Temporary reputation

Jeremy Sandford*
Federal Trade Commission

June 3, 2016

Abstract

Traditional models of reputation posit that a firm is one of two or more exogenously-imposed *types*, with a normal type being the most common. Normal types then choose effort without their type being known by their customers, incentivizing them to attempt to either pool with good types, or dissociate from bad types. We offer a model that generates an incentive to maintain reputation without imposing an exogenous type distribution. Outcomes are imperfectly correlated with not only contemporaneous effort, but previous effort levels as well. Firms willingly invest in unobservable high effort in the hopes that this will produce a favorable outcome, raising customer willingness to pay for the firm's service in the next period because they believe the firm more likely to have exerted high effort in a previous period. Our model can explain several salient features of markets with imperfectly-observed effort.

1 Introduction

Under what conditions should we assume that past performance is predictive of future results? It is not obvious, for example, that a customer's favorable or unfavorable history with a firm is an essential determinant of current outcomes; employees and management change and budgets may be cut or increased, implying that the link between past and future results is tenuous. Still, employee turnover and changes in business strategy are gradual processes, implying at least a transitory link between past and present. This paper studies a model of temporary reputation, in which a firm's actions in

*jsandford@ftc.gov. The views expressed in this article are those of the author and do not necessarily reflect those of the Federal Trade Commission.

the immediate past is informative regarding current outcomes. This is a departure from much of the literature on reputation, which posits that firms are exogenously endowed with one of two or more "types," some of which are committed to particular actions such as high effort.[1] This model generates reputational incentives without uncertainty over firm type, by assuming that the effects of effort are persistent across a finite number of periods. This model can explain customers assigning higher weight to recent events than to decades of good (bad) service[2] since knowledge of high effort in the distant past is rightly discarded by customers. However, our model of two-period persistent effort can even explain extreme persistence of effort states in which a firm has long periods of sustained high (low) quality, followed by long periods of sustained low (high) quality, as might reasonably characterize the historical arc of companies such as GM, Apple, Microsoft, or Kodak.[3]

This paper adds to a small literature on the reputational effects of persistent effort without types (cf. Board and Meyer-Ter-Vehn (2013), Bohren (2013), Jarque (2010), and Dilme (2012)). Models of persistent effort offer several advantages over type-based models. One, they are more realistic; the literature has little to say on the genesis of a type distribution (see Sandford (2010) for one exception). Two, they generate interesting reputational dynamics. As Cripps et al. (2004) show, reputation in a broad class of type-based models is an impermanent phenomenon, as eventually customers will learn a firm's type with a high degree of certainty, and his reputational incentives will dissipate. While this difficulty can be circumvented in type-based models to some extent,[4] the implausibility of these models' assumptions and their inability to generate real world features such as some firms consistently playing high effort and some consistently playing low effort suggest that alternative models of reputation are worth considering. Finally, models of persistent effort (including this one) often allow for complete characterizations of equilibrium behavior, in contrast to the folk theorem-like limiting results from the type-based literature.

Our model posits that a firm exerts either high or low effort in each discrete time period. Effort is unobserved by customers, who instead experience a good or bad outcome. Current outcome is determined by a probabilistic function of both the firm's current effort and the effort in the most

[1] See Milgrom and Roberts (1982), Tadelis (1999), Mailath and Samuelson (2001), Horner (2002), Liu (2011), and Faingold and Sannikov (2011) for prominent examples of this literature.

[2] See "Malaysia Airlines warns of further losses, 8/28/14, bbcnews.com, which states that bookings on Malaysia Airlines fell 33% and the continued viability of the company is in question after two black swan plane crashes, despite their having an excellent safety record for decades.

[3] See "Why companies fail" by Megan McArdle, March 2012, *The Atlantic* for an entertaining look at the shifting fortunes of large American companies.

[4] Mailath and Samuelson (2001) posits a constant churning of firms unobservable to customers, while Horner (2002), Rob and Sekiguchi (2006), and Sandford (2014) construct reputational models of customers continually switching between firms.

recent k prior periods (most of the paper considers the comparatively tractable case of $k = 1$). A customer with a good outcome then increases her posterior belief that the firm played high effort in that period, and this higher belief in turn increases her willingness to pay in the following period. This creates a reputational incentive for firms, even with no uncertainty over firm type. This incentive is also transitory, as a customer's belief that high effort was played k periods ago no longer directly affects her willingness to pay in the current period. In the case of $k = 1$, the firm's problem has two state variables: the customer belief that high effort was played in the previous period, and the actual effort level played (known only to the firm). The former influences contemporary profit and the manner in which customers update their posterior after observing the outcome at the end of the period. The latter influences the probability of a good or bad outcome for each possible current effort level. Models with an imperfect correlation between effort (private information) and outcome (public information) can be used to study the incentives of mechanics, teachers,[5] Realtors,[6] physicians,[7] and even economists.[8]

This paper uses a simple, discrete-time model similar to those used in virtually all of the type-based reputation literature cited above, to generate reputational effects when there is no uncertainty about firm type. In contrast, other models of persistent effort use more complicated continuous time setups; while these allow for rich characterizations of behavior, they are sometimes difficult to compare with the extant reputational literature. Board and Meyer-Ter-Vehn (2013) model the quality of a firm's product as endogenously switching between "high" and "low", with the former more likely when high effort is played. Customers occasionally receive signals correlated with current quality, and because after a signal, customers receive no new information until the next signal, firms have a reputational incentive to exert high effort to make favorable signals more likely. Bohren (2013) uses a continuous time model similar to the discrete time model used here, in which past quality investments have an effect on current quality which decays exponentially. Dilme (2012) uses a continuous time model in which firms must pay a switching cost when changing from high to low effort, or vice-versa. This switching cost generates persistence in effort across periods, as generically the firm will not want to pay the cost for small changes in state variables from moment to moment. Because of this persistence, firms have a reptuational incentive to generate good outcomes, as customers will correctly infer upon receiving a good outcome that future prospects look good. Jarque (2010) and Fernandes and

[5]See Figlio et al. (2013), which points out that students in introductory economics courses learn more when taking the course from non-tenure track faculty.

[6]See *New York Times*, "Realtors agree to stop blocking web listings," May 28, 2008, for an interesting example of Realtors acting against their customers' interests.

[7]See*New York Times*, "Caesarean births are at a high in U.S.," March 23, 2010, for a discussion of the high rates of Casearian births and possible causes which include physician moral hazard.

[8]See Posner (1999) for a discussion of the incentives of economists testifying as expert witnesses.

Phelan (2000) study models of repeated moral hazard with persistent effects of effort that resemble this model. Like these papers, this paper can explain the decaying reputational effects that come from persistent effort, but, in contrast to these papers, can also explain long periods of high (low) quality followed by long periods of low (high) quality without needing exogenous shocks.

This paper also generates different results than the type-based literature. Here, as shown in proposition 1, a firm's incentive for high effort is generally *greatest* for high and low existing reputations, and lowest for intermediate reputations. Firms prefer to invest in high effort both to maintain already-high reputations and to bolster horrible reputations, but prefer to coast on established reputations that are just so-so, leading to "work-shirk-work" equilibria. In contrast, in type-based models, reputational effects erode as beliefs approach 1 or 0. The reason for the distinction is that in this model, if a customer is very sure that a firm played high effort in the previous period (established reputation close to 1), a failure most likely means the firm unobservably shirked in the current period, so future reputation will take a large hit. On the other hand, in the event of a current success, the customer's current belief is confirmed, and reputation stays more or less the same. Similarly, when reputation is very poor, a successful outcome is far more likely given high effort, and so the firm's reputation will greatly increase. However, for an intermediate reputation, customers are uncertain what effort level was played int he previous period, and so neither a success nor a failure are particularly helpful for determining current effort.

Section 4 proposes a number of numerical examples, most of which share the feature that firms with current high reputation are more likely to play high effort than firms with intermediate reputations. Interestingly, one of the examples demonstrates that this model is capable of generating extreme persistence of effort; firms will play low (or high) effort for literally thousands of periods, before switching to high (or low) effort, which they will then play for a similarly long time. Thus, even this model of temporary reputation generated by effort persisting for two periods is capable of explaining apparently very-long-run behavior of firms, as well as scenarios in which a firm with a well-established record of high quality service appears to switch to shoddy quality, or a company overcomes a long period of low quality to establish itself as a high quality producer.

Section 2 describes the paper's model, while section 3 provides the main analytical result for the special case of two-period effort persistence and then demonstrates how the model can be numerically solved using value function iteration. Section 4 describes by example several different types of equilibrium behavior that the model can generate. Section 5 discusses extending the model to effort persisting longer than two periods. Finally, section 6 concludes.

2 Model

One firm serves a series of customers. Assume that each period, a firm interacts with one customer, and that each customer-firm interaction lasts for exactly one period. Customers are myopic, while the firm maximizes the discounted sum of its profits. In every period, the firm exerts either high or low effort, but this effort is not observable, even *ex post*. Instead, the outcome of each customer-firm interaction is either a *success* or a *failure* with successes being relatively more likely under high effort, and failures under low effort. Assume that the firm's history of successes and failures are publicly available, e.g. via an Internet review site or word of mouth.

Firms charge customers a price that is increasing in the customer's belief that a success will occur. Firms thus face an intertemporal tradeoff: while high effort lowers current profits, it may increase future profits by raising the price future customers are willing to pay. Each customer's payoff depends on the price and whether her interaction with the firm was a success or failure, with a success yielding a higher payoff than a failure.

This modeling setup is appropriate for credence good markets, in which customers are unsure of the quality of the service for which they pay, even *ex post*. For example, a customer advised that her car needs expensive repairs when she brought it in for an oil change will view the visit as a failure, but will not be entirely sure whether the repairs are actually efficient, saving her money down the road, or whether the mechanic is trying to cheat her. Similarly, a patient who recovers quickly after a doctor's visit will regard the visit as a success, but has no way of knowing whether she would have recovered just as quickly without the visit. Instead, a customer's willingness to pay is derived from the firm's prior outcomes. We study the effectiveness of reputation in providing an incentive to the firm to exert high effort.

2.1 Information

Let $E_t = [e_t, e_{t-1}, e_{t-2}, ...]$ denote the vector of effort levels played by the firm in periods up to and including t, with $e_{t-k} \in \{L, H\}$ for all k denoting either low or high effort. E_{t-1} comprises the firm's private information at the beginning of time period t. Further, E_t determines the outcome experience by the customer in time period t. Specifically, let $f(E_t)$ denote the probability of a success in time period t, with complimentary probability of a failure. Assume that if vectors E_t and $\tilde{E}_t$ are identical in all but one element (without loss of generality, suppose that $e_{t-k} = H$ while $\tilde{e}_{t-k} = L$), then $f(E_t) \geq f(\tilde{E}_t)$.[9] This means that higher effort in any time period up to and including the current

[9]One simple example is $f(E_t) = \frac{e_t + e_{t-1}}{2}$, so that the probability of a success in the current period is equal to 1 if high effort was played in periods t and $t-1$, equal to 0 if low effort was played in both periods, and equal to $\frac{1}{2}$ if high effort

period makes a success weakly more likely. The key modeling assumption differentiating this paper from much of the literature on reputation is that $f(E_t)$ depends non-trivially on e_{t-k} for $k \geq 1$; most of the reputation literature assumes that only current period effort affects the current outcome.

Let the public history of outcomes up to and including time t be given by $S_t = [s_t, s_{t-1}, ...]$, where each $s_{t-k} \in \{S, F\}$, and S denotes a success, F a failure. The information in S_{t-1} is public knowledge at the beginning of time period t. As the outcome vector S_t is correlated with the vector of effort levels E_t, customers employ Bayes' rule to estimate the probability that high effort was played in any given period, given outcome vector S_t.

Let μ describe a customer's prior belief that a firm plays high effort in any given time period. A customer newly matched with a firm will assign probability μ to the firm playing high effort in the current period. Let μ^t_τ denote the belief customers hold at the beginning of time t that the firm played high effort in period τ, and let $M^t = [\mu^t_t, \mu^t_{t-1}, \mu^t_{t-2}, ...]$ denote the vector of such beliefs. Trivially, $\mu^t_t = \mu$ for all t, while μ^t_{t-1}, for example, depends on both the prior μ, the outcome of the previous period, s_{t-1}, and possibly other beliefs in M^t.[10]

Throughout the paper, we treat the prior μ as an exogenous parameter; in particular, we do not impose any equilibrium requirements on μ, such as that it equal the probability a firm plays high effort in the current period.[11] This is both for tractability and because, regardless of how μ is defined, the model's reputational dynamics derive from the effect of current effort on future customer beliefs (that is, from how e_t affects M^{t+1}). The belief vector M^t is a state variable for the firm's dynamic optimization problem; from here on, we will also refer to M^t as a firm's *reputation*.

2.2 The firm's dynamic optimization problem

Suppose that utility of a customer visiting the firm at time t is given by $U_t = A * I\{s_t = S\} - P$, where $I\{s_t = S\}$ is equal to one if the outcome in the current period is a success, and equal to zero if it is a failure. A success thus increases a customer's utility by A relative to a failure. We assume that

was played in exactly one of the two periods. This could be generalized to $f(E_t) = \alpha\Sigma_{j=0}^{k}\delta_j e_{t-j}$ where $\alpha \leq 1, \Sigma\delta_j = 1$, and $\delta_j \in [0, 1]$ for each j. Section **??** considers an extension of the function $f(E_t) = \frac{I\{e_t=e_H\}+I\{e_{t-1}=e_H\}}{2}$ that bounds the probability of a success between ρ and $(1-\rho)$, for some $\rho \in (0, \frac{1}{2})$.

[10]To see why μ^t_{t-1} may depend on past beliefs, consider the following example. Suppose that $f(E_t) = \frac{I\{e_t=e_H\}+I\{e_{t-1}=e_H\}}{2}$, and that $\mu = \frac{1}{2}$. Suppose the outcome of period t is a success. If customers believe at the beginning of time t that effort in time period $t-1$ was certainly low (so that $\mu^t_{t-1} = 0$), then the prior μ updates to 1 (in other words, $\mu^{t+1}_t = 1$). On the other hand, if customers believe at the beginning of time t that effort in time period $t-1$ was high with probability $\frac{3}{4}$, then their posterior belief about time period t upon observing a success in period t is $\mu^{t+1}_t = \frac{7}{10}$.

[11]Although nothing rules out defining μ so that it depends on the public history, including in such a way that it does equal the probability of high effort in the current period.

the firm is able to set prices so as to extract all expected surplus from customers. Price P_t is thus given by $P_t(M^t) = E[I\{s_t = S\}|M^t]$.[12] High effort is costly for the firm relative to low effort; suppose that $c(e_H) = c > 0$ while $c(e_L) = 0$.

Each period, the firm decides whether to exert high or low effort. High effort yields lower contemporaneous utility, but is more likely to produce a successful outcome, which may increase the price received from future customers. The firm's choice in period t depends on two state variables: the vector of effort levels played prior to time t. E_{t-1}, and the vector of customer beliefs, M^t.

To model this intertemporal tradeoff, given state variables (M^t, E_{t-1}) define the following one-period ahead state variables:

$$E_t(e_t) = (e_t, E_{t-1})$$
$$M^{t+1}(s_t) = [\mu, \mu_t^t(s_t), \mu_{t-1}^t(s_t), \mu_{t-2}^t(s_t), ...]$$

where (e_t, E_{t-1}) is the vector with e_t in position 1, and $E_{t-1}(j)$ in position $j+1$, for $e_t \in \{e_L, e_H\}$. $\mu_\tau^t(s_t)$ is the posterior probability that high effort was played in period τ given outcome s_t. $\mu_\tau^t(s_t)$ is calculated by Bayes' rule, and is described by equation (1) below:

$$\mu_\tau^t(s_t) = \Pr(e_\tau = H|s_t) = \frac{\Pr(s_t|e_\tau = H)\mu_\tau^t}{\Pr(s_t|e_\tau = H)\mu_\tau^t + \Pr(s_t|e_\tau = L)(1-\mu_\tau^t)} \tag{1}$$

As an example of how to calculate $\mu_\tau^t(s_t)$, if $f(E_t) = \frac{I\{e_t=e_H\}+I\{e_{t-1}=e_H\}+I\{e_{t-2}=e_H\}}{3}$, and prior beliefs are μ_t, μ_{t-1}, and μ_{t-2}, the posterior probability that high effort was played in period $t-1$ is:

$$\Pr(e_{t-1} = H|s_t = S) = \frac{\frac{\mu_t+1+\mu_{t-2}}{3}\mu_{t-1}}{\frac{\mu_t+1+\mu_{t-2}}{3}\mu_{t-1} + \frac{\mu_t+\mu_{t-2}}{3}(1-\mu_{t-1})} = \frac{\mu_{t-1}(1+\mu_t+\mu_{t-2})}{\mu_t+\mu_{t-1}+\mu_{t-2}}$$

Given the customer's Bayesian updating, the firm then chooses high or low effort to maximize its discounted sum of short term profits, $p(M) - c$. To describe the firm's optimization problem, we suppress time subscripts and superscripts, so that (M, E) is taken to mean (M^t, E_{t-1}). Then, define $M'_S(j)$ ($M'_F(j)$) to be the j^{th} element of the one-period-ahead belief vector, given a success (failure) occurs in the current period, with M'_S and M'_F referring to the entire vector of beliefs, given a success and failure, respectively.

Suppressing time subscripts, the firm's optimization problem is described by equation (2). Note that $E(j)$ and $M_S(j)$ refer respectively to the j^{th} element of vectors E and M_S. The firm discounts

[12] We assume that each customer has an outside option of 0.

future payoffs geometrically, at rate $\beta \in (0,1)$.

$$\begin{aligned} V[M^t, E_{t-1}] = \max_{e_t \in \{e_H, e_L\}} & p(M^t) - c(e_t) \\ & + \beta \left(f(E_t(e_t)) V[M_S^{t+1}, E_t(e_t)] + (1 - f(E_t(e_t)) V[M_F^{t+1}, E_t(e_t)] \right) \\ \text{subject to: } & E_t(e_t) = (e_t, E_{t-1}) \\ & \mu_\tau^{t+1}(s_t) = \Pr(e_\tau = H | s_t) \text{ (calculated by Bayes rule, from equation (1))} \\ & \mu_{t+1}^{t+1} = \mu \end{aligned} \tag{2}$$

3 Two-period effort

In this section, we consider a simple application of the model, namely that of effort which affects quality for two periods. Specifically, we assume that the probability of a success in any period is given by equation 3, below.

$$f(e_t, e_{t-1}) = \begin{cases} (1-\rho) & \text{if} \quad e_t = e_{t-1} = e_H \\ \frac{1}{2} & \text{if} \quad e_t = e_H \text{ and } e_{t-1} = e_L \text{ or } e_t = e_L \text{ and } e_t = e_H \\ \rho & \text{if} \quad e_t = e_{t-1} = e_L \end{cases} \tag{3}$$

The parameter $\rho \in (0, \frac{1}{2})$ measures the amount of noise with which customers observe outcomes. In the extreme case $\rho = 0$, two consecutive periods of high (low) effort always produce a success (failure), while in the case of $\rho \in (0, \frac{1}{2})$, high effort increases the chance of a success, but failures are possible regardless of effort played. Note that equation (3) assumes that the previous two effort levels are of equal importance in determining the likelihood of a success; it would also be natural to assume that current effort has a greater effect on the probability of success than does past effort.

Since the probability of a success depends only on current effort e_t and effort in the previous period e_{t-1}, the state variables in the firm's problem (2) reduce to two scalars, μ_{t-1}^t and e_{t-1}. Given our assumption that $\mu_t^t = \mu$ for all t, the customer's perceived probability of a success in any given period is a function of only μ_{t-1}^t. Specifically, the customer perceives the probability of a success as $\rho + (1-2\rho)\frac{\mu + \mu_{t-1}^t}{2}$,[13] and thus the price charged by the firm is given by:

$$P(\mu_{t-1}^t) = A \left(\rho + (1-2\rho) \frac{\mu + \mu_{t-1}^t}{2} \right) \tag{4}$$

Suppressing time subscripts for simplicity, let $\tilde{\mu} = \mu_{t-1}^t$ and $\tilde{e} = e_t - 1$ describe the firm's state variables, while $\tilde{\mu}' = \mu_t^{t+1}$ and $\tilde{e}' = e_t$ describe the one-period-ahead state variables. Of course, $\tilde{e}'$ is

[13] Note that an alternate way of writing equation (3) is $f(e_t, e_{t-1}) = \rho + (1-2\rho)\frac{I\{e_t = e_H\} + I\{e_{t-1} = e_H\}}{2}$.

also the firm's choice variable in the current period, and $\tilde{\mu}'$ depends on whether or not the outcome in the current period is a success or a failure. Hence, let $\tilde{\mu}'_S$ and $\tilde{\mu}'_F$ denote the one-period-ahead value of $\tilde{\mu}$ given a success and a failure, respectively. Both are defined by equation (1).

The firm's optimization problem is then given by:

$$V[\tilde{\mu},\tilde{e}] = \max_{\tilde{e}'\in\{e_L,e_H\}} P(\tilde{\mu}) - c(\tilde{e}') + \beta\left(f(\tilde{e},\tilde{e}')V[\tilde{\mu}'_F,\tilde{e}'] + (1-f(\tilde{e},\tilde{e}'))V[\mu'_S,\tilde{e}']\right) \tag{5}$$

$$\text{subject to:}\quad \tilde{\mu}'_S = \frac{\mu\left(\tilde{\mu}(1-\rho)+(1-\tilde{\mu})\frac{1}{2}\right)}{\mu\left(\tilde{\mu}(1-\rho)+(1-\tilde{\mu})\frac{1}{2}\right)+(1-\mu)\left(\tilde{\mu}\frac{1}{2}+(1-\tilde{\mu})\rho\right)}$$

$$\tilde{\mu}'_F = \frac{\mu\left(\tilde{\mu}\rho+(1-\tilde{\mu})\frac{1}{2}\right)}{\mu\left(\tilde{\mu}\rho+(1-\tilde{\mu})\frac{1}{2}\right)+(1-\mu)\left(\tilde{\mu}\frac{1}{2}+(1-\tilde{\mu})(1-\rho)\right)}$$

While we characterize the solutions to the firm's optimization problem (5) numerically, we first present a key analytical result. The firm's incentive to exert high effort is to reap the benefits of a higher reputation tomorrow. Thus, the marginal benefit of high effort is related to the difference in next period's state variable $\tilde{\mu}'$ after a success, relative to that after a failure. This difference in turn depends on accumulated reputation $\tilde{\mu}$. In particular, this difference is largest for extreme values of $\tilde{\mu}$, those closest to 0 and 1. The reason is straightforward; as customers become increasingly sure that the firm played high (low) effort in the previous period, a failure (success) in the current period leaves them increasingly sure that low (high) effort was played in that period. Hence, the firm's incentive for high effort will be greatest when it has an established low or high reputation, and will be muted for intermediate reputations. This difference is the smallest for intermediate reputations, specifically when established reputation $\tilde{\mu}$ equals the customer's prior belief μ. The firm's incentives are thus non-monotonic in established reputation. Lemma 1 demonstrates this formally.

Lemma 1. *$\tilde{\mu}'_S-\tilde{\mu}'_F$ is strictly decreasing in $\tilde{\mu}$ for $\tilde{\mu}<1-\mu$ and strictly increasing in $\tilde{\mu}$ for $\tilde{\mu}>1-\mu$.*

Proof: $\tilde{\mu}'_S$ and $\tilde{\mu}'_F$ are calculated using Bayes' rule as follows:

$$\tilde{\mu}'_S = \frac{P(\text{high effort and success})}{P(\text{success})} \quad \tilde{\mu}'_F = \frac{P(\text{high effort and failure})}{P(\text{failure})}$$

Note that P (success) and P (failure) are given by the denominators in the expressions for $\tilde{\mu}'_F$ and $\tilde{\mu}'_S$ in equation (5). Solving for the least common denominator of the expressions for $\tilde{\mu}'_S$ and $\tilde{\mu}'_F$ in (5), we have that:

$$\tilde{\mu}'_S-\tilde{\mu}'_F = \frac{\frac{1}{2}\mu(1-\mu)(1-2\rho)}{P(S)P(F)} \tag{6}$$

Equation (6) and the quotient rule imply that the sign of $\frac{\partial}{\partial\tilde{\mu}}(\tilde{\mu}'_S - \tilde{\mu}'_F)$ is the same as the sign of $-\frac{\partial}{\partial\tilde{\mu}}(P(S)P(F))$. Using the expressions for $P(S)'$ and $P(F)$ from (5), we have that:

$$\frac{\partial}{\partial\tilde{\mu}}P(S) = \frac{\mu(1-2\rho)}{2} + \frac{(1-\mu)(1-2\rho)}{2} = \frac{1}{2} - \rho$$
$$\frac{\partial}{\partial\tilde{\mu}}P(F) = -\mu\left(\frac{1}{2}-\rho\right) - (1-\mu)\left(\frac{1}{2}-\rho\right) = -\left(\frac{1}{2}-\rho\right)$$

And so, by the product rule,

$$-\frac{\partial}{\partial\tilde{\mu}}P(S)P(F) = \left(\frac{1}{2}-\rho\right)(P(S)-P(F)) \geq 0$$
$$\iff \mu\left(\frac{1}{2}-\rho-\left(\frac{1}{2}-\rho\right)\tilde{\mu}-\left(\frac{1}{2}-\rho\right)(1+\tilde{\mu})\right)$$
$$\iff -\mu\tilde{\mu}(1-2\rho) + (1-\mu)(1-\tilde{\mu})(1-2\rho) \geq 0$$
$$\iff \tilde{\mu} \geq 1-\mu$$

Since $\frac{\partial}{\partial\tilde{\mu}}(\tilde{\mu}'_S - \tilde{\mu}'_F)$ and $-\frac{\partial}{\partial\tilde{\mu}}(P(S)P(F))$ have the same sign, the claim follows. ■

An implication of proposition 1 is that the firm has a greater incentive to play high effort for accumulated reputations above or below $1-\mu$. This result differs from the reputation literature, in which a high accumulated reputation often destroys a firm's incentive for high effort.[14]

Since the firm's incentive constraint depends heavily on the term $\tilde{\mu}$, it is worth noting that $\tilde{\mu}$ depends not only on the outcome in the previous period, but also on the outcomes in all prior periods. For example, if customers experienced a successful outcome in period t and an unsuccessful outcome in period $t+1$, in period $t+2$, $\tilde{\mu}$ will be relatively higher than it would be were the outcome in period t were also unsuccessful. This is because given unsuccessful outcomes in the recent past, a customer is more likely to assign "blame" for a bad outcome in the current period to lower past effort, meaning that posterior beliefs on current effort will suffer a relatively smaller 'penalty'. On the other hand, given successful outcomes in the recent past, a failure in the current period will prompt the customer to strongly suspect that the firm played low effort only the current period, and so state variable $\tilde{\mu}$ will be low. Thus, we view $\tilde{\mu}$ as a firm's reputation, accumulated over all previous periods.

3.1 Numerically solving the model

We employ value function iteration to numerically solve the firm's optimization problem given in (5). In particular, we determine the optimal policy function $\tilde{e}'(\tilde{\mu}, \tilde{e})$ and the implied value function $V[\tilde{\mu}, \tilde{e}]$. In doing so, we prove by example that the following types of equilibria can occur:

[14]While $V[\tilde{\mu}, \tilde{e}]$ is an increasing function of $\tilde{\mu}$ and $\tilde{e}$, it is not necessarily true that $V[\tilde{\mu}'_S, \tilde{e}] - V[\tilde{\mu}'_F, \tilde{e}]$ is also U-shaped in $\tilde{\mu}$.

1. The firm always plays low effort
2. The firm always plays high effort
3. The firm usually plays low effort, but switches to high effort whenever its accumulated reputation $\tilde{\mu}$ is unusually high or unusually low.
4. The firm generally plays high effort, but switches temporarily to low effort when the customer is especially uncertain ($\tilde{\mu}$ close to $\frac{1}{2}$) what effort level was recently played
5. There are two states, one in which the firm plays high effort, with only occasional deviations to low effort, and one in which the firm plays mainly low effort, with only occasional deviations to high effort. The firm remains in one state for long periods of time, with a low probability of switching states.
6. Rapid switching between high and low effort.

Equilibria of type 1 exist in any reputation model. However, equilibria of type 2, while also common in reputational models, often suffer from the firm's incentives eroding as reputation approaches 1; nothing of the sort happens here, where the customer knows to discount outcomes occurring in the distant past. The straightforwardness of high-effort equilibria (firm plays high effort so long as the cost of doing so is sufficiently low) is a desirable feature of our model.

Equilibria of type 3 demonstrate that the firm's incentive for high effort is greatest for very low and very high reputations, a feature novel to the reputation literature. Equilibria of type 4 demonstrates the same principle, that the incentive for high effort can be weakest when the customer is most uncertain about which effort level was played recently.

Finally, equilibria of type 5 are entirely novel to the reputation literature. As we will see, the firm can play high (low) effort seemingly indefinitely, with mostly successes (failures) as outcomes. However, certain essentially random sequences of outcomes trigger a reputational shift extreme enough to cause the firm's incentives to change, and thus the state variable $\tilde{e}$ to switch either from e_H to e_L or vice-versa. Once the state variable $\tilde{e}$ has switched, the firm plays a new effort level for an extended period, until another, different, essentially random series of outcomes changes its incentives again. Such equilibria are not generally seen in the reputational literature, and may help explain phenomenon such as brands or companies with apparently radically different efforts towards customer satisfaction cross time, such as department stores. Equilibria of type 6 are essentially sped-up versions of those of type 5, and see the firm near its indifference point in most time periods.

3.2 Value function iteration

First, note that $V : [0,1] \times \{e_L, e_H\} \Rightarrow \mathbb{R}$ is a bounded function over a subset of $\mathbb{R}^2$ (representing e_L and e_H with arbitrary real numbers, such as 0 and 1, respectively). Second, define the operator T mapping the set of bounded functions on $[0,1] \times \{e_L, e_H\}$ into itself as:

$$TV[\tilde{\mu}, \tilde{e}] = \max_{e \in \{e_L, e_H\}} p(\tilde{\mu}) - c(e) + \beta(f(\tilde{e}, e)V[\mu'_F(\tilde{\mu}), e] + (1 - f(\tilde{e}, e))V[\mu'_S(\tilde{\mu}), e])$$

A standard argument shows that T satisfies Blackwell's sufficiency conditions for a contraction (Stokey and Lucas, 1989, pg. 54). Third, by the Contraction Mapping Theorem, T has exactly one fixed point, and for any bounded function W, $T^n W$ converges to the value function V as defined in equation (5) (Stokey and Lucas, 1989, pg. 50). Given that T is a contraction, we numerically solve for the value function V in (5) by iteratively applying the operator T to an arbitrary initial "guess" W until $|T^{n+1}W - T^n W|$ is below some threshold.

We discretize the interval $[0,1]$ into a grid of $n+1$ uniformly-spaced points, $\{\tilde{\mu}_1, \tilde{\mu}_2, ..., \tilde{\mu}_{n+1}\}$ such that $\tilde{\mu}_i = \frac{i-1}{n}$, for $i = 1$ to $i = n+1$. Then, define a $(2n+2) \times 1$ vector V describing the future payoff to all $2n+2$ states $(\tilde{\mu}, \tilde{e})$; the first $n+1$ values correspond to state variable $\tilde{e} = e_L$, while the subsequent $n+1$ correspond to $\tilde{e} = e_H$. Given the discretized model's $2n+2$ state variables and 2 control variables, we compute two $(2n+2) \times 2$ matrices tracking current and future payoffs respectively for each effort choice, given any possible combination of state variables. For each matrix, the row determines the current state, and the column the control variable.

The first matrix, tracking current payoffs, consists of two $(n+1) \times 2$ matrices stacked on top of each other, $\begin{bmatrix} \Pi_L \\ \Pi_H \end{bmatrix}$, with Π_L corresponding to $\tilde{e} = e_L$ and Π_H to $\tilde{e} = e_H$. Applying the firm's optimization problem (5) to the discretized problem, we have $\Pi_L(i,j) = \Pi_H(i,j) = p(\tilde{\mu}_i) - c(e_j)$, where e_1 denotes low current effort and e_2 high current effort.[15]

The second matrix, describing the expected future payoff, is given by $\begin{bmatrix} F_{LL}V_L & F_{LH}V_H \\ F_{HL}V_L & F_{HH}V_H \end{bmatrix}$, where the four $n+1 \times n+1$ matrices F_{LL}, F_{HL}, F_{LH}, and F_{HH} give probability distributions over $\tilde{\mu}'$ given, respectively, $(\tilde{e} = e_L, e = e_L)$, $(\tilde{e} = e_H, e = e_L)$, $(\tilde{e} = e_L, e = e_H)$, and $(\tilde{e} = e_H, e = e_H)$. The $2n+2 \times 1$ vector V is divided into two $(n+1) \times 1$ vectors, so that $V = \begin{bmatrix} V_L \\ V_H \end{bmatrix}$, where $V_L(i)$ is the continuation value of state $(\tilde{\mu}_i, e_L)$, and $V_H(i)$ the continuation payoff for $(\tilde{\mu}_i, e_H)$.

For example, row 7 of matrix F_{LH} gives a probability distribution over next period's reputation state variable $\tilde{\mu}$, given current state $\tilde{\mu}_7$, previous effort $\tilde{e} = e_L$, and current effort $e = e_H$. From (3), there is a $\frac{1}{2}$ probability of a success, and a $\frac{1}{2}$ probability of a failure, and so the model calls for row

[15] Since current profit $p(\tilde{\mu}_i) - c(e_j)$ does not depend on state variable $\tilde{e}$, we have that $\Pi_L(i,j) = \Pi_H(i,j)$.

7 of F_{LH} to have two non-zero elements, each equal to $\frac{1}{2}$, located in columns j and k, respectively, where $\mu_j = \mu'_F(\mu_7)$ and $\mu_k = \mu'_S(\mu_7)$. However, as the discrete nature of the model means that such j and k generically do not exist, each probability is split across two adjacent columns. For example, if $n = 100$, and we have that $\mu'_s(\mu_7) = .8463$, the 7^{th} row of F_{LH} would have a value of $\frac{1}{2} * (1 - .63)$ in column 85, and a value of $\frac{1}{2} * .63$ in column 86.

In matrix notation, the discrete version of the model is given by equation (7):

$$TV = max\left\{\begin{bmatrix} \Pi_L \\ \Pi_H \end{bmatrix} + \beta\begin{bmatrix} F_{LL}V_L & F_{LH}V_H \\ F_{HL}V_L & F_{HH}V_H \end{bmatrix}\right\} \tag{7}$$

The $\max$ operator applied to the $(2n+2)\times 2$ matrix inside of the brackets in (7) returns a $(2n+2)\times 1$ vector whose i^{th} element is the maximal element in row i of P. Let $X_L(i)$ describe the column of the maximizing element in row i, $i = 1, 2, ..., n+1$, and $X_H(j)$ describe the column of the maximizing element in row j, $j = n+2, ..., 2n+2$. Following the discussion at the beginning of this section, for any initial guess of the value function V, iteratively applying the operator T until $|T^{n+1}V - T^nV| < \epsilon$, for some $\epsilon > 0$ generates the solution to (7). The vectors $\begin{bmatrix} V_L \\ V_H \end{bmatrix}$ and $\begin{bmatrix} X_L \\ X_H \end{bmatrix}$ generated on the final iteration of (7) then describe, respectively, the firm's value function and optimal policy function over the state space.

Table 1 summarizes the matrices defined for the discretized model.

matrix name	dimension	generic element
Π_L	$(n+1)\times 2$	$\Pi_L(i,j) = p(\tilde{\mu}_i) - c(e_j)$
Π_H	$(n+1)\times 2$	$\Pi_L(i,j) = p(\tilde{\mu}_i) - c(e_j)$
F_{LL}	$(n+1)\times(n+1)$	$F_{LL}(i,j) = \Pr(\text{next state is } \tilde{\mu}_j \mid \tilde{\mu}_i, \tilde{e} = e_L, e = e_L)$
F_{HL}	$(n+1)\times(n+1)$	$F_{HL}(i,j) = \Pr(\text{next state is } \tilde{\mu}_j \mid \tilde{\mu}_i, \tilde{e} = e_H, e = e_L)$
F_{LH}	$(n+1)\times(n+1)$	$F_{LH}(i,j) = \Pr(\text{next state is } \tilde{\mu}_j \mid \tilde{\mu}_i, \tilde{e} = e_L, e = e_H)$
F_{HH}	$(n+1)\times(n+1)$	$F_{HH}(i,j) = \Pr(\text{next state is } \tilde{\mu}_j \mid \tilde{\mu}_i, \tilde{e} = e_H, e = e_H)$
V_L	$(n+1)\times 1$	$V_L(i) = V(i)$
V_H	$(n+1)\times 1$	$V_H(i) = V(n+1+i)$
X_L	$(n+1)\times 1$	$X_L(i)$ is equal to 1 if firm optimally plays effort e_L given state $(\tilde{\mu}_i, \tilde{e} = e_L)$, equal to 2 if high effort e_H is optimal
X_H	$(n+1)\times 1$	$X_H(i)$ is equal to 1 if firm optimally plays effort e_L given state $(\tilde{\mu}_i, \tilde{e} = e_H)$, equal to 2 if high effort e_H is optimal

4 Equilibrium scenarios

We now examine several possible scenarios for equilibrium firm behavior. Each scenario illustrates a type of equilibrium behavior, and the existence of each scenario is proven by example, using the discrete version of the model given by equation (7).

Scenario 1: the firm plays high effort if $\tilde{\mu}$ is either high or low. Suppose that $\beta = .99$, $\rho = .05$, $A = 5$, $c = .76$ and $\mu = .5$. In this case, in state $\tilde{e} = e_L$, the firm plays high effort for $\tilde{\mu} \in [0, .397] \cup [.813, 1]$, while for $\tilde{e} = e_H$, the firm plays high effort for $\tilde{\mu} \in [0, .361] \cup [.701, 1]$. In this case, the firm exerts high effort either to maintain a high reputation or to bolster a low reputation, but is content to coast on an intermediate level of accumulated customer good will. This scenario is the clearest illustration of proposition 1 offered in this section.

Scenario 2: the firm plays the same effort level in every period Suppose that $\beta = .99$, $\rho = .1$, $A = 5$, and $\mu = .7$. So long as the cost of high effort $c \leq .57$ the firm plays $e = e_H$ in every period, regardless of history. Unlike in traditional reputation models, there is no potential for a firm's incentives for high effort to be eroded upon repeated good outcomes. Similarly, if $c \geq .57$ the firm plays $e = e_L$ in every period, regardless of history. Here, the reputational gain from having a more favorable distribution over customer outcomes is not enough to compensate the firm for the short-term effort cost.

Scenario 3: the firm plays e_H whenever $\tilde{e} = e_L$ and for low and for high values of $\tilde{\mu}$ when $\tilde{e} = e_H$. Suppose that $\beta = .99$, $\rho = .05$, $A = 5$, $c = .75$, and $\mu = .45$. Here, a numerical implementation of (7) reveals that the firm always plays e_H in the period after playing e_L, and will play e_H in the period after e_H so long as the firm's reputation $\tilde{\mu}$ is either below .350 or above .761.

The firm's incentives in state $\tilde{e} = e_H$ illustrate the result of proposition 1; its incentive for high effort is greatest for both low and high values of the accumulated reputation $\tilde{\mu}$. When the model is simulated, the firm often plays high effort for two consecutive periods, then one period of low effort, and then repeats. Behavior following the 5% of periods in which high effort produces a failure or low effort a success may differ. Generally, should a failure occur in the same period as the firm plays low effort, it's accumulated reputation $\tilde{\mu}$ will fall to about $\tilde{\mu} = .21$. Following the period of low effort, the firm choose to play high effort regardless of $\tilde{\mu}$, and should a success occur in that period, it's reputation increases to about .78. Given the state $(\tilde{\mu} = .78, \tilde{e} = e_H)$, the firm again prefers high effort in the following period, and if this results in another success, its reputation falls to $\tilde{\mu} = .636$, at which point it prefers low effort in the following period.

Longer periods of high effort can be sustained only by occasional failures. For example, in a simulated model that ran for $100,000$ periods, the longest consecutive number of periods of high effort was 9. Here, in the first, third, fifth, and seventh periods, despite high effort being played, the outcome was a failure. Thus, accumulated reputation $\tilde{\mu}$ alternated between the range $[.761, 1]$ (following a success) and $[0, .350]$ (following a failure). Periods 8 and 9 both produces successes, driving accumulated reputation down to .635 in period 10, spurring the firm to exert low effort.

This scenario exhibits an important feature of the model: accumulated reputation may increase or decrease following a success (or a failure). Since two previous effort levels, plus luck, determine current outcome, a very high value of $\tilde{\mu}$ means that the customer is virtually certain that high effort was played in the previous period. Given this, a failure indicates virtual certainty that low effort was played contemporaneously, but a success tells the customer little, as even low effort produces a 50% chance of a success. Consequently, given a success in a period with a high value of $\tilde{\mu}$, reputation is pushed downward, towards the prior belief μ. The opposite happens following a failure in a period with a very low value of $\tilde{\mu}$; accumulated reputation will increase towards the prior μ.

Scenario 4: the firm plays $e_t = e_H$ in periods in which accumulated reputation $\tilde{\mu}$ is high, and $e_t = e_L$ in periods in which $\tilde{\mu}$ is low. Suppose that $\beta = .99$, $\rho = .1$, $A = 5$, $c = .553$ and $\mu = .55$. Here, a solution to (7) yields that in state $\tilde{e} = e_L$, current effort $e_t = e_H$ if and only if $\tilde{\mu} \geq .759$, while if $\tilde{e} = e_H$ current effort $e_t = e_H$ if and only if $\tilde{\mu} \geq .718$. In a simulated run of 100,000 periods of the model, the firm plays high effort in 8.627% of periods, and never two or more times consecutively. This implies that the only time the firm plays high effort is when accumulated reputation $\tilde{\mu}$ is high in state e_L. This happens after several failures are followed by a success; the success causes the customer to think (incorrectly) that high effort was likely played, causing the firm's reputation to shoot up to $\tilde{\mu} = .763$, incentivizing high effort. However, after two consecutive successes, the firm's reputation falls to $\tilde{\mu} = .708$, which is insufficient to incentivize further high effort in state e_H.

Why does the firm invest in high effort only in periods in which it already has a high reputation $\tilde{\mu}$? After all a success in such a period is relatively worthless to the firm, and will likely even cause its reputation to *decrease*. The reason is that given a high accumulated reputation $\tilde{\mu}$, a failure is particularly damaging to the firm's prospects, as it will cause customers to update to a very low beliefs that the firm played high effort, damaging future revenue. The logic is the same as that of proposition 1: the largest difference in $\tilde{\mu}'_S$ and $\tilde{\mu}'_F$ comes when $\tilde{\mu}$ is either very low or very high, and a success with low reputation or a failure with high reputation is quite informative.

Scenario 5: the firm plays high effort only when reputation $\tilde{\mu}$ is sufficiently low Suppose that $\rho = .1$, $A = 5$, $c = .48$, $\beta = .99$, and $\mu = .3$. Here, for $\tilde{e} = e_L$, the firm plays high effort if and only if $\tilde{\mu} \leq .505$, while if $\tilde{e} = e_H$, the firm plays high effort if and only if $\tilde{\mu} \leq .337$. The firm plays low effort for reputations above these cutoffs. In this scenario, a firm views it as worthwhile to try to buttress a truly horrible reputation, while being content to coast on a mediocre or better accumulated reputation with customers.

Scenario 6: extreme persistence of the effort state In this scenario, the firm plays high effort almost exclusively for long periods of time, punctuated only very occasionally by low effort. However, after some unlikely sequence of outcomes, the firm will eventually shift to low effort, and once in this state, will play low effort almost exclusively for a long period of time.

Suppose that $\beta = .99$, $A = 5$, $c = .8795$, $\rho = .02$, and $\mu = .65$. In this case, the firm will prefer high effort in state e_L for $\tilde{\mu} \in [.882, 1]$ and in state e_H for $\tilde{\mu} \in [0, .211] \cup [.750, 1]$. Given a prior of $\mu = .65$, the firm plays low effort initially. Successive failures drive $\tilde{\mu}$ down to a value of about $.416$, where it remains until a success occurs. Given that $\rho = .02$, a success occurs in only 2% of periods.

Given $\tilde{\mu} = .416$, a success, when it does happen, pushes the firm's accumulated reputation to $.855$. However, given that the success was the result of luck, and not of high effort, the firm continues to prefer low effort in the state $(\tilde{\mu} = .855, \tilde{e} = e_L)$; the accumulated reputation must top $.882$ for the firm to prefer high effort. For the firm to acquire such a high reputation, its customers must think it very likely that it played low effort in the previous period, and a success must occur in the current period. Since a long string of failures only drives reputation down as low as $\tilde{\mu} = .416$, the only way this can happen is for the sequence success, failure, success to occur in three consecutive periods. The first success raises $\tilde{\mu}$ to $.855$, the failure lowers it to $.225$, and a second success raises $\tilde{\mu}$ all the way to $.898$, at which point the firm switches to high effort. The probability of such a sequence is $.02 * .98 * .02 = .000392$.

Once the firm makes the switch to high effort, its reputation stabilizes at $\tilde{\mu} = .802$ given continual successes. However, as soon as there is a failure, its reputation falls to $.264$, which is low enough for the firm to prefer low effort in the following period. So, once in the high effort state, the firm switches back to the low effort state with probability $.02$. Once in the low effort state, his accumulated reputation again stabilizes at $.416$, and he remains in this state until the sequence success, failure, success reoccurs.

5 More persistent effort

Sections 3 and 4 posit a model in which two periods of effort affect current outcomes. Here, we comment on the effect of more persistent effort. Suppose that current outcome is a function of ρ, current effort, and effort in the previous k periods. How does k influence firm incentives for high effort?

First, consider the direct extension of equation (3), which gives equal weight to 2 effort levels in determining current outcome. Specifically, assume that the probability of a success is determined by $f(e_t, e_{t-1}, ..., e_{t-k})$:

$$f(e_t, e_{t-1}, ..., e_{t-k}) = \rho + (1 - 2\rho)\frac{e + \Sigma_{j=1}^{k} e_{t-j}}{k+1} \tag{8}$$

What effects does this change have on the firm's optimization problem (2)? First, it increases the number of state variables, as the customer must track beliefs for effort not only in periods t and $t-1$, but effort in periods $t-k$ through $t-2$ as well. Second, and more importantly, it diminishes the effect of any one period's effort on outcome, including the current period's. By spreading the relevant effort levels over a larger number of periods, any one effort level becomes less important to future outcomes, and so the firm's incentive to exert high effort is decreased. Consequently, for any parameters ρ, c, β, and A, there exists some cutoff $\overline{k}$ such that $k > \overline{k}$ implies that the firm exerts low effort in every period.

On the other hand, if the determinants of outcome depend on current effort and the previous k effort levels, but the the importance of effort levels decays over time, then the model is qualitatively similar to the two-period model studied above. For example, consider a model in which current outcome is determined by an infinite history of effort levels as follows:

$$f(\cdot) = (1 - \rho)\Sigma_{t=1}^{\infty} \left(\frac{1}{2}\right)^{t} * e_t$$

Since, much as in equation (3) in the two-period model, current effort determines half of the non-probabilistic portion of the current outcome, this model will behave qualitatively similar to the two-period model, despite being much more difficult to solve.

6 Conclusion

This paper studies the reputational effects of persist effort. The result is a model which is arguably more intuitive than type-based reputation models. Rather than relying on a black box exogenous type distribution, firms have an additional incentive to exert high effort because good outcomes today are correlated with good outcomes tomorrow through the persistent effects of effort. Thus, hiring decisions, decisions affecting corporate culture, or capital purchases affecting quality all have effects that last for at least a couple of periods, but not forever. A version of the model in which effort persists for two periods can explain not only why firms might exert high effort despite having no short-run incentive to do so, but also extreme persistence of effort levels, that is firms playing high effort for long stretches before switching to low effort for a similarly long stretch, and then possibly back again. Such extreme persistence could explain the gradual changes in fortune of certain companies.

References

S. Board and M. Meyer-Ter-Vehn. Reputation for quality. *Econometrica*, 81(6):2381–2462, 2013.

J. A. Bohren. stochastic games in continuous time: Persistent actions in long-run relationships. 2013. manuscript.

M. W. Cripps, G. J. Mailath, and L. Samuelson. Imperfect monitoring and impermanent reputations. *Econometrica*, 72(2):407–432, 2004.

F. Dilme. Building (and milking) trust: Reputation as a moral hazard phenomenon. 2012. manuscript.

E. Faingold and Y. Sannikov. Reputation in continuous-time games. *Econometrica*, 79(3):773–876, 2011.

A. Fernandes and C. Phelan. A recursive formulation for repeated agency with history dependence. *Journal of Economic Theory*, 91(2):223–247, 2000.

D. N. Figlio, M. O. Schapiro, and K. B. Soler. Are tenure track professors better teachers. 2013. manuscript.

J. Horner. Reputation and competition. *American Economic Review*, 92(3):644–663, 2002.

A. Jarque. Repeated moral hazard with effort persistence. *Journal of Economic Theory*, 145:2412–2423, 2010.

Q. Liu. Information acquisition and reputation dynamics. *The Review of Economic Studies*, 78(4): 1400–1425, 2011.

G. J. Mailath and L. Samuelson. Who wants a good reputation. *Review of Economic Studies*, 68(2): 415–441, 2001.

P. Milgrom and J. Roberts. Predation, reptuation, and entry deterrence. *Journal of Economic Theory*, 27(2):280–312, 1982.

R. A. Posner. The law and economics of the economic expert witness. *Journal of Economic Perspectives*, 13(2):91–99, 1999.

R. Rob and T. Sekiguchi. Reputation and turnover. *RAND Journal of Economics*, 37(2):341–361, 2006.

J. A. Sandford. Experts and quacks. *RAND Journal of Economics*, 41(1):199–214, 2010.

J. A. Sandford. Competition and endogenous impatience in credence good markets. *Journal of Institutional and Theoretical Economics*, 169:531–565, 2014.

S. Tadelis. What's in a name? reputation as a tradeable asset. *American Economic Review*, 89(3): 548–563, 1999.