Games and Economic Behavior 91 (2015) 45–59

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Plausible cooperation ☆

Olivier Compte [a], Andrew Postlewaite [b,*]

[a] Paris School of Economics, 48 Bd Jourdan, 75014 Paris, France
[b] Department of Economics, University of Pennsylvania, 3718 Locust Walk, Philadelphia, PA 19104-6207, United States

ARTICLE INFO

Article history:
Received 4 April 2013
Available online 30 March 2015

JEL classification:
D01
D70

Keywords:
Repeated games
Private monitoring
Bounded rationality
Cooperation

ABSTRACT

There is a large repeated games literature illustrating how future interactions provide incentives for cooperation. Much of the earlier literature assumes public monitoring. Departures from public monitoring to private monitoring that incorporate differences in players' observations may dramatically complicate coordination and the provision of incentives, with the consequence that equilibria with private monitoring often seem unrealistically complex or fragile. We set out a model in which players accomplish cooperation in an intuitively plausible fashion. Players process information via a mental system – a set of psychological states and a transition function between states depending on observations. Players restrict attention to a relatively small set of simple strategies, and consequently, might learn which perform well.

1. Introduction

Cooperation is ubiquitous in long-term interactions: we share driving responsibilities with our friends, we offer help to relatives when they are moving and we write joint papers with our colleagues. The particular circumstances of an agent's interactions vary widely across the variety of our long-term relationships but the mechanics of cooperation are usually quite simple. When called upon, we do what the relationship requires, typically at some cost. We tend to be upset if our partner seems not to be doing his part and our willingness to cooperate diminishes. We may be forgiving for a time but stop cooperating if we become convinced the relationship is one-sided. We sometimes make overtures to renew the relationship when opportunities arise, hoping to restart cooperation. Incentives to cooperate stem from a concern that the relationship would temporarily break down, while incentives to be less cooperative when the relationship feels one-sided stem from the fear of being taken advantage of by a noncooperative partner. Such simple behavior seems to be conducive to cooperation under a broad range of circumstances, including those in which we get only a noisy *private* signal about our partner's efforts in the relationship, that is, when our partner does not always know if we are less than satisfied with his effort.

Despite the fundamental importance of cooperation in understanding human interaction in small or large groups, the theory of repeated games, while providing important insights about repeated interactions, does not capture the simple intuition in the paragraph above. When signals are private, the quest for "stable" rules of behavior (or equilibria) typically

☆ This paper is a substantially revised version of earlier working papers: "Repeated Relationships with Limits on Information Processing", "Effecting Cooperation" and a previous version of the current paper. We thank Ted Bergstrom, Larry Samuelson, Alvaro Sandroni and the participants at many seminars at which this paper was presented for helpful comments. A large part of this work was done while Postlewaite visited the Paris School of Economics; their hospitality is gratefully acknowledged. The authors thank the Gould Foundation for financial support, and Postlewaite thanks the National Science Foundation for financial support (grant SES-1260753).

* Corresponding author.
E-mail addresses: compte@enpc.fr (O. Compte), apostlew@econ.upenn.edu (A. Postlewaite).

http://dx.doi.org/10.1016/j.geb.2015.03.010

produces strategies that are finely tuned to the parameters of the game (payoffs, signal structure),[1] or to the assumed sequencing/timing of actions and signals.[2] These "rules of behavior" fail to remain stable when the parameters of the game are changed slightly.[3] Their robustness to changes in timing is typically not addressed, nor is their plausibility. We propose an alternative theory/description of how strategic players accomplish cooperation via realistic and intuitively plausible behavior.

A descriptively plausible theory of cooperation should have realistic strategies. A strategy is a complex object that specifies behavior after all possible histories, and the number of possible histories increases exponentially with the number of interactions. If I and my spouse alternate cooking dinner and whoever cooks can either shirk or put in effort each time he cooks, there will be approximately a billion possible histories after one month. For each of these billion histories, both I and my spouse will have received imperfect signals about the effort put in by the other on the nights they cooked, and for each of the histories, I must decide whether or not to put in effort the next time I cook. It is inconceivable that I recall the precise history after even a month let alone after several years.

A more realistic description is that I rely on some summary statistic in deciding whether or not to put in effort – the number of times it seemed effort was put in over the past several times my spouse cooked, for example. In this way, histories are catalogued in a relatively small number of equivalence classes or states.[4] The pooling of histories into classes is intuitively plausible – not just a mathematically convenient pooling, and my action today depends only on the equivalence class containing the history. The strategies we consider will conform to these desiderata.

A second property of a plausible theory of cooperation is that the equilibrium behavior should be consistent with agents coming to that behavior. In the standard approach to repeated games there is typically no realistic story of how players would arrive at the proposed equilibrium strategies. It seems extremely implausible that players could compute appropriate strategies through introspection in repeated games with private signals.[5] Equilibrium strategies in such a setting typically rely on my knowing not only the distribution of signals I receive conditional on the other player's actions, but also on the distribution of his signals given my actions, something I never observe. Even if one entertains the possibility that players compute equilibrium strategies through introspection there is the question of how the players might know these signal distributions. Alternatively, one might posit that players could "learn" the equilibrium strategies. However, the set of strategies is huge and it is difficult to see how a player might learn which strategies work well. Even if one restricted attention to strategies that are deterministic functions of histories, finding an optimal strategy amounts to finding an optimal partition of the histories among all possible partitions.

Our view is that the second property – that agents might come to behave in accordance with equilibrium predictions – calls for a restriction on the strategies that agents choose among. The nature of the restriction one imposes is a critical issue. In this paper, we propose a priori constraints on how agents process signals, motivated by plausible psychological considerations or cognitive limitations, and ask when the restricted family of strategies so generated is conducive to cooperation.[6]

To summarize, our goal is to find sets of strategies that have the following desirable properties: (i) the number of strategies in the set should be small enough that players might ultimately learn which perform well; (ii) the strategies should be based on a cataloging of histories that is intuitively plausible; (iii) the sets of strategies allow agents to cooperate under a broad set of circumstances; and (iv) equilibrium cooperation obtains in a way that is robust to the parameters of

---

[1] See in particular the belief free literature in repeated games (Piccione, 2002; Ely and Välimäki, 2002). See Compte and Postlewaite (2013) for a critique.

[2] Repeated relationships are typically modeled as a stage game played repeatedly, with the players choosing actions simultaneously in the stage game. In reality, the players may be moving sequentially and the signals they get about others' actions may not arrive simultaneously. The choice to model a repeated relationship as simultaneous play is not based on a concern for realism, but for analytic convenience. A plausible theory of cooperation should not hinge on the fine details of the timing of actions: we should expect that behavior that is optimal when play is simultaneous to be optimal if players were to move sequentially.

[3] Fundamental to the standard approach to repeated games with private signals is the analysis of incentives of one party to convey to the other party information about the private signals he received, either directly (through actual communication), or indirectly (through the action played). Conveying such information is necessary to build punishments that generate incentives to cooperate in the first place.

Incentives to convey information, however, are typically provided by making each player indifferent between the various messages he may send (see (Compte, 1998; Kandori and Matsushima, 1998)), or the various actions he may play (belief free literature). There are exceptions, and some work such as Sekiguchi (1997) or Compte (2002) *does* have players provided with *strict* incentives to use their observation. But, these constructions rely on fine tuning some initial uncertainty about the opponent's play (as in the work of Bagwell, 1995).

Finally, when public communication is allowed and signals are not conditionally independant, strict incentives to communicate past signals truthfully may be provided (Kandori and Matsushima, 1998), but the equilibrium construction relies on simultaneous communication protocols.

[4] Aumann (1981) suggested modeling agents as having a finite number of states of mind. This has led to the study of repeated game equilibria in which players are *constrained* to using finite automata (Rubinstein, 1986; Neyman, 1985, 1998; Ben-Porath, 1993), and to the study of repeated game equilibria in which strategies can be *implemented* by simple automata (Abreu, 1986), or approximated by finite automata (Kalai and Stanford, 1988).

[5] In a different context (that of repeated games with perfect monitoring), Gilboa (1988) and Ben-Porath (1990) have expressed a related concern, distinguishing between the complexity associated with *implementing* a repeated game strategy, and the complexity associated with *computing* best response automata. Our concern is not computational complexity per se, but rather the ability to perform relevant computations without precise knowledge of distributions.

[6] Although Aumann (1981) is not motivated by learning considerations, he mentions that assuming a bound on the number of states of mind would "put a bound on the complexity a strategy can have, and enable an analysis in the framework of finite games." In particular, in the context of a repeated prisoners' dilemma with perfect observations, he reports an example in which only few strategies are compared. Although Kalai et al. (1988) argue that the example lacks robustness, the path we follow is in the spirit of Aumann's example.