**T. Offerman: Beliefs and decision rules in public good games. Theory and experiments.**
Dordrecht: Kluwer, 1997, 248 pages, ISBN 0-792-9991-9

*On expectations, preferences, and learning*

This book presents an interdisciplinary approach to the study of individual decisions in public good games. There are various ways to separate the various elements found in this book. First, as Offerman asserts, this book is interdisciplinary in the sense that it considers theories, ideas and hypotheses advanced from the fields of economics as well as social psychology. Second, it is explicitly divided into three parts; a theoretical part, an experimental part, and finally a summary and interpretation of the results. But in addition, there appears to be another important partition in this study. As the title explains, the study focuses on public good games, and this is one perspective from which to consider it. But at the same time the kind of questions addressed by the author are

*"Do players appreciate the strategic nature of the game? Do they learn from their experience? Are they concerned about the payoffs of other players? Do systematic differences exist between the decision rules of various players?"* (back cover)

Obviously, these are questions that have a much more general meaning than merely in the context of public good games. In other words, this is a broad study of individual behavior in general. And it is this last aspect that is one of its main strengths.

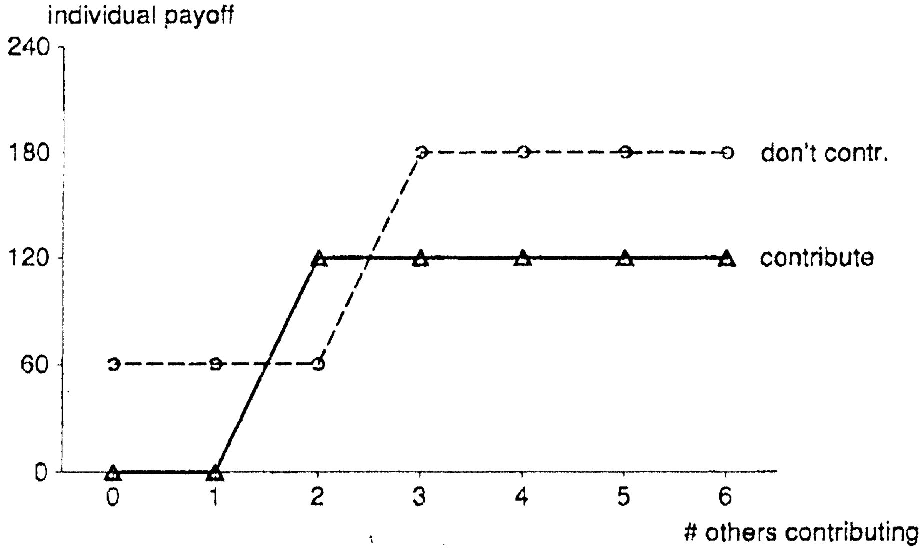


**Fig. 1.** Payoff individual player in step-level public good game, with 7 players, c = 60, and f = 180

What is the exact issue we are dealing with here? That is, what is the public good game considered, and why does it deserve so much attention? Suppose there is a group of people that might be made happy collectively with some public good, for example some roads, an army, public lighting, or a dike. Next, suppose that these people individually but simultaneously decide whether to contribute or not to such a public good. Due to fixed startup costs, there is a minimum aggregate contribution needed before the public good can be provided, that is, only if enough people contribute will the public good be realized. In the specific case of a step-level public good considered in this study, when the threshold is not reached, all contributions made are money thrown into the air, and also all contributions beyond the threshold are completely useless. Contributions made are never refunded to the individuals, no matter that they were useless or redundant.

In the experiments played, in each round (the standard game is played repeatedly) an individual player would get an endowment of c cents, and then face the binary choice whether to contribute c or not. If at least 3 players in the group contributed, the payoff would be f − c (with f ≥ 3c in all experiments) for each player, whereas it would be 0 otherwise. Hence, if the public good is realized, a contributor will end up with c − c + (f − c) cents, whereas a free rider would have c + (f − c) cents. If the public good fails to be realized, a contributor ends up with c − c cents, whereas the other players would still have their c cents. Hence, it seems better not to contribute. The structure is very similar to a n-person Prisoner's Dilemma or social dilemma game. Figure 1 highlights the structure of this game.[1]

[1] The author stresses that the game resembles more the game Chicken than a Prisoner's Dilemma because there are Nash equilibria that coincide with the Pareto optimum, namely those where exactly the critical number of players contributes, i.e., at the second intersection of the two curves in Figure 1. Although this is correct as such, in Chicken all three Nash equilibria are Pareto efficient. Moreover, for the game Chicken the analogue of Figure 1 would correspond only to the right-hand side, from 2 other players contributing onwards, and the left-hand part would disappear.

Notice that if all players contribute, their payoff (120 cents) is higher than if all players do not contribute (60 cents). But notice also that the payoff for a non-contributing player is always above that for a contributing player, no matter how many other players contribute. There is exactly one exception, the special case of the critical contributor. Suppose two other people contribute. In this case an additional individual player's decision would be critical for the realization of the public good. If such a critical contributor does not add his support, the project will not materialize, and he will keep his c cents. However, if the critical contributor does help to realize the project, he will end up with $c-c+(f-c)$ cents. Since $f-c>c$ in all experiments, the critical contributor would be better off contributing. From a public good point of view, the critical contributor seems somewhat artificial, since the probability of an individual player being the critical player approaches zero in a large society. But from an individual behavior point of view, this feature makes for an intriguing game, for one of the crucial questions for an individual player might be how many other players he expects to contribute. Being critical really pays here. Part of this study is focused on measuring these beliefs or expectations.

But even for given beliefs, different players might make different contribution choices, as they might apply different decision rules.[2] Some people care only about their own payoff, while other people may care about the joint payoff with other people. Psychologists have paid a lot of attention to measuring these motivations, whereas economists tend to equate preferences with maximizing one's own expected value as measured by a monetary payoff (but as Offerman shows one can formulate expected payoffs including utility assigned to the outcomes of others). The players' underlying motivations are measured and various hypotheses in this respect are tested.

As to beliefs, these may surely change during the periods played. Whenever the number of other players contributing does not correspond to what a player has expected, his beliefs are likely to be affected. The author considers various theories to model this (both explicitly and implicitly).

Whereas positive theories of learning were rarely considered until a decade ago, nowadays it seems an intrinsic element of experimental economics. No matter what the economic context, the issue to be analyzed, and the hypotheses to be tested, the experimental design is such that much of the analysis can be focused on the learning behavior of the individual players. An alternative would be to present players with an economic decision problem, and to observe how they decide in a one-off situation. The problem with this latter possibility, as observed by the author, is that the data in such a single period may contain relatively more noise because some subjects only fully understand the rules of a game after some practice. Clearly, this might be a problem, but it could be largely offset by a good, thorough explanation, some exercises or a quiz, and some trial periods

---

[2] This term 'decision rule' in the title might be somewhat confusing. What the author means by *"decision rules"* is more 'decision criteria', or preferences in the economics jargon. That is, the word 'rules' does not imply a procedural approach to bounded rationality here.

without payoff. Two other possible explanations are that the use of computers makes it so easy to play a large number of periods in a relatively short period, and that, because of the fixed costs for the subjects of an experiment, the average costs of recruitment decrease as the number of periods played increases. Whereas these motives are slightly less noble than the one advanced by the author, the great advantage of this all is that the fascinating and still relatively little understood issue of learning gets so much more attention.

Part I contains three chapters: the first one on experimentation in the social sciences, a second chapter on some theoretical foundations of individual behavior in public good games, and finally a chapter on basic experimental tools.

Chapter 1 starts by comparing and contrasting the approaches of economists and psychologists. As Offerman concludes, these are not as different as sometimes asserted, and *"experimental economics provides an incentive for cooperation between economists and psychologists"* (p. 5). The two main differences concern incentives and manipulation of the subjects. Psychologists often do not believe incentives are important, while they accept manipulation. While some experimental evidence would seem to suggest that incentives do not matter for individual behavior, what those people must mean is that monetary payoffs do not always completely control preferences, and that there might be other incentives besides the monetary reward related to the experiment. Since economics is about agents seeking to do the best they can, arguing that incentives do not matter would be asserting that economics does not exist. The author suggests that the argument that thinking is costly does not imply that rewards should be of significant magnitude, because attributing costs to thinking will inevitably lead to an infinite regress: deciding how to decide how to decide how ... While it is undeniable that thinking is costly, I am pretty sure that no player will get caught in such an infinite regress. The problem only appears to occur in a procedurally bounded rationality approach. But Lipman (1991) has shown that such a regress may converge. This is an important mathematical confirmation of the following well-known fact: In reality, agents *do* reach conclusions, choosing what they perceive to be best for them. Lipman's result does not tell us *how and what* people decide, but for the question of incentives in experiments, the infinite regress argument is irrelevant.

As to manipulation of experimental subjects, this is completely out of order in economics. Although Offerman notes that there exists a couple of exceptions, he should also mention that one of those exceptions was omitted, explicitly giving the reason of manipulation, from the Ledyard (1995) survey. Probably this further weakens incentives for economists to use manipulation. Though the brief sketch of the history of the experimental method is somewhat superficial (among other things reinforcing the myth that Vernon Smith basically initiated experimental economics), the chapter as a whole is a good starting point for anybody new to experimental economics.

In chapter 2 the public good game is presented and analyzed in detail. A game-theoretic analysis shows that many Nash equilibria exist, both in pure and mixed strategies, many of which are asymmetric: nobody contributing, a critical number of people contributing, all contributing with some probability, etc. Hence, the players face a double coordination problem. Not only do they need to adjust their actions to each other such that an equilibrium is played, but they also can coordinate on different equilibria with different people contributing and with different payoff consequences (in particular, they might want to avoid the worst equilibrium where nobody contributes). The author also provides some discussion of criteria like symmetry, and payoff-dominance which make some equilibria more likely than others.

Allowing for heterogeneity in the players' preferences greatly complicates matters. Psychologists distinguish, for example, individualists, cooperators, altruists, and aggressors, who might use different decision criteria. Two classic psychological hypotheses, the triangle hypothesis and the false consensus hypothesis, are discussed. The economic approach focuses on expected monetary reward, but the utility function can be modified to allow, for example, for warm-glow cooperators, or material cooperators. Equilibrium conditions for such situations are presented.

Since this is a strategically complicated game, one might assume making errors is natural. One way to deal with this is Quantal Response theory, which assumes that players are in equilibrium even with respect to the errors they make. More natural might be to follow a dynamic, non-equilibrium approach, in which players make errors, but learn as they gain experience. The author focuses on two classes of relatively simple learning models. With naive Bayesian learning, agents update their beliefs about the actions of others using Bayes' rule, and they choose a best-response to the thus given belief. This is naive in the sense that it neglects the strategic aspects of the interaction, in particular the aspect that other players might be updating and hence changing their actions as well. In fact, this is just a form of best-reply dynamics mediated through explicit beliefs. The other class of learning models upon which Offerman focuses is reinforcement learning. Reinforcement learning is based on the principle that players try actions, and actions that lead to better outcomes are more likely to be repeated. It is a minimal form of modeling learning, in the sense that one does not need to make many assumptions about the reasoning procedures followed by the agents. In particular, no explicit expectations or beliefs are modeled.

In chapter 3 the experimental tools are presented. To elicit the players' beliefs as to what other player might do, and their value orientations or preferences, the following incentive compatible mechanisms are used. (Notice that this latter aspect differs in an important sense from many other studies in the field. It is important because otherwise people can say what they want – usually something simple – without any need to report truthfully.) The decomposed game technique is used to assess an individual's preferences towards outcomes for himself or others. A player goes through a set of 24 pairwise choices in which each choice allocates a certain amount of money to himself and to an anonymous player. Since the chosen allocations are actually carried out, a player has no reason not to act in accord with his preferences. A scoring rule is used to elicit beliefs about the actions of other players in the public good game. Since payoffs in this stage depend on the difference between the stated beliefs and the actual outcome, it is in the interest of a player to report his true beliefs as best as he can.

In general, there seems to be a problem with the use of the argument that some players' actions can be explained by the fact that they care so much about others. As mentioned by various economists, these explanations are somewhat suspect because, if a player really cares about others, why would he not take as much money as possible in an experiment, and then give it outside the lab to those who really need it? After all, there is some tension between the altruism argument and the way experiments are organized in economics, with the effort made to implement anonymity as strictly as possible. In the experiments reported in this book this problem largely disappears because the preferences are elicited not only in an incentive compatible way, but also in the same experimental fashion (with anonymity) as the ensuing public good game.

Part II reports on two different sets of experiments; chapter 4 focuses on decision rules or criteria with given beliefs, and chapter 5 concentrates on beliefs and learning.

In chapter 4 many different tests of specific hypotheses from both psychology and economics are presented. We will not go into the details of all separate findings, because no general definite conclusion presents itself. While this might be disappointing to some, it seems inherent to contributing to a positive theory of learning. There is so much that we do not know about individual behavior in economic contexts, and it seems so unlikely that one great organizing principle governs all such human behavior, that one might be better off carefully investigating experimental data, and looking for regularities, as is done in this study.

Typical average contribution levels are 40 to 50%. With 7 players, the minimum number of contributors needed would be 43% (3 out of 7), and with 5 players this would be 60%. There are strong period to period fluctuations, but not much of a trend is evident. One treatment stands out. With group size 7, and $f = 180$ and $c = 60$, we almost converge to the 'no contributions at all' equilibrium. In fact, we see a clear downward trend, and in period 18, the average contribution by cooperators is 0, while only 10% of the individualist contribute.

The tested preferences or value orientations show that individual propensities to contribute vary substantially: 61% of the players are classified as individualist, 30% as cooperator. In the public good game the cooperators contribute more than the individualists. As to the transformed utility functions (allowing for more than just the player's own monetary payoff), Maximum Likelihood estimates appear to favor the hypothesis of warm-glow cooperators. Concerning the influence of the players' beliefs about the actions of others, it turns out that players are more inclined to contribute if they believe their contribution is critical. Also, the effects of increasing the payoff of the public good, changing the group size, and having players interacting either with the same group or with a changing group of players is documented and analyzed. Interesting is the following framing effect. The same public good experiment can be presented as a public bad experiment. Take, for example, a common pool, like an ocean with whales. If too many contribute to fishing, the whales disappear. The experimental outcomes are significantly better (reaching the threshold of cooperative actions more often, leading to higher average payoffs for the subjects) in the public good game than in the public bad game. The analysis in chapter 4 is static. All observations are averaged or aggregated from period 1 to period 20, and related to the given preferences and expectations at the beginning of the experiment, as if no interaction with changing beliefs and learning took place.

In chapter 5 the possibility of learning and changing beliefs is considered. This is done with the given preferences of the players, but it would have been interesting as well to see whether these value orientations change over time. It might very well be that these are not so much a given psychological feature of individual players, but more a socially learned characteristic.

Every period the players report their expectations. The updating of these beliefs seems to be largely consistent with naive Bayesian hypotheses. To check how far the players do indeed appear to abstract from strategic considerations, some control treatments with games against nature were run, in which the players faced an investment problem that closely resembles the public good game. (One could imagine the other players replaced by automata that make the decision to contribute or not with given, fixed probabilities). The results in these games against nature seem to support the idea that the players are non-strategic, naive Bayesian updaters also in the standard public good game. One anomaly reported by the author is that the players need considerably more time in the simpler, non-strategic, games against nature. An explanation might be that the public good game, while being a strategic game, is also a game with a double coordination task. Hence, in the standard game a player might see other players as helping to 'solve' the game, and he just needs to adapt to the actions of the others. In the game against nature, a player has to do the job all by himself. Similar evidence was found with respect to experiments reported in Egidi (1994).

Finally, some simulations are presented with the two learning models upon which this study focuses: naive Bayesian learning and reinforcement learning. These simulations are intended to answer the following two questions: Do they reproduce the stylized facts of the experimental data? And do we observe convergence of the players' actions in the long run, possibly long after the human players would have been exhausted. Simulations of both learning models seem to trace the dynamics of the experiments quite well. Most simulations eventually converge to some Nash equilibrium, usually those most similar to what was observed in the experiments.

The concluding part III evaluates the results in chapter 6. Following a similar approach by Roth and Erev (1995), Offerman proposes to classify games according to the kind of rules used by the players. The author conjectures that in simple games players might reason enough about the game to reach a Nash equilibrium, and that difficult games with multiple equilibria might be best characterized by belief learning, while in really complicated games reinforcement learning might be best. Although this seems a very useful and promising approach, and this book contributes significantly to such a project, some remarks are in order.

It is not entirely clear how Offerman would interpret such a classification, for some alternative forms of learning like imitation, or learning direction theory are mentioned already by the author elsewhere in this book. In fact there are so many learning possibilities (or to put it another way, there seems to be so little we know about how learning by humans actually takes place, and there seems to be so much about the human brain that we as economists are unlikely to know ever), that we might want to settle for modest modeling attempts, using simple learning models, not because we believe the cognitive processes used by the players are that simple, but because they are so complex. In other words, there might be good reasons to favor the old fashioned '*as if*' argument.

An example is reinforcement learning, cognitive processes modeled at the most basic level. The agents in reinforcement learning models do not need to know what payoff or utility is, but only recognize when they feel happier in the sense that they are more likely to repeat actions that triggered such a feeling in the future. But this does not necessarily imply that the players whose actions seem reasonably well described by such a model do not engage, consciously or unconsciously, in complicated mental processes. In fact, naive Bayesian learning and reinforcement learning are not that different. When reinforcement learners change the propensity to choose a certain action, this can be interpreted as a change in beliefs. In other words, reinforcement learning is a form of belief learning. [A recent integrated approach of these two types of learning models is carried out by Camerer and Ho (1996).]

To conclude, this is definitely recommended reading material, especially for those interested in the public goods issue, and social dilemmas in a more general sense. But since it is so rich in its overview of the recent experimental literature and methods, it will also prove very useful for all those interested in economic behavior in general, and in expectations, preferences, and learning in particular.

## References

Camerer C, Ho TH (1996) Experience-weighted attraction learning in games: a unifying approach. Mimeo

Egidi M (1994) Routines, hierarchies of problems, procedural behaviour: some evidence from experiments. IIASA (WP-94-58)

Ledyard JO (1995) Public goods: a survey of experimental research. In: Kagel JH, Roth AE (eds) Handbook of experimental economics. Princeton University Press, Princeton, NJ

Lipman B (1991) How to decide how to decide to ...: modeling limited rationality. Econometrica 59: 1105–1125

Roth AE, Erev I (1995) Learning in extensive-form games: experimental data and simple dynamic models in the intermediate term. Games and Economic Behavior 8: 164– 212

*Nicolaas J. Vriend, Department of Economics, Queen Mary and Westfield College, University of London, Mile End Road, London, E1 4NS, UK*
(e-mail: n.vriend@qmw.ac.uk)