Reputations: A Very Basic Introduction

Niccolò Lomys*

April 5, 2019

Preamble

These notes heavily draw upon Mailath (2019), Mailath and Samuelson (2015), and Mailath and Samuelson (2006), to which I refer if you want (or need) to learn more about reputations. All errors are my own. Please bring any error, including typos, to my attention.

1 Introduction

The word "reputation" appears throughout discussions of everyday interactions. Firms are said to have reputations for providing good service, professionals for working hard, people for being honest, newspapers for being unbiased, governments for being free from corruption, and so on. Reputations establish links between past behavior and expectations of future behavior—one expects good service because good service has been provided in the past, or expects fair treatment because one has been treated fairly in the past. These reputation effects are so familiar as to be taken for granted. One is instinctively skeptical of a watch offered for sale by a stranger on a subway platform, but more confident of a special deal on a watch from an established jeweler. Firms proudly advertise that they are fixtures in their communities, while few customers would be attracted by a slogan of "here today, gone tomorrow".

Repeated games allow for a clean description of both the myopic incentives that agents have to behave opportunistically and, via appropriate specifications of future rewards and punishments, the incentives that deter opportunistic behavior. As a consequence, strategic interactions within long-run relationships have often been studied using repeated games. For the same reason, the study of reputations has been particularly fruitful in the context of repeated games, the topic of these notes.

1.1 Adverse Selection Approach to Reputations

The adverse selection approach to reputations[1] considers games of incomplete information. The motivation typically stems from a game of complete information in which the players are "nor-

*Toulouse School of Economics, University of Toulouse Capitole; niccolo.lomys@tse-fr.eu.

[1] As opposed to the interpretative approach to reputations, see Mailath and Samuelson (2015).

mal", and the game of incomplete information is viewed as a perturbation of the complete information game. Attention is typically focused on games of "nearly" complete information, in the sense that a player whose type is unknown is very likely (but not quite certain) to be a normal type. For example, a player in a repeated game might be almost certain to have stage-game payoffs given by the prisoners' dilemma, but may with some small possibility have no other option than to play tit-for-tat. Again, consistent with the perturbation motivation, it is desirable that the set of alternative types be not unduly constrained.

The idea that a player has an incentive to build, maintain, or milk his reputation is captured by the incentive that player has to manipulate the beliefs of other players about his type. The updating of these beliefs establishes links between past behavior and expectations of future behavior. We say "reputations effects" arise if these links give rise to restrictions on equilibrium payoffs or behavior that do not arise in the underlying game of complete information. The basic goal is to identify circumstances in which reputation effects necessarily arise, imposing bounds on equilibrium payoffs that are in many cases quite striking.

2 Canonical Reputation Model: A Basic Introduction

In these notes, we study a sequential entry game which is a (simplified) version of the reputation model of Kreps and Wilson (1982) and Milgrom and Roberts (1982). A similar analysis was used by the same authors (see Kreps, Milgrom, Roberts and Wilson (1982)) to demonstrate that cooperation can be sustained in the finitely repeated prisoner's dilemma by introducing reputation effects through incomplete information.

2.1 Two Periods

Consider the entry game whose extensive form is in Figure 1. There is an incumbent and a potential entrant. First, the potential entrant decides whether to enter the market (play In) or not (play Out); if the entrant plays In, the incumbent decides whether to accommodate entry, or to fight. Payoffs are as in Figure 1. The game has two Nash equilibria: (In, Accommodate) and (Out, Fight). The latter violates backward induction, so (In, Accommodate) is the unique SPNE.

Consider next the sequential-entry version of the game, known as the *chain store game*. Now the game in Figure 1 is played twice, against two different entrants (E_1 and E_2), with the second entrant E_2 observing the outcome of the first interaction. The incumbent's payoff is the sum of payoffs in the two interactions. In this case, we have the *chain store paradox*: the only SPNE outcome is that both entrants enter, and the incumbent always accommodates. This is true for any finite chain store.[2]

Finally, suppose the incumbent is a long-lived player, playing the game at times $i = 0, 1, 2, \ldots$ with discount factor $\delta \in (0, 1)$. The incumbent faces a successions of short-lived entrants, with a new entrant in each period. The repeated game has a SPNE featuring (In, Accommodate) in every period. However, If δ is sufficiently large, then there are (many) other equilibria. Indeed,

[2]The chain store paradox originates from Selten (1978).

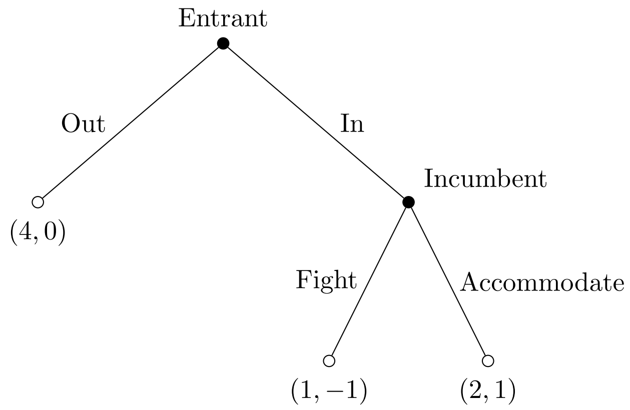

Figure 1: The stage game for the chain store. The first payoff is that of the incumbent, and the second is that of the entrant.

for sufficiently large δ, every payoff in the interval $[2, 4]$ is a SPNE payoff for the incumbent (more on this in Section 2.2). A payoff close to 4 for the incumbent seems intuitive—by fighting entry, the incumbent develops a *reputation* for being "tough", and entrants (at least eventually) stay out. However, there is nothing in the structure of the repeated game that captures this intuition.

Introducing Incomplete Information. Now we introduce incomplete information of a very particular kind. We suppose the incumbent could be *tough*, i.e. of type ω_t. The tough incumbent receives a payoff of 2 from fighting and only 1 from accommodating. The other incumbent is *normal*, i.e. of type ω_n, with payoffs as described in Figure 1. The entry game is still played twice, against two different entrants E_1 and E_2. Both entrants assign prior $\rho \in (0, 1/2)$ to the incumbent being ω_t.

Hereafter, for $i = 1, 2$, O_i stands for Out in period i, I_i stands for In in period i, F_i stands for Fight in period i, and A_i stands for Accommodate in period i.

Suppose first E_1 chooses O_1. Then, in any sequential equilibrium, the analysis of the second period is just that of the static game of incomplete information with E_2's beliefs on the incumbent given by the prior[3], and so E_2 optimally plays I_2, the normal type accommodates and the tough type fights.

We next analyze the behavior that follows E_1's choice of I_1, i.e. entry in the first market. Because in any sequential equilibrium, in the second period the normal incumbent accommodates and the tough incumbent fights, this behavior must be a PBE of the signaling game illustrated in Figure 2 (given the optimal play of the incumbent in the second period).

It is easy to verify that there are no pure strategy Nash equilibrium (make sure you understand this). There is, instead, a unique mixed strategy PBE (this is an example of hybrid PBE): type ω_n plays F_1 with probability α and A_1 with probability $1 - \alpha$, type ω_t plays F_1 for sure. Entrant E_2 enters for sure after A_1 and plays I_2 with probability β and O_2 with probability $1-\beta$ after F_1.

[3]This is so under the assumption that the extensive form of the two-period chain store is specified as first E_1 chooses I_1 or O_1, with each choice leading to a move of nature which determines the type of the incumbent.

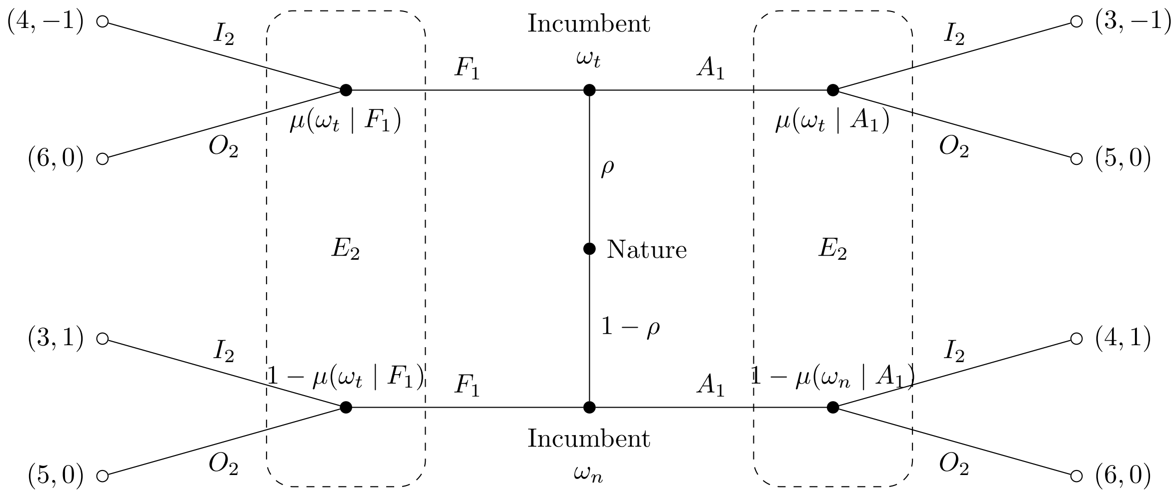

Figure 2: A signaling game representation of the subgame reached by E_1 playing I_1. The first payoff is that of the incumbent, and the second payoff is that of E_2 *assuming* the incumbent plays F_2 after a choice of I_2 by E_2.

Entrant E_2 is willing to randomize only if his posterior after F_1 that the incumbent is ω_t equals 1/2. Since that posterior is given by (applying Bayes' rule)

$$\begin{aligned}\mu(\omega_t \mid F_1) &= \frac{\mathbb{P}(F_1 \mid \omega_t)\mathbb{P}(\omega_t)}{\mathbb{P}(F_1 \mid \omega_t)\mathbb{P}(\omega_t) + \mathbb{P}(F_1 \mid \omega_n)\mathbb{P}(\omega_n)} \\ &= \frac{\rho}{\rho + (1-\rho)\alpha},\end{aligned}$$

solving

$$\frac{\rho}{\rho + (1-\rho)\alpha} = \frac{1}{2}$$

for α gives

$$\alpha = \frac{\rho}{1-\rho},$$

where $\alpha < 1$ since $\rho < 1/2$.

Type ω_n is willing to randomize if

$$4 = 3\beta + 5(1-\beta),$$

where the left-hand side is the payoff from playing A_1 and the right-hand side is the payoff from playing F_1; this gives

$$\beta = \frac{1}{2}.$$

It remains to determine the behavior of entrant E_1. This entrant faces a probability of F_1 given by

$$\rho + (1-\rho)\alpha = 2\rho.$$

Hence, if $\rho < 1/4$, E_1 faces F_1 with sufficiently small probability that he enters. However, if $\rho \in (1/4, 1/2)$, E_1 faces F_1 with sufficiently high probability that he stays out. For $\rho = 1/4$, E_1 is indifferent between O_1 and I_1, and so any specification of behavior is consistent with equilibrium.

Suppose $\rho < 1/4$, so that E_1 enters. This simple example shows the following. From the static viewpoint, the normal type would want to accommodate on the first period; however, by fighting (with some probability), he may convince E_2 that he is the tough type (i.e. the normal type *builds a reputation for being tough*), and thus convince E_2 to stay out and increase his payoff in the second period. Hence, accommodating in every period is no longer an equilibrium outcome even if the entry game is only repeated twice.

2.2 Infinite Horizon

Suppose now the time horizon is infinite with the incumbent discounting the future at rate $\delta \in (0, 1)$ and a new potential entrant in each period.

In the complete information game, the outcome in which all entrants enter and the incumbent accommodates in every period is an equilibrium. Moreover, the profile in which all entrants stay out and any entry is met with F is a SPNE, supported by the "threat" that play switches to the always-enter/always-accommodate equilibrium if the incumbent ever responds with A, provided δ is sufficiently high. In particular, the relevant incentive constraint for the incumbent is conditional on I (since the incumbent does not make a decision when the entrant chooses O), and takes the form

$$(1 - \delta) + 4\delta \geq 2 \Longleftrightarrow \delta \geq \frac{1}{3}.$$

Note that stage game is not a simultaneous move game, and so the repeated game does not have perfect monitoring. In particular, the incumbent's choice between F and A is irrelevant (not observed) if the entrant plays O (as the putative equilibrium requires). Subgame perfection, however, requires that the incumbent's choice of F be optimal, given that the entrant had played I. The one-shot deviation principle applies here: the profile is subgame perfect if, conditional on I, it is optimal for the incumbent to choose F, given the specified continuation play.

We now consider the *reputation game*, where the incumbent may be normal or tough. The profile in which all entrants stay out, any entry is met with F is a SPNE, supported by the "threat" that the entrants believe that the incumbent is normal and play switches to the always-enter/always-accommodate equilibrium if the incumbent ever responds with A.

Theorem 1. *Suppose the incumbent is either of type ω_n or type ω_t, and that type ω_t has prior probability less than $1/2$. Type ω_n must receive a payoff of at least $(1 - \delta) + 4\delta = 1 + 3\delta$ in any pure strategy NE in which ω_t always plays F.*

If type ω_t has prior probability greater than $1/2$, trivially there is never any entry and type ω_n has payoff 4 in any NE.

Proof. Pick any NE in pure strategies of the game. In this equilibrium, either the incumbent always plays F, (in which case, the entrants always stay out and the incumbent's payoff is 4), or there is a first period, say period τ, in which the normal type accommodates, revealing to future entrants that he is the normal type (since the tough type plays F in every period). In such an equilibrium, entrants stay out before τ (since both types of incumbent are choosing F), and there is entry in period τ. After observing F in period τ, entrants conclude the incumbent is the ω_t type, and there is no further entry. An easy lower bound on the normal incumbent's equilibrium payoff is then obtained by observing that the normal incumbent's payoff must be at least the payoff from mimicking the ω_t type in period τ. The payoff from such behavior is at least as large as

$$\begin{aligned}(1-\delta)\sum_{t=0}^{\tau-1}\delta^t 4 + (1-\delta)\delta^\tau + (1-\delta)\sum_{t=\tau+1}^{\infty}\delta^t 4 \\ &= (1-\delta^\tau)4 + (1-\delta)\delta^\tau + \delta^{\tau+1}4 \\ &= 4 - \delta^\tau(1-\delta)3 \\ &\geq 4-(1-\delta)3 \\ &= 1+3\delta,\end{aligned}$$

as was to be show. ■

For all $\delta \geq 1/3$, the outcome in which all entrants enter and the incumbent accommodates in every period is thus eliminated.

2.3 Infinite Horizon with Behavioral Types

In the reputation literature, it is standard to model the tough type as a *behavioral type*. In that case, the tough type is constrained to necessarily choose F. Then, the result is that in any equilibrium, $1 + 3\delta$ is the lower bound on the normal type's payoff. [The type ω_t from Sections 2.1 and 2.2 is an example of a *payoff type*.]

In fact, irrespective of the presence of other types, if the entrants assign positive probability to the incumbent being a tough behavioral type, for δ close to 1, the incumbent's payoff in any NE is close to 4. This is an example of a *reputation effect*.

Suppose there is a set of types Ω for the incumbent. Some of these types are behavioral. One behavioral type, denoted $\omega_0 \in \Omega$, is the *Stackelberg type*, who always plays F (i.e. the tough type). The *normal type* is $\omega_n \in \Omega$. Other types may include behavioral type ω_k, who plays F in every period before k and A afterwards. Suppose the prior beliefs over Ω are given by μ.

Lemma 1. *Consider the incomplete information game with types Ω for the incumbent. Suppose the Stackelberg type $\omega_0 \in \Omega$ receives positive prior probability $\mu(\omega_0) > 0$. Fix a NE. Let h^t be a positive probability period-t history in which every entry results in F. The number of periods in h^t in which an entrant entered is no larger than*

$$k^* := -\frac{\log \mu(\omega_0)}{\log 2}.$$

Proof. Denote by q_τ the probability that the incumbent plays F in period τ conditional on h^τ if entrant τ plays I. In equilibrium, if entrant τ does play I, then $q_\tau \leq 1/2$ (if $q_\tau > 1/2$, it is not a best reply for the entrant to play I). An upper bound on the number of periods in h^t in which an entrant entered is thus

$$k(t) := \#\{\tau \in \mathbb{N} : \tau < t \text{ and } q_\tau \leq 1/2\},$$

the number of periods in h^t where $q_\tau \leq 1/2$. (This is an upper bound, and not the actual number, since the entrant is indifferent if $q_\tau = 1/2$).

Let $\mu_\tau := \mathbb{P}(\omega_0 \mid h^\tau)$ be the posterior probability assigned to ω_0 after h^τ, where $\tau < t$ (so that h^τ is an initial "segment" of h^t). If entrant τ does not enter, $\mu_{\tau+1} = \mu_\tau$. If entrant τ does enter, then the incumbent fights (as h^t is a history in which every entry results in F) and[4]

$$\begin{aligned}\mu_{\tau+1} = \mathbb{P}(\omega_0 \mid h^\tau, F) &= \frac{\mathbb{P}(\omega_0, F \mid h^\tau)}{\mathbb{P}(F \mid h^\tau)} \\ &= \frac{\mathbb{P}(F \mid \omega_0, h^\tau)\mathbb{P}(\omega_0 \mid h^\tau)}{\mathbb{P}(F \mid h^\tau)} \\ &= \frac{\mu_\tau}{q_\tau}.\end{aligned}$$

Defining

$$\tilde{q}_\tau = \begin{cases} q_\tau & \text{if there is entry in period } \tau \\ 1 & \text{if there is no entry in period } \tau \end{cases},$$

we have, for all $\tau \leq t$,

$$\mu_\tau = \tilde{q}_\tau \mu_{\tau+1}.$$

Note that $\tilde{q}_\tau < 1 \implies q_\tau \leq 1/2$. Then,

$$\begin{aligned}\mu(\omega_0) = \tilde{q}_0\mu_1 &= \tilde{q}_0\tilde{q}_1\mu_1 \\ &= \cdots \\ &= \mu_t \prod_{\tau=0}^{t-1} \tilde{q}_\tau \\ &= \mu_t \prod_{\{\tau:\tau<t \text{ and } q_\tau \leq 1/2\}} \tilde{q}_\tau \\ &\leq \left(\frac{1}{2}\right)^{k(t)}.\end{aligned}$$

Taking logs,

$$\log \mu(\omega_0) \leq k(t) \log \frac{1}{2},$$

and so

$$k(t) \leq -\frac{\log \mu(\omega_0)}{\log 2},$$

as was to be shown. ■

The key intuition here is that since the entrants assign prior positive probability (albeit small)

[4] Since the entrant's action is a function of h^τ only, it is uninformative about the incumbent and so can be ignored in the conditioning.

to the Stackelberg type, they cannot be surprised too many times (in the sense of assigning low prior probability to F and then seeing F). Note that the upper bound is independent of t and δ, though it is unbounded in $\mu(\omega_0)$. The next result, established in much greater generality by Fudenberg and Levin (1989), follows.

Theorem 2. *Consider the incomplete information game with types Ω for the incumbent. Suppose the Stackelberg type $\omega_0 \in \Omega$ receives positive prior probability $\mu(\omega_0) > 0$. In any NE, the normal type's expected payoff is at least $1 + 3\delta^{k^*}$. Thus, for all $\varepsilon > 0$, there exists $\overline{\delta} \in (0,1)$ such that for all $\delta \in (\overline{\delta}, 1)$, the normal type's payoff in any NE is at least $4 - \varepsilon$.*

Proof. The normal type can guarantee histories in which every entry results in F by always playing F when an entrant enters. Such behavior yields payoffs that are no larger than the incumbent's NE payoffs in any equilibrium (if not, the incumbent would have an incentive to deviate). Since there is positive probability that the incumbent is the Stackelberg type, the history resulting from always playing F after entry has positive probability. Applying Lemma 1 yields a lower bound on the normal types payoff of

$$\sum_{t=0}^{k^*-1}(1-\delta)\delta^t + \sum_{t=k^*}^{\infty}(1-\delta)\delta^t 4 = 1 - \delta^{k^*} + 4\delta^{k^*} = 1 + 3\delta^{k^*}.$$

This can be made arbitrarily close to 4 by choosing δ sufficiently close to 1. ■

It is worth concluding with a few remarks.

- The payoff 4 is usually referred to as the *Stackelberg payoff*: this is the payoff the incumbent could secure if it were common knowledge he is the Stackelberg type. Theorem 2 says that a sufficiently patient incumbent gets arbitrarily close to the Stackelberg payoff. Even a small chance of being the Stackelberg type is just as good as being known to be that type.

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

- Theorem 2 makes few assumptions on the nature of incomplete information. In particular, the type space Ω can be infinite (even uncountable), as long as there is a gain of truth on the Stackelberg type ($\mu(\omega_0) > 0$).

- The result also holds for finite horizons. If the incumbent's payoff is the average of the flow (static) payoffs, then the average payoff is arbitrarily close to 4 for sufficiently long horizons.

2.4 Exercises

Exercise 1. Reconsider the two period reputation example illustrated in Figure 2. Suppose $\rho > 1/2$. Describe all of the equilibria. Which equilibria survive the Intuitive Criterion?

Exercise 2. Consider a stage game where player 1 is the row player and player 2 is the column player (as usual). Player 1 is one of two types, ω_n and ω_0. Payoffs are as follows.

		Player 2	
		L	R
Player 1	T	$2, 3$	$0, 2$
	B	$3, 0$	$1, 1$

ω_n

		Player 2	
		L	R
Player 1	T	$3, 3$	$1, 2$
	B	$2, 0$	$0, 1$

ω_0

The stage game is played twice, and player 2 is short-lived: a different player 2 plays in different periods, with the second period player 2 observing the action profile chosen in the first period. Describe all the equilibria of the game. Does the Intuitive Criterion eliminate any of them?

References

Fudenberg, Drew and David K. Levin (1989), “Reputation and Equilibrium Selection in Games with a Patirnt Player.” *Econometrica*, 57, 759–778.

Kreps, David M., Paul Milgrom, John Roberts, and Robert Wilson (1982), “Rational Cooperation in the Finitely Repeated prisoners’ Dilemma.” *Journal of Economic Theory*, 27, 245–252.

Kreps, David M. and Robert Wilson (1982), “Reputation and Imperfect Information.” *Journal of Economic Theory*, 27, 253–279.

Mailath, George J. (2019), “Modeling Strategic Behavior: A Graduate Introduction to Game Theory and Mechanism Design.” *World Scientific*.

Mailath, George J. and Larry Samuelson (2006), “Repeated Games and Reputations: Long-Run Relationships.” *Oxford University Press*.

——— (2015), “Reputations in Repeated Games.” H. Peyton Young and Shmuel Zamir eds., *Handbook of Game Theory with Economic Applications*, 165–238.

Milgrom, Paul and John Roberts (1982), “Predation, Reputation, and Entry Deterrence.” *Journal of Economic Theory*, 27, 280–312.

Selten, Reinhard (1978), “The Chain Store Paradox.” *Theory and Decision*, 9, 127–159.