Reputations

Larry Samuelson

Yale

January 30-31, 2012

I. Introduction

I.1 An Example: The Chain Store Game

Consider the chain-store game:

	Out	In
$Acquiesce$	5, 0	2, 2
$Fight$	5, 0	−1, −1

If played once, this game has a unique Nash equilibrium, $(Acquiesce, In)$.

What if the game is played (finitely) many times?

- One's intuition is that player 1 will fight early entry, in order to deter later entrants.

- However, the (finitely repeated) game has a unique subgame perfect equilibrium, in which $(Acquiesce, In)$ is played in every period. This is a simple backward-induction argument.

- This is known as the "chain store paradox."

- Similarly, it is intuitive that players might cooperate in the early rounds of a finitely-repeated prisoners' dilemma, but here the only *Nash* equilibrium is that both players always defect.

Let's add some incomplete information:

– Now suppose that with probability $1 - \mu_0^*$, player 1 is the "normal" type.

– With probability μ_0^* (small), player 1 is a "commitment" type who always fights entry.

– Player 1's type is 1's private information.

– We can interpret the commitment type as having different payoffs, that make it optimal to fight, or as a type who is simply "committed" to fighting.

Some things that are no longer equilibrium outcomes:

- It is not an equilibrium for the normal type to acquiesce in every period.
- It is also not an equilibrium for the normal type to fight in every period. For example, the normal type will acquiesce in the last period.

The sequential equilibrium has the following properties (Kreps, Milgrom, Roberts and Wilson (1982)):

– There is an initial phase in which $(Fight, Out)$ is played, and 2's beliefs remain unchanged at the prior.

– There is a terminal phase in which 2 mixes between In and Out, and 1 mixes between $Acquiesce$ and $Fight$.

– If $(Acquiesce, In)$ occurs, then 1 is revealed to be normal, and $(Acquiesce, In)$ is played in every subsequent period.

– If Out is played, 2's beliefs remain unchanged. If $(Fight, In)$ is played, 2 revises upward the posterior probability that 1 is the commitment type.

- In the final period, player 2 mixes while player 1 plays $Acquiesce$.

As the number of repetitions grows:

- The length of the terminal phase remains fixed. This is determined by backward-induction calculations that do not depend on the length of the game.

- The initial phase grows long, consuming virtually all of the game.

- We thus see $(Fight, Out)$ most of the time.

I.2 The Reputation Literature

The idea of "reputation" has been modeled many different ways in the literature. Examples include:

- Models of expert advice, such as Morris (2001).
- Models of career concerns, such as Holmström (1982,1999).
- Models of bargaining, such as Abreu and Gul (2000).
- ⋮
- Models based on repeated games.

I.3 Reputations in Repeated Games

The idea of reputation is applied to repeated games in two ways:

– As an interpretation of equilibria in repeated games of complete information.

– By incorporating incomplete information.

I.4 The Objective

Reputation ideas were originally applied to finitely-repeated games, such as the chain-store game or the prisoners' dilemma, as a way of expanding the set of equilibrium outcomes to avoid seemingly counterintuitive implications.

More recently, reputation ideas have been applied to infinitely-repeated games, as a way of putting bounds on equilibrium payoffs.

The techniques are much the same.

This is commonly interpreted as a robustness exercise, but this interpretation must be adopted with care.

II. Reputations in Repeated Games of Perfect Monitoring

II.1 The Model

Players include:

- The normal player 1, a long-run, optimizing player.
- One or more commitment types of player 1, mechanical players.
- Player 2, a short-run player.

II.2 Stackelberg Types and Actions

Define:

$$
\begin{aligned}
v_1^* &= \max_{a_1 \in A_1} \min_{\alpha_2 \in BR(a_1)} u_1(a_1, \alpha_2) \\
a_1^* &\in \arg\max_{a_1 \in A_1} \min_{\alpha_2 \in BR(a_1)} u_1(a_1, \alpha_2) \\
v_1^{**} &= \sup_{\alpha_1 \in \Delta A_1} \min_{\alpha_2 \in BR(\alpha_1)} u_1(\alpha_1, \alpha_2).
\end{aligned}
$$

For example, consider the product-choice game:

	h	ℓ
H	$2,3$	$0,2$
L	$3,0$	$1,1$

II.3 The Basic Reputation Result

Proposition 1 *Let* A_2 *be finite and let* μ_0 *attach positive probability to both the normal type and the pure Stackelberg type. Then there exists a (parameter dependent)* k *such that in every Nash equilibrium, player 1's payoff is at least*

$$\delta^k v_1^* + (1-\delta^k)\min_a u_1(a).$$

Fudenberg and Levine (1989).

We'll first interpret this result, and then sketch the proof.

II.4 Interpretation

- The implication is that a patient player 1 gets arbitrarily close to the Stackelberg payoff. A chance of being the Stackelberg player is just as good as being that player.

- It is familiar that results in repeated games require patience. However, in the folk theorem, patience is important in strengthening incentives, while here is is important in reducing the cost of reputation building.

- Despite this interpretation of the role of patience, this result and its proof says nothing about equilibrium behavior.

- A_2 can be infinite, at the cost of a more cumbersome argument.

– We can also accommodate multiple short-run players. We simple replace "best response" with "Nash equilibrium in the induced short-run player game" in the Stackelberg definitions.

II.5 Should We Expect Stackelberg Types?

This result is sometimes criticized for requiring the fortuitous appearance of (only) the Stackelberg type.

The commitment type was considered intuitive in the chain store game, but this may not always be the case.

- The result still holds if the commitment type is not the Stackelberg type, but it may give a lower payoff bound. Whatever commitment type is there, player 1 can do as well as being known to be that commitment type.
- There may be many commitment types, such as countably many. In this case, player 1 "gets to choose his favorite type."

II.6 The Argument

There are three parts to the proof:

– A characterization of player 2's beliefs on a useful set of states.

– A characterization of player 2's behavior, given these beliefs.

– The combination of these two elements to get the proof.

II.7 Beliefs

Some notation:

Ω	Set of states
Ω'	States in which a_1^* is chose in every period
Ω''	States in Ω' whose histories all have positive probability
μ_0^*	Prior probability of Stackelberg type
q_t	Probability player 2 attaches to a_1^* in period t
η_z	Number of periods $q_t < z$.

The belief lemma:

Lemma 1

$$\mathbb{P}\left\{\eta_z > \frac{\ln \mu_0^*}{\ln z} | \Omega'\right\} = 0.$$

The argument: On Ω'',

$$\mathbb{P}(\text{Stackelberg type}|h_{t+1}) = \frac{P(\text{Stackelberg type}|h_t)}{q_t}$$

This follows immediately from Bayes' rule if only 1's actions are observed. A straightforward technical argument extends to the case in which 2's actions are observed as well.

Then we note that

- Either q_t is large or the probability attached to the Stackelberg type increases.

- A high probability of the Stackelberg type implies that q_t is large.

II.8 Best responses

The best response lemma:

Lemma 2 *There exists* $z^* \in (0, 1)$ *such that*

$$\alpha_1(a_1^*) > z^* \Rightarrow BR(\alpha_1) \subset BR(a_1^*).$$

This is obvious when A_2 is finite, and otherwise requires a continuity argument.

II.9 The Reputation Result

- Fix $z > z^*$. Let $k = \frac{\ln \mu_0^*}{\ln z^*}$.
- Let the normal player 1 play a_1^* in each period. Then in all but k periods, 2 plays a best response to a_1^*.
- Then 1's payoff is at least

$$\delta^k v_1^* + (1-\delta^k) \min_a u_1(a).$$

III. Reputations in Repeated Games of Imperfect Monitoring

We are interested in imperfect monitoring for two reasons:

- Monitoring may often be imperfect.
- We can then consider mixed commitment types in perfect monitoring games.

The arguments are much the same, but require two changes:

- We must revisit the Stackelberg payoff notion.
- We need a new belief lemma.

III.1 The Model

Player i now has an action set A_i and a signal set Z_i.

Special cases:

- Perfect monitoring, mixed commitment type.
- Public signals, player 2 cannot see previous player-2 actions.
- Public signals, public short-run actions.
- Private signals.

III.2 Stackelberg Payoffs

Let α_1^* be the commitment action.

The bound on player 1's payoff will be something like

$$\inf_{\alpha_2 \in BR(\alpha_1^*)} u_1(\alpha_1^*, \alpha_2).$$

However, we need to be more precise about "$BR(\alpha_1^*)$".

Let

$$BR_\varepsilon(\alpha_1) = \{\alpha_2 | \exists \alpha_1' \ s.t. \ \alpha_2 \in BR(\alpha_1'), |\rho(\cdot|\alpha_1, \alpha_2) - \rho(\cdot|\alpha_1', \alpha_2)| < \varepsilon\}.$$

$$BR_\varepsilon^*(\alpha_1) = \{\alpha_2 : supp(\alpha_2) \subset BR_\varepsilon(\alpha_1)\}$$

III.3 The Reputation Result

Proposition 2 *For any $\varepsilon > 0$, there exists k such that in any Nash equilibrium, player 1's payoff is at least*

$$(1-\varepsilon)\delta^k \inf_{\alpha_2 \in B^*_\varepsilon(\alpha_1^*)} u_1(\alpha_1^*, \alpha_2) + (1-(1-\varepsilon)\delta^k) \min_{a\in A} u_1(a).$$

Fudenberg and Levine (1992).

An example. Consider the product choice game.

	h	ℓ
H	$2,3$	$0,2$
L	$3,0$	$1,1$

Suppose h and ℓ are public, but H and L give rise to public signals $\overline{y}$ and $\underline{y}$ with probabilities $(p, 1-p)$ and $(q, 1-q)$ $(p > q)$.

If we set $\alpha_1^* = \frac{1}{2}H + \frac{1}{2}L$, then for all ε, $BR_\varepsilon(\alpha_1^*) = \Delta A_2$, and we have a lower bound on player 1's payoff of $\frac{1}{2}$.

If we shade the commitment type just slightly toward H, then $BR_\varepsilon(\alpha_1^*) = \{L\}$, and we get a bound close to $\frac{3}{2}$.

This exceeds the public-monitoring complete information payoff.

III.4 Sorin's Belief Lemma

Let Ω be a Borel space with filtration $\mathcal{F}_t$. Let $\hat{\mathbb{P}}$ and $\tilde{\mathbb{P}}$ be measures on Ω, and let $\mathbb{P}$ be their (nontrivial) mixture $\mu\hat{\mathbb{P}}+(1-\mu)\tilde{\mathbb{P}}$. Then for

Lemma 3 *For any ε and $\psi > 0$, there exists k such that*

$$\hat{\mathbb{P}}\left\{|\{t : d_t(\mathbb{P},\hat{\mathbb{P}}) \geq \psi\}| > k\right\} \leq \varepsilon,$$

where

$$d_t(\mathbb{P},\hat{\mathbb{P}}) = \sup_{A\in\mathcal{F}_{t+1}} |\mathbb{P}(A|\mathcal{F}_t - \hat{\mathbb{P}}(a|\mathcal{F}_t)|.$$

IV. Temporary Reputations

IV.1 The Model

- Full support public or private monitoring.
- Player 1 is either normal or commitment type.
- Player 2 has unique best response to commitment action, that does not give Nash equilibrium of the stage game.
- For every player-1 pure action, the signal distributions induced by player 2's pure actions are linearly independent.

IV.2 The Temporary Reputations Result

Proposition 3

$$\mathbb{P}(\textit{Commitment type}|\mathcal{F}_{2t}) \to 0 \; \tilde{\mathbb{P}} \; almost \; surely.$$

Cripps, Mailath and Samuelson (2004,2007).

IV.3 The Intuition

- Suppose not. Player 2's beliefs are a martingale, and so must converge to an interior belief.

- 2's beliefs can converge to something in the interior only if play converges to α_1^*.

- Then 2 will play a best response to α_1^*.

- Then the normal player 1 will deviate from α_1^*, a contradicting 2's beliefs.

This is not a limiting result in δ, but the limits must be interpreted carefully.

V. Discussion

There are many extensions we might consider.

What are reputation models missing?

VI. Two Long-Run Players

VI.1 The Model

Suppose player 2 is also a long-run player. There will be two sources of asymmetry in the model:

- There is incomplete information about player 1's type.
- We will fix δ_2 and let $\delta_1 \to 1$.

We should in general expect to work with some asymmetry.

Assume monitoring is perfect.

VI.2 A First Intuition

Then it seems we should expect our previous results to go through in a pretty straightforward way. In particular, if player 1 plays a_1 in each period, we again have

$$\mathbb{P}\left\{\eta_z > \frac{\ln \mu_0^*}{\ln z} | \Omega'\right\} = 0.$$

In particular, this result depends only on belief updating, and is unaffected by whether player 2 is short-run or long-run.

Then we need only check that player 2 will play a best response.

VI.3 The Difficulty

Consider the game:

	L	C	R
T	$10, 10$	$0, 0$	$-z, 9$
B	$0, 0$	$1, 1$	$1, 0$

The (pure and mixed) Stackelberg action is T, with payoff 10.

Types for player 1:

- Normal, prior probability .8, in equilibrium plays trigger strategy that begins with T and plays T until first observing B, then switches to B.

- Stackelberg, prior probability .1, always plays T.

- Punishment, prior probability .1, plays T unless some out-of-equilibrium action observed, thereafter plays B.

Player 2's behavior:

– Alternate between L and R, as long as 1 plays T. Should 1 ever play B, conclude 1 is the punishment type and thereafter play C.

– Suppose 2 has deviated. Then 2 plays L and if 1 plays T, conclude 1 is not the punishment type, and play L. If one plays B, conclude 1 is the punishment type and play C.

Player 1's payoff is $\frac{10-z}{2}$, can be made arbitrarily close to his pure minmax payoff of 1.

Player 2's payoff is 9.5.

Is this an equilibrium?

- Consider player 2. 2's equilibrium payoff is 9.5. A deviation on the part of player 2 gives 2 a payoff of 10 if 1 is not the punishment type (probability .9), and 1 if player one is the punishment type (probability 01), for an expected payoff of 9.1.

- If 1 deviates, player 2's actions are clearly optimal, given her beliefs.

- What about player 1? If 1 deviates, he gets 1, which is suboptimal.

- Sequential rationality? If 1 deviates, subsequent play is a stage-game Nash equilibrium, which is OK. If 2 deviates, the normal player 1 gets 10, the maximum possible, and hence must be OK.

Comments:

- This example is adapted from Schmidt (1993). We can make all of player 1's types payoff types.
- Celantani, Fudenberg, Pesendorfer and Levine (1996) construct a similar example without a punishment type.

The key to this result is that player 2 learns to expect the Stackelberg action, but cannot be sure that this is the Stackelberg type. This is irrelevant for a short-run player 2, but the possibility of the punishment type can terrify a long-run player 2.

VI.4 Conflicting Interests

The stage game has conflicting interests if the pure Stackelberg action a_1^* mixed minmaxes player 2.

Examples:

	C	D
C	$2, 2$	$-1, 3$
D	$3, -1$	$0, 0$

This is the prisoners' dilemma. $a_1^* = D$, and so we have conflicting interests.

	h	ℓ
H	2, 3	0, 2
L	3, 0	1, 1

In the product choice game, $a_1^* = H$, and we do not have conflicting interests.

	Out	In
$Acquiesce$	5, 0	2, 2
$Fight$	5, 0	−1, −1

In the chain store game, $a_1^* = F$, and we have conflicting interests.

Consider a version of the ultimatum game:

- Simultaneous moves, proposal from player 1 (of how much goes to 2) and reservation value for 2.

- Finite set of offers, surplus of size 1.

- If offers are $\{0, 1/n, 2/n, \ldots, 1\}$, then $a_1^* = 1/n$, and we do not have conflicting interests.

- If offers are $\{1/n, 2/n, \ldots, 1\}$, then $a_1^* = 1/n$, and we do have conflicting interests.

	L	C	R
T	$3,2$	$0,1$	$0,1$
B	$0,-1$	$2,0$	$0,-1$

Here, $a_1^* = T$, which does not minmax player 2. Notice, however, that B does minmax player 2.

VI.5 The Reputation Result

Proposition 4 *Suppose there exists a positive-probability commitment type with action a_1' that mixed minmaxes player 2. Then in any Nash equilibrium, player 1's equilibrium payoff is at least*

$$\delta_1^k \min_{\alpha_2 \in BR(a_1')} u_1(a_1', \alpha_2) + (1 - \delta^k) \min_a u_1(a),$$

where k depends on δ_2 but not δ_1.

Schmidt (1993).

Comments:

- If the game has conflicting interests, a sufficiently patient player 1 gets very close to his Stackelberg payoff.

- Even if the game does not have conflicting interests, reputation may be valuable for player 1. Consider:

	L	C	R
T	$3,2$	$0,1$	$0,1$
B	$0,-1$	$2,0$	$0,-1$

VI.6 The Argument

The intuition:

- In our opening example, 2 feared best responding to the commitment action for fear of being punished.

- Now we have a case in which a best response to the commitment action minmaxes 2. No punishment can do worse than minmax 2. As a result, nothing can deter 2 from best-responding to commitment action.

The steps in the proof:

The intuition:

- Let Ω' be the outcomes with positive-probability histories that play the Stackelberg action in every period.

- Consider histories h_t consistent with Ω'. Show that there exists ε and L such that if player 2 expects less than her minmax payoff at h_t *conditional on* Ω', then there must exist a period with the next L periods at which player 2 attaches less than $1 - \varepsilon$ probability to the commitment type.

- Hence, for every L periods in which player 2 does not best respond to the commitment action, there must be a period in which the commitment action is expected with probability less than $1 - \varepsilon$.

– From our previous belief lemma, on Ω', there is a bounded number of times player 2 can expect t;he commitment action with probability less than $1-\varepsilon$, and hence a bounded number of times player 2 can not play a best response to the commitment action.

VI.7 Extensions

There are many extensions and variations on this basic result.

- Cripps, Schmidt and and Thomas (1996) show that one can get somewhat weaker reputation results for actions that do not minmax player 2. Intuitively, if player 1 is possibly committed to a_1', then player 1 can get a payoff at least

  $$\min_{\alpha_2 \in D(a_1')} u_1(a_1', \alpha_2),$$

  where

  $$D(a_1') = \{\alpha_2 | u_2(a_1', \alpha_2) \geq \underline{v}_2\}.$$

For an example, consider the battle of the sexes:

	L	R
T	$0,0$	$3,1$
B	$1,3$	$0,0$

The Stackelberg action is T, which does not minmax player 2, so we do not have conflicting interests. $D(T)$ is the set of actions that put probability at least $3/4$ on R, and so player 1 can get a payoff at least $9/4$.

– Celantani, Fudenberg, Pesendorfer and Levine (1996) extend this result to imperfect monitoring, in the process obtaining a quite high payoff bound.

– Evans the Thomas (1997) obtain a similar bound using "punishment" commitment types.

- Cripps, Mailath and Samuelson (2004,2007) once again establish a temporary-reputations result.

- We can establish some results for equal discount factors, but these are quite limited. For example, consider the game:

	L	R
T	$1,1$	$0,0$
B	$0,0$	$0,0$

  Cripps and Thomas (1997) show that if player 1 is either normal or a commitment type playing T, there are equilibria with payoffs arbitrarily close to $(0,0)$.

VII. Reputation as Separation

VII.1 The Model

The goal is a model that

- allows us to characterize behavior, and
- allows us to examine reputations as assets.

The key will be to build a model of reputation based on separation and limited coordination.

The players and actions:

- Long-lived player 1.
- Continuum of short-lived, player 2's with idiosyncratic, private signals.
- Player 1 can be normal (probability μ_0) or inept (probability $1-\mu^0$).
- The normal player 1 chooses from $\{L, H\}$. Inept player one always chooses L.
- L is free, H costs c.
- There are two signals $\overline{z}$ and $\underline{z}$. H gives $\overline{z}$ with probability ρ_H, L gives $\overline{z}$ with probability $\rho_L < \rho_H$.

Expectations:

- F is a cumulative distribution of player 2's expectation of high effort.
- Player 1's payoff is $p(F)$ minus cost (if any), where p is continuous (weak convergence) and increasing (first-order stochastic dominance).
- One example: perfect price discrimination.

In each period, the existing player 1 is replaced with probability λ, with a new player whose type is drawn anew.

VII.2 Equilibrium

There always exists an L equilibrium.

Proposition 5 *Fix $\lambda > 0$. Then for sufficiently small c, there exists a high-effort equilibrium.*

The intuition is that the normal type chooses H to convince consumers his is not the inept type.

VII.3 The Role of Replacements

Replacements are important in this result:

Proposition 6 *If $\lambda = 0$, then the is a unique pure sequential equilibrium, in which the normal type always plays L.*

The ides is that without replacements, too many good signals do "too good" a job of convincing consumers, in the process destroying incentives.

Holmström (1982,1999) studies a related phenomenon, in that incentives deteriorate with success, but there is always updating in his model, even with low effort, and his is a model of symmetric information.

VII.5 The Role of Idiosyncratic Consumers

Suppose there is only one consumer, willing to pay a price equal to the probability of high effort?

To what extent are Markov equilibria of this common consumer model analogous to the equilibria of the idiosyncratic consumer model?

Proposition 7 *Fix $\lambda > 0$. Then there exists a sufficiently small c such that there exists a Markov equilibrium of the common-consumer model with high effort.*

Proposition 8 *Let $\lambda = 0$ and suppose $\rho_H = 1 - \rho_L$. Then there is unique pure Markov equilibrium, in which normal types always exert low effort*

VII.6 Continuity

This suggests a parallel between idiosyncratic consumers and common consumers, given that we consider Markov equilibria in the latter. But, the assumption $\rho_H = 1 - \rho_L$ is a bit troubling. Let us explore that.

Proposition 9 *Let $\lambda = 0$ and suppose there are no integers m and n with $(1-\rho_H)^m \rho_H^n = (1-\rho_L)^n \rho_L^m$. Then*

- *There exists a pure Markov equilibrium in which the normal type initially plays H with probability 1.*
- *Once a bad signal appears, L is played thereafter.*
- *H lasts for an exponentially distributed number of periods.*

– $0 = \liminf \alpha(\mu) < \limsup \alpha(\mu) = 1$.

Proposition 10 *Let* $\rho_H = 1 - \rho_L$ *and* $\lambda = 0$. *Let*

$$\frac{\rho_H + c(1 - \delta\rho_H)}{\delta(1 - 2\rho_H)} < 1.$$

Then there exists a mixed equilibrium in which the normal type initially exerts effort with high probability. In this equilibrium, $0 = \liminf \alpha(\mu) < \limsup \alpha(\mu) = 1$.

Hence, we think of Markov equilibria coupled with a continuity requirement as capturing the essence of idiosyncratic consumers in the common consumer model.

VII.7 Discrete Actions

Can we simplify further? Consider the game:

	h	ℓ
H	$3-c, 3$	$1-c, 2$
L	$3, 0$	$1, 1$

Proposition 11 *Let $\lambda > 0$. Then for sufficiently small c, there exists a pure Markov equilibrium in which the normal firm exerts high effort, and consumers choose h, if and only the posterior probability of the normal type exceeds $1/2$.*

VII.8 Lost Consumers

What if consumers abandon the firm when they get too pessimistic? Consider the game:

	d	b
H	$0,0$	$2-c,1$
L	$0,0$	$2,-1$

Proposition 12 *If $\lambda = 0$, then ever pure Markov equilibrium features low effort. If $\lambda > 0$, then for small λ and c, there exists a pure Markkov equilibrium in which (H,b) is played for posteriors exceeding $1/2$ and (L,d) is played otherwise.*

VII.9 Bad Reputations

Consider the "bad reputation" game of Ely and Välimäki (2003).

A consumer decides whether to hire a firm. The firm (privately) observes the state and then decides whether to provide H or L service.

The consumer cannot observe what service is provided.

Payoffs:

If the state is θ_H, payoffs are:

	$hire$	not
H	v, v	0, 0
L	$-w, -w$	0, 0

If the state is θ_L, payoffs are

	$hire$	not
H	$-w, -w$	0, 0
L	v, v	0, 0

States are equally likely and $w > v$.

Equilibrium:

There is a unique equilibrium in the stage game, with payoffs (v, v).

The repeated game of complete information also has an equilibrium with payoffs (v, v).

Suppose there is incomplete information. The firm is likely to be normal, but with some probability is a bad type that always chooses H.

Proposition 13 *As $\delta \to 1$, the largest equilibrium payoff for the firm converges to zero.*

- The consumer will hire the firm if and only if the probability the firm is normal is at least some threshold μ^*.

- Consider a posterior just above μ^*. Then the next H signal pushes the firm into the "never hire" region.

- As a result, the firm will choose L no matter what the state. But the consumer will not hire such a firm, a contradiction.

VIII. Discussion