Reputation in Repeated Interaction

Jonathan Levin

May 2006

We now investigate the idea that if there is some uncertainty about strategic motives, a player who plays a game repeatedly may be able to capitalize on it by building a reputation for a certain sort of behavior. As we will see, these effects can be quite powerful in some situations — even a small amount of uncertainty about a player's intentions can allow the player to completely dictate how the game will be played.

1 Examples

Reputation for Quality A firm that produces an experience good faces a series of customers. Each decides whether or not to purchase; upon purchase, the firm decides on high or low quality $q \in \{0, 1\}$. If a consumer doesn't purchase, both firm and customer get 0. If the customer buys, he gets $vq-p$, while the firm gets $p - cq$. We assume that $v > c$, so that high quality is efficient.

Our intuition says that the firm should be able to establish a reputation for high quality by actually producing high quality a few times. However, in the one-shot game or in a finite repetition, backward induction tells us that the firm will always produce low quality, and hence that the customer won't buy the good. And in the infinite horizon game, the folk theorem says that essentially anything can happen.

Chain Store Game For this model, imagine a drug company being sued by patients who claim to have been harmed by a drug. There is a series of potential litigants. If sued, the drug company can either fight of settle. If a litigant doesn't sue, she gets zero, while the company gets 2. If she sues and the company concedes, both get 1. If she sues and the company fights, both get -1.

Again, it seems reasonable that the drug company will fight to establish a reputation for being hard to sue. But in a finite version of this game, it will settle every suit from the start. And in an infinitely repeated version, anything can happen.

Monetary Policy Reputational models are often used to capture central bank behavior. The canonical problem is that a central bank wants agents to believe it will have a tight money supply (in order to keep inflation down), but once agents have set prices, the bank is tempted to raise money supply to boost output. Thus, central bankers need to establish a reputation for being tough on inflation.

Credible Advice. Consultants, political advisors, and others would often like to build a reputation for giving unbiased advice. However, they may have certain biases, or may just be worried about being labeled as biased or extremist. The question is whether such advisors can successfully acquire a reputation for giving unbiased credible advice.

2 Reputation as Commitment

We start with a fairly general model with one long-run player facing a sequence of small opponents. The development here follows Fudenberg and Levine (1989). I focus on finitely-repeated games, but the same argument applies to infinitely repeated games if the long-run player has a sufficiently high discount factor.

Let G be a two-player simultaneous move game. Suppose that players $1, 2$ have action sets A_1, A_2, and sssume that $BR_2(a_1)$ is unique for all $a_1 \in A_1$. We consider a T-period repetition of this game with incomplete information. Prior to play, nature chooses a type $\theta \in \Theta$ for player 1, and then player 1 meets a sequence of player 2, in periods $t = 1, 2, ..., T$. Suppose that Θ consists of:

- A "rational" player 1, who maximizes

$$\frac{1}{T} \sum_{t=1}^{T} u_1(a_1, a_2)$$

- For each $a_1 \in A_1$, a type θ^{a_1}, who plays a_1 in every period.

Assume that Nature chooses $\theta \in \Theta$ with probability $\mu(\theta) > 0$ for all θ.

Proposition 1 *For all $\varepsilon > 0$, there exists $\underline{T}$ such that for all $T \geq \underline{T}$, all perfect bayesian equilibria of the T-period game with incomplete information give the rational player a payoff of at least:*

$$\max_{a_1 \in A_1} u_1 \left(a_1, BR_2(a_1)\right) - \varepsilon.$$

Proof. Define

$$a_1^* = \arg \max_{a_1 \in A_1} u_1(a_1, BR_2(a_1)).$$

Pick $\alpha > 0$ such that if $\sigma(a_1^*) \geq 1 - \alpha$, then $BR_2(\sigma) = BR_2(a_1^*)$. Also, note that following any history h_t, the following identity applies:

$$\mu_2 \left(\theta^{\text{Rat}} | h_t\right) + \mu_2 \left(\theta^{a_1^*} | h_t\right) = 1$$

Now consider some PBE where player 2 doesn't play $BR_2(a_1^*)$ following a history h_t. It must be that $\sigma\left(a_1^* | h_t, \theta = \theta^{\text{rat}}\right) < 1 - \alpha$. Consider the updating that player 2 would do if she saw h_t followed by a_1^*. Using Bayes rule,

$$\frac{\mu_2 \left(\theta^{\text{rat}} | h_{t+1}\right)}{\mu_2 \left(\theta^{a_1^*} | h_{t+1}\right)} < (1 - \alpha) \frac{\mu_2 \left(\theta^{\text{rat}} | h_t\right)}{\mu_2 \left(\theta^{a_1^*} | h_t\right)}$$

Let k be such that

$$\mu \left(\theta^{\text{rat}}\right) (1 - \alpha)^k < \alpha \mu \left(\theta^{a_1^*}\right).$$

If player 1 follows the strategy of playing a_1^* in every period, there can be at most k periods in which player 2 does not play $BR_2(a_1^*)$ because after that point player 2 will assign probability at least $1 - \alpha$ to type $\theta^{a_1^*}$, and then will play $BR_2(a_1^*)$ as a best-response.

Now, pick $\underline{T}$ such that:

$$\left(\underline{T} - k\right) u_1 \left(a_1^*, BR_2(a_1^*)\right) + k \min_{a_2} u_1 \left(a_1^*, a_2\right) > \underline{T} \left[u_1 \left(a_1^*, BR_2(a_1^*)\right) - \varepsilon\right]$$

Thus, if $T > \underline{T}$, player 1 can always ensure a payoff $u_1 \left(a_1^*, BR_2(a_1^*)\right) - \varepsilon$ by deviating, and hence must get at least this in equilibrium. *Q.E.D.*

Note that when the game is played, player 1 does not actually "build" a reputation. That is, as the game goes on, opponents do not put more weight on player 1 being committed. Rather, they know that even if he is rational, player 1 will act as if he is committed. One point to note here is that the when $\Pr(\theta^{a_1^*})$ is small, then we may need to pick $\underline{T}$ very large.

You might think that it is important here for the short-run players to be able to observe perfectly what the long-run player is doing. In fact, that need not be the case. Fudenberg and Levine (1992) show that a version of their "bounds" argument applies even if the long-run player's behavior is imperfectly observed. Basically the idea is that short-run players will be actively learning about the long-run player's behavior over time, with some outcomes making it more likely that the long-run player is a Stackelberg type and some making it less likely. A "normal" type will have an incentive to play non-Stackelberg actions, but Fudenberg and Levine show that as the horizon becomes sufficiently long, the normal types will still mimic the Stackelberg type to get a payoff nearly equal to the Stackelberg payoff. Cripps, Mailath and Samuelson (2003) study the long-run, or asymptotic, pattern of behavior in this model.

3 Many Long-Run Players

Fudenberg and Levine's result shows that if there is a single long-run player and a sequence of short run players, then the possibility of building a reputation completely determines how the game will be played (or at least the payoffs). Fudenberg and Maskin (1986) show that if there are two long-run players and both are potentially "crazy" then anything can happen. Their result extends Kreps et al.'s (1982) famous paper showing that cooperation was possible in a repeated prisoner's dilemma where there was some chance one player was committed to tit-for-tat.

Proposition 2 *Let G be a two-player game. Suppose that $(v_1, v_2) \in V^*$. Then for any $\varepsilon > 0$, there exists $\underline{T}$ and a form of behavior for crazy types, θ_1^c, θ_2^c such that in the T-period game with $\Pr\left(\theta_1 = \theta_1^{rational}\right) = \Pr\left(\theta_2 = \theta_2^{rational}\right) = 1 - \varepsilon$, and $\Pr\left[\theta_1 = \theta_1^c\right] = \Pr\left[\theta_2 = \theta_2^c\right] = \varepsilon$ there exists a PBE in which the rational players' average payoffs are within ε of (v_1, v_2).*

Proof. (with some loss of generality) Suppose that (v_1, v_2) pareto dominate some Nash equilibrium s^* of G with payoffs (e_1, e_2). And assume that there is some (a_1, a_2) such that $u_i(a_1, a_2) = v_i$. Suppose the crazy types play a_i until someone plays something other than (a_1, a_2), and then plays s_i^* ever after. Consider the strategy for rational types:

I. In periods $1, 2, ..., T - \hat{T}$, play a_1 so long as no one deviates.

II. If someone deviated from phase I, play s^* for the rest of the game.

III. If no one deviated from phase I, then from $T - \hat{T} + 1$ on, play some PBE σ^* of the $\hat{T}$-period game, which has the property that s^* is played whenever players 1 and 2 are known to be rational.

We want to show this is a PBE. Clearly, there are no profitable deviations in phase III for rational types since this is a PBE, and similarly in phase II since this is a NE in every period. So consider phase I.

Before $T - \hat{T}$, if rational player i deviates, he gains at most $\overline{v} - v_i$ for one period, but loses $v_i - e_i$ in periods $t+1, ..., T - \hat{T}$, and loses something more in the last $\hat{T}$ periods. How much does he lose in the last $\hat{T}$ periods? In these last $\hat{T}$ periods, if the rational player has not deviated before $T - \hat{T}$, he could follow the strategy: play a_i until his opponent plays something other than a_{-i}, then play something random to reveal his rationality, then play s_i^* until the end. This gives a payoff of at least:

$$(1-\varepsilon)\left(2\underline{v} + (\hat{T}-2)e_i\right) + \varepsilon \hat{T} v_i$$

On the other hand, if rational i deviates in the first $T - \hat{T}$ period, he will get $\hat{T}e_i$ in the last $\hat{T}$ periods. The loss in the last $\hat{T}$ periods from deviating in Phase I is thus at least:

$$\varepsilon \hat{T}(v_i - e_i) - (1-\varepsilon)2(e_i - \underline{v})$$

Taking $\hat{T}$ large enough, we can ensure that this is greater than or equal to $\overline{v} - v_i$. Then rational i will not deviate in phase I. Finally, since $\hat{T}$ is now fixed, we can simply take T large enough to make the payoffs from this equilibrium be within ε of (v_1, v_2). *Q.E.D.*

1. The result says that even in long finite games, essentially "anything can happen" if there is a small chance players aren't rational. It suggests that (i) the backward induction solution to these games may not be "robust" to the introduction of particular types of irrationality, and also that long finite games may behave much like infinitely repeated games.

2. Of course, the full strength of the result depends on being able to find just the right kind of crazy type. The result depends crucially on this. FM suggest that some kinds of craziness may be more reasonable than others.

3.1 Extensions

The two results we have shown have an interesting contrast in that the equilibrium payoffs are sharply constrained if there is just a single long-run player who may be committed, but not constrained at all if there are two long-run players who may be committed. A natural question is whether there are situations where payoffs can be pinned down by reputation even with two long-run players.

- Schmidt (1993) observes that with two long-run players, it is hard to use reputation to pin down payoffs. The reason is that once both players have revealed to be rational, the folk theorem kicks in and anything can happen. The problem spills over to situations where one player has revealed rationality. Schmidt identifies a class of games with conflicting interest payoffs where a Fudenberg-Levine type result applies when one player is much more patient than the other.
- Abreu and Pearce (2003) consider infinitely repeated two-player games, where at the outset players can announce a repeated game strategy and with $\varepsilon \to 0$ probability remain committed to it. They ask whether in such a game, there exists a payoff profile (v_1, v_2) such that if, in all subgames where rationality has been revealed, play yields (v_1, v_2), then (v_1, v_2) will be the payoff that results from the start of play. They show that there is exactly one such payoff vector, which remarkably, coincides with the solution to the Nash demand game with endogenous threats.

4 Bad Reputation

Our analysis suggests that a long-run player interacting with a sequence of a short-run players will typically benefit from reputation effects since he has the option of acting committed to a certain strategy. But this need not be the case if there is imperfect monitoring! We now consider an example of Ely and Valimaki (2003) in which reputation has perverse consequences.

The main character is a mechanic who interacts with a motorist. The motorist's car needs either a tune-up or an engine replacement with equal probability. Denote these possibilities as $\theta \in \{\theta_t, \theta_e\}$. The motorist can't tell which is needed, but if he hires the mechanic, the mechanic will be able to tell. The mechanic will then choose a repair $a \in \{t, e\}$. The motorist's payoff depends on the treatment and the state:

	θ_t	θ_e
t	u	$-w$
e	$-w$	u

We assume that $w > u > 0$ and that the motorist has some outside option that gives a payoff zero.

Let (β_t, β_e) denote the probability hte mechanic will perform the right repair in each of two states, i.e. β_a is the probability of repair a in state θ_a. The motorist's payoff from hiring is then:

$$-w + \frac{1}{2}(\beta_t + \beta_e)(u + w)$$

The motorist can get ensure zero by not hiring, so a *necessary* condition for hiring is that $\beta_t, \beta_e \geq \beta^* = \frac{w-u}{w+u}$.

In the benchmark case, we assume the mechanic is "good" and has the same preferences as the motorist. In the unique sequential equilibrium of the one-shot interaction, the motorist will hire the mechanic and he'll do the right repair. This remains essentially correct even if there is a small probability μ that the mechanic is a "bad" type who always replaces the engine. Given that the good mechanic does the right thing, the motorist's expected payoff is:

$$-w + (1 - \mu)(u + w) + \mu\frac{1}{2}(u + w) = u - \mu\frac{1}{2}(u + w).$$

Thus if $\mu \leq p^* = \frac{2u}{u+w}$, the motorist will hire the mechanic in a one-shot interaction and the good mechanic will do the right thing.

We now investigate the idea that even if motorists assign only a small probability of the mechanic being bad, this can distort the reputational incentives of a good mechanic in such a way that the motorist may not want to hire. To model this, we imagine an infinite sequnce of motorists who decide in turn whether to hire the mechanic after observing earlier repairs (but not what was actually needed).

In this game, a bad mechanic always chooses $a = e$. The good mechanic maximizes his average payoff using discount factor $\delta \in (0, 1)$. His strategy specifies for each date k and history h_k the probabilities $\beta_t^k(h^k)$ and $\beta_e^k(h^k)$ of doing the right repair.

Since each motorist is a short run player, motorist k will want to hire given history h^k if she expects the right repair with sufficient probability. Her decision will be based on the probability $\mu^k(h^k)$ she assigns to the mechanic being bad and the expected behavior β_t^k, β_e^k of the good mechanic. Of course if $\mu^k(h^k) > p^*$, she will certainly choose not to hire, and if $\mu^k(h^k) \leq p^*$ a necessary condition for her to hire is that $\beta_t^k, \beta_e^k \geq \beta^*$.

EV's surprising result is that if the mechanic is sufficiently patient, then he will be unable to realize positive profits over the course of the game even if μ is quite small. Specifically, let $\overline{V}(\mu, \delta)$ denote the supremum of the mechanic's Nash equilibrium average payoffs in the game with discount factor δ and prior μ.

Proposition 3 *For any $\mu > 0$, $\lim_{\delta \to 1} \overline{V}(\mu, \delta) = 0$.*

Proof. To begin, note that if $\mu > p^*$, there is a unique Nash equilibrium in which the mechanic is never hired. The first motorist does not hire, so beliefs are not updated. The second motorist then doesn't hire and so on.

Suppose that $\mu \leq p^*$ amd consider a Nash equilibrium in which the mechanic is hired in the first period. The updated probability of being bad depends on the first period behavior. In particular,

$$\mu^1(t) = 0$$

and

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

Recall that the mechanic will only be hired if $\beta_t \geq \beta^* > 0$. So $\mu^1(e) > \mu > \mu^1(t) = 0$. Note that $\mu^1(e)$ is lower (i.e. the motorist has a better opinion of the mechanic following an engine replacement) when either β_t is low or β_e is high. Letting $\beta_t = \beta^*$ and $\beta_e = 1$, define

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

to be the smallest possible posterior probability of a bad mechanic given an engine replacement and prior μ. We know that $\Upsilon(\mu) > \mu$ for all $\mu \in (0, p^*)$. Also Υ is continuous and strictly increasing in μ.

Now, let $p_1 = p^*$ and define p_m such that $\Upsilon(p_m) = p_{m-1}$. Under this definition, if the prior $\mu \geq p_{m+1}$, and an engine replacement is observed, then the posterior will be *at least* p_m. The sequence $p_1, p_2, ...$ is strictly decreasing and $\lim_{m \to \infty} p_m = 0$. We will now use an induction argument on m to bound the mechanic's Nash equilibrium payoffs as $\delta \to 1$.

We have already seen that if the prior μ exceeds p^*, the mechanic gets zero payoff. For the induction step assume that if the prior μ exceeds p_m, the mechanic's payoff is bounded above by some $\overline{V}_m(\delta)$ with $\lim_{\delta\to 1}\overline{V}_m(\delta) = 0$. To complete the induction argument, assume that the prior μ exceeds p_{m+1} and consider a Nash equilibrium where the mechanic is hired in the first period (this is wlog since if he isn't hired until the second period, the game starting in the second period has the same prior and a higher payoff).

Since in the first period the mechanic must be choosing the correct action with probability at least β^* in each state (otherwise the motorist wouldn't hire him), his payoff is bounded above by his payoff from choosing the correct action with probability 1 in each state (either this is his optimal strategy or he mixes in which case it's one of his best responses). Letting $z(a|\tau)$ denote the continuation payoff if the state is τ and the good mechanic chooses action a:

$$\overline{V}(\mu, \delta) \leq (1-\delta)u + \delta \frac{z(e|e) + z(t|t)}{2} \tag{1}$$

We have assumed that $\mu > p_{m+1}$, so if the mechanic chooses e, then $\mu^1(e) > p_m$ and hence by the induction assumption:

$$z(e|e) \leq \overline{V}_m(\delta).$$

Now, consider the incentive compatibility constraint for the mechanic. He must be willing to choose e when the state is actually e rather than deviating to t. So

$$(1-\delta)u + \delta z(e|e) \geq -(1-\delta)w + \delta z(t|e)$$

or, re-arranging:

$$z(t|e) \leq \frac{1-\delta}{\delta}(u+w) + z(e,e) \tag{2}$$

Combining the inequalities:

$$\overline{V}(\mu, \delta) \leq (1-\delta)\frac{3u+w}{2} + \delta\overline{V}_m(\delta) = \overline{V}_{m+1}(\delta).$$

Since $\lim_{\delta\to 1}\overline{V}_m(\delta) = 0$, then evidently $\lim \overline{V}_{m+1}(\delta) = 0$. By induction, this holds for all m, so we have shown that $\lim_{\delta\to 1}\overline{V}(\mu, \delta) = 0$ for any μ that is greater than some p_m. Since $\inf_m p_m = 0$, the proof is complete. *Q.E.D.*

Intuitively, the problem is that once there is a sufficiently high belief that the mechanic is bad, motorists will stop hiring and the game will effectively end. This means that if beliefs are such that the mechanic is only one engine replacement away from this region, and he cares about future payoff enough, he will be exceedingly averse to doing a replacement today even if one is needed — *unless the continuation payoff from a tune-up is also exceedingly low*. The dilemma is that if the mechanic is not willing to do a replacement, the motorist will anticipate this and refuse to hire him because she only wants to hire if he's willing to do engine replacements when they're needed. Thus the only way the mechanic will be hired when she's one e away from being fired forever is if the continuations from a tune-up are also close to zero. If δ is high, this means the overall expected payoff in this region must be very low. But now, this expands the region of beliefs that the mechanic wants to avoid and creates a new region where he's just one engine replacement away from a bad belief region. Sadly for the mechanic, this unraveling continues until we've shown that his payoffs must be low for any prior belief!

4.1 Comments

- The result doesn't imply the motorist will never get hired. For instance, if $\mu < p^*$, there is a Nash equilibrium in which the first motorist hires the mechanic, but no future motorist ever hires him regardless of what he does with the car. Since future motorists ignore his behavior, the mechanic will do the right thing in the first period. The problem with this equilibrium is that if the second motorist see t, she'll believe with probability 1 that the mechanic is good. So it seems quite unreasonable for her not to hire.

- A key problem here is that each motorist is a short-run player. She does not internalize the benefits she creates for later motorists if she hires and gives the mechanic a chance to signal his goodness by performing a tune-up. EV show that if there is a single long-run motorist, there is an equilibrium that essentially gets the first-best outcome, even if the bad mechanic may be a strategic player who can try to imitate a good mechanic rather than automatically choosing e each period.

References

[1] Abreu, D. and D. Pearce (2003) "Bargaining, Reputation and Equilibrium Selection in Repeated Games," Working Paper.

[2] Cripps, M., G. Mailath and L. Samuelson (2003) "Imperfecct Monitoring and Impermanent Reputations," *Econometrica*.

[3] Ely, J. and J. Valimaki (2003) "Bad Reputation," *Quarterly J. Economics.*

[4] Ely, J., D. Fudenberg and D. Levine (2003) "When is Reputation Bad?" Working Paper.

[5] Fudenberg, D. and D. Levine (1989) "Reputation and Equilibrium Selection in Repeated Games," *Econometrica.*

[6] Fudenberg, D. and D. Levine (1992) "Maintaining a Reputation when Strategies are Imperfectly Observed," *Rev. Econ. Studies.*

[7] Fudenberg, D. and E. Maskin (1986) "The Folk Theorem with Discounting or with Incomplete Information," *Econometrica.*

[8] Kreps, D., P. Milgrom, J. Roberts and R. Wilson (1982) "Cooperation in the Finitely Repeated Prisoners' Dilemma," *J. Econ. Theory.*

[9] Schmidt, K. (1993) "Reputation and Equilibrium Characterization in Repeated Games with Conflicting Interests," *Econometrica.*