

# Reputations

Niccolò Lomys\*

March 29, 2020

These notes heavily draw upon [Mailath \(2019\)](#), [Mailath and Samuelson \(2015\)](#), [Levin \(2006\)](#), and [Mailath and Samuelson \(2006\)](#). All errors are my own. Please bring any error, including typos, to my attention.

If you wish to learn more about Reputations, Part IV in [Mailath and Samuelson \(2006\)](#) and [Mailath and Samuelson \(2015\)](#) are excellent starting points.

## 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.

---

\*Toulouse School of Economics, University of Toulouse Capitole; [niccolo.lomys@tse-fr.eu](mailto:niccolo.lomys@tse-fr.eu).

## 1.1 Adverse Selection Approach to Reputations

The *adverse selection* approach to reputations<sup>1</sup> considers *games of incomplete information*. The motivation typically stems from a game of complete information in which the players are “normal”, 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 rewards 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 the 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 *reputation 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

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. Rewards 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

---

<sup>1</sup>As opposed to the interpretative approach to reputations, see [Mailath and Samuelson \(2015\)](#).

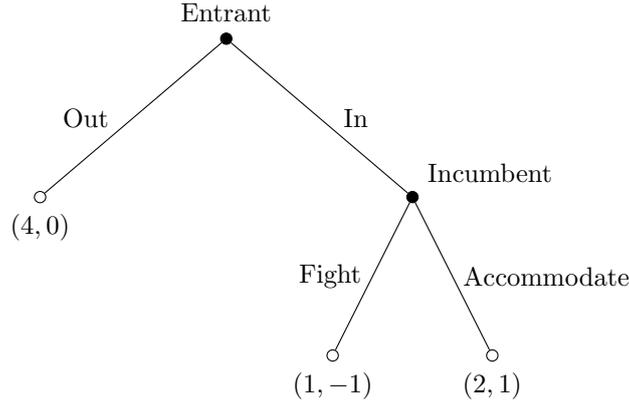


Figure 1: The stage game for the chain store. The first reward is that of the incumbent, and the second is that of the entrant.

$E_2$ ), with the second entrant  $E_2$  observing the outcome of the first interaction. The incumbent’s payoff is the sum of rewards 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.<sup>2</sup>

Finally, suppose the incumbent is a long-lived player, playing the game at times  $i = 0, 1, 2, \dots$  with discount factor  $\delta \in (0, 1)$ . The incumbent faces a succession 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  $\delta$  is sufficiently large, then there are (many) other equilibria. Indeed, for sufficiently large  $\delta$ , 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  $\omega_t$ . The tough incumbent receives a reward of 2 from fighting and only 1 from accommodating. The other incumbent is *normal*, i.e. of type  $\omega_n$ , with rewards 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  $\omega_t$ .

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<sup>3</sup>, and so  $E_2$  optimally plays  $I_2$ , the normal type

<sup>2</sup>The chain store paradox originates from Selten (1978).

<sup>3</sup>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

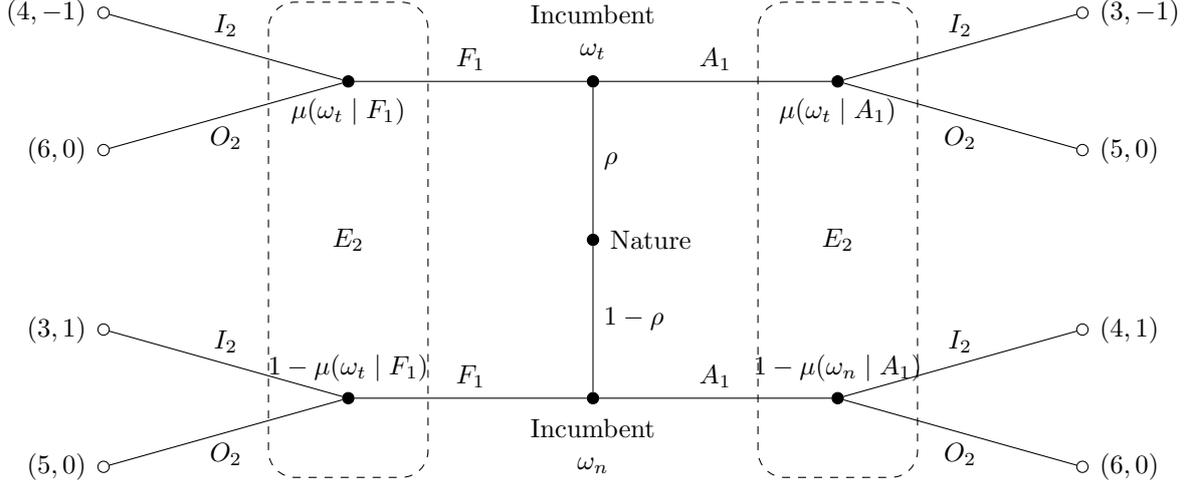


Figure 2: A signaling game representation of the subgame reached by  $E_1$  playing  $I_1$ . The first reward is that of the incumbent, and the second reward is that of  $E_2$  assuming the normal incumbent plays  $A_2$  and the tough incumbent plays  $F_2$  after a choice of  $I_2$  by  $E_2$ .

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 equilibria (make sure you understand this). There is, instead, a unique mixed strategy PBE (this is an example of hybrid PBE): type  $\omega_n$  plays  $F_1$  with probability  $\alpha$  and  $A_1$  with probability  $1 - \alpha$ , type  $\omega_t$  plays  $F_1$  for sure. Entrant  $E_2$  enters for sure after  $A_1$  and plays  $I_2$  with probability  $\beta$  and  $O_2$  with probability  $1 - \beta$  after  $F_1$ .

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

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

solving

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

for  $\alpha$  gives

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

---

incumbent.

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

Type  $\omega_n$  is willing to randomize if

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

where the left-hand side is the reward from playing  $A_1$  and the right-hand side is the reward 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 in 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 reward 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  $\delta$  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

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

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

If type  $\omega_t$  has prior probability greater than  $1/2$ , trivially there is never any entry and type  $\omega_n$  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  $\tau$ , 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  $\tau$  (since both types of incumbent are choosing  $F$ ), and there is entry in period  $\tau$ . After observing  $F$  in period  $\tau$ , entrants conclude the incumbent is the  $\omega_t$  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  $\omega_t$  type in period  $\tau$ . 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* (also called *action* or *commitment 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  $\omega_t$  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  $\delta$  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  $\Omega$  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  $\omega_k$ , who plays  $F$  in every period before  $k$  and  $A$  afterwards. Suppose the prior beliefs over  $\Omega$  are given by  $\mu$ .

**Lemma 1.** *Consider the incomplete information game with types  $\Omega$  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  $q_\tau$  the probability that the incumbent plays  $F$  in period  $\tau$  conditional on  $h^\tau$  if entrant  $\tau$  plays  $I$ . In equilibrium, if entrant  $\tau$  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  $\mu_\tau := \mathbb{P}(\omega_0 | h^\tau)$  be the posterior probability assigned to  $\omega_0$  after  $h^\tau$ , where  $\tau < t$  (so that  $h^\tau$  is an initial "segment" of  $h^t$ ). If entrant  $\tau$  does not enter,  $\mu_{\tau+1} = \mu_\tau$ . If entrant  $\tau$  does enter, then the incumbent fights (as  $h^t$  is a history in which every entry results in  $F$ ) and<sup>4</sup>

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

---

<sup>4</sup>Since the entrant's action is a function of  $h^\tau$  only, it is uninformative about the incumbent and so can be ignored in the conditioning.

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,

$$\begin{aligned} \mu(\omega_0) &= \tilde{q}_0 \mu_1 = \tilde{q}_0 \tilde{q}_1 \mu_2 \\ &= \dots \\ &= \mu_t \prod_{\tau=0}^{t-1} \tilde{q}_\tau \\ &= \mu_t \prod_{\{\tau: \tau < t \text{ and } q_\tau \leq 1/2\}} \tilde{q}_\tau \\ &\leq \left(\frac{1}{2}\right)^{k(t)}. \end{aligned}$$

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) 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  $\delta$ , 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  $\Omega$  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  $\bar{\delta} \in (0, 1)$  such that for all  $\delta \in (\bar{\delta}, 1)$ , 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  $\delta$  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  $\Omega$  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) rewards, then the average payoff is arbitrarily close to 4 for sufficiently long horizons.
- You might think that it is important here for the short-lived players (in our example, entrants) to be able to observe perfectly what the long-lived player (in our example, the incumbent) is doing. In fact, that need not be the case. [Fudenberg and Levin \(1992\)](#) show that a version of their “bounds” argument applies even if the long-lived player's behavior is imperfectly observed. The idea is that short-lived players will be actively learning about the long-lived player's behavior over time, with some outcomes making it more likely that the long-lived 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 Levin \(1992\)](#) 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 \(2004\)](#) study the long-run, or asymptotic, pattern of behavior in this model.

**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,  $\omega_n$  and  $\omega_0$ . Rewards are as follows.

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

		Player 2	
		$L$	$R$
Player 1	$T$	3, 3	1, 2
	$B$	2, 0	0, 1
		$\omega_0$	

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?

### 3 Bad Reputations

The previous analysis suggests that a long-lived player interacting with a sequence of short-lived players will typically benefit from reputation effects since he has the option of acting committed to a certain strategy. This need not be the case if there is imperfect monitoring. The following example, due to [Ely and Välimäki \(2003\)](#), shows that reputation concerns can have perverse consequence. In particular, we will show that the reputational concern of the long-lived player to look good in the current period undermines commitment power and results in the loss of all surplus. These are called “bad reputations”.

There are two players: a motorist and a mechanic. The motorist’s car needs either an engine replacement or a mere tune-up with equal probability. Denote these possibilities as  $\theta \in \{\theta_e, \theta_t\}$ . The motorist lacks the expertise to determine which repair is necessary. The motorist therefore considers bringing the car to a mechanic who possesses the ability to diagnose the problem and perform the necessary repair. The mechanic, if hired, privately observes the state  $\theta$  and, conditional on this information, chooses a repair  $a \in \{e, t\}$ , indicating engine replacement or tune-up. Let  $(\beta_e, \beta_t)$  be the mechanic’s repair strategy, where  $\beta_a$  is the probability that the mechanic performs repair  $a$  in state  $\theta_a$  (i.e. the correct repair).

The motorist’s reward depends on the repair and the state as in the following table:

	$\theta_e$	$\theta_t$
$e$	$u$	$-w$
$t$	$-w$	$u$

Assume that  $w > u > 0$  and that the motorist has an outside option that gives him a reward normalized to zero. Given  $(\beta_e, \beta_t)$ , the motorist’s expected reward from hiring

the mechanic is

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

Since motorist can secure a reward equal to zero by not hiring, a necessary condition for hiring the mechanic is that

$$\min\{\beta_e, \beta_t\} \geq \beta^* := \frac{w - u}{u + w} > 0.$$

To see this, if  $\beta_e < \beta^*$ , then even if  $\beta_t = 1$ , the expected reward to the motorist from hiring the mechanic would be negative.

Consider the benchmark scenario in which the mechanic is known to be *good*; that is, his rewards are identical to those of the motorist. In this case, there is no incentive problem, and the first-best outcome is the unique sequential equilibrium of the one-shot interaction between the motorist and mechanic: the mechanic does the correct repair if hired, and the motorist optimally hires.

This conclusion remains true even when the motorist believes with small probability  $\mu > 0$  that the mechanic is *bad*, i.e. a mechanic who always replaces the engine independently of the state. Given that the good mechanic always does the correct repair and the bad mechanic always replaces the engine, the motorist's expected reward from hiring the mechanic is now

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

Thus, if

$$\mu \leq p^* := \frac{2u}{u + w},$$

the good mechanic does the correct repair if hired, and the motorist optimally hires. If  $\mu \in (p^*, 1]$ , the motorist is too pessimistic about the mechanic's type and strictly prefers not to hire even if the good mechanic can be expected to perform the correct repair.

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 sequence of motorists who decide in turn whether to hire the same mechanic after observing earlier repairs but not what was actually needed.

The good mechanic chooses a strategy, which in this dynamic game is a pair of probabilities  $(\beta_e^k(h^k), \beta_t^k(h^k))$  specifying the probability of changing the engine and performing a tune-up, respectively, at date  $k$  as a function of his previous history  $h^k$ . The mechanic's private history records for each previous date  $l$ , the pair  $(\theta^l, a^l)$  consisting of the state and the chosen repair. The good mechanic maximizes the expected discounted average payoff where  $\delta \in (0, 1)$  is the discount factor. The bad mechanic replaces the engine whenever hired, regardless of the history.

Since each motorist is a short run player, motorist  $k$  will want to hire given history  $h^k$  if he expects the correct repair with sufficient probability. His decision will be based on

the probability  $\mu^k(h^k)$  he assigns to the mechanic being bad and the expected behavior  $\bar{\beta}_a^k := \mathbb{E}(\beta_a^k | h^k)$ ,  $a \in \{e, t\}$ , of the good mechanic. Of course, if  $\mu^k(h^k) > p^*$ , the motorist will certainly choose not to hire, and if  $\mu^k(h^k) \leq p^*$ , a necessary condition for him to hire is that  $\min\{\bar{\beta}_e^k, \bar{\beta}_t^k\} \geq \beta^*$ .

In this repeated interaction the good mechanic's incentives at date  $k$  are a mix of the short-run desire to choose the correct repair for the current customer, and the long-run strategic objective to maintain a good reputation. We show next that this reputational incentive distorts the otherwise good intentions of the mechanic, and that the severity of this distortion increases with the discount factor. In fact, when the good mechanic is sufficiently patient, his services become essentially worthless. Formally, let  $V(\mu, \delta)$  be the supremum of discounted average Nash equilibrium payoffs for the good mechanic when  $\mu$  is the prior probability that the mechanic is bad and  $\delta$  is the discount factor. Since the good mechanic's rewards are identical to those of the motorists, this also measures the value of the mechanic's services to the population of motorists when  $\delta$  is close to 1. The next theorem shows that this average discounted equilibrium payoff is small for high discount factors.

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

**Proof.** If  $\mu > p^*$ , there is a unique Nash equilibrium outcome in which the mechanic is never hired. The first motorist does not hire, so beliefs are not updated. The second motorist then does not hire and so on. In this case, the result trivially holds.

Now suppose  $\mu \in (0, p^*]$  and consider a Nash equilibrium in which the mechanic is hired. The updated probability of a bad mechanic will depend on the repair chosen. When the mechanic chooses to replace the engine, Bayes' formula gives the updated probability as follows:

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

Since the mechanic is hired only if  $\beta_t \geq \beta^* > 0$ , we have  $\mu^1(e) > \mu$ . That is, observing an engine replacement is always bad news. For each  $\mu \in [0, 1]$ , let

$$\Upsilon(\mu) := \frac{\mu}{\mu + (1 - \mu)\left[\frac{1}{2} + \frac{1}{2}(1 - \beta^*)\right]}$$

be the smallest possible posterior probability of a bad mechanic when an engine replacement has been observed. The preceding argument implies that  $\Upsilon(\mu) > \mu$  for every  $\mu \in (0, p^*]$ . It is also easy to see that  $\Upsilon$  is strictly increasing and continuous.

We construct a decreasing sequence of cutoff points  $p_m$  defined by  $p_1 := p^*$  and  $p_m := \Upsilon^{-1}(p_{m-1})$  for  $m > 1$ . We will use induction on  $m$  to bound the payoffs across all Nash equilibria when the prior exceeds  $p_m$ . For the induction hypothesis, suppose that there exists a bound  $V_m(\delta)$  with  $\lim_{\delta \rightarrow 1} V_m(\delta) = 0$  and  $V(\mu, \delta) \leq V_m(\delta)$  for all  $\mu > p_m$ .

Note that we have already shown this for  $m = 1$ .<sup>5</sup> Fix  $\mu > p_{m+1}$ . It suffices to consider a Nash equilibrium in which the mechanic is hired in the first period.<sup>6</sup>

Since in the first period the good mechanic must be choosing the correct action with probability at least  $\beta^* > 0$  in each state (otherwise the motorist would not 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 is one of his best responses). Thus, we must have

$$V(\mu, \delta) \leq (1 - \delta)u + \delta \left[ \frac{z(e | \theta_e) + z(t | \theta_t)}{2} \right], \quad (1)$$

where  $z(a | \theta_a)$  denotes the expected continuation payoff in state  $\theta_a$  from choosing action  $a$ . Since we have taken  $\mu > p_{m+1}$ , we know that the posterior  $\mu^1(e)$  is at least  $\Upsilon(p_{m+1})$  which by definition is larger  $p_m$ . Thus, by the induction hypothesis,

$$z(e | \theta_e) \leq V_m(\delta). \quad (2)$$

Now, consider the incentive compatibility constraint for the good mechanic. He must be willing to choose  $e$  when the state is  $\theta_e$  rather than deviating to  $t$ . That is, we must have

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

or, re-arranging,

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

and so by (2),

$$z(t | \theta_e) \leq \frac{1 - \delta}{\delta}(u + w) + V_m(\delta). \quad (3)$$

Now since the motorist's behavior is conditioned only on the public history, the mechanic's continuation value from choosing  $t$  cannot depend on the mechanic's private history. Thus,  $z(t | \theta_e) = z(t | \theta_t)$ . Now we can substitute the bound (2) and the bound obtained from incentive compatibility (given by (3)) into (1) and rearrange to obtain

$$V(\mu, \delta) \leq (1 - \delta) \frac{3u + w}{2} + \delta V_m(\delta).$$

By the induction hypothesis, the limit of the right-hand side is zero as  $\delta$  approaches one.

By induction, it follows that the desired result holds for each  $\mu > \inf_m p_m$ . Since the sequence is decreasing,  $\inf_m p_m = \lim_m p_m$ , and we can conclude the proof by observing

---

<sup>5</sup>Quiz: Do you see why?

<sup>6</sup>Take any Nash equilibrium in which the mechanic is hired with zero probability until date  $k > 1$ . Then the continuation play beginning at date  $k$  is a Nash equilibrium with prior  $\mu$ . Since the mechanic's reward was zero in the first  $k - 1$  periods and the mechanic's minmax reward is zero, the continuation payoff in the equilibrium that begins at date  $k$  can be no smaller than the payoff to the original equilibrium.

that  $p_m \rightarrow 0$ . Indeed, since  $\Upsilon$  is continuous,  $\lim_m p_m = \lim \Upsilon(p_{m+1}) = \Upsilon(\lim_m p_m)$ , and therefore  $\lim_m p_m$  is a fixed point of  $\Upsilon$ . Since  $\Upsilon(\mu) > \mu$  for every  $\mu \in (0, p^*]$ , that 0 is the only fixed point (smaller than  $p^*$ ) of  $\Upsilon$  follows. ■

Intuitively, the problem with bad reputations is the following. Once there is a sufficiently high belief that the mechanic is bad, motorists will stop hiring. 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 is willing to do engine replacements when they are needed. Thus, the only way the mechanic will be hired when he is one engine replacement away from being fired forever is if the continuation payoff from a tune-up is also close to zero. If  $\delta$  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 is just one engine replacement away from a bad belief region. This unraveling continues until the mechanic's payoffs must be low for any prior belief.

Theorem 3 does not imply that the motorist will never be hired. The repeated game has many equilibria, including equilibria in which the mechanic is hired many times on the equilibrium path. However, the theorem implies that, as  $\delta$  is close to 1, the frequency of hiring converges to zero so that the average value also declines to zero. A typical equilibrium has the following structure: the mechanic is hired for sure up to date  $k$  and never hired thereafter.

One could argue that these equilibria have an implausible feature: if the mechanic has ever performed a tune-up prior to date  $k$ , he will have perfectly revealed himself to be good. Nevertheless, it must be the case that even after these histories, the mechanic will not be hired frequently. If he were hired, then the incentive to separate would again be too strong, and the mechanic would perform a tune-up with probability 1.

A key problem here is that each motorist is a short-lived player. Thus, he does not internalize the benefits he creates for later motorists if he hires and gives the mechanic a chance to signal his goodness by performing a tune-up. [Ely and Välimäki \(2003\)](#) show that if there is a single long-lived 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  $e$  in each period.

**Exercise 3.** Consider a variant of [Ely and Välimäki \(2003\)](#)'s model where only the mechanic's action in the first period ( $e$  or  $t$ ) is observed by subsequent short-lived motorists. This means that a period- $k$  public history,  $k \geq 1$  only includes the actions in the first period.

- (a) Suppose that  $\mu \in (u/w, 2u/(u+w)]$  and that  $\delta > (u+w)/(2u+w)$ . Show that in the unique sequential equilibrium of the game the good mechanic always chooses  $e$  and the motorist never hires the mechanic in the first period.
- (b) Now suppose that  $\mu \in (0, u/w]$ . Show that there is a sequential equilibrium of the game where the good mechanic always chooses  $e$  and the motorist never hires the mechanic in the first period. Show that there is a sequential equilibrium of the game where the good mechanic chooses the correct repair in the first period and the motorist hires the mechanic in the first period.

## References

- Cripps, Martin W., George J. Mailath, and Larry Samuelson (2004), “Imperfect Monitoring and Impermanent Reputations.” *Econometrica*, 72, 407–432.
- Ely, Jeffrey C. and Juuso Välimäki (2003), “Bad Reputation.” *The Quarterly Journal of Economics*, 118, 785–814.
- Fudenberg, Drew and David K. Levin (1989), “Reputation and Equilibrium Selection in Games with a Patient Player.” *Econometrica*, 57, 759–778.
- (1992), “Maintaining a Reputation when Strategies are Imperfectly Observed.” *The Review of Economic Studies*, 59, 561–579–778.
- 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.
- Levin, Jonathan (2006), “Reputation in Repeated Interaction.” *Lecture Notes*.
- 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.” *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.