Tulane Economics Working Paper Series

Bad Reputation under Bounded and Fading Memory

Benjamin Sperisen
Department of Economics
Tulane University
bsperise@tulane.edu

Working Paper 1527
October 2015

Abstract

I relax the full memory assumption in Ely and Välimäki's (2003) mechanic game, where reputation is bad for all players. First I consider "bounded memory," where only finitely many recent periods are observed. For long memory, reputation is still bad. Shortening memory avoids bad reputation but only by making it "useless." There is no "happy middle:" reputation is either useless or reduces equilibrium payoffs for any memory length. I find a qualitatively different result for "fading memory," where players randomly sample past periods with probabilities "fading" toward zero. Unlike bounded memory, reputation is not bad but remains useful under sufficiently fast fading. This result extends to a more general class of both good and bad reputation games, suggesting reputation leaves long-run player behavior unaffected in some realistic word-of-mouth environments.

Keywords: reputation; bounded memory; fading memory; mechanic game
JEL codes: C73, D8

Bad Reputation under Bounded and Fading Memory

Benjamin Sperisen*

October 19, 2015

Abstract

I relax the full memory assumption in Ely and Välimäki's (2003) mechanic game, where reputation is bad for all players. First I consider "bounded memory," where only finitely many recent periods are observed. For long memory, reputation is still bad. Shortening memory avoids bad reputation but only by making it "useless." There is no "happy middle:" reputation is either useless or reduces equilibrium payoffs for any memory length. I find a qualitatively different result for "fading memory," where players randomly sample past periods with probabilities "fading" toward zero. Unlike bounded memory, reputation is not bad but remains useful under sufficiently fast fading. This result extends to a more general class of both good and bad reputation games, suggesting reputation leaves long-run player behavior unaffected in some realistic word-of-mouth environments.

JEL Classification: C73; D8
Keywords: Reputation; Bounded memory; Fading memory; Mechanic game

1 Introduction

The reputation literature has shown that even very small uncertainty about a player's type can have dramatic effects on equilibrium behavior and payoffs. Building on the seminal work of Kreps and Wilson (1982) and Milgrom and Roberts (1982), Fudenberg and Levine (1989; 1992) show that in a general class of games, introducing such uncertainty assures the long-run player of a payoff arbitrarily close to the payoff achieved by a credible commitment to an action of her choice. Ely and Välimäki (2003) (henceforth EV) construct a model where reputation has a similarly dramatic but negative effect on payoffs of all players, a phenomenon they call "bad reputation."

These models typically assume that short-run players see the full history of past outcomes. In reality, agents often perceive reputation through only limited excerpts of the past, raising the question: is it possible to avoid bad reputation by limiting how much of the past is seen?

*Department of Economics, Tulane University. Email address: bsperise@tulane.edu. I am especially grateful to Thomas Wiseman for his extensive advice and guidance on this paper, as well as to Stefano Barbieri, V. Bhaskar, Florian Kuhn, Dale Stahl, Maxwell Stinchcombe, Caroline Thomas, and numerous seminar participants for their helpful comments.

The focus of this paper is answering that question using two different relaxations of the full memory assumption: "bounded memory" (modeling a public list of recent reviews) and "fading memory" (modeling word-of-mouth). For the bounded memory case, I find that EV's bad reputation result holds when memory is sufficiently long (but still bounded). Shortening the memory sufficiently avoids the bad reputation result by allowing the repeated one-shot equilibrium, but in doing so it makes reputation "useless:" short-run players do not learn enough from the history to make decisions that improve their payoff. I find that there does not exist a "happy middle" where reputation is neither bad nor useless.

By contrast, fading memory can achieve this happy middle. When memory "fades" sufficiently quickly, bad reputation is avoided. But the history is still sometimes useful to the short-run players, and their expected payoffs are strictly greater than for the one-shot game. When fading is sufficiently slow — agents talk to each other very frequently — the bad reputation result returns.

EV's model, the mechanic game, has the feature that the long-run player is clearly harmed by reputation, as are the short-run players. EV point out these dynamics could be a concern in a number of asymmetric information settings involving expert sellers, such as auto mechanics, lawyers, management consultants, and medical doctors. These markets generally involve consumers who are not perfectly informed about the seller's past. Consider motorists who solicit information about an auto mechanic through word-of-mouth, patients who choose to see a dentist after reading the first few reviews listed on Yelp, or a consultant who provides prospective clients with a list of only her most recent references on her resumé. Behavior in the mechanic game within settings where consumers have a limited view of the past may shed light on both positive questions, such as when and why experts are hired in real markets with this feature, and normative issues, such as the optimal design and welfare effects of review websites.

In the mechanic game, the "good" mechanic (rational long-run player) and the motorist (short-run player) have coinciding interests in the stage game: the motorist's car has a problem, and both want the problem fixed correctly. Motorists do not know which repairs their cars need (either a cheap tune-up or an expensive engine replacement), but the mechanic does. The motorist would like to hire the good mechanic instead of an outside option, if she does the right repair. However, the introduction of even a tiny probability that the mechanic is a commitment type (the "bad" mechanic, who performs an expensive engine replacement no matter what problem the car has) impedes the ability of the good mechanic and motorists to cooperate.

When motorists can see the entire history of hiring decisions and repairs, a history with sufficiently many engine replacements and no tune-ups yields a belief that the mechanic is likely enough to be bad that not hiring is necessarily preferable. All subsequent motorists avoid the mechanic, "freezing" the bad belief and preventing the mechanic from ever being hired again. At a "critical" history, where the mechanic is just one engine replacement away from a frozen bad belief, a sufficiently patient good mechanic is inevitably tempted into performing a tune-up, even when an engine replacement is needed, to signal that she is good. Such signaling behavior is harmful to the motorist receiving the unnecessary tune-up, whose best response is to not hire. Before the critical history, the mechanic's anticipation of the critical motorist's decision not to hire must lead to a certain (possibly unnecessary) tune-up even earlier, so this previous motorist also does not hire, and so on by backward induction, leading to a complete unraveling of the market and no hiring on the equilibrium path of all renegotiation-proof Nash equilibria.

However, real mechanics likely do not expect that any particular action is certain to be observed by every subsequent customer. Even if a mechanic knew with certainty that a customer would immediately report her actions to the world in a review on Yelp, she knows that many future potential customers may not see the review, either because they do not check Yelp at all, or because with time, the review is eventually pushed out of sight on the first page. Word-of-mouth seems particularly unlikely to yield fully informed customers due to its decentralized, random nature.

The online review example motivates the bounded memory model, where each motorist sees a finite number of periods into the past, but no further.[1] I find that when memory is long enough, exactly the same equilibrium behavior as the full memory model is obtained because the events in the beginning periods "echo" through the participation decisions of the motorists. If the mechanic signals that she is good with a tune-up, all motorists who see that first tune-up hire her, while the next "generation" of motorists, who do not see that tune-up but see all the hiring that followed, infer her type and also hire even if they see only engine replacements, and so on forever. By contrast, performing only engine replacements early on eventually leads to the mechanic not being hired, and subsequent motorists can infer that the mechanic never performed a tune-up, even in periods they no longer observe, leading to a critical history and thus bad reputation.

Making the memory too short for an individual motorist to learn much about the mechanic allows an equilibrium that avoids the bad reputation result, with the one-shot equilibrium played repeatedly. But reputation is also rendered "useless:" it does not help players because it is too uninformative to have any effect on behavior at all. In this case, customers would not bother paying any cost (such as the time to open a webpage) to view the mechanic's history, as it never provides valuable information that would change their decision — they are just as happy going into the decision "blind."[2]

Under bounded memory, is there a way to provide valuable information to motorists without doing more harm than good? I show that the answer is "no." Any equilibrium for any memory length either gives the motorists (as well as the mechanic) an expected payoff lower than that of the one-shot game, or motorists have the same expected payoff and are no better off paying attention to the history than ignoring it. Providing useful information to future customers unavoidably tempts the mechanic into harmful signaling, and so reputation is either bad or useless under bounded memory.

The second type of history observation I consider is "fading memory," where motorists see the

[1] In addition to the standard assumption of stationary strategies, I also use a robustness refinement (purifiability in the sense of Harsanyi, 1973) to rule out unrealistically fragile equilibria. This refinement requires equilibria be robust to tiny, private perturbations in payoffs, motivated by the fact that models are only approximate, rather than exact, descriptions of reality. Section 3.2 has a detailed discussion.

[2] This also raises the question of why, for example, the owner of an online review site would expend resources providing such memory to customers.

last period with probability $\lambda \in (0, 1)$, the second-to-last period with probability λ^2, and so on for all past periods. This can be thought of as modeling the decentralized randomness of word-of-mouth. Imagine a town where residents leave randomly with some probability each period (and new residents take their place), swapping stories about the mechanic with those who happen to be in town at the time but not with those who have left; a motorist asking around about the mechanic will more likely hear about the most recent experiences, while ancient history is nearly forgotten.[3]

Like long bounded memory, fading memory yields the bad reputation result for high λ. Unlike bounded memory, when λ is small enough, reputational incentives are too weak to cause bad reputation, but reputation is still useful. This result seems reasonably realistic: the good mechanic is not diverted from serving customers by extremely strong reputation incentives, good and bad mechanics are both sometimes hired, and some of the more discerning customers avoid the bad type.

Why does fading memory avoid the "bad or useless" result for bounded memory? Bounded memory allows an unbroken chain of observations from the first customer through all subsequent customers: what happens in period 0 is seen in period 1, what happens in period 1 is seen in period 2, and so on. Therefore, what happens in any period t can affect what happens at any future period even after period t is forgotten. When the memory is long enough, this "echo" is exactly what happens, causing bad reputation. By contrast, the gradual, random forgetting of fading memory bounds the probabilities of such observation chains, thereby limiting reputational incentives.

In fact, this result for fading memory with low λ applies to a more general class of reputation games: when the long-run player has a strictly dominant action in the stage game, the long-run player behaves myopically and always chooses that dominant action (as though reputation did not exist), while the payoffs of short-run players are often greater than under the static game because they are sometimes well-informed. This class includes many good reputation games, such as the chain store game of Selten (1978), variants of which have been widely used to study reputation.[4]

The assumption of "low λ" does not mean trivially small: for a patient mechanic, it corresponds to a given motorist talking to an average of $\frac{1}{2}$ future potential customers about their experience, high enough to cover scenarios with significant (but not totally ubiquitous) word-of-mouth communication.[5] The result is also robust to any correlation between observations, allowing applications to more centralized communication like online forums where public messages "fade" over time. This suggests reputation may be too weak to affect long-run player behavior in many real word-of-mouth situations.

Though fading memory is intended as a model of word-of-mouth communication, it differs

[3]Note that this model means a resident telling one future customer is correlated with the same resident telling another customer. All the fading memory results are robust to such correlation as long as the departures of residents are independent, and the result for low λ is robust to any correlation between observations.

[4]See, for example, Pitchik (1993), Aoyagi (1996), and Wiseman (2008).

[5]Section B.2 of the Appendix gives a higher upper bound that allows talking to an average of $\frac{2}{3}$ future motorists, given some reasonable restrictions on equilibria, which can be even higher depending on the parameterization of the stage game.

importantly from existing work on word-of-mouth (e.g. Ellison and Fudenberg, 1995 and Banerjee and Fudenberg, 2004) where players randomly sample some fixed number of past events. Key differences include that under fading memory, the "sample size" is random, and that players are more likely to observe the recent past than distant past. The fading of past events is critical for ruling out the never-ending echo that occurs in the bounded memory case.

This paper also relates to other (full memory) extensions of the mechanic game. EV show that when the motorist is also a long-run player, an equilibrium exists where the mechanic and motorist are able to interact. Mailath and Samuelson (2006) consider the possibility of random "captive consumers," who hire no matter the history. Ely, Fudenberg, and Levine (2008) extend bad reputation to a broader class of games, illustrating the difference between bad and good reputation. Though this paper only considers bounded and fading memory as applied to the original EV mechanic game, there is no apparent reason why similar results would not apply to such generalizations.

Bounded memory repeated games without reputation have been the topic of a number of papers (for example, see Sabourian, 1998; Mailath and Olszewski, 2011; and Bhaskar, Mailath, and Morris, 2013) but such games with reputation are relatively unexplored. Liu and Skrzypacz (2014) study a bounded memory product choice game with reputation, finding cyclical behavior, the "riding of reputation bubbles." Liu (2011) also finds cyclical equilibria in an environment where short-run players incur a cost to observe limited records of past long-run player actions. These papers show reputation being continually accumulated, exploited, and then replenished.

A key assumption behind these cyclical results is that only the long-run player's actions are observed, which Liu and Skrzypacz call "limited records." This assumption makes models more tractable, and Liu and Skrzypacz point out that it is realistic in applications like the Better Business Bureau, which does not show how much business a firm gets but does show complaints. Importantly, this gives the long-run player the power to "clean" her history regardless of what the short-run players do, and therefore unilaterally control her reputation.

In many other settings, however, the long-run player needs the cooperation of short-run players to send the signals she wants. For example, a consultant or lawyer needs clients to hire her in order to provide references to future clients; she cannot generate a "high quality" signal through sheer effort alone. Instead, future prospective clients observing that she was not hired may interpret that as a bad signal, leading to persistent unemployment. This lack of total control over one's reputation is critical to the bad reputation effect: the mechanic *needs* the cooperation of motorists to establish a reputation. While Liu and Skrzypacz find cyclical behavior that is qualitatively different from both the one-shot and full memory cases, this paper instead finds *non*-cyclical behavior because the reputation gets "stuck." To what extent this applies to other bounded memory reputation games remains an interesting open question.

The rest of the paper is organized as follows. Section 2 defines EV's mechanic game and summarizes their result. Section 3 presents the bounded memory model and the corresponding results. Section 4 does the same for fading memory. Section 5 concludes. Omitted proofs are relegated to the Appendix.

2 The Mechanic Game

In the mechanic game, reputation leads to a lower payoff for both the long-run and short-run players than in the static game. A long-lived car mechanic faces a different short-lived motorist each period. Each motorist's car is in one of two states, each requiring a different repair: either a cheap tune-up c or an expensive engine replacement e. The states are drawn iid each with probability $\frac{1}{2}$. The motorist does not know which repair is needed, but the mechanic does.

In each period, the motorist first chooses to either hire the mechanic or choose an outside option N with payoff zero. If hired, the mechanic observes the state of the car, either θ_c or θ_e. The motorist benefits if the mechanic performs the correct repair, receiving payoffs according to the following table:

	θ_c	θ_e
c	u	$-w$
e	$-w$	u
N	0	0

Assume $w > u > 0$. This insures that if the mechanic chooses the repair independent of the state, the motorist will prefer the outside option.

The mechanic can be one of two types: good (g) and bad (b). The mechanic's type is denoted s. The good mechanic has the same stage payoff as the motorist (in the table above), and wants to maximize her expected discounted average payoff, discounted at rate $\delta \in (0, 1)$. The bad mechanic is non-strategic and simply performs engine replacements, regardless of the state.[6] Motorists observe the full history of repair and hiring decisions, but not the previous motorists' states (i.e., it is not known whether the repairs were correct), as public knowledge. Beginning at period 0, the first motorist has prior belief μ^0 that the mechanic is bad, and subsequent motorists update their beliefs about the mechanic's type according to Bayes' rule.

In the one-shot game, the motorist's expected payoff for hiring is simply $\frac{1}{2}\mu(u - w) + (1 - \mu)u$ where μ is the belief that the mechanic is bad (doing the right repair is strictly dominant for the good mechanic). She will hire the mechanic only if this expected payoff is nonnegative, which is clearly false when the belief μ is greater than critical value p^*, where p^* is defined so that $p^* \frac{u-w}{2} + (1 - p^*)u = 0$.

EV prove that the supremum of the mechanic's Nash equilibrium payoffs must converge to zero for δ close enough to one, so that equilibria where the mechanic is hired must have the mechanic hired only infrequently. They point out that equilibria with such infrequent hiring have an implausible feature: once the mechanic performs a single tune-up, she reveals herself to be good (with certainty) to all future motorists. After a tune-up, it makes sense that all subsequent motorists (knowing the mechanic is good) will want to hire, and the mechanic will want to perform

[6]EV also show their result holds for a strategic bad mechanic who receives u for performing an engine replacement , $-w$ for a tune-up, and 0 when not hired.

correct repairs for them. For this reason, EV use the following renegotiation-proofness assumption to rule out such dubious behavior.

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

EV then find the following dramatic result.

Proposition 2.1. *Let $\mu^0 > 0$ be given. When δ is close enough to one, any Nash equilibrium satisfying Assumption 1 has a unique equilibrium outcome where the mechanic is never hired.*

Without reproducing the proof here, the intuition behind it is that in any equilibrium, if the mechanic performs some number L engine replacements and no tune-ups, the motorists' beliefs must rise above p^* and they do not hire. If a motorist hires at any history, the mechanic must perform an engine replacement with sufficient probability (otherwise the motorist's expected payoff from hiring would be negative), and this means that, with positive probability, the mechanic performs L consecutive engine replacements on the equilibrium path. After performing $L - 1$ engine replacements (and no tune-ups), an engine replacement gives a continuation payoff of 0, compared with a continuation payoff of u for a tune-up (since all future motorists hire). When she is sufficiently patient, the mechanic always performs a tune-up, so she cannot be hired at a "critical" history (i.e., after $L - 1$ engine replacements without tune-ups); backwards induction leads to the result of no hiring on the equilibrium path.

3 Bounded Memory

The disaster of Proposition 2.1 is driven by the fact that if the mechanic were ever hired, with positive probability she would arrive at a critical history where providing a needed engine replacement results in her never being hired again. This is because the evidence leading to a belief greater than p^* (i.e., the mechanic is too likely to be bad to hire) is never forgotten, and because the mechanic is not hired, she never has another chance to change beliefs, and so a sufficiently patient mechanic cannot resist the temptation to show she is good at such a history.

Can this disastrous outcome be avoided if the damning evidence is eventually forgotten? This section considers the case where memory is bounded: motorists view only the finite T most recent periods.

I make the following two assumptions. First, I follow the example of Rosenthal (1979) and Liu (2011) in assuming that the short-run players (motorists) do not know when the game starts, and have steady-state beliefs induced by the play of the game, meaning they update beliefs according to Bayes' rule as if they have an improper uniform prior over the period at which they enter.

Second, I restrict attention to equilibria which are "purifiable" in the sense of Harsanyi (1973). This robustness refinement rules out unrealistically fragile equilibria that may exist in such bounded memory environments, as shown in a number of papers.[7] A purifiable equilibrium is one that

[7]See Bhaskar (1998), Bhaskar, Mailath, and Morris (2008), and Bhaskar, Mailath, and Morris (2013) for examples.

survives the addition of tiny, independent private shocks to player payoffs. The motivation for this refinement is that, as games are only approximations of reality, some private payoff information always exists, and equilibria that do not survive arbitrarily close approximations of payoffs are therefore unrealistic. Section 3.2 gives the formal definition.

3.1 Preliminaries

I introduce some useful notation. A full history $\mathfrak{h} = (\mathfrak{h}_0, \mathfrak{h}_1, ..., \mathfrak{h}_t)$ is a t-length sequence of events, where h_k is the outcome at period k, either a repair (c or e) or a no-hiring decision (N). Let $H \equiv \{c, e, N\}$. The game begins in period 0, and so motorists arriving at period $t \in \{0, ..., T-1\}$ observe a history $h \in H^t$, which is the same as the full history. Motorists arriving in all periods $t' \geq T$ observe a bounded history $h \in H^T$, rather than the full history $\mathfrak{h} \in H^{t'}$. Since bounded histories are dealt with more often than full ones, I often refer to bounded histories simply as "histories." The set of all (bounded) histories is $\mathcal{H} \equiv \bigcup_{k=0}^{T} H^k$, and the set of full histories is $\bar{\mathcal{H}} \equiv \bigcup_{k=0}^{\infty} H^k$.

The strategies for the motorists can be represented as a function $\rho : \mathcal{H} \to [0, 1]$, where $\rho(h)$ gives the probability that a motorist hires at history h. The mechanic's strategy $\sigma \equiv (\sigma_c, \sigma_e)$ is given by two functions $\sigma_x : \bar{\mathcal{H}} \to [0, 1]$ where $x \in \{c, e\}$ denotes a needed repair (for the hiring motorist) and $\sigma_x(\mathfrak{h})$ gives the probability the mechanic performs the needed repair x at full history $\mathfrak{h}$. However, Section 3.2 shows that the purifiability requirement implies the mechanic plays a stationary strategy dependent only on the history seen by the motorist, i.e. we can write $\sigma_x : \mathcal{H} \to [0, 1]$. I denote the beliefs of motorists about the mechanic with the function $\mu : \mathcal{H} \to [0, 1]$, where $\mu(h)$ is the posterior belief that the mechanic is the bad type at history h.

For bounded histories, I use negative subscript $-k$ to indicate the event k periods ago, i.e. $h = (h_{-T}, ..., h_{-2}, h_{-1})$. It will also be convenient to use notation appending and prepending outcomes onto histories. Let a bounded history $h \in H^T$ be given. Let hx denote the bounded history consisting of the $T-1$ most recent outcomes of h followed by $x \in H$, i.e. $hx = (h_{-(T-1)}, ..., h_{-1}, x)$. Let xh denote the bounded history consisting of the x followed by the $T-1$ oldest outcomes of h,i.e. $xh = (x, h_{-T}, ..., h_{-2})$. I also use superscripts to denote the repetition of an outcome in a history. For example, e^T is the history with all T periods containing an engine replacement, and Ne^{T-1} is the history with an engine replacement T periods ago followed by all engine replacements.

If a motorist at history h hires a mechanic, the expected payoff of hiring must be nonnegative (otherwise the motorist chooses the outside option):

$$\mu(h)\left(\frac{u-w}{2}\right) + (1-\mu(h))(\beta(h)u - (1-\beta(h))w) \geq 0,$$

where $\beta(h) \equiv \frac{1}{2}\sigma_c(h) + \frac{1}{2}\sigma_e(h)$ is the probability that the good mechanic performs the correct repair at h. Solving for $\beta(h)$ gives

$$\beta(h) \geq \frac{1}{u+w}\left[w + \frac{\mu(h)}{1-\mu(h)}\left(\frac{w-u}{2}\right)\right] \geq \frac{w}{u+w}. \tag{3.1}$$

Since $\beta(h) = \frac{1}{2}\sigma_c(h) + \frac{1}{2}\sigma_e(h) \leq \frac{1}{2} + \frac{1}{2}\sigma_c(h)$, I can substitute (3.1) to get the lower bound

$$\sigma_c(h) \geq \frac{w-u}{u+w} \equiv \beta^* \tag{3.2}$$

on the probability that a mechanic performs a needed tune-up.

For periods $t < T-1$ when motorists still observe the full history, the posterior evolves according to Bayes' rule in the same way as the full memory model. Suppose the belief at history $h \in H^t$ is $\mu(h)$. If the motorist does not hire, then $\mu(hN) = \mu(h)$; if the mechanic performs a tune-up at h, then $\mu(hc) = 0$; if the mechanic performs an engine replacement, then

$$\mu(he) = \frac{\mu(h)\rho(h)}{\mu(h)\rho(h) + (1-\mu(h))\rho(h)[\frac{1}{2}\sigma_c(h) + \frac{1}{2}\sigma_e(h)]}.$$

Define

$$\Upsilon(\mu) \equiv \frac{\mu}{\mu + (1-\mu)[\frac{1}{2} + \frac{1}{2}(1-\beta^*)]} \tag{3.3}$$

so that $\Upsilon(\mu(h))$ is a lower bound for $\mu(he)$ (again, assuming $h \in H^t$ for $t < T-1$), and inductively define $\Upsilon^1(\mu) \equiv \Upsilon(\mu)$ and $\Upsilon^{k+1}(\mu) \equiv \Upsilon(\Upsilon^k(\mu))$, so that $\Upsilon^k(\mu)$ is a lower bound for the posterior after observing k engine replacements and no tune-ups at $t \leq T-1$. Define

$$L(\mu^0) \equiv \min k \text{ such that } \Upsilon^k(\mu^0) > p^* \tag{3.4}$$

as an upper bound on the number of engine replacements that can be performed (without any tune-ups) in the first T periods before the posterior exceeds p^*. For $T > L(\mu^0) + 1$, it is straightforward to see that the mechanic cannot perform more than L engine replacements without any tune-ups within the first T periods on the equilibrium path. Motorists before period T who arrive after the Lth engine replacement will believe the mechanic is so likely to be bad that the payoff of hiring must be negative, preventing any hiring until at least period T.

Finally, I define terms similar to (3.3) and (3.4) for the case when the mechanic always performs the correct repair, and so performs an engine replacement with probability $\frac{1}{2}$ upon being hired. Given prior μ, the motorist's posterior after observing an engine replacement is

$$\bar{\Upsilon}(\mu) \equiv \frac{\mu}{\mu + \frac{1}{2}(1-\mu)} = \frac{2\mu}{1+\mu},$$

with $\bar{\Upsilon}^t(\mu)$ defined inductively like $\Upsilon(\cdot)$. Define

$$\bar{L}(\mu^0) \equiv \min t \text{ such that } \bar{\Upsilon}^t(\mu^0) > p^*.$$

Note that $\bar{L}(\mu^0) \leq L(\mu^0)$ because $L(\mu^0)$ is a lower bound that presumes the good mechanic performs engine replacements with maximum probability $\frac{1}{2} + \frac{1}{2}(1-\beta^*) > \frac{1}{2}$.

3.2 Purifiability

I follow the example of Bhaskar, Mailath, and Morris (2013) in constructing a sequence of perturbed games, which consist of the mechanic game with the addition of independent private shocks to each action for each player. Let Z be a compact subset of $\mathbb{R}^2$ and write $\Delta^*(Z)$ for the set of measures with support Z with strictly positive densities. At each full history $\mathfrak{h}$, the motorist draws payoff shocks $(z_r^Y, z_r^N) \in Z$ according to distribution $\psi_r \in \Delta^*(Z)$ and the mechanic draws payoff shocks $(z_g^c, z_g^e) \in Z$ from a distribution $\psi_g \in \Delta^*(Z)$. Denote $\psi \equiv (\psi_r, \psi_g)$. Let $\varepsilon > 0$ be some positive constant. The motorist's payoff to hiring is augmented by z_r^Y and the payoff to the outside option is augmented by z_r^N; the mechanic's payoff to doing a tune-up is augmented by z_g^Y and the payoff to an engine replacement is augmented by z_g^N. For example, when hired by a motorist needing an engine replacement, the mechanic receives payoff $u + \varepsilon z_g^e$ for performing an engine replacement and payoff $-w + \varepsilon z_g^c$ for performing a tune-up. Denote the perturbed game by $\Gamma(\varepsilon, \psi)$.

Strategies are mappings from both the public history (observed in the unperturbed game) and the private history of shocks. For the motorist, this consists of a single shock as they are only present for one period, so the strategies of all motorists can be represented as a function $\tilde{\rho} : H \times Z \to \{Y, N\}$. A mechanic strategy at period t in the perturbed game maps from a full history of t outcomes $\{N, c, e\}$, a private history of $t+1$ private shocks (including the current period's shocks), and $t+1$ motorist states, to a repair:

$$\tilde{\sigma} : \bigcup_{t=0}^{\infty} (H^t \times Z^{t+1} \times \Theta^{t+1}) \to \{c, e\}.$$

However, note that any outcomes before period $t - T$, as well as any private shocks and motorist states from all but the current period, are payoff irrelevant. This observation yields the following result, which states that any best response by the mechanic must ignore the private shocks and needed repairs from previous periods for almost all realizations of the current shock. The proof is a straightforward adaptation of Lemma 2 in Bhaskar, Mailath, and Morris (2013) and so is omitted.

Lemma 3.1. *For the perturbed game, let motorist strategies $\tilde{\rho}$ be given. For any any $k \in \{T, T+1, ...\}$, let any $h^{k-T} \in H^{k-T}, h^T \in H^T$; $\bar{z}, \bar{z}' \in Z^k$; $\bar{\theta}, \bar{\theta}' \in \Theta^k$ be given. If a mechanic strategy $\tilde{\sigma}$ is a best response to $\tilde{\rho}$, then $\tilde{\sigma}(h^{k-T}h^T, (\bar{z}, z_k), (\bar{\theta}, \theta_k)) = \tilde{\sigma}(h^{k-T}h^T, (\bar{z}', z_k), (\bar{\theta}', \theta_k))$ for any $\theta_k \in \{\theta_c, \theta_e\}$ and for almost all $z_k \in Z$.*

What this shows that is that in any equilibrium of the perturbed game, the mechanic's strategy is "essentially stationary:" at two periods $t, t' \geq T$, the mechanic plays identically conditional on having the same public history $h \in T$ with probability one. Hence, the probability that the mechanic performs a needed repair x if hired at public history $h \in \mathcal{H}$ and current shock $z \in Z$ is $\int \sigma_x(h, z)\, d\psi_g(z)$.

Turning back to the unperturbed game, an equilibrium of the perturbed game is weakly purifiable if there exists a sequence of equilibria in a sequence of perturbed games that converge in outcomes to the unperturbed equilibrium (as given in Definition 6 of Bhaskar, Mailath, and Morris, 2013).[8]

Definition 3.1. A strategy profile (ρ, σ) of the unperturbed game is *weakly purifiable* if there exists a sequence $(\psi^k, \varepsilon^k)_{k=1}^{\infty}$, with $\psi \equiv (\psi_r^k, \psi_g^k) \in (\Delta^*(Z))^2$ and $\lim_{k\to\infty} \varepsilon^k = 0$, such that there exists a sequence of strategy profiles $(\tilde{\rho}^k, \tilde{\sigma}^k)_{k=1}^{\infty}$ such that

$$\lim_{k\to\infty} \int \tilde{\rho}^k(h, z)\, d\psi_r^k(z) = \rho(h), \qquad \lim_{k\to\infty} \int \tilde{\sigma}_x^k(h, z)\, d\psi_g^k(z) = \sigma_x(h)$$

for all $h \in \mathcal{H}, x \in \{\theta_c, \theta_e\}$, and $(\tilde{\rho}^k, \tilde{\sigma}^k)$ are part of a sequential equilibrium of the perturbed game $\Gamma(\psi^k, \varepsilon^k)$.

Since Lemma 3.1 shows that any equilibrium in the perturbed game is "essentially stationary," purifiability requires that the mechanic play a stationary strategy as a best response to the motorists' stationary strategies.

3.3 Results

Can bounding the memory avoid the bad reputation outcome? The first result shows that if memory length is above a threshold (independent of the discount factor), bad reputation persists.

Proposition 3.1. *Let $\mu^0 > 0$ and $L(\mu^0)$ be given. If*

$$T > (L + 1) \left\lfloor \frac{u + w}{u} \right\rfloor, \tag{3.5}$$

then for any weakly purifiable sequential equilibrium satisfying Assumption 1 there is a unique equilibrium outcome where the mechanic is never hired when δ is close enough to one.

Why does bad reputation persist for a patient mechanic despite the fact that the "bad evidence" is erased? The first part of the proof shows that purifiability requires the mechanic always be hired at the history full of engine replacements (e^T). Purifiability means the mechanic cannot condition her strategy on payoff-irrelevant history (like the outcome T periods ago about to be erased), so the only way a "robust" equilibrium can have motorists hire at ce^{T-1} but not hire at e^T is if they have a sufficiently bad belief at e^T. Since not being hired then yields history $e^{T-1}N$, the bad belief at e^T requires that the bad mechanic eventually returns to history e^T sufficiently frequently (otherwise, motorists would have a good belief at e^T). The only way for the bad mechanic to reach e^T from $e^{T-1}N$ is through the history Ne^{T-1}. Yet for long memory, the good mechanic rarely reaches Ne^{T-1} because they are too likely to perform a tune-up in the meantime. Thus, the belief at Ne^{T-1} is so high that motorists do not hire, so the bad mechanic never returns to e^T, a contradiction.

[8] Bhaskar, Mailath, and Morris call this "weakly purifiable" as it only requires the existence of *some* sequence of perturbed games with equilibria converging to the equilibrium in question, in contrast to the stronger notion (that they call "Harsanyi purifiable") that for *any* sequence of perturbed games, there must exist such a corresponding sequence of equilibria that converge.

The second part of the proof deals with the case where the mechanic is always hired at e^T, which means performing a single tune-up ensures the mechanic is hired forever. Thus, a motorist knows whether the mechanic ever performed a tune-up since it will "echo" through the motorists' hiring decisions. If the mechanic is ever hired, she must be hired again soon even if she does an engine replacement (or else the history would be critical). Yet, a motorist who sees a history where the mechanic has L engine replacements, some no-hiring events (N), and no tune-ups can infer that a tune-up has never occurred in the full history, and so has a belief greater than the critical value p^*. Thus, the mechanic is not hired again at least until one of these engine replacements is erased. When memory is sufficiently long, this means the mechanic is not hired for a long time, and so the mechanic cannot resist the temptation to perform an incorrect tune-up to avoid this fate, yielding the bad reputation outcome.

What if memory is not so long? With a shorter bound on the memory, can we limit how high the motorist beliefs may rise and therefore avoid such a critical history? Indeed we can.

Proposition 3.2. *Let $\mu^0 > 0$ and $T < \bar{L}(\mu^0)$ be given. A sequential equilibrium exists for any δ where the good mechanic always performs the correct repair and the motorists always hire.*

Proof. It is easy to show that always hiring is a best response for motorists. At any period t, the motorists' posterior is less than or equal to $\bar{\Upsilon}^T(\mu^0) \leq p^*$, so the payoff of hiring is nonnegative. For the mechanic, the continuation payoff of a tune-up is equal to that of an engine replacement (she is always hired, no matter her strategy), so performing the correct repair always yields a greater payoff. ■

Proposition 3.2 avoids the disaster of Proposition 3.1, but it does so by preventing reputation from having *any* effect. Reputation neither tempts the mechanic into a costly tune-up that harms the motorist, nor does it help the motorists sort out a good mechanic from the bad, and the ex ante payoff for every motorist is equal to that of the one-shot game. Reputation is only useful to motorists if it sometimes gives them a posterior greater than p^*, so that they can avoid hiring the bad type; the myopic equilibrium avoids bad reputation precisely by ruling out that possibility.

Is there some memory length that avoids the bad reputation result but is also useful in the sense that motorists can achieve expected payoffs greater than those without any memory (i.e. the one-shot game), particularly between the lower bound of Proposition 3.1 and the upper bound of Proposition 3.2? The following result shows that such a "happy medium" is impossible.

Proposition 3.3. *Let $\mu^0 > 0$ and $T > 0$ be given. Let $\bar{v}^*(\delta)$ be the supremum of the expected motorist payoff for the set of weakly purifiable sequential equilibria satisfying Assumption 1, and let $\hat{v}^0$ be the maximum expected motorist payoff in a sequential equilibrium of the one-shot game.*[9] *For δ close enough to one,*

$$\bar{v}^*(\delta) \leq \hat{v}^0. \tag{3.6}$$

[9] Note that for there is a unique equilibrium of the one-shot game for all priors except p^*.

Proposition 3.3 shows that equilibria fall into one of the following two categories. Some equilibria yield lower expected payoffs for the motorists than under the one-shot game, as in Proposition 3.1. Reputation may be useful in the sense that it is individually beneficial to learn from the history, but the memory is doing a disservice to motorists — the market would work better without it. In the other category of equilibria, motorists have the same expected payoffs as under the one-shot game. Thus, the presence of reputation does not harm motorists, but it does not help them either. If accessing the history was at all costly to the motorists, they would prefer to hire the mechanic blindly. Hence, reputation becomes "useless."

The proof of Proposition 3.3 begins by showing that purifiability requires the mechanic always be hired in a "useful equilibrium" at the history e^T (all engine replacements), using a method similar to that used for Proposition 3.1. Since the mechanic is always hired at e^T, performing a tune-up ensures the mechanic is hired forever. A good mechanic who is hired always eventually ends up in the set of histories without any no-hire "N" events, and so, roughly speaking, histories containing Ns must have bad beliefs (if the mechanic is hired in equilibrium). Thus, the event of not being hired once "echoes," condemning the mechanic to never being hired again. If the equilibrium is useful, the memory must be informative enough that the motorist expects to not hire with positive probability. This implies the existence of a critical history where the mechanic cannot resist the temptation of an incorrect tune-up.

4 Fading Memory

The results of Section 3 show that bounded memory can only avoid bad reputation at the expense of making reputation "useless" — specifically, bounded memory cannot improve the motorists' expected payoff beyond that of the one-shot game. The information structure introduced in this section, called "fading memory," achieves the happy medium that bounded memory cannot.

Fading memory can be interpreted as roughly reflecting how word-of-mouth spreads. It is not certain that customers hear about previous experiences, nor is it certain that they do not, but it is more likely that they hear about recent history than the distant past. For example, this might model how information is spread in a community (perhaps a town or group of friends) whose members randomly leave and arrive each period. If experiences with the mechanic are shared among the current members, but potential customers do not talk to those who have left, the probability that they hear about an outcome decays exponentially with time.

Fading memory is defined as follows. Motorists observe each previous period with some positive probability. Let $p_t^{t'}$ denote the probability that the motorist in period t' observes the actions in period $t < t'$. By comparison, under full memory this probability is always one; under bounded memory, $p_t^{t'} = 1$ for $t' \in \{t+1, ..., t+T\}$ and $p_t^{t'} = 0$ for $t' > t+T$. Under fading memory, this probability is never one nor zero, instead starting relatively high right after the event and exponentially "fading" toward zero: $p_t^{t'} = \lambda^{t'-t}$ for $\lambda \in (0,1)$. I also allow motorists to observe the calendar date.[10] For the purpose of comparing the expected motorist payoff with the bounded memory case, I assume the motorists have the improper uniform prior before finding out the period in which they arrive.[11] It is assumed that observations are independent across periods-observed, but not necessarily across observers. More formally, let periods t, t', t'' be given such that $t < t' < t''$. The probability that motorist t'' observes both periods t and t' is assumed to be equal to $p_t^{t''} p_{t'}^{t''}$, but I allow the possibility that the probability that period t is observed by both motorists t' and t'' is *not* equal to $p_t^{t'} p_t^{t''}$. (The community model described above is consistent with these assumptions so long as the departures of community members are independent.)[12]

Low λ yields myopic behavior by the mechanic, thereby avoiding bad reputation. This result extends to a more general class of reputation games, defined in Section 4.1. This class includes many "good reputation" games like Selten's (1978) chain store game (see Section 4.3), as well as other participation and simultaneous-move games.[13]

4.1 Preliminaries

I define a class of reputation games where the long-run player has a strictly dominant action in the sense defined below. Let an extensive-form game between players 1 (the long-run player) and 2 (the short-run player) be given with finite action spaces A_1 and A_2, with player 2 moving first and where some subset of actions $\tilde{A}_2 \subset A_2$ lead to a decision node for player 1, moving second. Player 1 may be at any information set about player 2's action within the set $\tilde{A}_2$. (The special case of a simultaneous-move game is given by $\tilde{A}_2 = A_2$ and player 1's information set includes the entire set $\tilde{A}_2$.) The long-run player may observe a private signal $\theta \in \Theta$ drawn iid in each period, which may affect both players' payoffs. The payoff for player i of action profile $(a_1, a_2) \in A \equiv A_1 \times A_2$ is denoted $u_i(a_1, a_2, \theta)$. Since the stage game is extensive-form, strict dominance is needed at decision nodes for the long-run player; I define this notion as "strictly conditionally dominant" (in a simultaneous-move game, this notion is equivalent to a strictly dominant action).[14]

Definition 4.1. An action $a_d \in A_1$ for player 1 is *strictly conditionally dominant* if and only if $u_1(a_d, a_2, \theta) > u_1(a'_1, a_2, \theta)$ for all $a'_1 \in A_1 \backslash \{a_d\}$, all $a_2 \in \tilde{A}_2$, and all $\theta \in \Theta$.

I assume such a strictly conditionally dominant action $a_d \in A_1$ exists. Define $H \equiv (A_2 \backslash \tilde{A}_2) \cup (\tilde{A}_2 \times A_1)$ as the set of terminal nodes in the stage game; the long-run player observes all such

[10] This actually makes the exposition simpler, but the results also carry over to making the calendar date unobservable. Recall also that Section 3 restricted attention to (weakly) purifiable equilibria. The results in this section do not require the purifiability assumption, but it is also clear that the equilibria described by the results are purifiable for almost all priors $\mu^0 \in (0, 1]$ because they involve pure strategies with unique best responses.

[11] This comparison is made in Corollary 4.1.

[12] Proposition 4.1 is robust to any correlation between observations: see Section 4.4.

[13] Indeed, the argument behind the proof of Proposition 4.1 holds even without the presence of commitment types. Of course, the absence of commitment types also means the absence of useful reputation (or any reputation for that matter), which is the motivation for this result.

[14] The definition used here is equivalent to a unique strategy that is not "conditionally dominated" under Definition 4.2 in Fudenberg and Tirole (1991), who consider iterated conditional dominance as a solution concept for extensive-form games.

outcomes for previous periods.[15] If a short-run player does not observe a period, they see the signal U ("unobserved") in its place. Define $\tilde{H} \equiv H \cup \{U\}$. Short-run player strategies are given by a function

$$\rho : \bigcup_{t=0}^{\infty} \tilde{H}^t \to \Delta A_2.$$

Denote the set of short-run player histories as $\tilde{\mathcal{H}} \equiv \bigcup_{t=0}^{\infty} \tilde{H}^t$. Similarly to Section 3.1, I refer to elements $\tilde{h} \in \tilde{\mathcal{H}}$ simply as "histories," as opposed to "full histories." To avoid more complicated notation, I assume that the long-run player does not observe the histories observed by short-run players (i.e., which periods are hidden from them), though the results also hold without this assumption. The long-run player's strategy is a function

$$\sigma : \bigcup_{t=0}^{\infty} (H^t \times \Theta^t) \times A_2 \times \Theta \to \Delta A_1.$$

I refer to elements of $\mathcal{H} \equiv \bigcup_{t=0}^{\infty} H^t$ as "full histories." Player 1 is either a rational type ξ_0, with stage payoffs given by $u_1(a_2, a_1, \theta)$ and discount factor δ, or one of $N \in \mathbb{N}$ commitment types $\xi_1, ..., \xi_N$, where type j always plays some action $\alpha^j \in \Delta A_1$, and short-run players have prior beliefs $\mu^0(\xi)$ on each of these types. Define

$$z \equiv \max_{\theta \in \Theta} \left\{ \max_{a \in A} u_1(a, \theta) - \min_{a \in A} u_1(a, \theta) \right\} \tag{4.1}$$

$$z_d \equiv \min_{\theta \in \Theta} \min_{(a_1', a_2) \in (A_1 \setminus \{a_d\}) \times \tilde{A}_2} \{u_1(a_d, a_2, \theta) - u_1(a_1', a_2, \theta)\}. \tag{4.2}$$

z gives the maximum difference in stage payoffs between two action profiles, while z_d gives the minimum current period cost of choosing a different action over the strictly conditionally dominant action a_d.

4.2 Results

I first show that the mechanic plays a (unique) myopic strategy in all equilibria under fading memory when λ is below a threshold. For almost all priors $\mu^0 \in (0, 1]$, this also implies a unique equilibrium.[16] The type of equilibrium described by Proposition 4.1 is similar to that of Proposition 3.2 in that the mechanic always does the correct repair, but what differs is that the mechanic is *not* always hired, even when good, because motorists sometimes reach "useful" histories that yield beliefs greater than p^*.

[15] Proposition 4.1 is straightforward (if notationally cumbersome) to extend to imperfect monitoring, where each terminal node generates a signal drawn from a probability distribution.

[16] The mechanic always performs the correct repair in any equilibrium. For almost all priors, there exist no histories where the belief is exactly p^*, and so customers play a pure strategy of always hiring at histories where the belief is less than p^* and never hiring at histories with a belief greater than p^*. (At the Lebesgue measure zero set of priors where equilibria do have histories with exactly belief p^*, any mixed actions at such histories could be part of an equilibrium.)

This result is given for a more general set of games defined in Section 4.1, namely those where the long-run player has a strictly dominant action in the stage game. It is also stronger than Proposition 3.2 because it shows that *every* equilibrium has myopic behavior by the long-run player.

Proposition 4.1. *Consider the game defined in Section 4.1, where the rational player 1 has action* $a_d \in A_1$ *which is strictly conditionally dominant. If*

$$\lambda < \frac{z_d}{\delta(z + 2z_d)} \tag{4.3}$$

where z *and* z_d *are defined as in (4.1) and (4.2), then any sequential equilibrium has rational player 1 playing* a_d *at every history.*

One striking feature is that the upper bound on λ does not depend on δ or the prior beliefs. This is because the proof does not rely on calculating beliefs. The action at some period t affects the payoff at some later period $t' > t$ only if either t' observes t directly or there exists some sequence $(t_1, ..., t_n)$ such that motorist t_1 observes t, motorist t_j observes t_{j-1} for all $j \in \{2, ..., n\}$, and motorist t' observes t_n. The proof bounds the probability of such "observation chains." Of course, the motorists in these chains would also have to change their action in response to their observation to affect the payoff at t' (so the bound is not as high as it could be), but ignoring the actual beliefs allows using essentially the same technique for fading memory in other games. This is impossible in the bounded memory environment because such a chain always connects every period together even for the shortest memory case $T = 1$: period 1 observes period 0, period 2 observes period 1, and so on. In the case of the mechanic game, $z = z_d = u + w$, so the upper bound (4.3) on λ is $1/(3\delta)$; for $\lambda = \frac{1}{3}$ this corresponds to a customer sharing her experience with an average of $\frac{1}{3} + \frac{1}{3^2} + \cdots = \frac{1}{2}$ future customers.[17]

The following example illustrates why the upper bound on λ is sufficient for precluding the "echo" that occurs under bounded memory.

Example 4.1. Let $\lambda = \frac{1}{3}$. If the motorist hires at period 0, the difference in stage payoffs between doing the right repair and the wrong repair is $u + w$. The probability that motorist 1 observes the repair at period 0 is $\frac{1}{3}$. The probability that motorist 2 observes period 1 is $\frac{1}{9}$, and the probability that motorist 2 observes period 1 and period 1 observes period 0 is $\frac{1}{9}$, so the probability of a "chain" of observations between period 0 and period 2 is bounded by $\frac{2}{9}$. The probability that period 3 observes period 0 is $\frac{1}{27}$, the probability that 3 observes 1 and 1 observes 0 is $\frac{1}{9} \cdot \frac{1}{3} = \frac{1}{27}$, and the probability that 3 observes 2 and that period 0 "reaches" period 2 via an observation chain is less than or equal to $\frac{1}{3} \cdot \frac{2}{9} = \frac{2}{27}$, so the probability that such a chain exists between period 0 and period 3 is bounded by $\frac{4}{27}$. This pattern continues so that the probability of a chain from period 0 to period t is bounded from above by $2^{t-1}/3^t$. The maximum difference the period 0 repair can

[17] This upper bound does not require Assumption 1, and an improved upper bound for the mechanic game can be obtained using Assumption 1 and an intuitive restriction on equilibria: Appendix B.2 gives an upper bound between $2/(5\delta)$ and $1/(2\delta)$ (depending on the ratio w/u), corresponding to a customer talking to an average between $\frac{2}{3}$ and 1 future customer for δ close to one.

make in the stage payoff at any future period that has a chain of observations back to period 0 is $u+w$,[18] so the discounted sum over these differences is $\sum_{t=1}^{\infty} \delta^{t-1} \frac{2^{t-1}}{3^t}(u+w) = \frac{u+w}{3} \cdot \frac{1}{1-\frac{2\delta}{3}} = \frac{u+w}{3-2\delta}$, which is less than the period 0 benefit of doing the right repair $(u+w)$. Thus, doing the right repair is the only best response.

Because motorists sometimes receive information that is useful for avoiding the bad mechanic, the ex ante payoff for all but the first few motorists is strictly greater in this equilibrium than in the one-shot game. Furthermore, for comparison with the "bad or useless reputation" result of Proposition 3.3, the motorist's payoff is also greater than the one-shot payoff in expectation over the improper uniform prior on the calendar date.

Corollary 4.1. *Let $\lambda > 0$ satisfying (4.3) be given. Let $\hat{v}^0$ be the maximum expected motorist payoff in a sequential equilibrium of the one-shot game. There is a unique sequential equilibrium expected payoff $v^{*t}(\delta)$ for the motorist arriving in period t. Let the motorist expected payoff (taken over the improper uniform prior over the periods) be $v^*(\delta)$. For all δ, $v^{*t}(\delta) = \hat{v}^0$ for $t < \bar{L}$ and $v^{*t}(\delta) > \hat{v}^0$ for $t \geq \bar{L}$. Furthermore, $v^*(\delta) > \hat{v}^0$.*

When history "fades" too slowly (i.e. λ is high), the bad reputation outcome returns. I first present a result that is weaker than Proposition 3.1 because it uses a different order of taking limits: instead of holding λ fixed and letting $\delta \to 1$, it holds δ fixed and lets $\lambda \to 1$. I then give a stronger result with the same order of limits as Proposition 3.1, using a mild restriction on equilibria.

Proposition 4.2. *Let $\mu^0 > 0$ and $\delta > (u+w)/(2u+w)$ be given. Then for λ close enough to one, for any sequential equilibrium satisfying Assumption 1 there is a unique equilibrium outcome where the mechanic is never hired.*

The proof of Proposition 4.2 shows that as λ gets arbitrarily close to one, the mechanic who performs a tune-up at a critical history is hired with probability close to one for arbitrarily many periods, while doing an engine replacement yields arbitrarily many periods of not being hired. At some point, the "memory" of the repair will (at least directly) fade away, and this "premium" the mechanic receives for a tune-up will eventually go away (or at least the bounds used cannot rule that out). The proof does not rule out the possibility that this premium for the tune-up is eventually (at periods far in the future) replaced by an even greater premium for the engine replacement. Instead, it simply relies on λ being high enough that any such "reverse premium" is postponed long enough that it is discounted away.

Establishing a lower bound on λ independent of the discount factor clearly requires ruling out such a reverse premium, which seems intuitively implausible because it requires that motorists far into the future are somehow dissuaded from hiring because of a tune-up, rather than an engine replacement, that they never observe directly (if they observed it directly they would hire, of

[18] Intuitively, one would expect only a difference of only u because that is the maximum decrease in the stage payoff going from being hired to not being hired (the difference between the maximum payoff and the minmax payoff). Corollary B.1 shows this intuition holds given some natural restrictions, yielding a higher upper bound for λ.

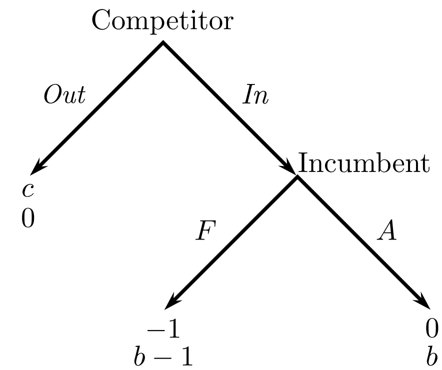

Figure 4.1: The chain store stage game, with payoffs at each node for the incumbent on top and for the competitor on bottom, with $b \in (0, 1)$ and $c > 1$.

course, because of Assumption 1). By restricting attention to equilibria where tune-ups do not, in expectation, dissuade future motorists from hiring, a lower bound on λ independent of δ is obtained that gives the bad reputation result.

Criterion 1. *Let a sequential equilibrium with strategy profile* (ρ, σ) *be given. Let* $\tilde{\sigma}^h$ *denote the strategy identical to* σ *except that the mechanic does a tune-up with certainty at full history* $h \in H^t$*. The equilibrium satisfies Criterion 1 if and only if doing a tune-up at* h *for* $h \in H^t$ *does not decrease the probability of being hired at any future period* k *given the motorists' strategies and beliefs, i.e.* $P_{\rho,\sigma}(\eta_k \neq N|h) \leq P_{\rho,\tilde{\sigma}^h}(\eta_k \neq N|h)$*, where* η_k *is the event at period* k*, for all* t*,* $k > t$*,* $h \in H^t$*, and* $\tilde{\sigma}^h$*.*

Criterion 1 is similar in spirit to the D1 Criterion (see, for example, Section 11.2 of Fudenberg and Tirole (1991)), but it is about actions instead of beliefs. (The proof of Proposition 4.1 shows that any equilibrium satisfying its assumptions must also satisfy Criterion 1.) The following result shows that the bad reputation result occurs when holding λ fixed and letting $\delta \to 1$.

Proposition 4.3. *Let* $\mu^0 > 0$ *and* $L(\mu^0)$ *be given. There exists* λ^* *such that for any* $\lambda \in (\lambda^*, 1)$*, for all sequential equilibria satisfying Assumption 1 and Criterion 1, there is a unique equilibrium outcome where the mechanic is never hired for* δ *close enough to one.*

4.3 The Chain Store Game

Among the class of games that Proposition 4.1 applies to is the infinitely repeated version of the chain store game, depicted in Figure 4.1, where the competitor is the short-run player and the incumbent is the long-run player. Let μ^0 be the prior belief that the incumbent is a "tough" commitment type that plays F every time entry occurs and probability $1 - \mu^0$ that the incumbent is a normal or "weak" (rational) type. The chain store game is a typical example of a game where reputation is good for the long-run player: Fudenberg and Levine (1989) show that under full memory, there is a lower bound on long-run player payoffs that approaches the Stackelberg payoff (which in this game is c) as δ approaches one.

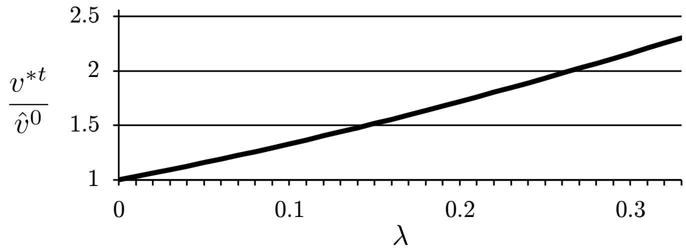

Figure 4.2: The ratio of the ex ante payoff v^{*t} for competitor $t = 20$ to the one-shot payoff $\hat{v}^0$ in any chain store game equilibrium with $\mu^0 = \frac{1}{5}, b = \frac{1}{4}$ (note that $v^{*t} = \hat{v}^0$ at $\lambda = 0$) is plotted for values of $\lambda \in [0, \frac{1}{3}]$, which satisfy (4.3) for $c = 2$ and any δ. These payoffs do not significantly change for periods past 20 (because $\psi(\lambda, 20) \approx \psi(\lambda, \infty)$).

For low λ, Proposition 4.1 shows that in any equilibrium the weak incumbent always accommodates (plays A). In this case, $z = c - 1$ and $z_d = 1$, so the upper bound (4.3) on λ is $1/[\delta(c+1)]$. As with the mechanic game, reputation increases the ex ante payoffs of the short-run players above that of the one-shot game. The increase here is more dramatic than in the mechanic game because a competitor need only observe one previous period to know if the incumbent is tough (all periods will be either F or Out) or weak (all periods are A). Thus, the ex ante payoff of competitor t is

$$v^{*t} = \psi(\lambda, t)(\mu^0(b-1) + (1-\mu^0)b) + (1 - \psi(\lambda, t))(1-\mu^0)b,$$

where $\psi(\lambda, t) \equiv \prod_{k=1}^{t}(1 - \lambda^k)$ is the probability of competitor t observing no history at all, which is strictly greater than the one-shot payoff for all competitors except at period 0.[19] This is straightforward to calculate and plotted in Figure 4.2 for some example parameters.

4.4 Correlation

Though this section generally assumes observations are independent across periods-observed, Proposition 4.1 (along with Corollary B.1) is robust to correlation between observations. This is because the proofs rely on Boole's inequality to bound the probability of the current repair being "chained" to any particular future period and then use the expected discounted sum of the effects of these chains at each future period for a bound on the continuation value. So long as the probabilities of these observations satisfy the "fading" definition, this correlation does not affect Proposition 4.1.

Correlation between observations can be interpreted as certain consumers being more connected to each other than others (a network of friends may be more likely to offer advice to each other than to strangers), and it does not have to be interpreted as decentralized communication. For example, consider messages posted on a centralized medium like an online forum, where it is visible to future customers while on the front page but with probability λ it disappears from view because other unrelated messages have pushed it off the front page. Others may publish replies underneath the post, sharing their own experiences; when the original post is pushed out of view, so are all

[19] $\psi(\lambda, t)$ converges absolutely as $t \to \infty$ to a value in the set $(0, 1)$ when $\lambda \in (0, 1)$ (Apostol, 1976).

these replies. In this case, if consumer t's post disappears at the end of period $t' > t$, then all consumers $t+1, ..., t'$ will see his message, but none of the consumers after period t' do. The ex ante probability that period t is observed by period $t+k$ is still λ^k, which is sufficient for the bounds used by Proposition 4.1 to ensure myopic long-run player behavior.

5 Conclusion

For the mechanic game with bounded memory, reputation is bad when short-run players have a long enough memory T. This is because early events that tarnish the mechanic's reputation "echo" for all following periods through the refusal of subsequent motorists to hire, which is observed by the following motorists who consequently also refuse to hire, and so on. When T is small enough, an equilibrium exists where the good mechanic always plays myopically. This equilibrium avoids bad reputation at the expense of making reputation irrelevant — motorists never see enough information to change their hiring decisions. It is impossible to shorten memory in a way that still provides valuable information to motorists but does not lower equilibrium payoffs, because the mechanic is too tempted to perform incorrect repairs: reputation is either bad or useless.

Under fading memory, when λ is less than a critical value, an equilibrium with myopic behavior by the mechanic exists, but reputation is still useful — sometimes motorists are informed enough that they do not hire. This increases the motorists' expected payoffs. The result holds more generally for reputation games where the long-run player has a strictly dominant action in the stage game: when λ is less than a critical value, the long-run player's equilibrium strategy is always to play the dominant action. This is because fading memory bounds the probability of an "observation chain" from the current period t to future period $\hat{t}$, where $t' > t$ observes t, $t'' > t'$ observes t', etc., which bounds the reputational payoffs of any signaling strategy. By contrast, such a chain always exists in bounded memory, even when $T = 1$, because period 1 always observes 0, period 2 always observes 1, etc. For high λ, the bad reputation result is recovered. The result for fading memory with small λ applies directly to the chain store game, for example, leading to a more dramatic increase in short-run player payoffs because they need only observe one past period to learn the long-run player's type.

While the above argument for fading memory carries over directly to the chain store game, the bounded memory arguments do not. For example, even for $T = 1$, there does not exist a "myopic" equilibrium in the chain store game. An interesting question is whether bounded memory equilibria in Stackelberg-type games like the chain store game can exhibit the cyclical behavior under limited records found by Liu and Skrzypacz (2014) or have the non-cyclical behavior of the mechanic game. More generally, behavior in other bounded memory reputation games remains largely unknown and is an interesting topic for future research.

References

Aoyagi, M. (1996): "Reputation and Entry Deterrence under Short-Run Ownership of a Firm," *Journal of Economic Theory*, 69(2), 411 – 430.

Apostol, T. M. (1976): *Introduction to Analytic Number Theory*. Springer-Verlag, New York.

Banerjee, A., and D. Fudenberg (2004): "Word-of-mouth learning," *Games and Economic Behavior*, 46(1), 1 – 22.

Bhaskar, V. (1998): "Information Constraints and the Overlapping Generations Model: Folk and Anti-Folk Theorems," *The Review of Economic Studies*.

Bhaskar, V., G. J. Mailath, and S. Morris (2008): "Purification in the infinitely-repeated prisoners' dilemma," *Review of Economic Dynamics*.

——— (2013): "A Foundation for Markov Equilibria in Sequential Games with Finite Social Memory," *The Review of Economic Studies*, 80(3), 925–948.

Ellison, G., and D. Fudenberg (1995): "Word-of-Mouth Communication and Social Learning," *The Quarterly Journal of Economics*, 110(1), 93–125.

Ely, J. C., D. Fudenberg, and D. K. Levine (2008): "When is reputation bad?," *Games and Economic Behavior*, 63, 498–526.

Ely, J. C., and J. Välimäki (2003): "Bad Reputation," *The Quarterly Journal of Economics*, 118, 785–814.

Fudenberg, D., and D. K. Levine (1989): "Reputation and Equilibrium Selection in Games with a Single Patient Player," *Econometrica*, 57, 251–268.

——— (1992): "Maintaining a Reputation when Strategies are Imperfectly Observed," *The Review of Economic Studies*, 59(3), pp. 561–579.

Fudenberg, D., and J. Tirole (1991): *Game Theory*. MIT Press, Cambridge, MA.

Harsanyi, J. C. (1973): "Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points," *International Journal of Game Theory*, 2(1), 1–23.

Kreps, D. M., and R. Wilson (1982): "Reputation and Imperfect Information," *Journal of Economic Theory*, 27, 253–279.

Liu, Q. (2011): "Information Acquisition and Reputation Dynamics," *The Review of Economic Studies*, 78(4), 1400–1425.

Liu, Q., and A. Skrzypacz (2014): "Limited records and reputation bubbles," *Journal of Economic Theory*, 151(0), 2 – 29.

Mailath, G. J., and W. Olszewski (2011): "Folk theorems with bounded recall under (almost) perfect monitoring," *Games and Economic Behavior*, 71(1), 174 – 192.

Mailath, G. J., and L. Samuelson (2006): *Repeated Games and Reputations: Long-Run Relationships*. Oxford University Press, New York.

Milgrom, P., and J. Roberts (1982): "Predation, Reputation and Entry Deterrence," *Journal of Economic Theory*, 27, 280–294.

Pitchik, C. (1993): "Commitment, Reputation, and Entry Deterrence," *Games and Economic Behavior*, 5(2), 268 – 287.

Rosenthal, R. (1979): "Sequences of Games with Varying Opponents," *Econometrica*, 47, 1353–1366.

Sabourian, H. (1998): "Repeated games with M-period bounded memory (pure strategies)," *Journal of Mathematical Economics*, 30(1), 1 – 35.

Selten, R. (1978): "The Chain Store Paradox," *Theory and Decision*, 9(2), 127–159.

Wiseman, T. (2008): "Reputation and impermanent types," *Games and Economic Behavior*, 62(1), 190 – 210.

Appendix

A Proofs of Bounded Memory Results

A.1 Proof of Proposition 3.1

Suppose by contradiction that for any $\delta^* \in (0,1)$, there always exists $\delta \in (\delta^*, 1)$ such that a sequential equilibrium with strategies (ρ, σ) exists with positive probability of hiring on the equilibrium path.

I begin by introducing some useful notation in addition to that defined in Section 3.1. Define $s_b(h) \equiv 1, s_g(h) \equiv \frac{1}{2}(1-\sigma_c(h)) + \frac{1}{2}\sigma_e(h)$ so that $s_\xi(h)$ is the probability that type $\xi \in \{b, g\}$ does an engine replacement upon being hired at history h. Define $\phi^t_\xi(h)$ as the probability of type ξ reaching history h at period t. Define $Q^{\bar{t}}_\xi(h) \equiv \sum_{t=0}^{\bar{t}} \phi^t_\xi(h)$, and define

$$R_\xi(h) \equiv \lim_{\bar{t}\to\infty} \frac{1}{\bar{t}-(T-1)} Q^{\bar{t}}_\xi(h), \qquad R^*_\xi(h) \equiv \begin{cases} R_\xi(h) & R_g(h) + R_b(h) > 0 \\ Q^\infty_\xi(h) & \text{otherwise.} \end{cases}$$

$R_\xi(h)$ is the long-run average probability that type ξ is at history h. Note that $R_\xi(h) > 0$ only if $Q_\xi^\infty(h) = \infty$. By Bayes' rule, the belief at history h is given by

$$\mu(h) = \lim_{\bar{t}\to\infty} \frac{\mu^0 Q_b^{\bar{t}}(h)}{\mu^0 Q_b^{\bar{t}}(h) + (1-\mu^0) Q_g^{\bar{t}}(h)},$$

and for $R_b(h) + R_b(h) > 0$ this implies $\mu(h) = \frac{\mu^0 R_b(h)}{\mu^0 R_b(h) + (1-\mu^0) R_g(h)}$. I also abuse notation for $Q_\xi^{\bar{t}}(\cdot), R_\xi^*(\cdot), R_\xi(\cdot)$ by letting $Q_\xi^{\bar{t}}(\mathcal{A}) = \sum_{h\in\mathcal{A}} Q_\xi^{\bar{t}}(h), R_\xi^*(\mathcal{A}) = \sum_{h\in\mathcal{A}} R_\xi^*(h), R_\xi(\mathcal{A}) = \sum_{h\in\mathcal{A}} R_\xi(h)$ for a set $\mathcal{A} \subset \mathcal{H}$ of histories. Define $\zeta(h) \equiv \lim_{t\to\infty} \frac{Q_b^t(h)}{Q_g^t(h)}$ and define ζ^* such that $p^* = \mu^0\zeta^*/[\mu^0\zeta^* + (1-\mu^0)]$. Let $\mathcal{C}$ be the set of histories with c in them, and let $\mathcal{C}^+ \equiv \mathcal{C} \cup \{e^T\}$. Define $\mathcal{X} \equiv H^T \backslash (\mathcal{C}^+ \cup \{N^T\})$, and $\mathcal{X}^+ \equiv \mathcal{X} \cup \{N^T\}$.

The following lemma uses the fact that purifiability restricts players to playing the same strategy at histories where the decision problem is the same.

Lemma A.1. *Let a weakly purifiable sequential equilibrium with strategy profile (ρ, σ) be given. It must be that $\sigma_x(e^T) = \sigma_x(ce^{T-1}) = \sigma_x(Ne^{T-1})$ for either $x \in \{c, e\}$.*

Proof. The argument follows the logic of Bhaskar, Mailath, and Morris (2013). Since (ρ, σ) is weakly purifiable, there exists a sequence $(\psi^{\tilde{k}}, \varepsilon^{\tilde{k}})_{\tilde{k}=1}^\infty$ with $\psi^{\tilde{k}} \equiv (\psi_b^{\tilde{k}}, \psi_g^{\tilde{k}}) \in (\Delta^*(Z))^2$ and $\lim_{\tilde{k}\to\infty} \varepsilon^{\tilde{k}} = 0$, such that there exists a sequence of strategy profiles $(\tilde{\rho}^{\tilde{k}}, \tilde{\sigma}^{\tilde{k}})_{\tilde{k}=1}^\infty$ such that $(\tilde{\rho}^{\tilde{k}}, \tilde{\sigma}^{\tilde{k}})$ is a sequential equilibrium in the $\Gamma(\psi^{\tilde{k}}, \varepsilon^{\tilde{k}})$ perturbed game and

$$\lim_{k\to\infty} \int \tilde{\rho}^k(h,z)\, d\psi_r^k(z) = \rho(h), \qquad \lim_{k\to\infty} \int \tilde{\sigma}_x^k(h,z)\, d\psi_g^k(z) = \sigma_x(h) \tag{A.1}$$

for all $h \in \mathcal{H}, x \in \{\theta_c, \theta_e\}$. Let some $\tilde{k}$ be given, and consider the mechanic's decision problem at a history h after being hired by a motorist needing repair x and current shock z. For any $h \in \{e^T, ce^{T-1}, Ne^{T-1}\}$, the mechanic's decision problem is the same since the continuation payoff for performing repair $x' \in \{c, e\}$ is the same in all three cases (since the histories are identical except for the oldest period about to be erased): $V(e^{T-1}x')$. For almost all shocks $z \in Z$, the mechanic's best response is unique, and so $\tilde{\sigma}_x^{\tilde{k}}(e^T, z) = \tilde{\sigma}_x^{\tilde{k}}(ce^{T-1}, z) = \tilde{\sigma}_x^{\tilde{k}}(Ne^{T-1}, z)$ for almost all $z \in Z$. Thus, for all $\tilde{k}$,

$$\int \tilde{\sigma}_x^{\tilde{k}}(e^T, z)\, d\psi_g^{\tilde{k}}(z) = \int \tilde{\sigma}_x^{\tilde{k}}(ce^{T-1}, z)\, d\psi_g^{\tilde{k}}(z) = \int \tilde{\sigma}_x^{\tilde{k}}(Ne^{T-1}, z)\, d\psi_g^{\tilde{k}}(z),$$

and so by (A.1), $\sigma_x(e^T) = \sigma_x(ce^{T-1}) = \sigma_x(Ne^{T-1})$. ∎

The following lemmas characterize the case where $\rho(e^T) < 1$.

Lemma A.2. *Suppose that $\rho(e^T) < 1$ and $T > L$. Then $\zeta^* \geq \zeta(Ne^{T-1})$ and $\zeta(e^T) \geq \zeta(Ne^{T-1})$ in any weakly purifiable equilibrium where the mechanic is hired on the equilibrium path.*

Proof. Suppose not, so then $\mu^* < \mu(Ne^{T-1})$ or $\mu(e^T) < \mu(Ne^{T-1})$. The former case clearly implies $\rho(Ne^{T-1}) = 0$, and in the latter case, Lemma A.1 implies that $\rho(Ne^{T-1}) = 0$. Thus, upon reaching a history in $\mathcal{X}^+$, the bad mechanic never returns to a history outside $\mathcal{X}^+$. Note also that there are never more than L engine replacements performed in the first T periods, so the bad mechanic always arrives in period T at a history in $\mathcal{X}^+$. Thus, $Q_b^t(e^T) = 0$ for all t. If the good mechanic is hired at any history with positive probability in equilibrium, she must perform a tune-up with positive probability. Since Assumption 1 requires she be hired immediately following a tune-up, and she must perform a tune-up with positive probability following each hiring, she must reach ce^{T-1} with positive probability and be hired. Thus, the good mechanic reaches e^T with positive probability in equilibrium, and so $\lim_{t\to\infty} Q_g^t(e^T) > 0$ which gives $\mu(e^T) = 0$. Lemma A.1 then requires that $\rho(e^T) = 1$, a contradiction. ■

Lemma A.3. *Suppose that $\rho(e^T) < 1$ and $T > L$. Then $\zeta(Ne^{T-1}) > \zeta(e^{T-1}N)$ in any weakly purifiable sequential equilibrium.*

Proof. Suppose not, so $\zeta(Ne^{T-1}) \le \zeta(e^{T-1}N)$.

Let any $hN \in \mathcal{X}^+$ be given. If $\zeta(hN) \le \tilde{\zeta}$ for some $\tilde{\zeta} \in (0,1)$, then it must be that $\zeta(Nh) \le \tilde{\zeta}$ or $\zeta(eh) \le \tilde{\zeta}$. Suppose by contradiction that $\zeta(Nh) > \tilde{\zeta}$ and $\zeta(eh) > \tilde{\zeta}$. Then

$$\begin{aligned}
\zeta(hN) = \frac{R_b^*(hN)}{R_g^*(hN)} &= \frac{(1-\rho(Nh))R_b^*(Nh) + (1-\rho(eh))R_b^*(eh)}{(1-\rho(Nh))R_g^*(Nh) + (1-\rho(eh))R_g^*(eh)} \\
&> \frac{(1-\rho(Nh))\tilde{\zeta}R_g^*(Nh) + (1-\rho(eh))\tilde{\zeta}R_g^*(eh)}{(1-\rho(Nh))R_g^*(Nh) + (1-\rho(eh))R_g^*(eh)} = \tilde{\zeta}.
\end{aligned}$$

A similar argument shows that for any $he \in \mathcal{X}^+$, $\zeta(he) \le \tilde{\zeta}$ implies that $\zeta(Nh) < \tilde{\zeta}$ or $\zeta(eh) < \tilde{\zeta}$ (the strictness of the inequalities is due to the fact that the good mechanic plays c with positive probability at Nh and eh conditional on playing e with positive probability).

The supposition that $\zeta(Ne^{T-1}) \le \zeta(e^{T-1}N)$ implies that either $\zeta(N^2e^{T-2}) < \zeta(e^{T-1}N)$ or $\zeta(eNe^{T-2}) < \zeta(e^{T-1}N)$, and continuing the argument by backward induction, there exists at least one sequence of histories $\{\tilde{h}^k\}_{k=1}^K$ (where K may be ∞), such that $\tilde{h}^1 = Ne^{T-1}$, $\tilde{h}^{k+1} \in \{N\tilde{h}^k, e\tilde{h}^k\}$ for each k, and $\zeta(\tilde{h}^{k+1}) < \zeta(e^{T-1}N)$. Let $\mathcal{B}^+$ denote the set of all histories in all such possible sequences, and let $\mathcal{B} \equiv \mathcal{B}^+ \backslash \{Ne^{T-1}\}$. Define $\tau_\xi(h)$ as the probability that type ξ at history $h \in \mathcal{B}$ switches to a history $hx \notin \mathcal{B}^+$ in the next period, and define $\hat{\tau}_\xi(h)$ as the probability that type ξ at history $h \in \mathcal{B}$ switches to Ne^{T-1} in the following period. Define $\omega_\xi(h)$ as the probability that type ξ at history $h \notin \mathcal{B}$ switches to a history $hx \in \mathcal{B}$ in the next period. We can then write

$$\phi_\xi^{t+1}(\mathcal{B}) = \sum_{h\in\mathcal{B}} \left[1 - (\hat{\tau}_\xi(h) + \tau_\xi(h))\right] \phi_\xi^t(h) + \sum_{h\in\mathcal{H}\backslash\mathcal{B}} \omega_\xi(h)\phi_\xi^t(h),$$

and summing over t and taking the limit gives

$$\begin{aligned}
R_\xi(\mathcal{B}) &= \sum_{h\in\mathcal{B}} \left[1 - (\hat{\tau}_\xi(h) + \tau_\xi(h))\right] R_\xi(h) + \sum_{h\in\mathcal{H}\backslash\mathcal{B}} R_\xi(h)\omega_\xi(h) \\
&= R_\xi(\mathcal{B}) - \sum_{h\in\mathcal{B}} (\hat{\tau}_\xi(h) + \tau_\xi(h))R_\xi(h) + \sum_{h\in\mathcal{H}\backslash\mathcal{B}} R_\xi(h)\omega_\xi(h)
\end{aligned}$$

Rearranging shows that long-run average probability of leaving $\mathcal{B}$ must be equal to the long-run average probability of entering $\mathcal{B}$:

$$\sum_{h\in\mathcal{B}} R_\xi(h)\left[\hat{\tau}_\xi(h)+\tau_\xi(h)\right] = \sum_{h\in\mathcal{H}\backslash\mathcal{B}} R_\xi(h)\omega_\xi(h).$$

Suppose that $R_b(Ne^{T-1})+R_g(Ne^{T-1})>0$. Note that $\hat{\tau}_g(h)+\tau_g(h)\geq\hat{\tau}_b(h)+\tau_b(h)$ because the only difference between $\hat{\tau}_g(h),\tau_g(h)$ and $\hat{\tau}_b(h),\tau_b(h)$, respectively, is that the good mechanic also has the additional possibility of leaving $\mathcal{B}$ by playing c. Since $\mathcal{B}$ is defined such that for all $h\in\mathcal{B}$, $R_b(h)/R_g(h)<\zeta(e^{T-1}N)$, then

$$\zeta(Ne^{T-1})=\frac{\sum_{h\in\mathcal{B}}R_b(h)\hat{\tau}_b(h)}{\sum_{h\in\mathcal{B}}R_g(h)\hat{\tau}_g(h)}\geq\frac{\sum_{h\in\mathcal{B}}R_b(h)\left[\hat{\tau}_b(h)+\tau_b(h)\right]}{\sum_{h\in\mathcal{B}}R_g(h)\left[\hat{\tau}_g(h)+\tau_g(h)\right]}=\frac{\sum_{h\in\mathcal{H}\backslash\mathcal{B}}R_b(h)\omega_b(h)}{\sum_{h\in\mathcal{H}\backslash\mathcal{B}}R_g(h)\omega_g(h)}. \tag{A.2}$$

However, note that for any $h\notin\mathcal{B}$, if $\omega_\xi(h)>0$, then $h\in\mathcal{X}^+\backslash\mathcal{B}$, since all histories in $\mathcal{X}^+\supset\mathcal{B}$ are only reachable within one period from other histories in $\mathcal{X}^+$ except for $e^{T-1}N$, and by definition of $\mathcal{B}$, $e^{T-1}N\notin\mathcal{B}$. Thus, $\omega_\xi(h)>0 \implies \zeta(h)\geq\zeta(e^{T-1}N)$ by the definition of $\mathcal{B}$. This implies that the right hand side of (A.2) is greater than or equal to $\zeta(e^{T-1}N)$, a contradiction.

Finally, suppose that

$$R_b(Ne^{T-1})+R_g(Ne^{T-1})=0. \tag{A.3}$$

Define $\hat{\tau}_\xi^{T-1}(h)$ as the probability that type ξ at history h reaches Ne^{T-1} in (exactly) $T-1$ periods. Let $\mathcal{T}\equiv\{h\in\mathcal{H}:\hat{\tau}_\xi^{T-1}(h)>0,\xi\in\{b,g\}\}$ as the set of histories h that can reach Ne^{T-1} in $T-1$ periods (note that any such h reachable by type b is also reachable by g and vice versa). Define $\mathcal{D}\equiv\mathcal{T}\cap H^T$ as the set of such non-initial histories. Note that $\mathcal{T}\backslash\mathcal{D}=\{N\}$ (the singleton set of the period 1 history "N"). Then $\phi_\xi^{t+T-1}(Ne^{T-1})=\sum_{h\in\mathcal{D}}\hat{\tau}_\xi^{T-1}(h)\phi_\xi^t(h)+\hat{\tau}_\xi^{T-1}(N)\phi_\xi^t(N)$. Summing from $t=1$ gives

$$Q_\xi^\infty(Ne^{T-1})=\sum_{h\in\mathcal{D}}\hat{\tau}_\xi^{T-1}(h)Q_\xi^\infty(h)+\hat{\tau}_\xi^{T-1}(N)\phi_\xi^t(N). \tag{A.4}$$

For all $h\in\mathcal{D}$ and $t>T$,

$$\phi_\xi^t(h)=\sum_{k=T}^{t-1}\tilde{\tau}_\xi^{h,t-k}(e^{T-1}N)\phi_\xi^k(e^{T-1}N)+\sum_{\tilde{h}\in\mathcal{X}^+}\tilde{\tau}_\xi^{h,t-T}(\tilde{h})\phi_\xi^T(\tilde{h})$$

where $\tilde{\tau}_\xi^{h,m}(\tilde{h})$ is the probability that type ξ at history $\tilde{h}$ reaches h in exactly m periods without reaching $e^{T-1}N$ in the intermediate periods . Summing over t starting from $t=T+1$ gives

$$Q_\xi^\infty(h)-\phi_\xi^T(h)=\sum_{t=T+1}^{\infty}\sum_{k=T}^{t-1}\tilde{\tau}_\xi^{h,t-k}(e^{T-1}N)\phi_\xi^k(e^{T-1}N)+\sum_{t=T+1}^{\infty}\sum_{\tilde{h}\in\mathcal{X}^+}\tilde{\tau}_\xi^{h,t-T}(\tilde{h})\phi_\xi^T(\tilde{h}). \tag{A.5}$$

Note that the first term on the right hand side is equal to the total probability that type ξ at $e^{T-1}N$

ever reaches h, multiplied by the sum over $\phi^l_\xi(e^{T-1}N)$ for all periods l, i.e.

$$\sum_{l=T}^{\infty}\phi^l_\xi(e^{T-1}N)\sum_{m=1}^{\infty}\tilde{\tau}^{h,m}_\xi(e^{T-1}N)=\sum_{m=1}^{\infty}\tilde{\tau}^{h,m}_\xi(e^{T-1}N)Q^\infty_\xi(e^{T-1}N).$$

Let $\tilde{P}^h_\xi(\tilde{h})=\sum_{m=1}^{\infty}\tilde{\tau}^{h,m}_\xi(\tilde{h})$ denote the probability that type ξ at $\tilde{h}$ ever reaches h without passing through $e^{T-1}N$. We can then rewrite (A.5) as

$$\begin{aligned}
Q^\infty_\xi(h) &= \sum_{l=1}^{\infty}\tilde{\tau}^{h,l}_\xi(e^{T-1}N)Q^\infty_\xi(e^{T-1}N)+\sum_{\tilde{h}\in\mathcal{X}^+}\phi^T_\xi(\tilde{h})\sum_{t=T}^{\infty}\tilde{\tau}^{h,t-T}_\xi(\tilde{h})+\phi^T_\xi(h)\\
&= \tilde{P}^h_\xi(e^{T-1}N)Q^\infty_\xi(e^{T-1}N)+\tilde{P}^h_\xi(\tilde{h})\sum_{\tilde{h}\in\mathcal{X}^+}\phi^T_\xi(\tilde{h})+\phi^T_\xi(h).
\end{aligned}$$

Note that $\tilde{P}^{h''}_b(h')\geq\tilde{P}^{h''}_g(h')$ for any histories $h',h''\in\mathcal{X}^+$ because the good type can only be more likely to reach h'' by doing a tune-up at some point, and then the only way to get back to histories within $\mathcal{X}^+$ (and therefore to h') is by passing through $e^{T-1}N$ (which is excluded by the definition of $\tilde{P}^{h''}_\xi(\cdot)$). Also note that $\phi^T_b(\tilde{h})\geq\phi^T_g(\tilde{h})$ for all $\tilde{h}\in\mathcal{X}^+$ since the good mechanic can only reach $\tilde{h}$ at period T by always performing engine replacements up to that point. Thus, we can see that

$$\zeta(h)=\frac{Q^\infty_b(h)}{Q^\infty_g(h)}=\frac{\tilde{P}^h_b(e^{T-1}N)Q^\infty_b(e^{T-1}N)+\tilde{P}^h_b(\tilde{h})\sum_{\tilde{h}\in\mathcal{X}^+}\phi^T_b(\tilde{h})+\phi^T_\xi(h)}{\tilde{P}^h_g(e^{T-1}N)Q^\infty_g(e^{T-1}N)+\tilde{P}^h_g(\tilde{h})\sum_{\tilde{h}\in\mathcal{X}^+}\phi^T_g(\tilde{h})+\phi^T_\xi(h)}\geq\min\left\{1,\zeta(e^{T-1}N)\right\}$$

for all $h\in\mathcal{D}$.

For any history $h\in\mathcal{T}$, in order to arrive at history Ne^{T-1} in $T-1$ periods requires performing $T-1$ consecutive engine replacements, and since the the good type must perform c with at least probability $\frac{1}{2}\beta^*$ when hired, it must be that $(1-\frac{1}{2}\beta^*)^{T-1}\hat{\tau}^{T-1}_b(h)\geq\hat{\tau}^{T-1}_g(h)$. From (A.4), we can write

$$\begin{aligned}
\frac{Q^\infty_b(Ne^{T-1})}{Q^\infty_g(Ne^{T-1})} &= \frac{\sum_{h\in\mathcal{D}}\hat{\tau}^{T-1}_b(h)Q^\infty_b(h)+\hat{\tau}^{T-1}_b(N)\phi^t_b(N)}{\sum_{h\in\mathcal{D}}\hat{\tau}^{T-1}_g(h)Q^\infty_g(h)+\hat{\tau}^{T-1}_g(N)\phi^t_g(N)}\\
&\geq \frac{1}{(1-\frac{1}{2}\beta^*)^{T-1}}\frac{\sum_{h\in\mathcal{D}}\hat{\tau}^{T-1}_g(h)Q^\infty_b(h)+\hat{\tau}^{T-1}_g(N)\phi^t_b(N)}{\sum_{h\in\mathcal{D}}\hat{\tau}^{T-1}_g(h)Q^\infty_g(h)+\hat{\tau}^{T-1}_g(N)\phi^t_g(N)}.
\end{aligned}$$

Above it was shown that for all $h\in\mathcal{D}$, $\zeta(h)\geq\min\left\{1,\zeta(e^{T-1}N)\right\}$. Note that $\zeta(N)=1$. Therefore

$$\zeta(Ne^{T-1})=\frac{Q^\infty_b(Ne^{T-1})}{Q^\infty_g(Ne^{T-1})}\geq(1-\tfrac{1}{2}\beta^*)^{-(T-1)}\min\left\{1,\zeta(e^{T-1}N)\right\}.$$

If $1\geq\zeta(e^{T-1}N)$, then the lemma is proven. If $1<\zeta(e^{T-1}N)$, then $\zeta(Ne^{T-1})\geq(1-\frac{1}{2}\beta^*)^{-(T-1)}$. The definitions of $\Upsilon(\cdot)$ and ζ^* give

$$\frac{\Upsilon^L(\mu^0)}{1-\Upsilon^L(\mu^0)}=\frac{\mu^0}{1-\mu^0}\frac{1}{(1-\frac{1}{2}\beta^*)^L}>\frac{\mu^0}{1-\mu^0}\zeta^*.$$

Then

$$\zeta(Ne^{T-1}) = \frac{Q_b^\infty(Ne^{T-1})}{Q_g^\infty(Ne^{T-1})} \geq \frac{1}{(1-\frac{1}{2}\beta^*)^{T-1}} \geq \frac{1}{(1-\frac{1}{2}\beta^*)^{L}} > \zeta^*,$$

but this is a contradiction of Lemma A.2. ■

The following lemma shows that the conclusions of the lemmas above lead to a contradiction when the mechanic is sufficiently patient.

Lemma A.4. *Let $T > L$ be given. For δ close enough to one, there does not exist a weakly purifiable sequential equilibrium (ρ, σ) such that the mechanic is hired in equilibrium, $\rho(e^T) < 1$, and $\zeta(e^T) \geq \zeta(Ne^{T-1}) > \zeta(e^{T-1}N)$.*

Proof. Suppose by contradiction that such an equilibrium does exist. Note that $\phi_\xi^{t+1}(e^{T-1}N) = (1-\rho(e^T))\phi_\xi^t(e^T) + (1-\rho(Ne^{T-1}))\phi_\xi^t(Ne^{T-1})$, so averaging over t and taking the limit $t \to \infty$ gives

$$R_\xi^*(e^{T-1}N) - \phi_\xi^T(e^{T-1}N) = (1-\rho(e^T))R_\xi^*(e^T) + (1-\rho(Ne^{T-1}))R_\xi^*(Ne^{T-1}).$$

Note that $\phi_\xi^T(e^{T-1}N) = 0$ because $e^{T-1}N$ contains $T-1 \geq L$ engine replacements. Then

$$\zeta(e^{T-1}N) = \frac{R_b^*(e^{T-1}N)}{R_g^*(e^{T-1}N)} = \frac{(1-\rho(e^T))R_b^*(e^T) + (1-\rho(Ne^{T-1}))R_b^*(Ne^{T-1})}{(1-\rho(e^T))R_g^*(e^T) + (1-\rho(Ne^{T-1}))R_g^*(Ne^{T-1})}. \tag{A.6}$$

Since we have supposed that $\zeta(e^T) \geq \zeta(Ne^{T-1}) > \zeta(e^{T-1}N)$, we have a contradiction because this implies that the right hand side of (A.6) is strictly greater than the left. ■

For the rest of the proof, suppose that $\rho(e^T) = 1$. Let $\hat{k} \equiv \lfloor \frac{u+w}{u} \rfloor$, and let $T > (L+1)\hat{k}$ be given. Then, $\mathcal{C}^+$ is a closed class, and so a motorist arriving at a history $h \notin \mathcal{C}^+$ knows that a c has not appeared before in the full history.

Lemma A.5. *Suppose that $\rho(e^T) = 1$. For δ close enough to one, if the mechanic is hired, there exists $\hat{\rho} > 0$ such that the mechanic is hired within $\hat{k}$ periods with probability greater than or equal to $\hat{\rho}$.*

Proof. If the mechanic is hired at history h in period t and performs c, then rehiring occurs the next period by Assumption 1, and because $\rho(e^T) = 1$, the mechanic is hired forever after as well, yielding a continuation payoff of u. Because the mechanic must weakly prefer the correct repair, incentive compatibility when an engine replacement is needed requires $(1-\delta)u + \delta V(he) \geq -(1-\delta)w + \delta u$. Suppose by contradiction that the mechanic is hired with probability zero for $k' > \hat{k}$ periods. Then $V(he) \leq \delta^{k'} u$, so incentive compatibility gives

$$(1-\delta)u + \delta^{k'+1}u \geq -(1-\delta)w + \delta u \implies \frac{u+w}{u} \geq \delta\frac{1-\delta^{k'}}{1-\delta}.$$

By l'Hôpital's rule, the limit of the right hand side as $\delta \to 1$ is $k' > \frac{u+w}{u}$, so for δ close enough to one, there is a contradiction. ■

I show that if the mechanic is hired on the equilibrium path at any history, there exists some $K \leq L$ and some history $h^0 \in H^T$ reached on the equilibrium path, without any tune-ups, and containing K engine replacements all within the $\hat{k}K$ most recent periods.

If the mechanic is not hired in equilibrium in the first T periods, then the mechanic is necessarily hired with positive probability at $h^0 = N^T$. Suppose instead that the mechanic is hired in equilibrium within the first T periods. Let $\underline{t}_0 < T$ be the earliest period at which the mechanic is hired with positive probability in some weakly purifiable sequential equilibrium. Lemma A.5 implies that there must be a history where the mechanic is hired with positive probability within $\hat{k}$ periods of $\underline{t}_0$, and performs an engine replacement. Furthermore, there can be at most L engine replacements in the first T periods. Thus, there exists a history $h^0 \in H^T$, reached in equilibrium at period T, where the mechanic is hired some $K \leq L$ times and performs an engine replacement, and all of those engine replacements happen at or after period $T - K\hat{k}$ and before period T.

Having established the existence of h^0, suppose that the mechanic is hired at h^0 at some period $t \geq T$. Lemma A.5 also implies that the mechanic must be hired and perform an engine replacement with positive probability at or before period $t+\hat{k}$, $t+2\hat{k}$ and so on. Thus, for δ close enough to one, the mechanic must be hired with positive probability in equilibrium at some sequence of histories $(\tilde{h}^l)_{l=0}^{\infty}$ where $\tilde{h}^0 = h^0 N^{k_0}$ and $\tilde{h}^{l+1} = \tilde{h}^l e N^{k_{l+1}}$, where $k_0, k_1, k_2, ...$ are all non-negative integers less than or equal to $\hat{k}$. Thus, in equilibrium with positive probability, the mechanic is hired L times, beginning at some period $\underline{t}_0 \geq T - K\hat{k}$ and before period $T + (L-K)\hat{k}$. Since there are L engine replacements in the (bounded) history and at least one N occurrence, the motorist knows a tune-up has not occurred in the full history, and so $\mu(\tilde{h}^{L-K}) > p^*$. Since $T > (L+1)\hat{k}$, all of the histories $\tilde{h}^{L-K}N, \tilde{h}^{L-K}N^2, ..., \tilde{h}^{L-K}N^{\hat{k}+1}$ still contain L engine replacements (none of the engine replacements have yet been "erased") and so have beliefs greater than p^*. Yet Lemma A.5 implies that the mechanic is hired with positive probability at at least one of these histories, so there is a contradiction.

A.2 Proof of Proposition 3.3

First, if $\mu^0 > p^*$, the result is trivial because the mechanic is never hired, both in the one-shot and repeated games, because the prior is too high to hire even if the mechanic does the correct repair with certainty, so $\bar{v}^*(\delta) = \hat{v}^0 = 0$ for all δ. The rest of the proof considers the case $\mu^0 \in (0, p^*]$, and so $\hat{v}^0 = \mu^0 \frac{u-w}{2} + (1-\mu^0)u$. I reuse the notation from the proof of Proposition 3.1.

The following lemma shows that a "useful reputation" equilibrium requires a "useful history" where the belief exceeds μ^* and occurs with positive long-term probability.

Lemma A.6. *Let v^* be the motorist's expected payoff. If $v^* > \hat{v}^0$, there must exist $h \in \mathcal{H}$ such that $\mu(h) > p^*$ and $R_b(h) + R_g(h) > 0$.*

Proof. Let v^*_ξ be the motorists' expected payoff conditional on type ξ, and let $v^*_\xi(h)$ be the expected payoff conditional on history h. Then

$$\begin{aligned}
v^* &= \mu^0 v_b^* + (1-\mu^0) v_g^* = \mu^0 \sum_{h\in\mathcal{H}} R_b(h) v_b^*(h) + (1-\mu^0) \sum_{h\in\mathcal{H}} R_g(h) v_g^*(h) \\
&= \sum_{h\in\mathcal{H}} \left[\mu^0 R_b(h) v_b^*(h) + (1-\mu^0) R_g(h) v_g^*(h)\right] \\
&\leq \sum_{h\in\mathcal{H}} \left[\mu^0 R_b(h)\rho(h)\frac{u-w}{2} + (1-\mu^0) R_g(h)\rho(h)\left(\tfrac{1}{2}u + \tfrac{1}{2}\sigma_e(h)\right)\right] \\
&= \sum_{h\in\mathcal{H}} \rho(h)\left[\mu^0 R_b(h)\frac{u-w}{2} + (1-\mu^0) R_g(h)\left(\tfrac{1}{2}u + \tfrac{1}{2}\sigma_e(h)\right)\right]. \qquad \text{(A.7)}
\end{aligned}$$

Let $R(h) \equiv R_b(h) + R_g(h)$. For any h such that $R(h) > 0$, Bayes' rule gives

$$\mu(h) = \frac{\mu^0 R_b(h)}{R(h)}, \qquad 1-\mu(h) = \frac{(1-\mu^0)R_g(h)}{R(h)}$$

$$R_b(h) = \frac{\mu(h)R(h)}{\mu^0}, \qquad R_g(h) = \frac{(1-\mu(h))R(h)}{1-\mu^0},$$

and substituting into (A.7) yields

$$\begin{aligned}
v^* &\leq \sum_{h\in\mathcal{H}} \rho(h)\left[\mu(h)R(h)\frac{u-w}{2} + (1-\mu(h))R(h)\left(\tfrac{1}{2}u + \tfrac{1}{2}\sigma_e(h)\right)\right] \\
&= \sum_{h\in\mathcal{H}} R(h)\rho(h)\left[\mu(h)\frac{u-w}{2} + (1-\mu(h))\left(\tfrac{1}{2}u + \tfrac{1}{2}\sigma_e(h)\right)\right] \\
&\leq \sum_{h\in\mathcal{H}} R(h)\rho(h)\left[\mu(h)\frac{u-w}{2} + (1-\mu(h))u\right].
\end{aligned}$$

Suppose that $v^* > \hat{v}^0$. Since $\hat{v}^0 = \mu^0\frac{u-w}{2} + (1-\mu^0)u$ for $\mu^0 \in (0, p^*]$, a necessary condition is that there exist h such that $R(h) > 0$, $\rho(h) < 1$, and $0 > \mu(h)\frac{u-w}{2} + (1-\mu(h))u$, which implies $\mu(h) > p^*$. ■

Suppose that $\rho(e^T) < 1$. To rule this case out, I adapt analogous lemmas from the proof of Proposition 3.1. Lemma A.1 holds here, but replacements are needed for Lemmas A.2 and A.3, since they require $T > L$ instead of arbitrary T.

Lemma A.7. *Let a weakly purifiable equilibrium with strategy profile (σ, ρ) that contradicts (3.6) be given. If $\rho(e^T) < 1$, then $\zeta(e^T) \geq \zeta(Ne^{T-1})$.*

Proof. Suppose not, so then $\mu(e^T) < \mu(Ne^{T-1})$. Lemma A.1 implies that $\rho(Ne^{T-1}) = 0$. Then $\phi_\xi^{t+1}(e^T) = \rho(ce^T)s_\xi(ce^T)\phi_\xi^t(ce^{T-1}) + \rho(e^T)s_\xi(e^T)\phi_\xi^t(e^T)$ and summing over t and taking the limiting average gives

$$R_\xi(e^T) = \rho(ce^T)s_\xi(ce^T)R_\xi(ce^{T-1}) + \rho(e^T)s_\xi(e^T)R_\xi(e^T).$$

Note that upon reaching a history in $\mathcal{X}^+$, the bad mechanic never returns to a state outside $\mathcal{X}^+$

(because $\rho(Ne^{T-1}) = 0$). Since $\rho(e^T) < 1$, for the bad mechanic e^T is a transient state for the bad type, so $R_b(e^T) = 0$. If $R_g(e^T) > 0$, then $\rho(e^T) = 1$ by Assumption 1, a contradiction. If $R_g(e^T) = 0$, then it must be that $R_g(ce^{T-1}) = 0$, which similarly requires $R_g(ece^{T-2}) = R_g(cce^{T-2}) = R_g(Nce^{T-2})$. A backward induction argument shows that $R_g(\mathcal{C}) = 0$ for the set $\mathcal{C}$ of all histories containing c, which implies $v^* = 0 < \hat{v}^0$ since motorists expect a tune-up with probability zero, a contradiction. ■

Lemma A.8. *Suppose that $\rho(e^T) < 1$. Then $\zeta(Ne^{T-1}) > \zeta(e^{T-1}N)$ in any weakly purifiable sequential equilibrium violating (3.6).*

Proof. The proof of Lemma A.3 does not rely on the assumption that $T > L$ until (A.3), so we need only rule out the case where

$$R_b(Ne^{T-1}) + R_g(Ne^{T-1}) = 0. \tag{A.8}$$

Let $R(h) \equiv R_b(h) + R_g(h)$. $R(Ne^{T-1}) = 0$ implies that $R_b(e^T) = 0$. If $R_g(e^T) > 0$, then necessarily $\mu(e^T) = 0$ and by Lemma A.1, $\rho(e^T) = 1$, a contradiction. Thus, $R_g(e^T) = 0$. By the same arguments as in the proof of Lemma A.7, $R_g(e^T) = 0$ leads to a contradiction. ■

Lemmas A.7 and A.8 then give the following lemma ruling out the $\rho(e^T) < 1$ case, whose proof is essentially the same as that of Lemma A.4 and so is omitted.

Lemma A.9. *For δ close enough to one, there does not exist a weakly purifiable sequential equilibrium (ρ, σ) contradicting (3.6) such that $\rho(e^T) < 1$ and $\zeta(e^T) \geq \zeta(Ne^{T-1}) > \zeta(e^{T-1}N)$.*

For the remainder of the proof, suppose that $\rho(e^T) = 1$. Then $\mathcal{C}^+$ is an closed class set of states (once the mechanic enters $\mathcal{C}^+$, she never leaves).

Suppose that $\rho(N^T) > 0$. Note that conditional on being at a history in $\mathcal{X}^+$, the mechanic is hired with some minimum probability $q > 0$ within some $\tilde{k} \leq T + 1$ periods, either because she is hired at some history $h \neq N^T$ with positive probability or because she is not hired for T periods, thereby reaching history N^T and then being hired with positive probability $\rho(N^T) > 0$. Since the good mechanic performs c with at least probability $\frac{1}{2}\beta^*$ upon being hired, $\phi_g^{t+\tilde{k}}(\mathcal{X}^+) \leq (1 - \frac{1}{2}\beta^*\tilde{\rho})\phi_g^t(\mathcal{X}^+)$ for all $t \geq T$. This implies that $R_g(\mathcal{X}^+) = 0$. Suppose there exist a history $h \in \mathcal{X}$ such that $R_b(h) > 0$. This means that $\mu(h) = 1$, and so $R_b(hN) \geq R_b(h)$. By induction, this means that $R_b(N^T) > 0$, and so $\mu(N^T) = 1$. But then a motorist at N^T will not hire, a contradiction of $\rho(N^T) > 0$.

Finally, suppose that $\rho(N^T) = 0$, so N^T is an absorbing state. If the good mechanic is at a history in $\mathcal{X}$, she will be at a history not in $\mathcal{X}$ within T periods (and never return) with some minimum probability $q > 0$, because either she is not hired for T periods and reaches N^T (an absorbing state) or she is hired at some history $h' \in \mathcal{X}$, and performs a tune-up with at least probability $\frac{1}{2}\beta^*$. Thus, $\phi_g^{t+T}(\mathcal{X}) \leq (1 - q)\phi_g^t(\mathcal{X})$, and a similar argument to that in the previous paragraph shows that $R_g(\mathcal{X}) = R_b(\mathcal{X}) = 0$.

Thus, Lemma A.6 shows that an equilibrium violating (3.6) requires $\mu(N^T) > p^*$. Note that $\phi_\xi^{t+1}(N^T) = \phi_\xi^t(N^T) + (1-\rho(eN^{T-1}))\phi_\xi^t(eN^{T-1})$. We can then write

$$\begin{aligned} R_\xi(N^T) &= \lim_{\bar{t}\to\infty} \frac{1}{\bar{t}-(T-1)} \sum_{t=T}^{\bar{t}} \phi_\xi^t(N^T) \\ &= \lim_{\bar{t}\to\infty} \frac{1}{\bar{t}-(T-1)} \sum_{t=T}^{\bar{t}} \left\{ \phi_\xi^T(N^T) + \sum_{t'=T}^{t-1} (1-\rho(eN^{T-1}))\phi_\xi^{t'}(eN^{T-1}) \right\} \\ &= \phi_\xi^T(N^T) + (1-\rho(eN^{T-1})) \lim_{\bar{t}\to\infty} \frac{1}{\bar{t}-(T-1)} \sum_{t=T}^{\bar{t}} \sum_{t'=T}^{t-1} \phi_\xi^{t'}(eN^{T-1}) \\ &= \phi_\xi^T(N^T) + (1-\rho(eN^{T-1})) \lim_{\bar{t}\to\infty} \frac{1}{\bar{t}-(T-1)} \sum_{t=T}^{\bar{t}} Q_\xi^{t-1}(eN^{T-1}). \end{aligned}$$

Since $\mu(N^T) > p^* \implies \zeta(N^T) > \zeta^*$, we have

$$\zeta^* < \frac{R_b(N^T)}{R_g(N^T)} = \lim_{\bar{t}\to\infty} \frac{\phi_b^T(N^T) + (1-\rho(eN^{T-1}))\frac{1}{\bar{t}-(T-1)}\sum_{t=T}^{\bar{t}} Q_b^{t-1}(eN^{T-1})}{\phi_g^T(N^T) + (1-\rho(eN^{T-1}))\frac{1}{\bar{t}-(T-1)}\sum_{t=T}^{\bar{t}} Q_g^{t-1}(eN^{T-1})}. \quad \text{(A.9)}$$

Note that $\phi_b^T(N^T) = \phi_g^T(N^T)$ (since the probability of the mechanic not being hired T times consecutively at the beginning of the game is independent of type). Since $\zeta^* > 1$, (A.9) requires that $Q_b^\infty(eN^{T-1}) > 0$ and

$$\lim_{\bar{t}\to\infty} \frac{\sum_{t=T}^{\bar{t}} Q_b^{t-1}(eN^{T-1})}{\sum_{t=T}^{\bar{t}} Q_b^{t-1}(eN^{T-1})} = \zeta(eN^{T-1}) > \zeta^*,$$

and so $\mu(eN^{T-1}) > p^*$.

Now, let some $h \in \mathcal{X}$ be given such that the following hold:

1. the $k \geq 0$ most recent outcomes are N (i.e., $h_{-k} = h_{-(k-1)} = \cdots = h_{-1} = N$), and the $(k+1)$th last outcome is e (i.e., $h_{-(k+1)} = e$),
2. $\mu(hN^l) > p^*$ for all $l \geq 0$ (and hence $V(h) = 0$), and
3. $Q_b^\infty(h) > 0$.

I show that this implies $\mu(eh) > p^*$ or $\mu(Nh) > p^*$. Note that $\phi_\xi^{t+1}(h) = (1-\rho(eh))\phi_\xi^t(eh) + (1-\rho(Nh))\phi_\xi^t(Nh)$ where eh and Nh denote the histories with the oldest element being e and N, respectively, followed by the $T-1$ oldest outcomes of h. Summing over $t = T+1$ to some $\bar{t}$ gives

$$Q_\xi^{\bar{t}}(h) - \phi_\xi^T(h) = (1-\rho(eh))Q_\xi^{\bar{t}-1}(eh) + (1-\rho(Nh))Q_\xi^{\bar{t}-1}(Nh).$$

We can see that

$$\zeta^* < \zeta(h) = \lim_{\bar{t}\to\infty} \frac{Q_b^{\bar{t}}(h)}{Q_g^{\bar{t}}(h)} = \lim_{\bar{t}\to\infty} \frac{\phi_b^T(h) + (1-\rho(eh))Q_b^{\bar{t}-1}(eh) + (1-\rho(Nh))Q_b^{\bar{t}-1}(Nh)}{\phi_g^T(h) + (1-\rho(eh))Q_g^{\bar{t}-1}(eh) + (1-\rho(Nh))Q_g^{\bar{t}-1}(Nh)}.$$

I now show that for δ close enough to one, either

$$\lim_{\bar{t}\to\infty} \frac{Q_b^{\bar{t}-1}(eh)}{Q_g^{\bar{t}-1}(eh)} > \zeta^* \qquad \text{or} \qquad \lim_{\bar{t}\to\infty} \frac{Q_b^{\bar{t}-1}(Nh)}{Q_g^{\bar{t}-1}(Nh)} > \zeta^*. \tag{A.10}$$

Suppose not. Then $\phi_b^T(h)/\phi_g^T(h) > \zeta^*$. Note that $\phi_\xi^T(h) = (1-\rho(h^{T-1}))\phi_\xi^{T-1}(h^{T-1})$ where h^{T-1} is the $(T-1)$-length history of the $T-1$ oldest outcomes in h. Then

$$\zeta^* < \frac{\phi_b^T(h)}{\phi_g^T(h)} = \frac{(1-\rho(h^{T-1}))\phi_b^{T-1}(h^{T-1})}{(1-\rho(h^{T-1}))\phi_g^{T-1}(h^{T-1})},$$

and so $\mu(h^{T-1}) > p^*$. By induction, this argument shows that $\mu(h^{T-k'}) > p^*$ for all $k' \in \{1, ..., k\}$. This implies that $\rho(h^{t-k'}) = 0$ for all $k' \in \{1, ..., k\}$, and since $V(h) = 0$, we have $V(h^{T-k}) = 0$. Since $\phi_b^{T-k}(h^{T-k}) > 0$, the mechanic is hired at $h^{T-(k+1)}$ on the equilibrium path, and so the mechanic must perform a needed engine replacement; incentive compatibility requires

$$(1-\delta)u + \delta V(h^{T-k}) = (1-\delta)u \geq -(1-\delta)w + \delta V(h^{T-(k+1)}c) = -(1-\delta)w + \delta u,$$

which is a contradiction for δ close enough to one. Thus, $\phi_b^T(h)/\phi_g^T(h) \leq \zeta^*$, and so (A.10) is proven.

By induction, there exists some history h, such that the most recent outcome is e, $\mu(hN^l) > \mu^*$ for all $l \geq 0$ and $Q_b^\infty(h) > 0$. Hence, the mechanic is hired at xh for some $x \in \{N, e, c\}$ on the equilibrium path, and so the good mechanic must weakly prefer performing a needed engine replacement at xh. Since $V(h) = 0$ and the continuation payoff to playing c is u (since the mechanic is hired forever after), incentive compatibility at xh requires $(1-\delta)u + \delta \cdot 0 \geq -(1-\delta)w + \delta u$, which clearly gives a contradiction for δ close enough to one.

B Proofs of Fading Memory Results

B.1 Proof of Proposition 4.1

Let any sequential equilibrium with strategies (ρ, σ), and a full history h at period t be given. Suppose player 1 faces a decision node, i.e. she is at some information set over player 2's action $a_2 \in \tilde{A}_2$ and observes private signal $\theta \in \Theta$. Let $V(h(a_2, a_1))$ denote player 1's continuation payoff following the action profile (a_2, a_1).[20] Player 1 plays a_d with certainty if

$$E_{a_2}[(1-\delta)u_1(a_d, a_2, \theta) + \delta V(h(a_2, a_d))] > E_{a_2}[(1-\delta)u_1(a_1', a_2, \theta) + \delta V(h(a_2, a_1'))] \qquad \text{(B.1)}$$

for all $a_1' \in A_1 \backslash \{a_d\}$, where $E_{a_2}[\cdot]$ is the expectation over player 2's actions $a_2 \in \tilde{A}_2$ given player 1's information set (in the mechanic game, of course, player 1 knows a_2 (hiring) because $\tilde{A}_2$ is a singleton). (B.1) can be rearranged as

$$E_{a_2}[V(h(a_2, a_d)) - V(h(a_2, a_1'))] < \frac{1-\delta}{\delta} E_{a_2}[u_1(a_d, a_2, \theta) - u_1(a_1', a_2, \theta)]. \qquad \text{(B.2)}$$

The left hand side of (B.2) is equal to the discounted sum of the expected differences in stage payoffs at every future period. Denoting player 1's stage payoff at some period $\hat{t} > t$ as $\nu(\hat{t})$, let $\bar{\nu}_{a_2,a_1}(h, \hat{t}) \equiv E[\nu(\hat{t})|h(a_2, a_1)]$ be the expected stage payoff at $\hat{t}$ conditional on reaching full history $h(a_2, a_1)$. An upper bound on the change in the expected stage payoff at $\hat{t}$ due to choosing an action different from a_d at full history h is $\Delta\bar{\nu}(h, \hat{t}) \equiv \max_{(a_2, a_1') \in \tilde{A}_2 \times A_1} \{\bar{\nu}_{a_2, a_1'}(h, \hat{t}) - \bar{\nu}_{a_2, a_d}(h, \hat{t})\} \leq z$. The action at t can only affect player 1's payoff at $\hat{t}$ if period $\hat{t}$ observes t directly, or observes some period $t' \in \{t+1, ..., \hat{t}-1\}$ that observes t, etc.; otherwise, player 2's action at period $\hat{t}$ is necessarily independent of the events of period t. This notion of an "observation chain" is formalized as "t reaches $\hat{t}$" in the following definition.

Definition B.1. Let two periods t' and $t'' > t'$ be given. Inductively define the relation "t' k-*reaches* t''" as follows. If period t'' observes period t', then t' is said to 0-reach t''. If period t'' observes some period $\tilde{t} \in \{t'+1, ..., t''-1\}$ and $\tilde{t}$ k-reaches t', then t' is said to $(k+1)$-reach t''. More simply, if (and only if) period t' k-reaches t'' for some $k \in \{0, 1, ...\}$, then t' is said to *reach* t''.

Let $\phi(t, \hat{t})$ denote the probability that t reaches $\hat{t}$, which gives the upper bound $\Delta\bar{\nu}(h^t, \hat{t}) \leq \phi(t, \hat{t})z$. The following lemma gives an upper bound for $\phi(t, \hat{t})$.

Lemma B.1. *For any two periods t and $\hat{t} > t$, $\phi(t, \hat{t}) \leq 2^{\hat{t}-t-1}\lambda^{\hat{t}-t}$.*

Proof. The proof is by induction. For $\hat{t} = t+1$, $\phi(t, \hat{t}) = \lambda$ is trivially true. Now suppose that for some $\hat{t} > t$, $\phi(t, t') \leq 2^{t'-t-1}\lambda^{t'-t}$ for all $t' \in \{t+1, ..., \hat{t}\}$. The probability that t reaches $\hat{t}+1$ is the probability that short-run player $\hat{t}+1$ observes either t or some period $t' \in \{t+1, ..., \hat{t}-1\}$ such that t reaches t'. Then Boole's inequality gives

$$\begin{aligned} \phi(t, \hat{t}+1) &\leq \lambda^{\hat{t}+1-t} + \sum_{t'=t+1}^{\hat{t}} \lambda^{\hat{t}+1-t'}\phi(t, t') \leq \lambda^{\hat{t}+1-t} + \sum_{t'=t+1}^{\hat{t}} \lambda^{\hat{t}+1-t'}(2^{t'-t-1}\lambda^{t'-t}) \\ &= \lambda^{\hat{t}+1-t}\left(1 + \sum_{k=0}^{\hat{t}-t-1} 2^k\right) = \lambda^{\hat{t}+1-t}\left(1 + \frac{1-2^{\hat{t}-t}}{1-2}\right) = 2^{\hat{t}-t}\lambda^{\hat{t}+1-t}. \end{aligned}$$

[20] "$h(a_2, a_1)$" denotes the full history consisting of (a_2, a_1) appended to h.

■

I can now write an upper bound for the left hand side of (B.2):

$$E_{a_2}[V(h(a_2,a_d)) - V(h(a_2,a_1'))] \leq (1-\delta)\sum_{k=1}^{\infty}\delta^{k-1}\Delta\bar{\nu}(h,t+k)$$

$$\leq \quad (1-\delta)\sum_{k=1}^{\infty}\delta^{k-1}\phi(t,t+k)z \leq (1-\delta)z\sum_{k=1}^{\infty}\delta^{k-1}(2^{k-1}\lambda^k).$$

Since $\delta\lambda < \frac{1}{2}$,

$$E_{a_2}[V(h(a_2,a_d)) - V(h(a_2,a_1'))] \leq (1-\delta)\lambda z\sum_{k=0}^{\infty}(2\delta\lambda)^k = \frac{(1-\delta)\lambda z}{1-2\delta\lambda}. \tag{B.3}$$

The right hand side of (B.3) is a strictly increasing function of λ for $\lambda \in (0, 1/(2\delta))$. Substituting (4.3) into λ gives

$$\begin{aligned} E_{a_2}[V(h(a_2,a_d)) - V(h(a_2,a_1'))] \quad &< \quad \frac{(1-\delta)z\left(\frac{z_d}{\delta(z+2z_d)}\right)}{1-2\delta\left(\frac{z_d}{\delta(z+2z_d)}\right)} = \frac{(1-\delta)}{\delta}z_d \\ &\leq \quad \frac{1-\delta}{\delta}E_{a_2}[u_1(a_d,a_2,\theta) - u_1(a_1',a_2,\theta)] \end{aligned}$$

for any $a_1' \in A_1, a_2 \in \tilde{A}_2$, so (B.1) is true.

B.2 A Higher Upper Bound for λ for Myopic Equilibria in the Mechanic Game

Proposition 4.1 assumes that the "worst case" when an "observation chain" reaches a future period is the stage payoff decreasing by the maximum feasible amount z; in the mechanic game, this difference is $u + w$. A tighter bound that seems natural is the difference between the highest feasible payoff and the minmax payoff (u). The following corollary uses that bound on the stage payoff difference to give a higher upper bound on λ, using Assumption 1 and Criterion 1. For δ close to one, as w/u approaches 1 the bound (B.4) approaches $\frac{2}{5}$ (corresponding to motorists talking to an average of $\frac{2}{3}$ future motorists) and as w/u approaches ∞, (B.4) approaches $\frac{1}{2}$ (corresponding to an average of 1 future motorist).

Corollary B.1. *Consider the fading memory mechanic game with*

$$\lambda < \frac{1}{\delta\left(2+\frac{u}{u+w}\right)}. \tag{B.4}$$

Then the action outcome of any sequential equilibrium satisfying Assumption 1 and Criterion 1 has the good mechanic doing the correct repair when hired.

Proof. Let σ denote the equilibrium strategy of the good mechanic, and let $\bar{\sigma}$ be the strategy identical to σ except that at any full history containing a tune-up, the mechanic does the right repair with certainty (it may be that $\sigma = \bar{\sigma}$). The following result allows a simplification of the continuation payoffs for a tune-up.

Lemma B.2. *Let a sequential equilibrium under fading memory given λ satisfying Assumption 1 and Criterion 1 be given. At any full history $h \in \mathcal{H}$ on the equilibrium path containing a tune-up, it is a best response for the mechanic to perform the correct repair.*

Proof. Any motorist observing the full history (which occurs with positive probability at every full history) must hire due to Assumption 1. This is only possible if the mechanic performs the correct repair with at least positive probability β^* no matter the car's state, so it must be a best response. ■

Thus, deviating to $\bar{\sigma}$ must be a best response at any full history containing a tune-up. Calculation of the expected stage payoffs following a tune-up (simply the probability of being hired times u) is simpler for $\bar{\sigma}$ and allows them to be used as upper bounds on the expected stage payoffs following an engine replacement because of Criterion 1.

For the remainder of the proof, suppose that the mechanic deviates to $\bar{\sigma}$, which has the same continuation payoffs as σ at every full history, an implication of Lemma B.2. Let the notation and arguments in the proof of Proposition 4.1 (Appendix B.1) up to and including Lemma B.1 be given, except that all notation is with respect to the strategy $\bar{\sigma}$ (not σ) and a_2 is omitted from subscripts (because in the mechanic game, at a mechanic's decision node, a_2 is known to be "hire"). At any full history h, let $\pi_a^{\hat{t}}(h)$ be the probability that the mechanic is hired at period $\hat{t} > t$ conditional on repair a at h.

Criterion 1 implies that the continuation payoff for a tune-up is greater than or equal to that of an engine replacement because $\bar{\nu}_c(h, \hat{t}) = \pi_c^{\hat{t}}(h)u \geq \pi_e^{\hat{t}}(h)u \geq \bar{\nu}_e(h, \hat{t})$. Therefore, when the motorist at h^t needs a tune-up, performing a tune-up strictly dominates an engine replacement.

What remains to be shown is that performing a needed engine replacement strictly dominates doing an incorrect tune-up. Let $\breve{\sigma}$ be the strategy identical to $\bar{\sigma}$, except that any full history following he (i.e. any full history that begins with h followed by e at period t) the mechanic always does the right repair. Let $\breve{V}, \breve{\pi}_a^{\hat{t}}, \breve{\nu}_a$ be the analogues of $V, \pi_a^{\hat{t}}, \bar{\nu}_a$ (which are defined for $\bar{\sigma}$) for a deviation to $\breve{\sigma}$. Note that $\breve{\nu}_e(h, \hat{t}) = \breve{\pi}_e^{\hat{t}}(h)u$, $\breve{\pi}_c^{\hat{t}}(h) = \pi_c^{\hat{t}}(h)$, $\breve{V}(hc) = V(hc)$, and $\breve{V}(he) \leq V(he)$. The fact that $\breve{\nu}_c(h, \hat{t}) - \breve{\nu}_e(h, \hat{t}) = (\breve{\pi}_c^{\hat{t}}(h) - \breve{\pi}_e^{\hat{t}}(h))u \leq \phi(t, \hat{t})u$ yields

$$\begin{aligned} V(hc) - V(he) &\leq \breve{V}(hc) - \breve{V}(he) = (1-\delta)\sum_{k=1}^{\infty} \delta^{k-1}(\breve{\nu}_c(h, \hat{t}) - \breve{\nu}_e(h, \hat{t})) \\ &\leq (1-\delta)\sum_{k=1}^{\infty} \delta^{k-1}\phi(t, t+k)u. \end{aligned}$$

By Lemma B.1,

$$V(hc) - V(he) \leq (1-\delta)u \sum_{k=1}^{\infty} \delta^{k-1}(2^{k-1}\lambda^k) = \frac{(1-\delta)\lambda u}{1-2\delta\lambda}.$$

Then substituting (B.4) into λ gives

$$V(hc) - V(he) < \frac{(1-\delta)u\left(\frac{1}{\delta(2+u/(u+w))}\right)}{1-2\delta\left(\frac{1}{\delta(2+u/(u+w))}\right)} = \frac{(1-\delta)u}{\delta\left(2+u/(u+w)\right)-2\delta} = \frac{1-\delta}{\delta}(u+w).$$

Therefore, doing an incorrect tune-up is not a best response. ■

B.3 Proof of Corollary 4.1

The mechanic has a unique equilibrium strategy, and the only case where the motorists do not have a unique equilibrium strategy is if they are indifferent; hence the equilibrium payoffs are unique. Motorist t observing history $\tilde{h} \in \tilde{H}^t$ for some t does not hire if $\mu(\tilde{h}) > p^*$. Note that by definition of $\bar{L}$, for periods $t < \bar{L}$, there are no histories yielding a belief greater than p^*, and so hiring is always a best response, yielding the one-shot equilibrium payoff for the motorist. The expected payoff of a motorist conditional on period t and the mechanic's type s is

$$v^{*t}(\tilde{h}; s) = \begin{cases} \rho(\tilde{h})\frac{u-w}{2} & s = b \\ \rho(\tilde{h})u & s = g, \end{cases}$$

where $\rho(\tilde{h})$ is the probability of hiring given $\tilde{h}$. Thus, the expected motorist t payoff is

$$v^{*t} = \sum_{\tilde{h}\in\tilde{H}^t} [\mu^0 P_{\rho,\sigma}(\tilde{h}|b)v^*(\tilde{h}; b) + (1-\mu^0)P_{\rho,\sigma}(\tilde{h}|g)v^*(\tilde{h}; g)] \tag{B.5}$$

where $P_{\rho,\sigma}(\tilde{h}|s)$ is the probability that history $\tilde{h}$ occurs at period t conditional on type s given ρ, σ. Bayes' rule gives

$$\mu(\tilde{h}) = \frac{\mu^0 P_{\rho,\sigma}(\tilde{h}|b)}{P_{\rho,\sigma}(\tilde{h})} \implies \mu^0 P_{\rho,\sigma}(\tilde{h}|b) = P_{\rho,\sigma}(\tilde{h})\mu(\tilde{h})$$

$$1-\mu(\tilde{h}) = \frac{(1-\mu^0)P_{\rho,\sigma}(\tilde{h}|g)}{P_{\rho,\sigma}(\tilde{h})} \implies (1-\mu^0)P_{\rho,\sigma}(\tilde{h}|g) = P_{\rho,\sigma}(\tilde{h})(1-\mu(\tilde{h})),$$

where $P_{\rho,\sigma}(\tilde{h})$ is the probability of $\tilde{h}$ at t (unconditional on the mechanic's type). Substituting the right hand side expressions into (B.5) gives

$$\begin{aligned} v^{*t} &= \sum_{\tilde{h}\in\tilde{H}^t} P_{\rho,\sigma}(\tilde{h})[\mu(\tilde{h})v^*(\tilde{h}; b) + (1-\mu(\tilde{h}))v^*(\tilde{h}; g)] \\ &= \sum_{\tilde{h}\in\tilde{H}^t} P_{\rho,\sigma}(\tilde{h})\rho(\tilde{h})\left[\mu(\tilde{h})\frac{u-w}{2} + (1-\mu(\tilde{h}))u\right]. \end{aligned}$$

The term in brackets is the expected payoff of hiring at $\tilde{h}$, which is negative when $\mu(\tilde{h}) > p^*$ since the mechanic always does the right repair and therefore $\rho(\tilde{h}) = 0$. If there exists any $\tilde{h}$ such that $P_{\rho,\sigma}(\tilde{h}) > 0$ and $\mu(\tilde{h}) > p^*$, then I can write

$$\begin{aligned} v^{*t} &= \sum_{\tilde{h}\in\tilde{H}^t} P_{\rho,\sigma}(\tilde{h}) \left[\mu(\tilde{h})\frac{u-w}{2} + (1-\mu(\tilde{h}))u\right] \mathbf{1}\{\mu(\tilde{h}) \le p^*\} \\ &> \sum_{\tilde{h}\in\tilde{H}^t} P_{\rho,\sigma}(\tilde{h}) \left[\mu^0\frac{u-w}{2} + (1-\mu^0)u\right] = \hat{v}^0, \end{aligned} \quad \text{(B.6)}$$

where $\mathbf{1}\{\cdot\}$ is an indicator function (set to one if the statement is true and zero otherwise). The right hand side of (B.6) is the one-shot ex ante payoff. It is easy to see for every period $t > \bar{L}(\mu^0)$ that such a history $\tilde{h}$ exists (at a minimum, they observe the full history with only engine replacements with positive probability). Let some $t \ge \bar{L}(\mu^0)$ be given. Note that the unconditional probability that the outcome is e or N in any period k is greater than $\frac{1}{2}$.

Let q_t^k be the probability that motorist t observes exactly k periods. Let $\psi(r,m) \equiv \prod_{l=1}^{m}(1-r^k)$. The probability of observing the k most recent periods and none of the older ones is $\lambda^{k(k+1)}\psi(\lambda^k, t-k)$. The (unconditional) probability that the outcomes in periods $t-k-1, ..., t-1$ are all either e or N is greater than 2^{-k}. The probability that motorist t observes a history $\tilde{h}$ with $\mu(\tilde{h}, t) > p^*$ with k periods of either e or N is greater than $2^{-k}\lambda^{k(k+1)}\psi(\lambda^k, t-k)$. Since $\psi(\lambda, t)$ converges absolutely as $t \to \infty$ to some value $\chi \in (0,1)$ (Apostol, 1976), the probability that any motorist $t \ge \bar{L}$ observes a history $\tilde{h}$ such that $\mu(\tilde{h}, t) > p^*$ is bounded from below by $2^{-\bar{L}}\lambda^{\bar{L}(\bar{L}+1)}\chi$. Thus, there exists $\varepsilon > 0$ such that $v^{*t} > \hat{v}^0 + \varepsilon$ for all $t \ge \bar{L}(\mu^0)$, and so $v^* = \lim_{\bar{t}\to\infty} \frac{1}{\bar{t}+1}\sum_{t=0}^{\bar{t}} v^{*t} > \hat{v}^0$.

B.4 Proof of Proposition 4.2

Suppose by contradiction that for any $\lambda^* \in (0,1)$, there always exists $\lambda \in (\lambda^*, 1)$ such that a sequential equilibrium exists with positive probability of hiring on the equilibrium path. Let such an equilibrium be given. Without loss of generality, let period zero be the first period at which the mechanic is hired with positive probability. The following lemma establishes that if the mechanic is hired at some full history in equilibrium, she must be hired again sufficiently soon (or else the temptation to do a tune-up will be too great).

Lemma B.3. *Suppose the mechanic is hired at some full history $h \in H^t$ on the equilibrium path at period t with positive probability, such that*

- $t = 0$, *or*
- *the mechanic is hired with probability greater than $\lambda^{t(t+1)/2}$.*

Suppose the mechanic chooses e at h, and let $t' > t$ be the earliest future period at which the mechanic is again hired with probability greater than $\lambda^{t'(t'+1)/2}$. Define $K(\delta, u, w) \equiv \ln(\delta - (1 - \delta)(1 + w/u))/\ln\delta$. Then $t' \le t + K(\delta, u, w)$ for λ close enough to one.

Proof. For any period t, if the mechanic is hired at period 0 or at $h \in H^t$ with probability greater than $\lambda^{t(t+1)/2}$, then the mechanic must perform a needed engine replacement with positive probability; otherwise, the motorist who sees the full history $h \in H^t$ would not hire, since the probability that the full history is observed is $\prod_{k=1}^{t} \lambda^k = \lambda^{t(t+1)/2}$. Incentive compatibility gives

$$(1-\delta)u + \delta V(he) \geq -(1-\delta)w + \delta V(hc) \tag{B.7}$$

where $V(ha)$ is the continuation payoff of full history $ha \in H^{t+1}$. By definition, t' is the earliest period such that the mechanic is hired with probability greater than $\lambda^{t'(t'+1)/2}$ if she chooses e at period t, so an upper bound on her continuation payoff for e is

$$\begin{aligned} V(he) &\leq (1-\delta)\left[\sum_{k=t+1}^{t'-1} \delta^{k-t-1}(1-\lambda^{k(k+1)/2})u + \sum_{k=t'}^{\infty} \delta^{k-t-1}u\right] \\ &= (1-\delta)\sum_{k=t+1}^{t'-1} \delta^{k-t-1}(1-\lambda^{k(k+1)/2})u + \delta^{t'-t-1}u. \end{aligned}$$

Assumption 1 gives the following lower bound for the continuation payoff of c:

$$V(hc) \geq (1-\delta)\sum_{k=0}^{\infty} \delta^k \lambda^{k+1} u = \frac{(1-\delta)\lambda u}{1-\delta\lambda}.$$

Substituting these bounds into (B.7) gives

$$(1-\delta)u + (1-\delta)\sum_{k=t+1}^{t'-1} \delta^{k-t}(1-\lambda^{k(k+1)/2})u + \delta^{t'-t}u \geq -(1-\delta)w + \delta\frac{(1-\delta)\lambda u}{1-\delta\lambda}$$

$$\delta^{t'-t} \geq (1-\delta)\left[\frac{\delta\lambda}{1-\delta\lambda} - (1+w/u) - \sum_{k=t+1}^{t'-1} \delta^{k-t}(1-\lambda^{k(k+1)/2})\right]. \tag{B.8}$$

Let any $\varepsilon > 0$ be given. Taking the limit of the right hand side of (B.8) as $\lambda \to 1$, there exists λ^* such that for all $\lambda \in (\lambda^*, 1)$, $\delta^{t'-t} \geq \delta - (1-\delta)(1+w/u)/\exp(\varepsilon/(-\ln\delta))$ since $\exp(\varepsilon/(-\ln\delta)) > 1$. Solving for t' gives

$$(t'-t)\ln\delta \geq \ln(\delta - (1-\delta)(1+w/u)) - \frac{\varepsilon}{-\ln\delta} \implies t' \leq t + \frac{\ln(\delta-(1-\delta)(1+w/u))}{\ln\delta} + \varepsilon.$$

Since $\varepsilon > 0$ is arbitrary, we can pick $\varepsilon < \max\{1, \lceil K(\delta,u,w)\rceil - K(\delta,u,w)\}$. In that case, because t' is an integer, it must be that $t' \leq t + K(\delta,u,w)$ for λ close enough to one. ■

Lemma B.3 implies that for λ close enough to one, if the mechanic is hired at period 0, with positive probability she must be hired in at least $L+1$ periods (with greater than probability $\lambda^{t'(t'+1)/2}$ at each such period t') in the first $KL+1$ periods on the equilibrium path, which means there must exist full history $h' \in H^{\tilde{t}}$ at $\tilde{t} \leq KL+1$ on the equilibrium path that includes $L+1$

engine replacements and no tune-ups. This also implies that at each of these hirings, the mechanic must have performed a tune-up with at least probability β^* (see (3.2)). Yet this implies that the posterior of the motorist receiving the $(L+1)$th engine replacement at period t_{L+1} if he observes the full history must have been at least $\Upsilon^L(\mu^0) > p^*$. Thus, hiring was not a best response for that motorist with at least probability $\lambda^{t_{L+1}(t_{L+1}+1)/2}$, a contradiction.

B.5 Proof of Proposition 4.3

I begin by characterizing λ^*. For $t \in \{0, 1, ...\}$, $n \in \{2, 3, ...\}$, and $\lambda \in (0, 1)$, define

$$f(t, n; \lambda) \equiv \sum_{k=t+1}^{t+n-1} [\lambda^{k-t} - (1 - \lambda^{k(k+1)/2})].$$

Note the following useful properties about the function f.

Fact B.1. *$f(t, n; \lambda)$ is strictly decreasing in t, and strictly increasing in λ.*

Fact B.2. *Let any $\varepsilon > 0$ be given. For any $t \in \{0, 1, ...\}, n \in \{2, 3, ...\}$, there exists $\lambda' \in (0, 1)$ such that for any $\lambda \in (\lambda', 1)$, $f(t, n; \lambda) > n - 1 - \varepsilon$.*

Let n^* be an integer strictly greater than $1 + w/u$. Pick $\lambda^* \in (0, 1)$ such that

$$f(Ln^*, n^* + 1; \lambda^*) \geq 1 + w/u. \tag{B.9}$$

Let $\lambda \in (\lambda^*, 1)$ be given. Suppose by contradiction that for any $\delta^* \in (0, 1)$, there always exists $\delta \in (\delta^*, 1)$ such that there exists a sequential equilibrium with strategies ρ, σ where the mechanic is hired with positive probability on the equilibrium path. Without loss of generality, let the first such period be 0.

Let $\bar{\sigma}$ be the strategy identical to σ except that at any full history containing a tune-up, the mechanic does the right repair with certainty (it may be that $\sigma = \bar{\sigma}$). Lemma B.2, reproduced here as Lemma B.4 for convenience, shows that deviating to $\bar{\sigma}$ is a best response at any full history (the only histories at which $\bar{\sigma}$ may differ from σ are those with tune-ups, and for those histories doing the right repair is always a best response).

Lemma B.4. *Let a sequential equilibrium satisfying Assumption 1 and Criterion 1 be given. At any full history $h \in \mathcal{H}$ containing a tune-up, it is a best response for the mechanic to perform the correct repair.*

I use a technique here similar to the proof of Corollary B.1 to simplify calculation of continuation payoffs. For $\bar{\sigma}$, calculation of the expected stage payoffs following a tune-up is simple (due to Lemma B.4, it is the probability of being hired times u) and due to Criterion 1, they can be used as upper bounds on the expected stage payoffs following an engine replacement (shown below).

For the remainder of the proof, suppose that the mechanic deviates to $\bar{\sigma}$; since by Lemma B.4 such a deviation is a best response at any full history, the continuation payoffs are identical at all histories. At any full history $h \in H^t$, let $\pi_a^k(h)$ be the probability that the mechanic is hired at period $k > t$ conditional on doing repair $a \in \{c, e\}$ at h, and let $\bar{\nu}_a^k(h)$ denote the expected stage payoff at period k conditional on a. Criterion 1 requires that $\pi_c^k(h) \geq \pi_e^k(h)$. Since the mechanic performs all correct repairs following a tune-up, $\bar{\nu}_c^k(h) = \pi_c^k(h)u \geq \pi_e^k(h)u \geq \bar{\nu}_e^k(h)$.

Lemma B.5. *Let the assumptions of Proposition 4.3 and $\lambda \in (\lambda^*, 1)$ be given. For δ close enough to one, if there exists a sequential equilibrium where the mechanic is hired with positive probability at period 0, then there exists a full history $h^{\tilde{t}} \in H^{\tilde{t}}$ on the equilibrium path at some period $\tilde{t} \leq Ln^*$ where the mechanic is hired with probability greater than $1 - \lambda^{\tilde{t}(\tilde{t}+1)/2}$ and $h^{\tilde{t}}$ contains L engine replacements and no tune-ups such that the posterior after observing the full history is $\mu(h^{\tilde{t}}) \geq \Upsilon^L(\mu^0)$.*

Proof. The proof is by induction. Let $t_1 > 0$ be the first period after 0 at which the mechanic is hired with greater than probability $1 - \lambda^{t_1(t_1+1)/2}$, conditional on the mechanic doing an engine replacement in period 0.

I now show that $t_1 \leq n^*$ for δ close enough to one. The continuation payoff of a tune-up at 0 has a lower bound due to Assumption 1 given by

$$\frac{\delta}{(1-\delta)} V(c) = \sum_{k=1}^{\infty} \delta^k \bar{\nu}_c^k(h^0) \geq \sum_{k=1}^{t_1 - 1} \delta^k \lambda^k u + \sum_{k=t_1}^{\infty} \delta^k \bar{\nu}_e^k(h^0)$$

where h^0 is the empty history at period 0. Since the mechanic is hired, the incentive constraint $-(1-\delta)w + \delta V(c) \leq (1-\delta)u + \delta V(e)$ must hold (when an engine replacement is needed). A necessary condition for this incentive constraint is

$$-w + \sum_{k=1}^{n-1} \delta^k \lambda^k u + \sum_{k=n}^{\infty} \delta^k \bar{\nu}_e^k(h^0) \leq u + \sum_{k=1}^{n-1} \delta^k (1 - \lambda^{k(k+1)/2}) u + \sum_{k=n}^{\infty} \delta^k \bar{\nu}_e^k(h^0) \tag{B.10}$$

for any $n \leq t_1$. Suppose by contradiction that $t_1 > n^*$. After some rearrangement of (B.10), picking $n = n^* + 1$ gives

$$\sum_{k=1}^{n^*} \delta^k [\lambda^k - (1 - \lambda^{k(k+1)/2})] u \leq u + w. \tag{B.11}$$

Dividing by u and taking the limit of δ gives

$$\lim_{\delta \to 1} \sum_{k=1}^{n^*} \delta^k [\lambda^k - (1 - \lambda^{k(k+1)/2})] = f(0, n^* + 1) < 1 + w/u,$$

so (B.11) contradicts (B.9) for $\lambda > \lambda^*$ and δ close enough to one. Thus, the mechanic must be hired at period t_1 with probability greater than $1 - \lambda^{t_1(t_1+1)/2}$ and $t_1 \leq n^*$ at a full history $h^{t_1} \in H^{t_1}$ with one engine replacement and no tune-ups.

Now for some $j \geq 1$, let h^{t_j} be a full history at $t_j \leq jn^*$ on the equilibrium path where the mechanic is hired with probability greater than $1 - \lambda^{t_j(t_j+1)/2}$, such that h^{t_j} has j engine replacements and no tune-ups. I show that there exists period $t_{j+1} \leq t_j + n^*$ such that the mechanic is hired with probability greater than $1-\lambda^{t_{j+1}(t_{j+1}+1)/2}$. Since the mechanic is hired at h^{t_j} with greater than probability $1-\lambda^{t_j(t_j+1)/2}$, the mechanic must perform an engine replacement with positive probability when it is needed (by the same argument as above), and incentive compatibility gives the necessary condition

$$\begin{aligned} & -w + \sum_{k=t_j+1}^{t_j+n-1} \delta^{k-t_j} \lambda^{k-t_j} u + \sum_{k=t_j+n}^{\infty} \delta^{k-t_j} \bar{v}_e^k(h^{t_j}) \\ \leq \quad & u + \sum_{k=t_j+1}^{t_j+n-1} \delta^{k-t_j} (1-\lambda^{k(k+1)/2}) u + \sum_{k=t_j+n}^{\infty} \delta^{k-t_j} \bar{v}_e^k(h^{t_j}) \end{aligned}$$

for $n \leq t_{j+1} - t_j$. Suppose by contradiction that $t_{j+1} > t_j + n^*$. Picking $n = n^* + 1$ gives

$$\sum_{k=t_j+1}^{t_j+n^*} \delta^{k-t_j} [\lambda^{k-t_j} - (1-\lambda^{k(k+1)/2})] u \leq u + w. \tag{B.12}$$

Dividing by u and taking the limit of δ gives

$$\lim_{\delta \to 1} \sum_{k=t_j+1}^{t_j+n^*} \delta^{k-t_j} [\lambda^{k-t_j} - (1-\lambda^{k(k+1)/2})] = f(t_j, n^*+1) < 1 + w/u,$$

so (B.12) contradicts (B.9) for $\lambda > \lambda^*$ and δ close enough to one. Then there exists a full history $h^{t_{j+1}}$ following h^{t_j} on the equilibrium path for some $t_{j+1} \leq t_j + n^* \leq n^*(j+1)$ where the mechanic is hired with probability greater than $1 - \lambda^{t_{j+1}(t_{j+1}+1)/2}$ with $j+1$ engine replacements (where the good mechanic must have performed a tune-up with at least probability β^*) and no tune-ups. For δ close enough to one, this induction proves the existence of such a full history on the equilibrium path containing any number of engine replacements $j \leq L$ such the posterior upon observing the full history is at least $\Upsilon^j(\mu^0)$, if motorist 0 hires with positive probability. ■

Lemma B.5 shows the existence of some full history $h^{\tilde{t}} \in H^{\tilde{t}}$ at some $\tilde{t} \leq Ln^*$ on the equilibrium path whose full observation yields posterior $\mu(h^{\tilde{t}}) \geq \Upsilon^L(\mu^0) > p^*$ and the mechanic is hired with probability greater than $1 - \lambda^{\tilde{t}(\tilde{t}+1)/2}$, which requires that the motorist hire even if he observes the full history. Yet if he observes the full history, hiring cannot be a best response, a contradiction.