

Bad Reputation ^{*}

Jeffrey C. Ely[†] Juuso Välimäki[‡]

January 21, 2003

Abstract

We construct a model where the reputational concern of the long-run player to look good in the current period results in the loss of all surplus. This is in contrast to the bulk of literature on reputations where such considerations mitigate myopic incentive problems. We also show that in models where all parties have long-run objectives, such losses can be avoided.

1 Introduction

We construct a model where a long-lived agent's concern for his reputation undermines commitment power and results in the loss of all surplus. This provides a cautionary counterpoint to a pervasive idea in economic theory; namely, that the reputation of an agent with long-run interests provides implicit incentives to keep short-run commitments and can thereby substitute for explicit contractual enforcement.

The conventional wisdom that reputation enhances commitment power has its theoretical foundation in the game-theoretic literature on reputation effects pioneered by Kreps and Wilson (1982) and Milgrom and Roberts (1982) and extended in Fudenberg

^{*}This paper has previously been peer reviewed in NAJ Economics. See <http://dlevine.econ.ucla.edu/v4.htm>. We thank Gregory Pavlov for invaluable research assistance. We received helpful comments from Drew Fudenberg, David Levine, Thomas Mariotti, Stephen Morris, Michele Piccione, Tom Wiseman, the editor, and three anonymous referees.

[†]Department of Economics, Northwestern University. ely@nwu.edu. Financial support from NSF grant #9985462 is gratefully acknowledged.

[‡]Department of Economics, University of Southampton and Helsinki School of Economics. valimaki@hkkk.fi

and Levine (1989).¹ The general message of this literature is that reputation effects leave the agent at least as well-off as he would be in the complete absence of external incentives,² but typically raise long-run payoffs, often to the agent's first-best.

Just as in these traditional models of reputation, we consider a model where a sequence of short-lived players interact with the long-run agent. The short-lived players are uncertain about the type of the long-run player. A good type has incentives perfectly aligned with the short-run players while a bad type is committed to behavior that is harmful to the short-run players. In our model, the only decision for the short-run players is whether to engage the long-run player in an interaction. If the long-run player is sufficiently patient, then it will be in her best interest to choose actions that separate her from the bad type in order to avoid a bad reputation in the future. In many instances, such separating actions also hurt the short-run player. Hence as soon as the reputational incentive to separate from the bad type becomes sufficiently strong, the short-lived players find it in their best interest not to participate and the market fails altogether.

Our model is applicable in a variety of economic situations. In the market for credence goods, the seller first diagnoses the client's needs and then chooses a product to sell. Mechanics, medical doctors, management consultants and lawyers are prime examples of such sellers. Suppose the client knows that a bad consultant will suggest a major reorganization of the firm (to get larger fees) even when smaller changes would be appropriate. Then it may be in the best interest of a good consultant to suggest minor reforms to separate from the bad consultants, even if more drastic measures are necessary. But now from the client's point of view, even the good type is bad. Concerns of similar type are also present in the medical and legal profession. Other instances where our model might be applicable include the policy recommendations of the IMF, where too stringent conditions for a country applying for extra debt financing might be imposed in order to achieve separation from too lenient regime.

In order to give a clear and concrete application, our main model considers a long-

¹Further developments appeared in Fudenberg and Levine (1992), ?, ?, ?, and ?.

²The formal expression of this statement is the following. Without any external incentives the agent could expect at least his worst, and at most his best, Nash equilibrium payoff. Fudenberg and Levine (1992) show that reputation effects restrict the set of payoffs a long-run agent can expect and they provide explicit lower and upper bounds. The Fudenberg and Levine (1992) lower and upper bounds weakly exceed the worst and best Nash equilibrium payoff, respectively. In relation to these observations, an interpretation of our results is that a tighter upper bound is possible.

lived auto mechanic who wishes to maintain a reputation for being scrupulous, i.e. one who never inflates his diagnosis of the problem and always does the necessary repair. A sequence of short-lived motorists need to have a problem with their car diagnosed and repaired. Each motorist observes which repairs were done in the past, but only the mechanic knows which repairs were actually necessary. Thus the mechanic's reputation can be built only on the frequency with which various repairs are performed.

The logic of our bad reputation result is as follows. First, although the necessary repairs are i.i.d. across periods, there is always a chance that the most costly repair, say replacing the engine, may be necessary with disproportionately high frequency over many periods. In such an event, even a truly scrupulous mechanic will appear to his potential future customers to be exactly the opposite: a bad mechanic who ignores what is truly necessary and replaces the engine. Along such a history, the short-run reputation incentives soon become inconsistent with the mechanic's original objective of being scrupulous. Indeed, in order to maintain the belief that he is scrupulous, the mechanic must distort his behavior to bring the frequency of engine replacements back in line with what would be typical for a scrupulous mechanic. In extreme cases, the mechanic's reputation can only be salvaged by performing a minor repair in the current period, even if a more serious repair is in order. But once this is true, the current motorist will refuse to bring in his car.

The preceding argument suggests that unfortunate events can lead even a good mechanic to get stuck with a bad reputation. The main analysis in this paper concerns the effect of this future possibility on incentives in the early stages when reputation is still being formed. We show in fact that the incentive effects of these distant and unlikely histories multiply back even to the earliest stages of the reputation-building process, undermining the mechanic's otherwise good intentions and resulting in the loss of all surplus.

We demonstrate this result in two versions of the model. In our basic version of the model we formalize reputation in the classical way: incomplete information with types that are committed to certain stage-game strategies. Thus, the reputation to avoid is that of being committed to an *unscrupulous* strategy. This model gives the cleanest version of our result and also facilitates comparison with Fudenberg and Levine (1989), Fudenberg and Levine (1992) and the literature that followed.

Next we consider a version of the model in which the bad type of mechanic is an otherwise standard strategic player who simply has mis-aligned incentives. In this

case, the analysis is very different and considerably more subtle. In particular, now *both* types of the long-lived player act to balance short-run payoffs with long-run reputation maintenance. Indeed, even the bad mechanic is prepared to recommend a minor repair if that leads to a substantial improvement in reputation providing further opportunities to replace engines. In the typical equilibrium the bad mechanic randomizes and this slows down the reputation dynamics. From a technical standpoint this necessitates a new analytical approach because the reputation dynamics are endogenous.³ In terms of the results, this raises the possibility that bad reputation effects can be mitigated. If reputation loss from engine replacements is sufficiently slow relative to the discount factor, then the good mechanic can achieve high payoffs. Nevertheless, we are able to show that in all renegotiation-proof Nash equilibria the bad reputation effect is still predominant and a sufficiently patient mechanic is never hired.

The source of the problem in each of our versions of the model is evidently an informational externality among the short-run players. When a motorist hires the mechanic, he provides the mechanic with the opportunity to reveal additional information about his type. This additional information is valuable to potential future customers, but the current customer does not internalize this value when he makes his hiring decision. To illustrate this interpretation of the problem, we turn in the last section to a version of the model in which the motorist (the principal) is also a long-run player.

The principal's objective is to discover the type of mechanic (the agent) he is facing, fire the bad type and continue to hire the good type. As a long-lived player, the principal directly internalizes the benefits of experimenting with the agent. With this externality eliminated, we show that bad reputation is no longer a problem. We construct a class of sequential equilibria which have the following properties. First, no matter how high is the prior probability of a bad agent, when the discount factor is near enough to one, there will be perfect screening: the bad agent will eventually be discovered and terminated and the good agent will be hired in every period. Second, this screening is essentially costless for a patient principal: the equilibrium payoff to the principal converges to the value he would obtain if he were to observe the agent's

³In particular, we cannot bound in advance the absolute number of stages that are required before the mechanic's reputation is established or lost. In the analysis developed by Fudenberg and Levine (1989) for the case of commitment types, such a bound exists independently of the discount factor and this bound translates into a bound on the payoffs of the long-run player. With strategic types, the number of engine replacements which leads to the loss of reputation depends on the equilibrium strategy of the bad mechanic, and this in turn depends on the discount factor.

type before play begins.

Our model is most closely related to Morris (2001). Morris builds on the multi-stage cheap-talk model of Sobel (1985)⁴ in which an advisor has repeated opportunities to provide information for a decision-maker. The advisor, like our mechanic, has two possible types: one whose preferences are aligned with the decision-maker and one who would like to exploit the decision-maker. Morris builds reputational concerns into the model by supposing that the good advisor applies higher weight to tomorrow’s decision than today’s. When this weight is high enough, the unique outcome in the first period is the uninformative “babbling” outcome. In a two-stage cheap talk model, babbling arises because there is no payoff-relevant means of separating the two types in the first period if reputational concerns are sufficiently important for the second period. In our model, the mechanic takes a payoff-relevant action, so the logic underlying bad reputation is qualitatively different and in particular relies on a long horizon.⁵⁶ While Morris introduces an “instrumental” concern for reputation, there are many papers in which perverse incentives can be created when agents have intrinsic or market induced preferences for their reputation. These include Prendergast (1993), Prendergast and Stole (1996), Sharfstein and Stein (1990), and Zwiebel (1995).⁷

The remainder of this paper is organized as follows. In Section

2 The Basic Model

Consider the following situation. A motorist (the principal) has a car which is in need of repair. The motorist knows that the car requires one of two repairs with equal probability: an engine replacement or a mere tune-up; however the motorist lacks

⁴Another paper in this spirit is Bénabou and Laroque (1992) which looks at the ability of an insider to make manipulative recommendations to investors and profit from the movements in stock prices. Both Sobel (1985) and Bénabou and Laroque (1992) are “good reputation” models because the opportunistic advisor benefits from the presence of a committed honest type.

⁵With a small probability of the bad type and only two periods, there exist efficient renegotiation-proof equilibria for any relative weights on payoffs in the two periods.

⁶Morris also considers a parameterized infinite horizon version of the model in which the good advisor is infinitely more patient than the bad advisor. This model is closer to our model with commitment types since an impatient bad advisor will have no reputation concerns.

⁷On the other hand, models such as Holmström (1999) and Mailath and Samuelson (1998) are “good reputation” models. Although “career concerns” can induce inefficient effort, inefficient effort improves on no effort at all which is what would result without the reputational motive in these models.

the expertise to determine which repair is necessary. The motorist therefore considers bringing the car to a certain mechanic (the agent) who possesses the ability to diagnose the problem and perform the necessary repair. We will model this by supposing that the mechanic, if hired, privately observes the state $\theta \in \{\theta_e, \theta_t\}$, indicating respectively whether an engine replacement is necessary, or a tune-up will suffice. Conditional on this information, the mechanic then chooses a repair $a \in \{e, t\}$, indicating engine replacement, or tune-up. We can represent the mechanic's repair strategy by the pair of probabilities (β_e, β_t) where for each a , β_a is the probability that the mechanic performs repair a in state θ_a , i.e. when a is indeed the necessary repair.

The following table shows the payoffs to the motorist from the two possible repairs $\{t, e\}$ in the two different states $\{\theta_t, \theta_e\}$.

	θ_t	θ_e
t	u	$-w$
e	$-w$	u

The motorist also has an outside option which gives a constant payoff normalized to zero.⁸ We will assume $w > u > 0$. This implies that if the mechanic chooses the repair independently of the state, the payoff to the motorist is negative and the motorist would strictly prefer not to hire a mechanic who employs such a strategy. This captures the essential strategic feature of the problem we are studying in the current paper: the mechanic is an agent who possesses some expertise and (most importantly) expert information. The agent's services are valuable only if he can be expected to make use of this information in selecting a course of action. Specifically, the motorist would be willing to hire the mechanic only if $\beta_a \geq \beta^* \equiv \frac{w-u}{u+w} > 0$ for each a .

This paper is about the distortionary consequences of the incentive to avoid bad reputations. To see these effects in their most extreme form, we begin by considering the benchmark scenario in which the mechanic is known to be *good*; that is, his payoffs are identical to those of the motorist. In this case there is no incentive problem and the first-best outcome is the unique sequential equilibrium of the one-shot interaction between the motorist and mechanic: the mechanic will do the correct repair if hired, and the motorist therefore optimally hires.

⁸For example, we could assume that the payoffs u and w already incorporate the fixed diagnosis fee charged by the mechanic, and the outside option could be the motorist choosing a repair at random and delegating it to some other mechanic.

This conclusion remains true even when the motorist believes with small but positive probability μ that the mechanic is *bad*. A bad mechanic is any mechanic whose choice of repair is independent of the state. As argued above, the motorist's value for such a mechanic is strictly negative. Nevertheless, the (constrained) efficient outcome remains the unique sequential equilibrium when μ is less than some critical value $p^* < 1$. The good mechanic can still be expected to choose the appropriate repair and the motorist optimally hires given that the mechanic is sufficiently likely to be good. On the other hand, when $\mu \in (p^*, 1]$, the motorist is too pessimistic about the mechanic's type and strictly prefers not to hire even if the good mechanic can be expected to do the right thing.

Reputational incentives have the potential to distort the behavior of the good mechanic only when motorists base their hiring decisions on information about the service received by earlier customers. To model this effect, suppose now that a sequence of motorists decide in turn whether or not to hire the same mechanic. Each motorist can observe the repair performed for preceding customers but none can directly observe the repair that was actually necessary.⁹

The good mechanic chooses a strategy, which in this dynamic game is a pair of probabilities $(\beta_e^k(\tilde{h}_k), \beta_t^k(\tilde{h}_k))$ specifying the probability of changing the engine and performing a tune-up respectively at date k as a function of his previous history \tilde{h}_k . The mechanic's history records for each previous date l , the pair (θ^l, a^l) consisting of the state and the chosen repair,¹⁰ The good mechanic maximizes the expected discounted average payoff where $\delta \in (0, 1)$ is the discount factor. The bad mechanic replaces the engine whenever hired, regardless of the history.

Each motorist in the sequence will decide whether to hire the mechanic, and this decision will be based on the previous (public) history h^k of repairs. The motorists are "short-run" players: each cares only about his or her own repair. This means that a motorist's hiring decision is based solely on the updated probability $\mu^k(h^k)$ that the mechanic is bad, and the expected behavior $\bar{\beta}_a^k = \mathbf{E}(\beta_a^k|h^k)$ of the good mechanic, $a \in \{e, t\}$ ¹¹. In particular, when $\mu^k(h^k) > p^*$ the motorist will refuse to hire, and when $\mu^k \leq p^*$, the motorist will hire only if $\bar{\beta}_a^k \geq \beta^* > 0$ for each a .

In this repeated interaction, the good mechanic's incentives at date k are a mix

⁹This assumption can be relaxed to allow for a public imperfectly observed signal on the state.

¹⁰With a^l equal to the null action if the mechanic was not hired in period l .

¹¹The good mechanic is potentially conditioning his behavior on past states, which are not observed by the motorist.

of the short-run desire to choose the right repair for the current customer, and the long-run strategic objective to maintain a good reputation. We show next that this reputational incentive distorts the otherwise good intentions of the mechanic, and that the severity of this distortion increases with the discount factor. In fact, when the good mechanic is sufficiently patient, his services become essentially worthless. Formally, let $\bar{V}(\mu, \delta)$ be the supremum of discounted average Nash equilibrium payoffs for the good mechanic when μ is the prior probability that the mechanic is bad, and δ is the discount factor. Because the good mechanic's payoffs are identical to those of the motorists, this also measures the value of the mechanic's services to the population of motorists when δ is close to 1. Our first theorem shows that this average discounted equilibrium payoff is small for high discount factors.

Theorem 1 *When the bad mechanic is a commitment type, for any $\mu > 0$,*

$$\lim_{\delta \rightarrow 1} \bar{V}(\mu, \delta) = 0. \quad (1)$$

Proof: To begin with, recall that if $\mu > p^*$, then motorists are too pessimistic and there is a unique Nash equilibrium outcome in which the mechanic is never hired. Now suppose that $\mu \leq p^*$ and consider a Nash equilibrium in which the mechanic is hired. The updated probability of a bad mechanic will depend on the repair chosen. When the mechanic chooses to replace the engine, Bayes' formula gives the updated probability as follows:

$$\mu^1(e) = \frac{\mu}{\mu + (1 - \mu) \left[\frac{1}{2}\beta_e + \frac{1}{2}(1 - \beta_t) \right]}$$

Recall that the mechanic is hired only if $\beta_t \geq \beta^* > 0$. Therefore because $\mu \in (0, p^*)$, the posterior $\mu^1(e)$ strictly exceeds μ . That is, observing an engine replacement is always bad news. Let us thus define for each μ ,

$$\Upsilon(\mu) = \frac{\mu}{\mu + (1 - \mu) \left[\frac{1}{2} + \frac{1}{2}(1 - \beta^*) \right]}$$

i.e. the smallest possible posterior probability of a bad mechanic when an engine replacement has been observed. The preceding argument implies $\Upsilon(\mu) > \mu$ for every $\mu \in (0, p^*)$. It is also easy to see that Υ is strictly increasing and continuous.

We construct a decreasing sequence of cutoff points p_m defined by $p_1 = p^*$ and $p_m = \Upsilon^{-1}(p_{m-1})$ for $m > 1$. We will use an induction on m to bound the payoffs across all Nash equilibria when the prior exceeds p_m . For the induction hypothesis, suppose that there exists a bound $\bar{V}_m(\delta)$ with $\lim_{\delta \rightarrow 1} \bar{V}_m(\delta) = 0$ and $\bar{V}(\mu, \delta) \leq \bar{V}_m(\delta)$ for all

$\mu > p_m$. Note that we have already shown this for $m = 1$. Fix $\mu > p_{m+1}$. It suffices to consider a Nash equilibrium in which the mechanic is hired in the first period.¹²

Since the good mechanic is hired in the first period, he must be choosing the correct repair with positive probability in each state. In particular, this must be a best-response for the good mechanic, and this implies the following bound on the equilibrium payoff.

$$\bar{V}(\mu, \delta) \leq (1 - \delta)u + \delta \left[\frac{z(e|\theta_e) + z(t|\theta_t)}{2} \right] \quad (2)$$

where $z(a|\theta)$ represents the expected continuation payoff in state θ , from choosing action a . Because we have taken $\mu > p_{m+1}$ we know that the posterior $\mu^1(e)$ is at least $\Upsilon(p_{m+1})$ which by definition is at least p_m . Thus, the good mechanic can expect $z(e|\theta_e)$ to be no more than $\bar{V}_m(\delta)$ ¹³

Consider the following incentive compatibility constraint for the good mechanic:

$$z(t|\theta_e) \leq \frac{1 - \delta}{\delta} (u + w) + \bar{V}_m(\delta)$$

Since the motorist is hiring the mechanic, we must have $\beta_e > 0$, i.e. the mechanic chooses to replace the engine when that is the appropriate repair (at least with positive probability.) The above inequality is a necessary condition for the good mechanic to do so in equilibrium: the long-run payoff resulting from an engine replacement can be no less than that of a tune-up.

Now since the motorist's behavior is conditioned only on the public history, the mechanic's continuation value from choosing t cannot depend on the mechanic's private history. Thus $z(t|\theta_e) = z(t|\theta_t)$, and we can substitute the bound obtained from incentive compatibility into (

By induction it now follows that (

In the simple repeated game described in the previous section, there are many equilibria, including equilibria in which the mechanic is a large number of times on the equilibrium path. However, the bound in (

¹²Take any Nash equilibrium in which the mechanic is hired with zero probability until date $k > 1$. Then the continuation play beginning at date k is a Nash equilibrium with prior μ . Since the mechanic's payoff was zero in the first $k - 1$ periods, and the mechanic's minmax payoff is zero, the continuation payoff in the equilibrium that begins at date k can be no smaller than the payoff to the original equilibrium.

¹³We are using here the fact that the continuation play of any Nash equilibrium beginning with a history that is on the equilibrium path must itself be a Nash equilibrium of the continuation game whose prior is the updated posterior. The mechanic has some private history, but this is nothing more than an explicit randomization device for his own mixed strategy.

One could argue that these equilibria have an implausible feature: if the mechanic has ever performed a tune-up prior to date k , he will have perfectly revealed himself to be good. Nevertheless, it must be the case that even after these histories, the mechanic will not be hired frequently. If he were hired, then the incentive to separate would again be too strong, and the mechanic would perform a tune-up with probability 1.¹⁴ Since there is no incentive problem when the mechanic is known to be good, one would expect that in this case the motorists would continue hiring the mechanic. We can see this as an extremely weak form of a renegotiation-proofness argument: once revealed to be good, the mechanic would renegotiate with future motorists to realize their mutual gain.

While there is no established renegotiation-proofness refinement for repeated games with incomplete information¹⁵, any reasonable definition would entail the following minimal restriction.

Assumption 1 (Renegotiation-Proofness) *The mechanic is hired at any history on the equilibrium path at which he is known to be good.*

Matters are dramatically worse when we impose renegotiation-proofness.

Theorem 2 *When δ is close enough to 1, there is a unique renegotiation-proof Nash equilibrium outcome. In that outcome the mechanic is never hired.*

Proof: Let \bar{p} be the supremum of the set of posteriors at which the mechanic is hired with positive probability on the equilibrium path of some such equilibrium. Obviously $\bar{p} \leq p^* < 1$. We will show in fact that $\bar{p} = 0$. Suppose on the contrary that $\bar{p} > 0$. Then let \tilde{p} be any posterior greater than $\Upsilon^{-1}(\bar{p})$. Suppose the mechanic is hired on the equilibrium path of some Nash equilibrium at a history h^k where the posterior is $\mu^k(h^k) = \tilde{p}$. Then $\bar{\beta}_t(h^k) \geq \beta^*$ implying that if the mechanic replaces the engine, the new posterior will exceed $\Upsilon(\tilde{p}) > \bar{p}$. By the definition of \bar{p} , this will lead to a continuation value of zero. On the other hand, if the mechanic performs a tune-up, he

¹⁴It is possible to construct continuation sequential equilibria with no hiring even when the mechanic is known to be good. This is a consequence of the folk theorem for games with long-run and short-run players proved by Drew Fudenberg and Maskin (1990) and extended to the case of imperfect monitoring by Fudenberg and Levine (1994)

¹⁵The concepts in Abreu, Pearce, and Stachetti (1993) and Farrell and Maskin (1989) apply to strategic-form games with complete information. Applying either the “external consistency” of the former or the “internal consistency” of the latter to the complete information version (i.e. $\mu = 0$) of our extensive form game would yield Assumption

will reveal himself to be a good mechanic and will be hired in each subsequent period. This yields continuation value u . If δ is big enough to satisfy $-(1-\delta)w + \delta u > (1-\delta)u$, the mechanic will strictly prefer to perform a tune-up, even when an engine-replacement is necessary. Thus $\bar{\beta}_e(h^k) = 0$. This implies that the motorist at date k will refuse to hire. In fact we have shown that there is no renegotiation-proof Nash equilibrium in which the mechanic is hired on the equilibrium path when the posterior is greater than $\Upsilon^{-1}(\bar{p}) < \bar{p}$. But this contradicts the definition of \bar{p} , and therefore $\bar{p} = 0$. ■

3 Strategic Bad Type

Patrons of automobile garages are no doubt suspicious of the scruples of their mechanics. We have modeled this previously by introducing incomplete information about the mechanic in the form of a commitment type. While this is consistent with the traditional game-theoretic analysis of reputation, it may not be the most realistic way to capture the concerns of wary motorists. If the motorist is concerned that the mechanic has too much short-run incentive to replace engines, then the motorist should anticipate that such an unscrupulous mechanic would himself act to preserve his reputation and secure the opportunity to replace additional engines in the future. This of course would lead even the bad mechanic to perform tune-ups from time to time.

To model this possibility, we now introduce incomplete information about the *payoffs* of the mechanic. The good mechanic's payoffs are as before, but now that the stage payoffs to the bad mechanic are as follows.

	θ_t	θ_e
t	$-w$	$-w$
e	u	u

Thus the bad mechanic, absent reputation effects, always prefers to replace the engine. We assume that the bad mechanic does not know the necessary repair so that there are no gains from trade between the bad mechanic and the motorist. The bad mechanic maximizes the discounted sum of stage payoffs. In contrast to the previous section, the bad mechanic has now reputational concerns. He may have a strategic incentive to play t in order to pool with the good type of mechanic if this will increase the frequency with which he will be hired in the future. This effect has

the potential to improve the outcome because when the bad type is playing t with positive probability, a play of t is a weaker signal of a good mechanic. This reduces the incentive for the good mechanic to play t when the state is θ_e , improving the payoff for the motorist.

Nevertheless, the results of the previous subsection continue to hold with a strategic bad type. Here is a rough intuition. If the bad type were to play t with positive probability, it is because t will lead to a better “reputation” than e . Here reputation means continuation value, which is directly related to the frequency with which the mechanic will be hired in the future. A crucial observation is that at critical histories; those in which a play of e will lead to complete loss of reputation (never being hired again), the good type of mechanic values the improvement in reputation strictly more than the bad type of mechanic. A simple example illustrates the intuition. Suppose a critical history h has been reached and suppose that a choice of t allows the mechanic to “survive” for at least one more period, but gives only a small improvement in reputation. This means that a few more plays of e would again lead to a critical history. At h , the bad type is sacrificing today’s payoff for these few additional opportunities to play e . On the other hand, consider the good type with signal θ_e . By playing t , he is also sacrificing today’s payoff, but in return he gets not only the opportunity to do his preferred action a few more times, but in addition the opportunity to *costlessly* further improve his reputation by playing t if the state is θ_t tomorrow.

This argument shows that at critical histories, either the bad mechanic plays e with probability 1, in which case we are in a situation identical to the non-strategic model, where the good mechanic strictly prefers to play t and separate, or the bad mechanic mixes, in which case again the good mechanic strictly prefers to play t . In either case, the motorist would prefer not to hire. The formal proof appears in Appendix

Theorem 3 *In the infinite horizon model with a strategic bad type with prior probability greater than zero, the mechanic is never hired in any renegotiation-proof Nash equilibrium when the discount factor is close enough to 1.*

4 Long-Run Principal

We can identify two strategic themes that contribute to the bad-reputation effect. The first is an inability on the part of the mechanic to commit not to invest in his reputation. Because the motorists anticipate the mechanic’s reputation motive, they refuse to hire

when this motive is too strong. The second theme is an information externality among the short-run motorists. The motorists who refuse to hire do so because they care only about their own repair and they do not internalize the value of the information they would reveal to future motorists if they were to hire.¹⁶

When the private information of the mechanic is not verifiable there is no way to eliminate the commitment problem. In this section we explore whether the problem can be ameliorated in environments in which the externality can be internalized. In particular, we consider the case of two long-run players. The principal (corresponding to the motorist in the previous sections) must decide in each period whether or not to hire the agent (corresponding to the mechanic). The agent is either good or bad, the latter with prior probability μ . A good type of agent observes a signal from $\{\theta_e, \theta_t\}$ and chooses an action from $\{e, t\}$. The stage payoffs to the good type as a function of the signal and action are given in figure

The results in this section are positive. With discount factors close to 1, there exist sequential equilibria in which the principal achieves approximately his full-information value. That is, when the principal internalizes the full benefits of his experimentation with the agent, he can achieve the first-best. This is true despite the fact that the agent is also becoming increasingly patient and hence the temptation to manipulate reputation also increases. In fact, in the equilibria we construct, even the good agent will, from time to time, sacrifice short-run payoffs in order to separate from the bad type. The patient principal is willing to bear this short-run cost to reap the long-run benefits of learning.

Before describing the structure of the equilibria, let us illustrate the challenges involved by viewing the problem purely from a mechanism-design point of view. If the principal could commit to a mechanism and sought only to eventually separate the agent's types, the following employment contract would achieve this objective. During an initial evaluation phase of length K , the agent is asked to signal his type. The good agent should play t and the bad agent should play e in each of these initial periods. At the end of the evaluation period, the agent is permanently fired if he played e in at least one period, otherwise he is promoted and hired forever thereafter regardless of future actions. The length K is chosen sufficiently long in order to satisfy incentive compatibility.

¹⁶An informational externality is present in the earlier literature on reputation (e.g. in Fudenberg and Levine (1989)), but there it does work to the advantage of the long run player

There are three problems with this mechanism. First, although it is incentive compatible for the agent, it is not sequentially rational for the principal. After the first choice of e , the agent has been revealed to be bad and the principal strictly prefers not to hire the agent through the remainder of the evaluation period. Thus, we face the following challenge in constructing a sequential equilibrium which screens the agent in finite time. There necessarily arrives a stage in which the bad agent must voluntarily play e knowing that this will lead to immediate termination.

Second, the contract is not renegotiation-proof. After the first play of t , the principal knows that the agent is good and yet the principal asks the agent to play t independent of the state for the remainder of the evaluation period. Instead, the principal and agent would like to renegotiate the contract and promote the agent. In order to satisfy renegotiation-proofness, the equilibrium must promote the agent as soon as he is revealed to be good.

Finally, the principal's payoff in this contract is bounded away from the first-best even as the discount factor approaches one. To see this, take the incentive compatibility constraint for the bad agent

$$(1 - \delta^K)u \geq -(1 - \delta^K)w + \delta^K u$$

This implies that $\delta^K \leq \frac{u+w}{2u+w}$. But then the principal's payoff conditional on the good agent is

$$(1 - \delta^K) \left(\frac{u - w}{2} \right) + \delta^K u$$

and can be shown by simple algebra to be bounded away from u independent of the discount factor.

Our equilibria also have an evaluation period, but the length of the evaluation is not fixed but is repeatedly adjusted depending on the behavior of the agent. Specifically, the agent begins with a "score." This score is reduced each time e is played and increased each time t is played.¹⁷ When the score reaches zero, the agent is permanently fired. When the score achieves some fixed high threshold, the agent is promoted. The good agent will choose the appropriate action whenever the score is strictly greater than one. When the score is exactly one, the good agent will play t in order not to be fired. The bad agent will choose e in each period with probability close to 1. Thus,

¹⁷The actual increments are randomly chosen after some histories in order to satisfy some indifference constraints. For simplicity, we assume that these increments are determined by the outcome of a public randomization device and hence that the agent always knows his current score.

the bad agent will be fired in finite time. Moreover we show that the good agent will be playing the right action almost all of the time. It follows that a patient principal will achieve his full-information value. The formal proof is in Appendix

Theorem 4 *Given any prior μ , for any sequence of discount factors approaching one, there is a sequence of equilibria in which the principal's average payoff approaches its full-information value: $(1 - \mu)u$. Given any discount factor close enough to one, and any sequence of priors approaching zero, there is a sequence of equilibria in which the principal's average payoff approaches u .*

Our model in this section is an example of a repeated game with incomplete information on one side. There is already a well-established theoretical literature on such games studying many variations of the model¹⁸. However, none of the well-studied variations cover the specific environment we have here. An extensive literature deals with undiscounted games, the seminal references being Hart (1985) and Aumann and Maschler (1995). For the discounted case, the literature seems to have restricted attention to the case of perfect monitoring Cripps and Thomas (2001). In the game we study in this section there is discounting and imperfect monitoring. Thus our construction of efficient equilibria is new. In fact, the proof of Theorem

Finally, let us remark on a technical issue. We face a common conceptual difficulty that arises in repeated games with incomplete information. If the principal is observing his payoffs over time, then these should provide information about the agent's previous signals and this could help in distinguishing the good agent from the bad. In order to maintain an information structure consistent with the preceding sections, we will suppose that the principal gets no information about the agent's previous signals.

This can be motivated in a number of ways. There are many situations in which a principal has a preference for a certain rule of behavior by his agent, without ever seeing any physical "payoff." For example, the agent could be a judge and the principal the government official in charge of reappointing judges. The principal wants a judge who will rule "fairly." Here the signal would represent the outcome of the judge's deliberations. Without the judge's legal expertise, it would be impossible for the principal to infer the signal from the plain evidence brought to trial. Thus the principal would have to base reappointment decisions solely on the relative frequencies of verdicts.

Alternatively, the principal may have access to information about the signals of the agent, but this information may not be verifiable to a court. If the relationship between

¹⁸Sorin (1999) includes an illuminating survey.

the principal and agent is governed by an employment contract which can only depend on verifiable performance data, the situation would be equivalent to the one studied here.

In any case, allowing the principal to observe past payoffs would only improve his ability to monitor the agent and hence we obtain a stronger result with our formulation.

5 Conclusion

We have constructed a game where the desire to avoid bad reputation eliminates all possibilities for profitable interactions between a long-lived agent and a sequence of short-lived principals. In the construction, we have made use of a number of simplifying assumptions. First of all, we have assumed that no information beyond past actions is transmitted to future motorists. This assumption can be relaxed to allow for imperfect public signals on past states as well as actions. Second, we have assumed that there is a single commitment type. This assumption can also be relaxed as long as the prior probability of all types is small.

There are a number of possible directions for extending the analysis of the model. On the more substantive side, the analysis of the current model can be extended to other economic settings such as repeated delegation problems where the principals must choose between rules and discretion for the agent. From a more theoretical point of view, it seems that a theory is needed to decide when reputation is good in terms of payoffs for the long run player as in the previous literature, and when it is bad as in the current model.

A Proof of Theorem

We begin with some notation. A history h is a sequence of elements of $\{t, e, \emptyset\}$, where \emptyset means that the mechanic was not hired in that period, and either t or e shows the action choice in a period in which the mechanic was hired. The notation (h, a) refers to the history following h in which the mechanic was hired and chose action a , and h^k refers to the k th element of h . Given an equilibrium, let $V_g(h)$ and $V_b(h)$ denote the equilibrium continuation value at history h for the good (hereafter g) and bad (b) types respectively. The posterior probability of the bad type after history h is denoted $p(h)$. The equilibrium behavior strategy of type φ is β_φ .

We begin with a crucial preliminary result:

Lemma 1 $V_g(h) > V_b(h)$ for every history h on the equilibrium path at which $p(h) > 0$ and the mechanic is hired with positive probability.

Let h be such a history. Let y be a positive integer. We will analyze the play over the next y stages following h . Let $V_g^y(h)$ be the equilibrium expected payoff for g over these next y periods. Let Σ_b^y be the set of all y -horizon continuation strategies played with positive probability by b beginning at h . Denote by $V_b^y(\sigma_b^y|h)$ the expected payoff to b over the next y periods from playing $\sigma_b^y \in \Sigma_b^y$.

Let A^y be the set of all y -length sequences from $\{e, t\}$. To each $\bar{a} \in A^y$, we can associate an element $\sigma_b^y(\bar{a})$ of Σ_b^y , defined as follows. For each continuation subhistory \tilde{h} of length $l(\tilde{h})$ less than or equal to y , if hired at \tilde{h} , $\sigma_b^y(\bar{a})$ plays $\bar{a}^{l(\tilde{h})}$ whenever $\beta_b(h, \tilde{h})(\bar{a}^{l(\tilde{h})}) > 0$ (otherwise it plays the unique action which b plays at history (h, \tilde{h}) .)

Let Θ^y be the set of all y -length sequences from $\{\theta_e, \theta_t\}$. Define σ_g^y to be the y -horizon strategy for g which chooses the appropriate repair in each period. Let $V_g^y(\sigma_g^y|h, \bar{\theta})$ be the expected payoff to g over the next y periods from playing σ_g^y conditional on $\bar{\theta}$ being the realized sequence of states. We will now prove

$$\mathbf{E}_{\bar{\theta}} V_g^y(\sigma_g^y|h, \bar{\theta}) \geq \min_{\sigma_b^y \in \Sigma_b^y} V_b^y(\sigma_b^y|h) \quad (3)$$

where $\mathbf{E}_{\bar{\theta}}$ denotes expectation with respect to the distribution over Θ^y .

To do so, we will compare the payoffs $V_g^y(\sigma_g^y|h, \bar{\theta})$ and $V_b^y(\sigma_b^y(\bar{a}(\bar{\theta}))|h)$ where $\bar{a}(\bar{\theta})$ is the sequence of actions that would be played by σ_g^y when the sequence of signals is $\bar{\theta}$. Consider any subhistory \tilde{h} of length $y - 1$. In the next period, b earns no more from playing according to $\sigma_b^y(\bar{a}(\bar{\theta}))$ than g earns conditional on $\bar{\theta}$ from playing according to σ_g^y . This is because by doing so, g earns u if he is hired, which is the maximum stage payoff. Now suppose that it is true of any subhistory of a given length $k < y$ that b earns no more over the next $y - k$ periods from playing according to $\sigma_b^y(\bar{a}(\bar{\theta}))$ than g earns conditional on $\bar{\theta}$ from playing according to σ_g^y .

Let \tilde{h} be a history of length $k - 1$. For the same reason as above, in the next period, b earns no more from playing according to $\sigma_b^y(\bar{a}(\bar{\theta}))$ than g earns conditional on $\bar{\theta}$ from playing according to σ_g^y . If these two strategies play the same action at \tilde{h} (i.e. $\sigma_b^y(\bar{a}(\bar{\theta}))(\tilde{h}) = [\bar{a}(\bar{\theta})]^k$), or if the motorist does not hire, then each strategy leads to the same successor history. By the induction hypothesis, we obtain the desired conclusion for \tilde{h} in this case. On the other hand, suppose $\sigma_b^y(\bar{a}(\bar{\theta}))(\tilde{h}) \neq [\bar{a}(\bar{\theta})]^k$. Then by

construction of $\sigma_b^y(\bar{a}(\bar{\theta}))$ it must be that in equilibrium, b plays $[\bar{a}(\bar{\theta})]^k$ with probability zero at history (h, \tilde{h}) . Thus, the posterior probability is zero that the mechanic is bad after the play of $[\bar{a}(\bar{\theta})]^k$, and so by renegotiation-proofness, the continuation payoff to g from playing σ_g^y conditional on $\bar{\theta}$ is the highest possible $\sum_{z=k}^y u\delta^z$. Again, the desired conclusion follows.

By induction we have that beginning with history h , b earns no more over the next y periods from playing according to $\sigma_b^y(\bar{a}(\bar{\theta}))$ than g earns conditional on $\bar{\theta}$ from playing according to σ_g^y . Since $\sigma_b^y(\bar{a}(\bar{\theta})) \in \Sigma_b^y$, this implies (

Now

$$V_g(h) = \lim_{y \rightarrow \infty} V_g^y(h) \geq \lim_{y \rightarrow \infty} V_g(\sigma_g^y|h) = \lim_{y \rightarrow \infty} \mathbf{E}_{\bar{\theta}} V_g^y(\sigma_g^y|h, \bar{\theta})$$

and

$$V_b(h) = \lim_{y \rightarrow \infty} \min_{\sigma_b^y \in \Sigma_b^y} V_b^y(\sigma_b^y|h)$$

which together with (

We can now use this to show that in fact $V_g(h) > V_b(h)$ for any history h on the equilibrium path such that $p(h) > 0$ and the mechanic is hired with positive probability. Indeed, let h be such a history. First suppose that the bad type of mechanic mixes. Then,

$$\begin{aligned} V_b(h) &= -w + \delta V_b(h, t) \\ &= u + \delta V_b(h, e) \end{aligned}$$

which together imply $V_b(h, t) > V_b(h, e)$. This allows us to bound $V_g(h)$:

$$\begin{aligned} V_g(h) &\geq u + \delta \left[\frac{V_g(h, t) + V_g(h, e)}{2} \right] \\ &\geq u + \delta \left[\frac{V_b(h, t) + V_b(h, e)}{2} \right] \\ &> u + \delta V_b(h, e) \\ &= V_b(h) \end{aligned}$$

If on the other hand, the bad type plays a pure action a , then

$$\begin{aligned} V_g(h) &\geq \frac{1}{2} \left(\frac{u}{1-\delta} \right) + \frac{1}{2} \left(\max \left\{ -w + \frac{u}{1-\delta}, u + \delta V_g(h, a) \right\} \right) \\ &\geq \frac{1}{2} \left(\frac{u}{1-\delta} + u + \delta V_b(h, a) \right) \end{aligned}$$

Now we claim that $V_b(h, a) < \frac{u}{1-\delta}$ and hence that $V_g(h) > u + \delta V_b(h, a) \geq V_b(h)$. For (h, a) is on the equilibrium path, and hence if $V_b(h, a) = \frac{u}{1-\delta}$, then beginning with history (h, a) , the mechanic must be hired for sure in every subsequent period even though type b is always playing e with probability one. But then the best-response of the the good type is to do the right repair in every period. This means that there is a finite number of consecutive plays of e after which the posterior exceeds $\frac{2u}{u+w}$ and the motorists will refuse to hire, a contradiction. ■

The proof is concluded analogously to the non-strategic case. Let us take p^* to be the supremum of all posteriors in which the mechanic is hired with positive probability on the equilibrium path. Because the bad mechanic does not know the necessary repair, $p^* < 1$. We will show that $p^* = 0$.

Given a prior p that the agent is bad, let $\Upsilon_a(\alpha_g, \alpha_b; p)$ denote the updated probability that the agent is bad after observing action a in some period in which the good type is playing e with total probability α_g and the bad type is playing e with total probability α_b . Note that Υ_a is increasing in p for $a \in \{t, e\}$, and Υ_e is decreasing in α_g , and increasing in α_b , while Υ_t has the opposite monotonicities. To ease notation, write $\Upsilon_a(h) = \Upsilon_a(\beta_g(h)(e), \beta_b(h)(e); p(h))$. Write also $\Upsilon_t^k(h)$ for the posterior probability of a bad agent given that t was observed in k consecutive periods following h .

Say that a history h is critical if there is an action $\bar{a}(h)$ such that $\Upsilon_{\bar{a}(h)}(h) > p^*$.

Lemma 2 *If $p^* > 0$ then there exists $\bar{p} < p^*$ such that $V_b(h) \leq u$ for all h on the equilibrium path such that $p(h) \in (\bar{p}, p^*]$.*

We introduce some additional notation. Recall that either there exist probabilities $0 < \underline{\alpha}_g < \bar{\alpha}_g < 1$ such that the mechanic will be hired in equilibrium at h only if $\underline{\alpha}_g \leq \beta_g(h)(e) \leq \bar{\alpha}_g$ or the bad mechanic is choosing both actions with positive probability at h . Let \underline{p} satisfy $\Upsilon_e(\bar{\alpha}_g, 1; \underline{p}) = p^*$, and define

$$\begin{aligned} f(p) &= \min_{\alpha_g, \alpha_b, a} \Upsilon_a(\alpha_g, \alpha_b; p) \\ \text{subject to } & \Upsilon_{a'}(\alpha_g, \alpha_b; p) \leq p^* \quad a' = e, t \\ & \underline{\alpha}_g \leq \alpha_g \leq \bar{\alpha}_g \end{aligned} \tag{4}$$

The function $f(\cdot)$ is continuous, increasing over $[0, p^*]$, and $f(p^*) = p^*$. Choose a sufficiently large integer K to satisfy

$$\frac{\delta^K}{1-\delta} u < u \tag{5}$$

By the continuity of f , we can find a \bar{p} less than, but close enough to p^* such that

$$f^K(p^0) > \underline{p} \text{ for all } p^0 \in (\bar{p}, p^*],$$

where $f^K(p^0)$ is the K -fold iteration of f . Let h be on the equilibrium path with $p(h) \in (\bar{p}, p^*]$. Let us classify continuation histories as follows. Let C denote the set of all continuation histories \hat{h} such that (h, \hat{h}) is critical. Let U be the set of all continuation histories \hat{h} satisfying

1. (h, \hat{h}) is on the equilibrium path
2. $l(\hat{h}) \leq K$
3. The mechanic is hired with positive probability at (h, \hat{h})
4. (h, \hat{h}) is not critical
5. There is no $k \leq l(K)$ such that (h, \hat{h}^k) is critical and $\hat{h}^k \neq \emptyset$.

We claim that $\beta_b(h, \hat{h})(\bar{a}(h, \hat{h})) > 0$ for $\hat{h} \in C$ and $\beta_b(h, \hat{h})(t) > 0$ for $\hat{h} \in U$. The first claim is immediate from the definitions. To prove the second, let $\hat{h} \in U$. For any subhistory \tilde{h} of \hat{h} , at which the mechanic was hired along \hat{h} , the history (h, \tilde{h}) is not critical (by

Since $l(\hat{h}) \leq K$, it follows that $p(h, \hat{h}) \geq f^K(p(h)) > \underline{p}$. Now suppose $\beta_b(h, \hat{h})(t) = 0$. Then $\Upsilon_e(h, \hat{h}) > \Upsilon_e(\beta_g(h, \hat{h}), 1; \underline{p})$ and since (h, \hat{h}) is on the path and the mechanic is hired with positive probability at (h, \hat{h}) , $\beta_g(h, \hat{h}) \leq \bar{\alpha}_g$, implying that $\Upsilon_e(h, \hat{h}) > p^*$. But this implies that (h, \hat{h}) is critical, a contradiction.

The claim implies that among the best-responses for b is a pure strategy which plays e at histories (h, \hat{h}) for $\hat{h} \in C$, and t at histories (h, \hat{h}) for $\hat{h} \in U$. We can thus conclude that $V_b(h)$ is bounded by the maximum payoff to any such pure strategy. This maximum is no greater than u . To see this, note that in the continuation, if the first critical history at which the mechanic is hired is reached within K periods, then this strategy earns a non-positive payoff in all periods prior to the critical history, then a payoff of u at the critical history and is never hired again. On the other hand, if no critical history is reached at which the mechanic is hired, then every reached continuation history of length less than or equal to K at which the mechanic was hired belongs to U . In this case, the strategy earns at most $\frac{\delta^K}{1-\delta}u$ which by equation (

Proof of Theorem

Since the mechanic is hired with positive probability at h , type g must be playing both e and t with positive probability. We claim that for δ close enough to 1, type b of mechanic must be playing t with positive probability. If b were playing e with probability 1, then $\bar{a}(h) = e$. A play of t leads to posterior 0, and by renegotiation-proofness, leads to gives payoff no less than $-w + \frac{\delta}{1-\delta}u$. But by playing e , the good type gets no more than u , which is strictly less for δ close enough to 1, a contradiction.

Since b is mixing, and since h is a critical history, we have $u = -w + \delta V_b(h, t)$. But by Lemma

B Proof of Theorem

The agent begins each period with a score τ . The set of possible scores is the set of non-negative integers, together with ∞ . The principal hires the agent when the score is $\tau \geq 1$, and does not hire the agent when the score is 0. Our interpretation of the score is as the level of reputation. All that remains to fully describe the principal's strategy is to specify a transition rule for the score, and the initial score. Starting with score $\tau > 1$, if the agent is hired (as is prescribed by the strategy) and the agent chooses action e , then the next period's score will be $\tau - 1$. In other words, performing an engine change reduces the reputational stock. (With the convention that $\infty - 1 = \infty$, we can thus interpret an agent with a score of ∞ as having been "promoted" since he will continue to be hired regardless of how often he plays e in the future.) On the other hand, if the agent plays action t , then he will begin the following period with a score that is determined by the outcome of a public randomization device. To define the probabilities, we introduce some additional notation. Define

$$V_b(\hat{\tau}) = u \sum_{k=0}^{\hat{\tau}-1} \delta^k$$

for all scores $\hat{\tau}$. For each τ , define (τ) to be any score which satisfies

$$u + \delta V_b(\tau - 1) < -w + \delta V_b((\tau)) \tag{6}$$

if such a score exists, otherwise set $(\tau) = \infty$. Let $\bar{\tau}$ be the greatest τ such that $(\tau) \neq \infty$. For each $1 \leq \tau \leq \bar{\tau}$, define $q(\tau)$ to satisfy

$$-w + \delta [q(\tau)V_b(\tau + 1) + (1 - q(\tau))V_b((\tau))] = u + \delta V_b(\tau - 1) \quad (7)$$

Note that $-w + \delta V_b(\tau + 1) < u + \delta V_b(\tau - 1)$, which together with equation (

The transition probabilities can now be defined. Suppose the agent plays action t when his score is τ . When $1 \leq \tau \leq \bar{\tau}$, the next period's score will be $\tau + 1$ with probability $q(\tau)$, and (τ) with the complementary probability. When $\tau > \bar{\tau}$, the next period's score will be ∞ with probability 1 (i.e. the agent is promoted.) Finally, in any period in which the agent is not hired, the next period's state will be the same as the previous. (Note that this implies that state 0 is an absorbing state in which the agent is never hired.) Since the transition among states depends only on the (public) history of the agent's actions and the outcome of the public randomization device, the current state is always common knowledge among the players.

We now describe the agent's strategy. The good type of agent plays the correct action whenever the score exceeds 1. With a score of 1, the good agent plays action t independent of his information. The bad agent plays action e whenever $\tau > \bar{\tau}$, and randomizes between t and e at any other score, playing e with probability $\alpha_b^* \in (0, 1)$. The precise value of α_b^* will be specified presently.

Given a prior p that the agent is bad, let $\Upsilon_a(\alpha_g, \alpha_b; p)$ denote the posterior probability that the agent is bad after observing action a in some period in which the good type is playing e with total probability α_g and the bad type is playing e with total probability α_b . Define $\Upsilon_a^2(\alpha_g, \alpha_b; p) = \Upsilon_a(\alpha_g, \alpha_b; \Upsilon_a(\alpha_g, \alpha_b; p))$ and recursively,

$$\Upsilon_a^n(\alpha_g, \alpha_b; p) = \Upsilon_a(\alpha_g, \alpha_b; \Upsilon_a^{n-1}(\alpha_g, \alpha_b; p)).$$

Let $V_P(\tau|\varphi)$ denote the principal's expected continuation value from the equilibrium strategies when the score is τ , conditional on the agent being type φ . Note that $V_P(\tau|g)$ depends only on the strategies of the principal and the good type of agent, both of which we have already specified. We will demonstrate below that for every score τ , $(1 - \delta)V_P(\tau|g)$ approaches u as the discount factor approaches one, and $(1 - \delta)V_P(\tau|b)$, which is negative,

approaches zero. Let δ be a discount factor such that $V_P(1|g) > 0$ (note that this restriction on the discount factor is independent of the prior), and let $p^* \in (0, 1)$ satisfy

$$(1 - p^*)V_P(1|g) - p^*V_P(1|b) = 0 \quad (8)$$

Note that p^* is independent of the prior and that p^* approaches 1 as $\delta \rightarrow 1$.

We set $\alpha_b^* < 1$ to satisfy

$$\Upsilon_e^{(\bar{\tau})}(1/2, 1; \Upsilon_t(0, \alpha_b^*; p^*)) < p^* \quad (9)$$

Since the left hand side is continuous in α_b^* , and is equal to zero when $\alpha_b^* = 1$, such an α_b^* exists. Note for future reference that α_b^* can be chosen arbitrarily close to 1.

Finally, the initial state τ^0 is defined to be the greatest integer τ satisfying $\Upsilon_e^{\tau-1}(1/2, 1; \mu) < p^*$, where μ is the prior probability of a bad agent. Note for future reference that τ^0 is independent of the discount factor and approaches ∞ as μ is allowed to go to zero.

The strategies have now been fully specified, we now demonstrate that these strategies form an equilibrium. Consider first optimality for the principal.

Lemma 3 *Suppose $\mu < p^*$. Consider any history in which the score at the beginning of period k is τ and p is the posterior probability of a bad agent.*

$$p \begin{cases} = 0 & \text{iff } \tau = \infty \\ \in (0, p^*) & \text{iff } 1 \leq \tau < \infty \\ = 1 & \text{iff } \tau = 0 \end{cases}$$

Proof: The first and third claims are obvious. We prove the second claim by two steps. First, consider a history where score $\tau < \infty$ has just been reached and the last play was t . Then, if the previous posterior was less than p^* , the new posterior will be less than $\Upsilon_t(0, \alpha_b^*; p^*) < p^*$ because Υ_t is increasing in its first and third arguments. Next, suppose $\tau > 0$ has been reached in period k by a play of e . Let y be the length of the most recent run of plays of e . Then in period $k - y - 1$, a score $\hat{\tau} < (\bar{\tau})$ was reached by a play of t , and hence the posterior was less than $\Upsilon_t(0, \alpha_b^*; p^*)$. That means that the posterior in period k is less than $\Upsilon_e^y(1/2, 1; \Upsilon_t(0, \alpha_b^*; p^*))$ since Υ_e is increasing

in its third argument. The latter is less than $\Upsilon_e^{(\bar{\tau})}(1/2, 1; \Upsilon_t(0, \alpha_b^*; p^*))$ since $\Upsilon_e(1/2, 1; p)$ is greater than p . By construction the latter is less than p^* .

It follows from these arguments that if the process starts at any score $0 < \tau < \infty$ with posterior less than p^* , then the posterior must remain below p^* forever. Thus, the second claim follows provided $\mu < p^*$ ■

It follows immediately from Lemma

To show optimality for the bad type of agent, we will show that V_b as defined above is the value function for the bad agent's strategy, and that it is the optimal value function. Obviously $V_b(0) = 0$ is the optimal value for the agent with score 0. Consider any score $\tau \geq 1$. With any such score, the bad agent is playing e with positive probability, hence we must show that $V_b(\tau)$ is achieved by playing e . This follows by definition: $V_b(\tau) = u + \delta V_b(\tau - 1)$ since the choice of e leads to successor score $\tau - 1$ for sure. With scores $1 \leq \tau \leq \bar{\tau}$, the bad agent is mixing, so we must show that in these states $V_b(\tau)$ is also achieved by action t . This is an immediate consequence of (

We have shown that V_b is the value function for the strategy of the bad agent. To show that it is the optimal value function, we will show that for any score τ in which a certain action is used with zero probability, that action yields a value no greater than $V_b(\tau)$. The only such scores and actions are scores $\tau > \bar{\tau}$ and action t . But for any $\tau > \bar{\tau}$, we have $-w + \delta V_b(\infty) \leq \sup_{\hat{\tau}} -w + \delta V_b(\hat{\tau}) \leq u + \delta V_b(\tau - 1)$ by the definition of $\bar{\tau}$. Since a choice of t would lead for sure to score ∞ , this establishes the claim.

The last step is to show optimality for the good type of agent. Denote by $V_g(\tau|\theta)$ the optimal continuation value for the good agent when he observes signal θ with score τ , and define $V_g(\tau) = \frac{1}{2}V_g(\tau|\theta_t) + V_g(\tau|\theta_e)$. Finally, let $\bar{V}_g(\tau) = q(\tau)V_g(\tau) + (1 - q(\tau))V_g(\tau - 1)$ be the expected continuation value after playing action t with score θ .

The first observation is that $V_g(\tau)$ is non-decreasing in τ , and hence playing action t upon observing signal θ_t is optimal at any score.

Next, for any score τ , $V_g(\tau) \geq V_b(\tau)$ (obviously with equality for $\tau = 0$). This is because $V_g(\tau)$ is bounded below by the payoff to the strategy which takes the appropriate action in every period. This strategy gives $\sum_{k=0}^{\tau-1} u\delta^k = V_b(\tau)$ for sure in the first τ periods, and the discounted continuation value thereafter. Since the discounted continuation value is non-negative, we

conclude that $V_g(\tau) \geq V_b(\tau)$. From this observation and (

Now consider a score $\tau > 1$ in which the good agent observes signal θ_e . Let W be the value to playing t and then continuing with an optimal strategy. We have

$$W = -w + \delta \bar{V}_g(\tau) \quad (10)$$

The optimal value is bounded by

$$V_g(\tau|\theta_e) \geq u, +\delta \left(\frac{u-w}{2} \right) + \delta^2 V_g(\tau) \quad (11)$$

because the right hand side is a lower bound for the payoff to the continuation strategy which plays e today and t after either signal tomorrow. Furthermore,

$$V_g(\tau) \geq \frac{1}{2} [W + u + \delta \bar{V}_g(\tau)] \quad (12)$$

because the right hand side is the payoff to playing t at score τ independent of the signal. Combining (

For each δ close enough to 1, we have constructed an equilibrium. In this equilibrium, the good agent is hired in every period with probability 1, and the bad agent is eventually terminated. To conclude the proof we calculate the equilibrium payoffs.

Proof of Theorem

To prove the second claim, let $\bar{\delta} < 1$ be the lower bound on the discount factor such that the strategies described above form an equilibrium. (Recall that this bound does not depend on the prior). Recall that the initial score, which we will here denote $\tau^0(\mu)$ to emphasize its dependence on μ , is independent of the discount factor and approaches infinity as the prior is taken to zero. Thus, for any discount factor greater than $\bar{\delta}$,

$$\lim_{\mu \rightarrow 0} (1 - \delta) V_P(\tau^0(\mu)|g) \rightarrow u$$

so that

$$\lim_{\mu \rightarrow 0} (1 - \delta) V_P(\tau^0(\mu)|\mu) = \lim_{\mu \rightarrow 0} (1 - \delta) \{ (1 - \mu) V_P(\tau^0(\mu)|g) + \mu V_P(\tau^0(\mu)|b) \} = u$$

since $(1 - \delta) V_P(\tau^0(\mu)|b) \geq -w$. ■

References

- Abreu, D., D. G. Pearce, and E. Stachetti (1993): “Renegotiation and Symmetry in Repeated Games,” *Journal of Economic Theory*, 60(2), 217–240.
- Aumann, R. J., and M. Maschler (1995): *Repeated Games with Incomplete Information*. MIT Press, Cambridge MA, with the collaboration of R. Stearns.
- Bénabou, R., and G. Laroque (1992): “Using Privileged Information to Manipulate Markets: Insiders, Gurus, and Credibility,” *Quarterly Journal of Economics*, 107, 921–958.
- Cripps, M., and J. Thomas (2001): “Some Asymptotic Results in Discounted Repeated Games with Incomplete Information,” <http://www.olin.wustl.edu/faculty/cripps/nash2001sent.pdf>.
- Drew Fudenberg, D. M. K., and E. Maskin (1990): “Repeated Games with Long-Run and Short-Run Players,” *Review of Economic Studies*, 57, 555–573.
- Farrell, J., and E. Maskin (1989): “Renegotiation in Repeated Games,” *Games and Economic Behavior*, 1, 327–360.
- Fudenberg, D., and D. K. Levine (1989): “Reputation and Equilibrium Selection in Games with a Single Patient Player,” *Econometrica*, 57, 251–68.
- (1992): “Maintaining a Reputation when Strategies are Imperfectly Observed,” *Review of Economics Studies*, 59, 561–579.
- (1994): “Efficiency and Observability with Long-Run and Short-Run Players,” *Journal of Economic Theory*.
- Hart, S. (1985): “Nonzero-Sum Two-Person Repeated Games with Incomplete Information,” *Mathematics of Operations Research*, 10, 117–153.
- Holmström, B. (1999): “Managerial Incentive Problems: A Dynamic Perspective,” *Review of Economics Studies*, 66, 169–182.

- Kreps, D. M., and R. Wilson (1982): "Reputation and Imperfect Information," *Journal of Economic Theory*, 27, 253–79.
- Mailath, G., and L. Samuelson (1998): "Your Reputation is Who You're Not, Not Who You'd Like to Be," U Penn CARESS Working Paper 98-11.
- Milgrom, P., and J. Roberts (1982): "Predation, Reputation, and Entry Deterrence," *Journal of Economic Theory*, 27, 280–94.
- Morris, S. (2001): "Political Correctness," *Journal of Political Economy*, 109(2), 231–265.
- Prendergast, C. (1993): "A Theory of "Yes Men"," *American Economic Review*, 83(4), 757–770.
- Prendergast, C., and L. Stole (1996): "Impetuous Youngsters and Jaded Old-Timers: Acquiring a Reputation for Learning," *Journal of Political Economy*, 104(6), 1105–1134.
- Sharfstein, D. S., and J. C. Stein (1990): "Herd Behavior and Investment," *The American Economic Review*, 80(3), 465–479.
- Sobel, J. (1985): "A Theory of Credibility," *Review of Economics Studies*.
- Sorin, S. (1999): "Merging, Reputation, and Repeated Games with Incomplete Information," *Games and Economic Behavior*, 29(1/2), 274–308.
- Zwiebel, J. (1995): "Corporate Conservatism and Relative Compensation," *Journal of Political Economy*.