Journal of Mathematical Economics 49 (2013) 345–354

Contents lists available at ScienceDirect

Journal of Mathematical Economics

journal homepage: www.elsevier.com/locate/jmateco

Bounded memory and permanent reputations☆

Daniel Monte *

Sao Paulo School of Economics - FGV, Brazil

ARTICLE INFO

Article history:
Received 10 February 2010
Received in revised form
9 May 2013
Accepted 28 June 2013
Available online 23 July 2013

Keywords:
Bounded memory
Reputation games
Learning

ABSTRACT

A recent result in repeated incomplete information games is that after an arbitrarily long history, any equilibrium of the continuation game must be an equilibrium of the complete information game. This result implies that reputation is a short-run phenomenon. We study a particular class of reputation games and show that bounded memory may lead to permanent reputations. In fact, for a particular range of parameters, a bounded memory player may never be able to learn anything at all.

1. Introduction

An important recent result in reputation games is that after an arbitrarily long history, any equilibrium of the continuation game must be an equilibrium of the complete information game (Cripps et al., 2004, CMS hereafter, prove this result for repeated games with imperfect public monitoring). Many recent papers have shown that this convergence result holds with an impressive generality, robust to different monitoring technologies and different underlying games.[1] This leaves open the question of how to explain reputation in long-term relationships.

We study a particular class of reputation games and show that bounded memory can explain long-run reputation. Our setting is a two-player infinitely repeated zero-sum game with one-sided incomplete information and perfect monitoring. One uninformed player faces a player that, with some exogenous probability, is committed to a specified mixed strategy. If both players have full memory, we show that the learning result of CMS applies here: in the limit, the uninformed player's belief will converge to the true probability distribution of the underlying types. Reputation is, thus, a short-run phenomenon. In this paper, we study reputation in a world in which the uninformed player has bounded memory.

Bounded memory is modeled as a finite set of states. The strategy of a bounded memory player is to choose an action rule, which is a map from memory states to the set of actions, but also a transition rule from state to state. At the beginning of every stage game the only information that the player has about the history of the game is his current memory state. He can then compute a best response based on the beliefs about the actual history at that point, knowing that he is forgetful across periods. We say that a memory is finite if the set of beliefs that the player can hold in equilibrium is finite. We do, however, impose sequential rationality on the equilibrium strategies. This rules out the possibility of commitment ex-ante; in particular it rules out the automata models.[2]

We think of memory constraints as a categorization procedure. A bounded memory agent can categorize the world in only a finite number of ways. This captures the fact that in reality we may not be able to distinguish very similar histories. Our view is that people categorize histories and form coarse impressions, instead of precise beliefs.

We consider the case of public beliefs in this paper. This is done for technical reasons that will become clear later in the paper. This assumption can be motivated in at least two different ways. First, a slightly modified model in which the uninformed player takes an observable continuous action before the agent acts together with a single peaked utility function would give us the same results as in this paper—without the assumption of observable memory states. Second, one can think of this as literally being public beliefs, such as institutions that publish credit ratings.[3]

☆ I am very grateful to Dirk Bergemann, Stephen Morris and Ben Polak for their support throughout this project. I also thank John Geanakoplos, Pei-Yu Lo, George Mailath, and Joel Watson for comments and conversations. Finally, I thank two anonymous referees for very helpful and detailed comments.

* Tel.: +55 1137993727.
E-mail address: daniel.monte@fgv.br.

1 Benabou and Laroque (1992), Kalai and Lehrer (1993), Aumann et al. (1995), Jackson and Kalai (1999), Sorin (1999), and Cripps et al. (2007) prove important results on learning and reputation, but are less related to the underlying model of our paper.

2 In Monte (2007) we discuss the difference between bounded memory and finite automaton in detail.

3 There are no reasons to believe that our results would not extend to the case of private beliefs. However, the case of private beliefs greatly complicates the analysis, so we leave it for future work.

0304-4068/$ – see front matter
http://dx.doi.org/10.1016/j.jmateco.2013.06.006

Our main result is shown in Proposition 1 in the text. We compute an upper bound on the belief spread of the bounded memory player as a function of his memory size and the exogenous mixed strategy of the commitment type. If this mixed strategy is such that all actions are played with an almost uniform randomization, but such that it is still informative, we show that the bounded memory player will never be able to learn his opponent's true type. In fact, learning may be negligible if the memory size is small relative to the commitment type's mixed strategy.

The intuition for our result is that with bounded memory the agent can hold only a finite number of beliefs in equilibrium. And, these beliefs cannot be too far apart from each other, or else the sequential rationality constraints would not be satisfied: the player would not want to update his beliefs.

Cripps et al. (2007) extend the results of CMS by considering a game with private reputations and two long-lived players. Our class of games differs only slightly from the class of games studied in CMS and Cripps et al. (2007). First, we restrict attention to games in which the stage game between the normal type and the uninformed player is zero-sum. Second, CMS considers one long-lived player and a sequence of short-lived players (who are uninformed about the long-lived player's type), while in this paper both players are long lived and we make no assumption in their discount factor (a sequence of short-lived players is a special case here). Cripps et al. (2007) also consider two long-lived players, but our underlying class of games is significantly simpler, in that we have a specific commitment type that plays a stationary strategy, thus their results also apply here. Third, while CMS considers a game with imperfect monitoring (the long-lived player's actions are private) we consider a game with perfect monitoring but in which the commitment type plays a completely mixed strategy—as they pointed out (p. 408) the argument for long-run learning applies with very little change to this case. Finally, we restrict attention to games with two actions only.

Although we have restricted our attention to repeated zero-sum games, there are no reasons to think that our results would not hold in a general game. There are, however, two main reasons for considering zero-sum games. First, the question of updating beliefs under bounded memory is more clear cut here. In a repeated non-zero-sum game with incomplete information, a player with a bounded number of states faces two constraints: bounded complexity on implementing a strategy and bounded ability on updating beliefs about the actual type. The literature on automata has focused on the first issue, while we focus on the second.[4] In the setting of this paper, the bounded memory player is uninformed about the type of his opponent. Moreover, the complete information game has a unique equilibrium in the repeated game and, thus, the complexity of implementing a strategy is simple. Therefore, the issue is on updating beliefs and learning. The second reason for choosing this class of games is that we want to emphasize that an infinitely renewable reputation can happen in a world in which parties have completely opposite interests. The monitoring technology and the restriction to two-by-two games are done here for simplicity only. While considering imperfect monitoring is an easy extension, extending beyond two-by-two games greatly complicates the analysis, since the precise statement of the results will depend on the specific game considered.

A motivation for our underlying model can be based on the credibility models of Sobel (1985) and Benabou and Laroque (1992). An uninformed policy maker takes decisions based on the advice of an expert. The expert may be a "friend" or an "enemy" (Sobel's labels). In our context, if the policy maker is restricted to bounded memory, the enemy may explore the policy maker indefinitely, without ever deteriorating his reputation to zero.

Other authors have worked on alternative explanations for permanent reputations. In a game where types are continuously changing, permanent reputation can be sustained as shown by Holmstrom (1999), and Mailath and Samuelson (2001). Wiseman (2008) also has a result on impermanent reputations. In a related study, Bar-Isaac (2007) showed that a model of reputation in teams can endogenously introduce this type uncertainty and thus sustain reputation. Liu (2011) and Liu and Skrzypacz (2011) also show that reputation incentives can persist if record keeping is costly or exogenously restricted.

Ekmekci (2011) showed that if the memory of the uninformed player is restricted (in the form of a finite set of ratings) then there exists a rating system (set of ratings and a transition rule) that can explain permanent reputation. The main difference between our paper and Ekmekci's is that here memory is endogenous. It is part of the uninformed player's strategy and has to satisfy sequential rationality constraints. In Ekmekci (2011) the memory process is exogenous: designed by a third party.

The results in this paper do not rely on the discount factor of the players. Fudenberg and Levine (1989, 1992) studied reputation in repeated games with one long-run player facing a sequence of short-run players. In contrast, the literature on reputation with two long-run players is very recent (see Cripps et al., 2005, Atakan and Ekmekci, 2011, 2012, 2013)

We proceed as follows. Section 2 describes the model. In Section 3 we solve the model for the full memory case. In Section 4 we define memory and describe the game when the uninformed player is restricted to a bounded memory. Section 5 is the main part of the paper where we show that under bounded memory reputation will always be sustained. We conclude in Section 6.

2. Reputation game

We study a two-player game with incomplete information. We assume that at every period the game ends with probability $1-\delta$, where $\delta \in (0, 1)$.[5] One of the two players is informed about a choice of nature while the other player is uninformed. Before the first stage game, nature draws one of two possible types for the informed player: either a commitment type (c) or a normal type (r).[6] The commitment type is chosen with probability ρ whereas the normal type is the true type with probability $1-\rho$. The uninformed player is not aware of nature's choice.

The commitment type is playing a given mixed strategy known by both players. The normal type maximizes her repeated-game payoff. Given the exogenous stopping probability $(1-\delta)$, the normal type and the uninformed player maximize their expected payoffs which are discounted by an effective discount rate of δ.

The stage game is defined as a simultaneous game between the uninformed and the informed player. The action space is A_b for the uninformed player and A_r for the normal type. We assume $|A_b| = 2$ and $|A_r| = 2$. A public history at time t is denoted by h_t and is a sequence of actions played by both players. The set of all public histories at time t is $\mathcal{H}_t \equiv (A_r \times A_b)^t$. The commitment type is playing a given behavioral strategy, denoted by σ_c, which plays $\bar{q} \in \Delta(A_r)$ after any history, where $\bar{q}$ is common knowledge in

[4] See, for example, Neyman (1985), Rubinstein (1986), Abreu and Rubinstein (1988) and Kalai and Stanford (1988) for repeated games with automata.

[5] This is done to ensure that beliefs are consistent. In particular, given that a bounded memory player can keep track of time insofar as his memory allows him, the exogenous stopping rule is a natural way to ensure that the priors over the time periods are well defined. This will be more clear after the discussed on beliefs in Section 4.2.

[6] The normal type of a player is also denoted in the literature as a rational type, thus the notation r.