Evolution and Cooperation in Noisy Repeated Games by Drew Fudenberg and Eric Maskin

2010-05-18

Axel Bernergård[1]

1 Introduction

In repeated games where a given "stage game" is played repeatedly, rational players can cooperate in a game like the prionser's dilemma even though rational players who never play strictly dominated stragies could never cooperate in a one-shot prisoner's dilemma. This is achieved by threatening to punish in the future players who refuse to cooperate today. Though it is true that beneficial cooperation is possible in repeated play of a prisoner's dilemma, which gives the players a high payoff, the "folk theorems" of repeated games say that this is not the only possible outcome: if the players are sufficiently patient, then almost any feasible and individually rational payoff is possible. Like Mailath and Samuelson (2006) we can view these multiple equilibria of repeated games as a virtue that makes it possible to explain the virtue of richness in the behavior we observe around as. But, the multiplicity of equilibria also means an that the theory is unable to predict how repeated games are played. More importantly, the result that highly inefficient payoffs are possible in the repeated game runs counter to the widespread intuition that efficient equilibria should be more likely. Fudenberg and Maskin provide support for this intuition in a particular class of games by showing that in this class of games the assumptions that players make mistakes with a small probability and play evolutionary stable strategies that are not infinitely complex lead to a restriction of the set of payoffs that can occur in equilibrium in the direction towards efficiency. If there is a unique payoff pair that maximizes the sum of the players' payoffs then the restriction is so great that efficiency is predicted. So, for the particular case of a prisoner's dilemma which has such a unique efficient payoff pair, the assumptions that there is a small mistake probability and that players play evolutionary stable strategies that are not too complex are sufficient change the set of equilibrium payoffs from any payoff to the efficient payoffs.

The intuitive motivation for the result proceeds in two steps. The first step is to argue that punishments will be mild in evolutionary stable equilibria when the players make mistakes. When mistakes are possible, also histories which are off the equilibrium outcome path can occur. This means that the payoffs along the punishment paths will actually matter and the punishments are not just used to support the equilibrium outcome. If the punishments are harsh, then other strategies with milder punishments can invade. The second step is to take mild punishments as given and then argue that with mild punishments equilibrium strategies will be efficient. This is because the penalty for deviating to efficient play from a strategy profile that suggests inefficient play is mild. Fudenberg and Maskin note that these arguments are general and do not apply only in the class of games that the paper considers, and the results are therefore intended to be suggestive rather than definitive.

2 Model and Results

There is a symmetric finite two-player stage game $(\{1, 2\}, A, u)$ that is repeated infinitely many times. There are intended actions and realized actions. When a mistake occurs the realized action

[1] Axel.Bernergard@hhs.se

differs from the intended action. The players only observe realized actions. The players can only use finitely complex strategies; that is (I suppose), strategies that can be described by a machine with finitely many states. The players do not discount their payoffs but instead maximize their time-average payoffs so that a stream $(x_t)_{t=0}^{\infty}$ of instantaneous payoffs is worth $\lim_{T\to\infty} \frac{1}{T}\sum_{t=0}^{T} x_t$ in the repeated game.[2] The assumption that the probability for a mistake is small is captured by using lexiographic preferences for the repeated game of the following kind: Let σ and σ' be two strategy profiles for the repeated game. If σ gives a strictly higher time-average payoff than σ' when there are no mistakes, then σ is strictly preferred to σ' in the repeated game. If σ and σ' are equal when no mistakes are made, then compare σ and σ' conditional on there having been 1 mistake. If they are again equal, then continue with comparing them conditional on there having been two mistakes, and so on. The concept of evolutionary stability used is a generalization of neutral stability (the weak inequality evolutionary stability criteria) to this type of lexiographic preferences. A payoff is efficient if it maximizes the sum of players payoffs.

Theorem 1. *Let $\underline{u} = \min\{u :$ there exists u' such that (u, u') is efficient$\}$. If a finitely complex strategy s is ES, then the time-average payoff when both players use s is at least $\underline{u}$ if there are no mistakes*

Corollary. *Suppose that there is a unique efficient pair (u^*, u^*). If a finitely complex strategy s is ES, then the time-average payoff when both players use s is u^* if there are no mistakes.*

Proof Steps. The proof proceeds by letting s be a finitely complex ES strategy and letting h^* be the history after which s gives the worst continuation payoffs. Let v^* be this continuation payoff. Such a worst history exists by finite complexity. Then assume for a contraction that there is no v' such that (v^*, v') is efficient. It is then be possible to construct a strategy that can invade s, which contradicts that s is ES. Hence there is some v' such that (v^*, v') is efficient. It then follows from the definition of $\underline{u}$ that $v^* \geq \underline{u}$. Since the time-average payoff for s is weakly greater than v^* (the empty history is not worse than the worst history), the desired conclusion follows. $\square$

Theorem 2. *Let $\underline{u}$ be defined as in Theorem 1. If (v, v) is a strictly individually rational payoff pair such that $v > \underline{u}$, then there exists a finitely complex ES strategy s which gives the time-average payoff v if there are no mistaktes, provided that there exists a finitely complex strategy with these payoffs.*

[2] Fudenberg and Maskin do not discuss the existence of this limit, but intuitively it should exist for the instantaneous payoffs that occur for finitely complex strategies since such strategies will generate outcome paths that cycle.

3 References

Fudenberg, Drew, and Eric Maskin. 1990. "Evolution and Cooperation in Noisy Repeated Games." *The American Economic Review,* 80(2): 274-279.

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