Political Economy*, 97:828–862, 1989.
- [8] Drew Fudenberg, David Kreps, and Eric Maskin. Repeated games with long-run and short-run players. *Review of Economic Studies*, 57:555–574, 1990.

- [9] Drew Fudenberg and David K. Levine. Maintaining a reputation when strategies are imperfectly observed. *Review of Economic Studies*, 59:561–579, 1992.
- [10] Drew Fudenberg, David K. Levine, and Eric Maskin. The folk theorem with imperfect public information. *Econometrica*, 62:997–1040, 1994.
- [11] Douglas Gale and Robert W. Rosenthal. Price and quality cycles for experience goods. *Rand Journal of Economics*, 25:590–607, 1994.
- [12] Bengt Holmström. Managerial incentive problems—a dynamic perspective. In *Essays in Economics and Management in Honour of Lars Wahlbeck*, pages 209–230. Swedish School of Economics and Business Administration, Helsinki, 1982.
- [13] Benjamin Klein and Keith B. Leffler. The role of market forces in assuring contractual performance. *Journal of Political Economy*, 89:615–641, 1981.
- [14] David Kreps, Paul R. Milgrom, D. John Roberts, and Robert Wilson. Rational cooperation in the finitely repeated prisoner’s dilemma. *Journal of Economic Theory*, 27:245–252, 1982.
- [15] David Kreps and Robert Wilson. Reputation and imperfect information. *Journal of Economic Theory*, 27:253–279, 1982.
- [16] David M. Kreps. Corporate culture and economic theory. In James E. Alt and Kenneth A. Shepsle, editors, *Perspectives on Positive Political Economy*, pages 90–143. Cambridge University Press, Cambridge, 1990.
- [17] David K. Levine and César Martinelli. Reputation with noisy precommitment. *Journal of Economic Theory*, 78:55–75, 1998.
- [18] George J. Mailath and Stephen Morris. Repeated games with imperfect private monitoring: Notes on a coordination perspective. CARESS Working Paper #98-07, University of Pennsylvania, 1998.
- [19] George J. Mailath and Larry Samuelson. Who wants a good reputation? CARESS Working Paper #98-12, University of Pennsylvania, 1998.
- [20] Paul R. Milgrom and D. John Roberts. Predation, reputation and entry deterrence. *Journal of Economic Theory*, 27:280–312, 1982.
- [21] Tadashi Sekiguchi. Efficiency in repeated prisoner’s dilemma with private monitoring. *Journal of Economic Theory*, 76:345–361, 1997.

- [22] Joel Watson. Building a relationship. Mimeo, University of California, San Diego, 1995.
- [23] Joel Watson. Starting small and renegotiation. Discussion Paper 97-17, University of California, San Diego, 1997.