Best-experienced-payoff Dynamics and Cooperation in the Centipede Game

William H. Sandholm. University of Wisconsin-Madison. USA
Segismundo S. Izquierdo. Universidad de Valladolid. Spain
Luis R. Izquierdo. Universidad de Burgos. Spain

The Centipede game, the finitely repeated Prisoner's Dilemma, and related examples can be viewed as models of relationships in which each participant has repeated opportunities to take costly actions that benefit his partner, and where there is a commonly known date at which the interaction will end. In these settings, experimental and anecdotal evidence suggests that cooperative behavior often persists until close to the terminal date. However, when these games are analyzed in classical game theory, they are solved applying the logic of backward induction, and the obtained conclusion is clear-cut: there will be no cooperation at all. This discrepancy between the conclusions of backward induction reasoning and observed behavior (and most people's intuition about rationality) is a basic puzzle of game theory.

This paper studies a dynamic model of behavior in games that maintains the assumption that agents respond optimally to the information they possess. However, rather than imposing strong assumptions about agents' knowledge of opponents' behavior, we suppose instead that agents' information comes from direct experience playing the strategies available to them. We analyze this model in the Centipede game.

To be precise, following the standard approach of evolutionary game theory, we suppose that two populations of agents are recurrently randomly matched to play the two-player game. Individual agents receive an opportunity to switch strategies at random times. When the opportunity comes, the revising agent decides whether to continue to play his current strategy or to switch to an alternative one. To make this decision, the agent tests every strategy available to him by playing it against one agent drawn at random from the opposing population (with each strategy tested against a newly drawn opponent). The revising agent then switches to the strategy that achieved the highest payoff in the test, breaking ties in favour of the least cooperative strategy.

This model leads to conclusions dramatically different from the classical prediction, i.e. we obtain that there is an almost globally attracting state that exhibits high levels of cooperation.

To evaluate the robustness of these results, we alter our model by replacing a fraction of each population with "backward induction agents" who always stop at their initial decision node. Cooperative behaviour among the remaining agents always exhibits some persistence, with the exact degree depending on the proportion of "backward induction agents" introduced and the length of the game. For longer games (specifically, games with 20 decision nodes) cooperative behaviour is markedly robust, persisting even when two-thirds of the population always stops immediately.

Another robustness test we conduct consists in increasing the number of trials used to test each strategy. It seems clear that if the number of trials is made sufficiently large, so that the agents' information about opponents' behaviour is quite accurate, then the population's behaviour should come to resemble a Nash equilibrium, all of which entail stopping at the initial node. Our analysis shows, however, that stable cooperative behaviour can persist even for substantial numbers of trials (as large as 200).