Bad Reputation *

Jeffrey C. Ely[†] Juuso Välimäki[‡]

November 6, 2001

Abstract

We study the classical reputation model in which a long-run player faces a sequence of short-run players and would like to establish a reputation for playing a certain strategy. We present a model in which a long-run player's incentive to build a good reputation lowers the payoffs of all players in the game. This is in contrast with the standard reputation literature wherein reputation effects can only help the long-run player. The inefficiency in the present model can be thought of as arising from an information externality among the short-run players: short-run players do not have sufficient incentive to experiment and learn the long-run player's type. To demonstrate this, we also analyze the version of the model in which both players are long-lived, and show that efficiency is restored when the players are sufficiently patient.

1 Introduction

We study the classical reputation model in which a long-run player faces a sequence of short-run players and would like to establish a reputation for playing a certain strategy. We present a model in which a long-run player's incentive to build a good reputation lowers the payoffs of all players in the game. This is in contrast with the standard reputation literature wherein reputation effects can only help the long-run player. The inefficiency in the

*Gregory Pavlov provided invaluable research assistance.
†ely@nwu.edu Financial support from NSF grant #9985462 is gratefully acknowledged
‡juuso.valimaki@soton.ac.uk

present model can be thought of as arising from an information externality among the short-run players: short-run players do not have sufficient incentive to experiment and learn the long-run player's type. To demonstrate this, we also analyze the version of the model in which both players are long-lived, and show that efficiency is restored when the players are sufficiently patient.

In our model, the long-run player (agent) can provide a service to a sequence of short-run players (principals). The agent knows which of two services the principal requires, but the principal does not. Thus, the agent's service is valuable if and only if it is chosen contingent on the agent's private information. The agent is *good*: his preferences are perfectly aligned with the principal's, and hence in a one-shot version of the model there is a unique sequential equilibrium in which the agent chooses the appropriate service.

To study reputation effects, we extend the horizon and assume that the principals assign positive prior probability to a *bad* agent: one who cannot observe the necessary service, and has a preference for one service over the other, the bad service. Each principal in sequence observes the services performed for previous principals before deciding whether to hire the agent. We present a number of variations on this model. In each we find the following conclusion: the more patient the agent, and hence the stronger his reputational incentive, the lower the payoffs to all players.

The basic idea is simple: to maintain a reputation for being good, the agent should make sure that the empirical frequency of the services performed is close to the underlying probability distribution. This inevitably leads to histories in which the good agent must avoid the bad service even if it is is necessary in order to avoid loss of reputation. In particular, this would be necessary at histories in which the bad service has been required unusually frequently in the past and so the principals have become suspicious that the agent is bad. Because the principals can identify those histories after which this reputational incentive takes over, they refuse to hire at such histories. An induction argument now shows that in order to avoid reaching those histories, the good agent tries to bias his chosen service even earlier, the principal's forsee this, etc.

In the extreme versions of the model the above argument leads to a unique equilibrium in which the agent is never hired. More generally, we find that the payoffs in all Nash equilibria are uniformly bounded by a value that converges to the no-trade value as the principal becomes more and more patient. All of these results require only a positive probability, however small, of a bad agent.

It is apparent that the source of inefficiency is an information externality among the principals. At a history in which the good agent will try to separate from the bad agent by avoiding the bad service, a short-run principal will refuse to hire. The short-run principal does not internalize the benefit to future principals of the information generated about the agent's type. To illustrate this interpretation, we study in the last section a version of the model in which the principal is also a long-run player. 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 close 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 type before play begins.

2 Connections with the Literature

The standard reputation literature studies the effect of incomplete information on the payoffs to a long-run player in a repeated interaction with a sequence of short-run players. This literature began with the papers of Kreps and Wilson (1982) and Milgrom and Roberts (1982) which considered finite horizon models. The papers Fudenberg and Levine (1989) and Fudenberg and Levine (1992) provide a thorough analysis of the infinite horizon case.

Fudenberg and Levine (1989) analyzes games in which the long-run player's strategy is revealed at the end of each stage. They provide a lower bound for the payoffs to the long-run player in any Nash equilibrium of the game with incomplete information. This lower bound approaches the Stackelberg value as the discount factor approaches 1. The clear message of this paper is that when his strategy can be perfectly monitored, the long-run player can exploit reputation effects to obtain nearly his full-commitment payoff. Note that in our model, this payoff would correspond to the efficient outcome.

However, our model has imperfect monitoring (while the mechanic's action is perfectly observed in our model, his strategy is not), hence the Fudenberg and Levine (1989) bound does not apply. Nevertheless, our model in Section 3.3 does fit the framework of the Fudenberg and Levine (1992) paper which deals with imperfectly observed strategies. Fudenberg and Levine (1992) provide upper and lower bounds on the payoff to the long-run player and show that these bounds converge to the *generalized* Stackleberg values as the discount factor approaches 1. The generalized Stackleberg values roughly correspond to the best and worst *self-confirming* equilibria of the game in which the long-run player has commitment power. In other words, we imagine that the long-run player can choose any strategy and the short-run player must play a best-reply to some belief which is not dis-confirmed by the resulting path of play.

In our model, these bounds have no bite: the optimistic Stackleberg payoff is equal to the best possible payoff. The mechanic chooses the right repair, the motorist has the correct (hence self-confirmed) beliefs and optimally hires. The pessimistic Stackleberg payoff is equal to the worst possible payoff. Here the motorist does not hire because he believes that the mechanic will not choose the appropriate repair. Without hiring, the motorist cannot observe what the mechanic would have done and hence the motorist's belief cannot be dis-confirmed.

In the game we study, the Fudenberg and Levine (1992) upper bound can be tightened and in fact coincides with the lower bound in the limit as the discount factor approaches 1.

3 Section 1

Consider the following situation. A motorist's car has stopped working and is considering bringing it to a certain mechanic. The motorist knows that the car requires one of two repairs: an engine replacement or a tune-up. Each possibility is equally likely, but only the mechanic can determine which is the necessary repair.

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

	θ_t	θ_e
t	u	$-w$
e	$-w$	u

We will assume that $w > u > 0$. This implies that if the mechanic chooses the repair independently of what is required, the payoff to the motorist is negative. This will mean that the value to the motorist of hiring a mechanic who pursues such a strategy is less than the value of some outside option which we normalize to zero.[1] This is the essential strategic feature of the problem we are studying in this paper: the mechanic is an agent who possesses some expertise and (most importantly) expert information necessary for an optimal choice of the principal (the motorist in this example.)

Suppose that the mechanic is *good*, that is, his payoffs are identical to those of the motorist. Then the problem is trivial in the sense that the unique sequential equilibrium of this one-shot interaction is the first-best outcome in which the motorist hires the mechanic and the mechanic performs the appropriate repair.

In this paper we are going to analyze a model in which many motorists, each with broken-down cars, decide in sequence whether to hire the same mechanic. We will assume that the required repairs are drawn independently and that each motorist observes the repairs performed for all previous motorists before making the hiring decision.

In our basic model in which it is common knowledge that the mechanic is good, the problem is again trivial with any finite number of motorists. In the unique sequential equilibrium, the mechanic performs the correct repair for each motorist.

Things become interesting when we suppose that there is a positive probability μ that the mechanic is *bad*; that is, he always replaces the engine. Obviously if the mechanic was known to be bad, no motorist would ever hire him. This gives the good mechanic an incentive to avoid acquiring a reputation for being bad. We will show that this incentive has dramatic consequences.

Since for any finite number of motorists the efficient outcome is the unique sequential equilibrium, it is natural to assume that if the mechanic is ever revealed to be good, this continuation equilibrium will be played thereafter.

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

In the section 3.1, we show that in an infinite horizon model, in any Nash equilibrium satisfying Assumption 1, the mechanic is never hired. Things are not much better in the finite horizon version analyzed in section 3.2. We note

[1]For 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 this or some other mechanic.

that Assumption 1 is indeed an *assumption* in this model with incomplete information even with a finite number of motorists, as it is possible to construct equilibria in which the mechanic behaves differently because according to the equilibrium, a deviation by the mechanic to a zero probability action would lead future motorists to believe that the mechanic was bad and thus refuse to hire. Such equilibria seem implausible and in particular are vulnerable to renegotiation: after the mechanic has shown himself to be good he could propose a switch to the (Pareto superior) efficient continuation equilibrium and all motorists would agree. Nevertheless, in section 3.3, we characterize the best equilibrium within the set of *all* equilibria and show that the extreme inefficiency resulting from our assumption cannot be improved upon by much.

3.1 Infinite Horizon

The simplest and most extreme result is in the infinite horizon case. Suppose that there is an infinite sequence of agents and that the good mechanic discounts future payoffs by δ. Consider any Nash equilibrium in which Assumption 1 is satisfied. We begin by observing that there is a threshold posterior probability $p < 1$ such that no motorist who believes that the mechanic is bad with probability greater than or equal to p would hire the mechanic on the equilibrium path. (The payoff is strictly negative conditional on a bad mechanic). Let $p^* < 1$ be the supremum of all posteriors at which the mechanic is hired with positive probability on the equilibrium path. We claim that for δ close enough to 1, in fact $p^* = 0$ and hence that the mechanic will not be hired in any period.

Suppose the contrary, that $p^* > 0$. Let $0 < \tilde{p} < p^*$ satisfy

$$\Upsilon(\tilde{p}) := \frac{\tilde{p}}{\tilde{p} + (1 - \tilde{p})\frac{3u+w}{2(u+w)}} > p^* \tag{1}$$

The map Υ is continuous, strictly increasing, and $w > u$ implies that $\frac{3u+w}{2(u+w)} < 1$ and so $\Upsilon(\tilde{p}) > \tilde{p}$ for all $\tilde{p} < 1$. Thus we can find such a p. Since p^* is the supremum among posteriors in which the mechanic is hired, there exists a history on the equilibrium path at which the posterior p is in $(\tilde{p}, p^*)$ and the motorist hires with positive probability.

The motorist would only hire if the payoff is non-negative. This requires that the good mechanic replaces the engine with a total probability no greater than $\frac{3u+w}{2(u+w)}$. To see this, note that if the good mechanic replaces the engine with total probability greater than $\frac{3u+w}{2(u+w)}$, then the expected payoff to the motorist conditional on a good mechanic is no greater than

$$\left(\frac{1}{2} - \frac{3u+w}{2(u+w)}\right) w + \left(\frac{3}{2} - \frac{3u+w}{2(u+w)}\right) u = 0$$

and since the probability p of a bad mechanic is greater than zero, the unconditional expected payoff to the motorist would be strictly negative in this case.

Now by (1), the posterior following an engine replacement by the mechanic will be greater than p^*. This means that if the good mechanic chooses to replace the engine, he will not be hired by any subsequent motorist. On the other hand, suppose he chooses to perform a tune-up. Bayes' rule applies here because the good mechanic is performing a tune-up with positive probability (if not, the motorist would not have hired). So he will reveal himself to be good and, by Assumption 1, guarantee that he will be hired by *all* subsequent motorists. If the mechanic is sufficiently patient, then upon observing that an engine replacement is necessary, the good mechanic would optimally choose to instead perform a tune-up, incurring a payoff loss of $u + w$ today, in exchange for a gain of u in each future period.

Thus, at any history in which the posterior is in $(p, p^*]$, the good mechanic performs a tune-up independent of what is necessary. But this means that the expected payoff from hiring the mechanic at such a history is negative and that the mechanic cannot be hired in equilibrium. This contradiction proves that $p^* = 0$. We have shown

Proposition 1 *With an infinite sequence of motorists, and a prior probability $\mu > 0$ of a bad mechanic, the mechanic is never hired in any Nash equilibrium satisfying Assumption 1 when the discount factor is close enough to 1.*

We have assumed that motorists observe only the history of repairs and not the signal of the mechanic. Obviously if the mechanic's signal were perfectly observed the first-best could be achieved in equilibrium. On the other hand, if the motorists observe the signal of the mechanic with noise, then Proposition 1 should continue to hold.

3.2 Finite Horizon

For completeness, we present here the case of finitely many motorists. In this case, the good mechanic's long-run payoff is the sum of the payoffs in each period. We maintain the restriction to equilibria satisfying Assumption 1. In this model we have the following result.

Proposition 2 *There is an integer n, depending only on the prior, such that in any Nash equilibrium with any finite number of motorists, the mechanic can be hired in at most the last n periods.*

Proof: Set $p_0 = \frac{2u}{u+w}$ and define inductively p_k by

$$\Upsilon(p_k) = p_{k-1}$$

Observe that $p_k < p_{k-1}$ and $p_k \to 0$. We prove the claim by induction on k. Suppose $\mu \in (p_1, p_0]$. Let n_0 be the smallest integer greater than $\frac{u+w}{u}$. Suppose the mechanic is first hired in period τ. Then, as argued above, the good mechanic must be playing e with probability at most $\frac{3u+w}{2(u+w)}$ and thus the posterior following action e will exceed p_0. The expected payoff is negative to hiring the mechanic when the posterior is greater than $\frac{2u}{u+w}$, so the mechanic who plays e in period τ will not be hired again. By assumption, the mechanic who plays t will prove himself to be good and will be hired in all remaining periods. This implies that the number of remaining periods must be less than n_0, otherwise by playing t, the good mechanic will get at least $-w + n_0 u > u$ from playing t when the state is θ_e. Thus, the good mechanic would prefer to play t independent of the state, but this is a contradiction because the motorist would not hire in this case.

Now for the inductive step suppose that there is a number n_k satisfying the claim for all priors in $(p_k, p_{k-1}]$, and the prior belongs to $(p_{k+1}, p_k]$. Let τ be the first period in which the mechanic is hired. Arguing as above, the choice of e must lead to a posterior in $(p_k, 1]$ which means that the mechanic playing e in period τ will be hired a total of at most $\rho_k = \max_k n_k + 1$ times. The payoff to the bad mechanic from playing e is thus bounded above by $\rho_k u$. A choice of t on the other hand guarantees a payoff of at least $r_k u - w$ where r_k is the number of periods remaining. Because the mechanic is hired in period τ, the good mechanic must prefer to play e when the state is θ_e, and thus

$$r_k \leq \frac{\rho_k u + w}{u}$$

Since this bound does not depend on the total number of motorists, the result is proved. ■

3.3 Equilibria Without Assumption 1

Assumption 1, while intuitive, is nevertheless a restriction on the set of equilibria. There are sequential equilibria which do not satisfy the assumption which improve upon the extreme results of the preceding sections. For example, consider a strategy profile in which the mechanic is hired in the first k periods independent of history and never hired thereafter. This will be a sequential equilibrium for k small. But for any given prior, there will be an upper bound on the values of k for which this is an equilibrium, independent of the discount factor. Obviously, as the discount factor approaches 1, the average payoff from this equilibrium to the good type of mechanic approaches zero. We will show in this section that this class of equilibria cannot be improved upon by *any* Nash equilibrium. We return to the infinite horizon model, but now drop Assumption 1.

Theorem 1 *In the infinite horizon model, let $\bar{V}_g(\mu)$ be the supremum among all Nash equilibria for prior μ of the average payoff to the good type of mechanic. Then for all $\mu > 0$*

$$\lim_{\delta \to 1} \bar{V}_g(\mu) = 0$$

Proof: Define the sequence p_k as in the proof of Proposition 2. Recall that the mechanic cannot be hired on the path of any Nash equilibrium when the posterior exceeds p_0. Consider any history h on the path of some Nash equilibrium in which the posterior $p(h)$ following h is in $(p_k, p_{k-1}]$. Suppose that the mechanic is hired with positive probability at history h. Then as argued previously, the good type must play e with a total probability in $(0, \frac{3u+w}{2(u+w)}]$. This implies by the definition of p_k that $p(h, e) > p_{k-1}$, where the notation (h, e) refers to the successor history after a play of e.

Given a Nash equilibrium, for any history $\hat{h}$, let $V_g(\hat{h})$ denote the equilibrium discounted continuation value to the good type after history $\hat{h}$. When the mechanic is hired with positive probability at a history h on the equilibrium path, the good type must play e with positive probability when he observes signal θ_e. Thus,

$$u + \delta V_g(h, e) \geq -w + \delta V_g(h, t)$$

so that

$$V_g(h) \le u + \delta \max\{V_g(h,t), V_g(h,e)\} \tag{2}$$
$$\le 2u + w + \delta V_g(h,e) \tag{3}$$

Suppose $\mu \in (p_1, p_0]$. Obviously the initial history $\emptyset$ is on the path of any Nash equilibrium. If the mechanic is hired with positive probability in the first period, then we have $p(\emptyset, e) > p_0$, and hence $V_g(\emptyset, e) = 0$. By (3) we have $V_g(\emptyset) \le 2u + w$. This is an upper bound on the equilibrium value in any Nash equilibrium in which the mechanic is hired with positive probability in the first period. Since the value in any equilibrium is no greater than the value beginning in the first period in which the mechanic is hired with positive probability, and the posterior at that period must be the same as the prior, this upper bound in fact applies to all equilibria; i.e. $\bar{V}_g(\mu) \le (1-\delta)[2u + w]$ for all $\mu \in (p_1, p_0]$. Clearly

$$\lim_{\delta \to 1} \bar{V}_g(\mu) = 0 \quad \text{uniformly for all } \mu \in (p_1, p_0].$$

Now suppose the above claim holds for the interval $(p_k, p_0]$ for some $k \ge 1$, and let μ belong to $(p_{k+1}, p_k]$. Then in any Nash equilibrium in which the mechanic is hired with positive probability in the first period, (3) holds for $h = \emptyset$. Since in this case $(\emptyset, e)$ is on the equilibrium path (the good mechanic must play e with positive probability if the mechanic is hired), we have $V_g(\emptyset) \le 2u + w + \frac{\delta}{1-\delta}\bar{V}_g(p(\emptyset, e))$. We 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. Since $p(\emptyset, e) > p_k$, it follows that for every $\mu \in (p_{k+1}, p_k]$,

$$\bar{V}_g(\mu) \le (1-\delta)[2u + w] + \delta \sup_{p > p_k} \bar{V}_g(p)$$

and using the induction hypothesis, $\lim_{\delta \to 1} \bar{V}_g(\mu) = 0$, uniformly for each $\mu \in (p_{k+1}, p_k]$. The Theorem now follows by induction. ■

4 Strategic Bad Type

The arguments in the previous subsection make heavy use of the fact that a single play of t allows the good type to fully separate from the bad type. This gives the good type a powerful incentive to play t, undermining his usefulness to the current motorist. The results thus seem to depend on the assumption that the bad type is non-strategic, playing e in every period independent of history. In this section we show that the same conclusions hold when the bad type of mechanic is a strategic player.

Suppose 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 always prefers to replace the engine. The bad mechanic maximizes the sum of his stage payoffs in a finite horizon model, and the discounted sum of stage payoffs in an infinite horizon model. We continue to assume that the bad mechanic does not have the information necessary to make the correct repair, but now the bad mechanic has 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 will allow the mechanic to "survive" for at least one more period, but gives only a small improvement in reputation. This will mean 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.

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

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 kth 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 σ_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

$$\mathbf{E}_{\bar{\theta}} V_g^y(\sigma_g^y|h, \bar{\theta}) \geq \min_{\sigma_b^y \in \Sigma_b^y} V_b^y(\sigma_b^y|h) \tag{4}$$

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. $\sigma_b^y(\bar{a}(\bar{\theta}))(\tilde{h}) = \left[\bar{a}(\bar{\theta})\right]^k$), 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 $\sigma_b^y(\bar{a}(\bar{\theta}))(\tilde{h}) \neq \left[\bar{a}(\bar{\theta})\right]^k$. Then by construction of $\sigma_b^y(\bar{a}(\bar{\theta}))$ it must be that in equilibrium, b plays $\left[\bar{a}(\bar{\theta})\right]^k$ with probability zero at history $(h, \tilde{h})$. Thus, the posterior probability is zero that the mechanic is bad after the play of $\left[\bar{a}(\bar{\theta})\right]^k$, and so by Assumption 1, 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.

B y 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 (4).

Now

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

and

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

which together with (4) proves that $V_g(h) \geq V_b(h)$.

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,

$$V_b(h) = -w + \delta V_b(h,t) \tag{5}$$

$$= u + \delta V_b(h,e) \tag{6}$$

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

$$V_g(h) \geq u + \delta \left[\frac{V_g(h,t) + V_g(h,e)}{2}\right] \tag{7}$$

$$\geq u + \delta \left[\frac{V_b(h,t) + V_b(h,e)}{2}\right] \tag{8}$$

$$> u + \delta V_b(h,e) \tag{9}$$

$$= V_b(h) \tag{10}$$

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

$$V_g(h) \geq \frac{1}{2}\left(\frac{u}{1-\delta}\right) + \frac{1}{2}\left(\max\left\{-w + \frac{u}{1-\delta}, u + \delta V_g(h,a)\right\}\right) \tag{11}$$

$$\geq \frac{1}{2}\left(\frac{u}{1-\delta} + u + \delta V_b(h,a)\right) \tag{12}$$

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. 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 $\Upsilon_a(h) = \Upsilon_a(\beta_g(h)(e), \beta_b(h)(e); p(h))$.

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

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 there exist probabilities $0 < \underline{\alpha}_g < \bar{\alpha}_g < 1$ such that the mechanic will be hired in equilibrium at h only if $\underline{\alpha}_g \leq \beta_g(h)(e) \leq \bar{\alpha}_g$.

Let $\underline{p}$ satisfy $\Upsilon_e(\bar{\alpha}_g, 1; \underline{p}) = p^*$, and define

$$\begin{aligned} f(p) = \quad & \min_{\alpha_g, \alpha_b, a} \Upsilon_a(\alpha_g, \alpha_b; p) \\ \text{subject to} \quad & \Upsilon_z(\alpha_g, \alpha_b; p) \leq p^* \quad z = e, t \\ & \underline{\alpha}_g \leq \alpha_g \leq \bar{\alpha}_g \end{aligned} \tag{13}$$

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

$$\frac{\delta^K}{1-\delta} u < u \tag{14}$$

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

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

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 $\beta_b(h, \hat{h})(a_c(h, \hat{h})) > 0$ 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 5). That means that $\Upsilon_a(h, \tilde{h}) \leq p^*$ for $a = e, t$, i.e. the constraint in (13) is satisfied. $\Upsilon_a(h, \tilde{h}) \geq f(p(h, \tilde{h}))$ for $a = e, t$.

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 (14) is less than u. ■

Proof of Proposition 3 Suppose $p^* > 0$. Then by Lemma 2 there is a $\bar{p} < p^*$ such that $V_b(h) \leq u$ for all h on the equilibrium path such that $p(h) \in (\bar{p}, p^*]$. Let $p' < p^*$ satisfy $f(p') = \bar{p}$, where $f(\cdot)$ is defined in the proof of Lemma 2. By the definition of p^*, there exists a history h on the equilibrium path in which the mechanic is hired with positive probability such that $p(h) > p'$. We claim that h is a critical history. If not, then $\Upsilon_a(h) > f(p') = \bar{p}$ for both $a = t, e$. This implies that the continuation value for type b after choosing either action is no greater than u. Thus, a choice of t gives payoff no greater than $-w + \delta u$, strictly less than the payoff u guaranteed by a choice of e. So b must be playing e with probability one. This implies that $\Upsilon_e(h) \geq \Upsilon_e(\bar{\alpha}_g, 1; p') > p^*$ implying that h is a critical history after all, a contradiction.

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 $a_c(h) = e$. A play of t leads to posterior 0, and by Assumption 1, 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 1 $V_g(h, t) > V_b(h, t)$ so $u < -w + \delta V_g(h, t)$ so that the good type strictly prefers to play t at h, a contradiction. Thus, the mechanic cannot be hired in equilibrium at any history such that $p(h) > p'$. Since $p' < p^*$, this contradicts the definition of p^*, and thus $p^* = 0$. ■

5 Long-Run Principal

In this section 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 Θ 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 3. These are also the stage payoffs for the principal. A bad type of agent gets no signal and has payoffs given by figure 4. All discount future payoffs by the common discount factor δ.

We maintain the assumption that the principal receives no information about the past signals of the agent (and hence his past payoff realizations). 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 the principal and agent is governed by an employment contract under which can only depend on verifiable performance data, the situation would be equivalent to the one studied here.

In any case, the results of this section are positive. With discount factors close to 1, the first-best can be approximated even without monitoring on the basis of information about the agent's signal.

We will assume that the players have access to a public randomization device. We begin by describing the (sequential) equilibrium strategies that approximate the first-best. It is worth noting that the equilibrium constructed here satisfies Assumption 1, and in fact satisfy a strong form of renegotiation-proofness. For discount factors close to 1, the continuation equilibrium following every possible history (on and off the equilibrium path) approximates the first-best given the updated posterior at that history.

The principal's strategy is described as follows. The principal begins each period in some state ω. The set of possible states is the set of non-negative integers, together with ∞. The principal hires the agent in any state $\omega > 1$, and does not hire the agent in state 0. All that remains to fully describe the principal's strategy is to specify a transition rule among states, and the initial state. When in state $\omega > 1$, if the agent is hired (as is prescribed by the strategy) and the agent chooses action e, then the next period's state will be $\omega - 1$ (where the convention is that $\infty - 1 = \infty$). On the other hand, if the agent plays action t, then the next period's state will be determined by the outcome of a public randomization device. Define

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

for all states $\hat{\omega}$. For each ω, define $f(\omega)$ to be any state which satisfies

$$u + \delta V_b(\omega - 1) < -w + \delta V_b(f(\omega)) \tag{15}$$

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

$$-w + \delta\left[q(\omega)V_b(\omega + 1) + (1 - q(\omega))V_b(f(\omega))\right] = u + \delta V_b(\omega - 1) \tag{16}$$

Note that $-w + \delta V_b(\omega + 1) < u + \delta V_b(\omega - 1)$, which together with equation (15) implies that $q(\omega) \in (0, 1)$.

The transition probabilities can now be defined. Suppose the agent plays action t in state ω. When $1 \leq \omega \leq \bar{\omega}$, the next period's state will be $\omega+1$ with probability $q(\omega)$, and state $f(\omega)$ with the complementary probability. When $\omega > \bar{\omega}$, the next period's state will be ∞ with probability 1. 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 in every state except state 1. In state 1, the good agent plays action t independent of his information. The bad agent plays action e in every state $\omega > \bar{\omega}$, and randomizes between t and e in every other state, 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 $V_P(\omega|\varphi)$ denote the principal's expected continuation value from the equilibrium strategies when in state ω, conditional on the agent being type φ. Note that $V_P(\omega|g)$ 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 state ω, $(1 - \delta)V_P(\omega|g)$ approaches u as the discount factor approaches one, and $(1 - \delta)V_P(\omega|b)$, 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

$$(1-p^*)V_P(1|g) - p^* V_P(1|b) = 0 \tag{17}$$

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

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

$$\Upsilon_e^{f(\bar{\omega})}(1/2, 1; \Upsilon_t(0, \alpha_b^*; p^*)) < p^* \tag{18}$$

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 ω^0 is defined to be the greatest integer ω satisfying $\Upsilon_e^{\omega-1}(1/2, 1; \mu) < p^*$, where μ is the prior probability of a bad agent. Note for future reference that ω^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 state at the beginning of period τ is $\omega$and p is the posterior probability of a bad agent.*

$$p \begin{cases} = 0 & \textit{iff } \omega = \infty \\ \in (0, p^*) & \textit{iff } 1 \le \omega < \infty \\ = 1 & \textit{iff } \omega = 0 \end{cases}$$

Proof: The first and third claims are obvious. We prove the second claim by two steps. First, consider a history where state $\omega < \infty$ 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 $\omega > 0$ has been reached in period τ by a play of e. Let y be the length of the most recent run of plays of e. Then in period $\tau - y - 1$, a state $\hat{\omega} < f(\bar{\omega})$ 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 τ 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 $\Upsilon_e^{f(\bar{\omega})}(1/2, 1; \Upsilon_t(0, \alpha_b^*; p^*))$ 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 in any state $0 < \omega < \infty$ 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 3 that the principal optimally hires in state ∞, and does not hire in state 0. Consider any period in which the state is $1 \leq \omega < \infty$. By Lemma 3, the posterior must be less than p^*. Conditional on a good agent, the principal's value $V_P(\omega|g)$ is increasing in the state. Thus, letting $V_P(\omega|p)$ designate the principal's optimal continuation value in state ω, with posterior p, we have $V_P(\omega|p) \geq (1-p^*)V_P(\omega|g) - p^*V_P(\omega|b) \geq (1-p^*)V_P(1|g) - p^*V_P(1|b) = 0$ by (17). This implies $V_P(\omega|p) \geq \delta V_P(\omega|p)$. Since the latter is the principal's value from a one-stage deviation, we conclude that the principal's strategy is sequentially rational in state ω.

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 in state 0. Consider any state $\omega \geq 1$. In any such state, the bad agent is playing e with positive probability, hence we must show that $V_b(\omega)$ is achieved by playing e. This follows by definition: $V_b(\omega) = u + \delta V_b(\omega - 1)$ since the choice of e leads to successor state $\omega - 1$ for sure. In states $1 \leq \omega \leq \bar{\omega}$, the bad agent is mixing, so we must show that in these states $V_b(\omega)$ is also achieved by action t. This is an immediate consequence of (16).

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 in any state ω in which a certain action is used with zero probability, that action yields a value no greater than $V_b(\omega)$. The only such states and actions are states $\omega > \bar{\omega}$ and action t. But for any $\omega > \bar{\omega}$, we have $-w + \delta V_b(\infty) \leq \sup_{\hat{\omega}} -w + \delta V_b(\hat{\omega}) \leq u + \delta V_b(\omega - 1)$ by the definition of $\bar{\omega}$. Since a choice of t would lead for sure to successor state ∞, this establishes the claim.

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

The first observation is that $V_g(\omega)$ is non-decreasing in ω, and hence playing action t upon observing signal θ_t is optimal in every state.

Next, for any state ω, $V_g(\omega) \geq V_b(\omega)$ (obviously with equality for $\omega = 0$). This is because $V_g(\omega)$ is bounded below by the payoff to the strategy which takes the appropriate action in every period. This strategy gives $\sum_{k=0}^{\omega-1} u\delta^k = V_b(\omega)$ for sure in the first ω periods, and the discounted continuation value thereafter. Since the discounted continuation value is non-negative, we conclude that $V_g(\omega) \geq V_b(\omega)$. From this observation and (16), we have

$$-w + \delta \bar{V}_g(\omega) \geq u$$

and we conclude that playing action t is optimal in state $\omega = 1$ after signal θ_e.

Now consider a state $\omega > 1$ 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

$$W = -w + \delta \bar{V}_g(\omega) \tag{19}$$

The optimal value is bounded by

$$V_g(\omega|\theta_e) \geq u, +\delta \left(\frac{u-w}{2} \right) + \delta^2 V_g(\omega) \tag{20}$$

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,

$$V_g(\omega) \geq \frac{1}{2} \left[W + u + \delta \bar{V}_g(\omega) \right] \tag{21}$$

because the right hand side is the payoff to playing t in state ω independent of the signal. Combining (19), (20), and (21), we obtain

$$V_g(\omega|\theta_e) \geq u + \delta \left(\frac{u-w}{2} \right) + \delta^2 \left(W + \frac{u+w}{2} \right)$$

which shows that $V_g(\omega|\theta_e) > W$ for δ close enough to 1. Note for future reference that this restriction on the discount factor is independent of the prior.

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. It now remains to calculate the equilibrium payoffs. We are particularly interested in the payoffs when either the discount factor is close to one, or the probability of a bad agent is small.

Theorem 2 *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.*

Proof: For the first claim, fix the prior μ. The principal's average equilibrium value is

$$V_P(\omega^0|\mu) = (1-\mu)V_P(\omega^0|g) + \mu V_P(\omega^0|b)$$

Recall that the mixing probabilities α_b^* can be chosen independent of the discount factor and arbitrarily close to 1. This means, first, that for any z as close to 1 as desired, there is an equilibrium in which the probability bad agent reaches state 0 in period ω^0 is greater than z. This in turn implies $\lim_{\delta\to 1} V_P(\omega^0|b) = 0$. Consider now the stochastic process on states conditional on the agent being good. In equilibrium, whenever $\omega > 1$ conditional on the good agent, the successor state is either $\omega - 1$ or some state greater than ω, with equal probability. Since $V_P(\hat{\omega}|g)$ is increasing in $\hat{\omega}$, the equilibrium value $V_P(\omega^0|g)$ is greater than or equal to the value from the process in which the successor state is $\omega+1$ or $\omega-1$ with equal probability. But such a process is payoff equivalent to a symmetric random walk on the integers in which the flow payoff at state 0 is $\frac{u-w}{2}$, and the flow payoff at any other state is u. In a symmetric random walk, the mean time to return to the origin is infinite, hence the average value of such a process converges to u as the discount factor goes to 1. We have shown that $\lim_{\delta\to 1}(1-\delta)V_P(\omega^0|g) = u$, and this concludes the proof of the first claim.

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 state, which we will here denote $\omega^0(\mu)$ 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 $\bar{\delta}$,

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

so that

$$\lim_{\mu\to 0}(1-\delta)V_P(\omega^0(\mu)|\mu) = \lim_{\mu\to 0}(1-\delta)\left\{(1-\mu)V_P(\omega^0(\mu)|g) + \mu V_P(\omega^0(\mu)|b)\right\} = u$$

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

References

Fudenberg, D., and D. K. Levine (1989): "Reputation and Equilibrium Selection in Games with a Single Patient Player," *Econometrica*, 57.

——— (1992): "Maintaining a Reputation when Strategies are Imperfectly Observed," *Review of Economics Studies*, 59, 561–579.

Kreps, D. M., and R. Wilson (1982): "Reputation and Imperfect Information," *Journal of Economic Theory*, 27, 253–79.

Milgrom, P., and J. Roberts (1982): "Predation, Reputation, and Entry Deterrence," *Journal of Economic Theory*, 27, 280–94.