Bounded Memory Equilibrium*

Penélope Hernández[†] and Eilon Solan[‡]

March 29, 2009

Abstract

In the present paper we define a new solution concept for infinitely repeated games in which the players have bounded computational capacity, the *automaton equilibrium payoff*. We say that a payoff vector is an automaton equilibrium payoff if (1) it can be supported by a pair of strategies that can be implemented by automata, (2) no player has a smaller automaton that, on the equilibrium path, yields the same (or higher) long-run average payoff, and (3) to profit, a player has to switch to an automaton with significantly larger memory.

The second condition reflects the intuition that players have a lexicographic utility (as Abreu and Rubisntein 1988): subject to maximizing their payoff they would like to minimize the size of the automaton that they use. An implication of this requirement is that the punishment phase must be part of the equilibrium path, since otherwise a player may reduce the number of states in its automaton while still being able to implement the equilibrium path.

The third condition says that as soon as a player does not significantly increase its memory size, he cannot profit. This condition reflects the intuition that a significantly larger automaton is not feasible, so that players need not to worry about deviations to such automata. Alternatively, since memory size is costly, one may think of deviations to a significantly larger automaton as too costly to implement.

Our main result if a folk theorem: every feasible and individually rational payoff vector is an automaton equilibrium payoff. Our construction suggests one way to implement a mixed automaton equilibrium. Any pure automaton which conforms the mixed automaton equilibrium has the same complexity and implememts the same equilibrium play but not in the same way. The play consists of three phases: a punishment phase, in which each player proves to the other that she can punish if necessary; a babbling phase, in which the players follow a pre-specified sequence of action pairs, and a regular phase, in which the players play repeatedly a sequence of action pairs that implements the desired equilibrium payoff.

The intuition behind the mixed automata that we construct is as follows. On the one hand, mixed strategy allows agents to reuse some states

*This work has been done while the second author was visiting Universidad de Valencia. The second author wishes to thank the Departamento de Análisis Económico at Universidad de Valencia for the hospitality during his visit.

[†]Departamento de Análisis Económico, Universidad de Valencia. Campus de Los Naranjos s/n, 46022 Valencia, Spain.

[‡]Department of Statistics and Operations Research, School of Mathematical Sciences, Tel Aviv University, Tel Aviv 69978, Israel.

in such a way the regular play is implemented by using the same states than in the punishment and babbling phases. On the other hand, each pure automaton has bounded memory, and he can be tricked. If the trick is detected, the agent punishes the deviator. If the trick is undetected, then once the trick is over the deviator does not know the memory state of the agent, and therefore she does not know what the agent expects her to play. At some point in the future she will therefore not stand in the agent's expectations, make some mistake and be punished. However, to learn the wiring of a non-negligible fraction of the agent takes a lot of memory, more than the other player has and it is here where the third requirement of the equilibrium concept acts.