

Discrete Public Goods: Contribution Levels and Learning as Outcomes of an Evolutionary Game

Christiane Clemens *Thomas Riechmann**

Diskussionspapier Nr. 263

September 2002

Abstract

This paper examines the learning dynamics of boundedly rational agents, who are asked to voluntarily contribute to a discrete public good. In an incomplete information setting, we discuss contribution games and subscription games, the latter including a money-back guarantee in case of provision failure. The theoretical results on myopic best response dynamics implying striking differences between strategies played in the two games are confirmed by simulations, where the learning process is modeled by an Evolutionary Algorithm. We show that the contribution game even aggravates the selective pressure leading towards the non-contributing equilibrium, thereby supporting results from laboratory experiments. In contrast to this, the subscription game removes the 'fear incentive', implying a higher percentage of successful provisions over time.

ISSN 0949 – 9962

JEL classification: C6 – C73 – D83 – H41

Keywords: bounded rationality, evolutionary games, evolutionary algorithms, learning, public goods

*Universität Hannover — Institut für Volkswirtschaftslehre, Abteilung Wachstum und Verteilung, Königsworther Platz 1, 30 167 Hannover, e-mail: clemens@vwl.uni-hannover.de and riechmann@vwl.uni-hannover.de

1 Introduction

The standard prediction on individual behavior when it comes to the question of voluntarily contributing to a public project is that strategic incentives to free-ride on the contributions of others lead towards a Nash-equilibrium characterized by a Pareto-inefficient underprovision of the public good.¹ This problem is even aggravated in incomplete information settings, since additionally there are no incentives to reveal individual preferences truthfully, a phenomenon Cornes and Sandler (1996) label as 'informational' free-riding.

Throughout the last two decades, attention has been drawn towards the analysis of discrete public goods, which are characterized by the feature that successful provision requires a certain amount of money to be raised or a minimum number of participants. Recent research on this subject shows that free-riding ceases to be the dominant strategy, but rather that discrete public good games are characterized by multiple equilibria, where efficient provision is one of the possible outcomes (see e. g. Palfrey and Rosenthal, 1984; Bagnoli and Lipman, 1989; Admati and Perry, 1991; Marx and Matthews, 2000).

In this context it is important to distinguish between two types of discrete public good games: the *contribution game* and the *subscription game*.² While contributions are lost in the first type, if the public good is not accomplished, they are refunded in the latter, thereby eliminating what van de Kragt *et al.* (1983) and Palfrey and Rosenthal (1984) call the 'fear incentive' adversely interfering with successful completion of the project. Nevertheless, efficient provision is a possible outcome for both types of games.

By summarizing the results from the literature, it is possible to identify three sets of problems which deserve particular treatment: Palfrey and Rosenthal (1984) emphasize the participation issue and point out that there is a problem concerning demand revelation. The latter arises especially in the incomplete information setting underlying the analysis of this paper. Last, due to the multiplicity of equilibria, there is a problem of equilibrium selection. So, do agents really coordinate themselves and cooperate to obtain efficient outcomes, and if so, how? A common argument brought forward in this context is that learning processes — for instance, as modeled by evolutionary game theory — can serve as a means for equilibrium selection and will be a central aspect of our work (see Marimon, 1993; Vega-Redondo, 1996; Samuelson, 1997; Fudenberg and Levine, 1998).³

¹A comprise presentation of the theory of public goods can be found in Cornes and Sandler (1996).

²Admati and Perry (1991) were the first to use this terminology.

³Already Bliss and Nalebuff (1984) argue that in a dynamic context an agent has the opportunity to learn the response of other players. Palfrey and Rosenthal (1984, p. 191) state

More specific, this paper explores the question whether boundedly rational individuals, who possess little information on the structure of the economy, learn best responses when asked to voluntarily contribute to a discrete public good. We examine contribution and subscription games in the presence of incomplete information, (a) within a static framework, (b) by extending the analysis with the myopic best response dynamics of a repeated one-shot game for the special case of uniform behavior, and (c) by providing numerical simulations for the case of heterogeneous actions. In the latter, the decision process of agents is described by an Evolutionary Algorithm (EA), which belongs to the general class of adaptive learning algorithms and serves as a means to simulate the evolutionary game; see Riechmann (2001*a,c*).

Up to now there are only few approaches dealing with (threshold) public goods in incomplete information settings. Nitzan and Romano (1990) assume an uncertain cost structure and find that the Nash-equilibria lose their efficiency property such that free-riding strategies reappear. Marx and Matthews (2000) as well as Gradstein (1992, 1994) and Vega-Redondo (1995) employ models with *Bayesian learning*, where each agent has statistical information on preferences or donations of others which are periodically updated. The results are ambiguous, in that they find Bayesian equilibria with and without completion of the public project.⁴

Our approach is closely related to Menezes *et al.* (2001) but differs from their work to at least two major aspects: First, our model is embedded in a standard textbook decision problem of agents who split their endowments between consumption of a private good and a contribution to a public threshold good, where individual utility is of the quasilinear form. Second, we employ myopic best response learning in order to derive results regarding long-run dynamics and stochastic stability of equilibria, whereas Menezes *et al.* (2001) discuss dynamics for a single-object auction.

Our theoretical results as well as outcomes from the simulation of heterogeneous behavior show a general superiority of subscription games over contribution games regarding the possibility of attaining the efficient equilibrium, where the public project is completed. Yet, we also demonstrate for both types of games that the threshold public good is less likely to be

in their concluding remarks: ‘Because of the multiplicity of equilibria, learning or some form of coordination, is probably very important (possibly necessary) for the attainment of a Nash equilibrium’. Cavaliere (2000) solves the coordination problem by allowing for communication and correlated strategies.

⁴Models of Bayesian learning are often criticized for the assumptions they impose on individual information processing capacities. We drop the assumption of common knowledge, that is, every player in isolation has full knowledge of the relevant data and can costlessly figure out all equilibria. Instead, our analysis relies on the assumption of *boundedly rational* agents who are involved in an adaptive learning process.

accomplished throughout the learning process, if we introduce sources of randomness such as mistakes in strategy formation into the game. In this case, chances are notably biased towards the inefficient non-contribution equilibrium.

In this respect, our results support findings from laboratory experiments on public good provision.⁵ Isaac *et al.* (1989) and Cadsby and Maynes (1999) find that provision is encouraged in subscription games *vis-à-vis* contribution games,⁶ with the impact of the money-back guarantee increasing with the threshold level. That successful provision of the public good is negatively related to the threshold level, is supported by results of Suleiman and Rapoport (1992). Isaac *et al.* (1989) also see evidence for ‘cheap riding’, which describes a situation where agents have the incentives to obtain equilibrium outcomes with an unequal distribution of contributions.⁷ Moreover, they find that zero contributions occur far more often in the contribution game when compared to free-riding strategies in the continuous public good case, but they observe a general decay of contributions over time in all of the considered types of games. Bagnoli and McKee (1991) test subscription games. They focus on Nash-refinements, the efficiency of outcomes and group size effects and find the impact of the latter negligible.

Our analysis will proceed as follows. The analysis of section 2 is devoted to the theoretical analysis. We present the general assumptions of the model and discuss static as well as myopic best response equilibria for both, the contribution and the subscription game. Section 3 analyzes the learning dynamics, describes the basics of the model of EA-learning, and derives general results for the intertemporal performance of strategies. In section 4, we discuss the simulations. Section 5 concludes.

2 Private Provision of a Discrete Public Good under Incomplete Information

2.1 The Model

We consider a model of voluntary contribution to a discrete public good under incomplete information with $n \geq 2$ consumers, one private good x and one pure public good G . Each consumer i is endowed with exogenous wealth w_i , which he divides between private consumption x_i , and a con-

⁵For an excellent survey on public goods experiments, see Ledyard (1995).

⁶Dawes *et al.* (1986) present contradicting results, which can mainly be ascribed to the fact that they analyzed one-shot games without repetition, which is of secondary interest for us, since we are primarily concerned about learning and the evolution of strategies over time

⁷More specific: ‘If one’s own contribution is indispensable, then better stay on the cheap side’.

tribution to the public good b_i , such that the budget constraint $w_i = x_i + b_i$ is satisfied. For simplicity, all prices and marginal costs are normalized to unity. The utility function of agent i in general is assumed to be of the quasi-linear form

$$U(x_i, G) = x_i + \beta_i \ln G, \quad (1)$$

which allows us to abstract from income effects.⁸ The individual valuation of the public good, $\beta_i > 0$, is private information. In what follows, we refer to β_i as the true valuation in order to distinguish it from the actually reported value b_i .

The objective of agent i now is to determine the best-response function b_i^* , while taking the reports b_{-i} of other agents as given. To keep things simple, we will assume the n agents of the economy to be identical with respect to their preferences and endowments, that is $\beta_i = \beta$ and $w_i = w$.

The dynamic game is modeled as a simultaneous move, repeated one-shot game, where, in each period, player i receives an endowment $w(t)$ and chooses a report $b_i(t)$ from the discrete set of feasible contributions

$$\mathbb{B} = \{b^1 = 0, b^2 = \delta, b^3 = 2\delta, \dots, b^m = (m-1)\delta = \beta\}, \quad (2)$$

which is defined over a finite grid. The actions are equally spaced in the interval $[0, \beta]$, the distance between two neighboring actions given by $\delta > 0$. The number of different actions in the set is $\#(\mathbb{B}) = m$. By assuming β to be the largest contribution, we imply that a rational individual never contributes more to the public good than her maximum willingness to pay.⁹

The public good is assumed to be discrete, which means it is only provided, if contributions are collected to an amount sufficiently large to cover a given threshold level T .

$$G = \begin{cases} 0 & \text{if } \sum_{i=1}^n b_i < T \\ \sum_{i=1}^n b_i & \text{if } \sum_{i=1}^n b_i \geq T \end{cases} \quad (3)$$

The agents' contributions are regarded to be perfect substitutes. If aggregate contributions exceed the threshold, there is no rebate, thereby assuming that excess contributions are wasted. In general, we will assume $T > \beta$, which implies that no agent is able to complete the public project on her own. From (3) it becomes obvious that there exist n^m action profiles, i. e.

⁸From the standard textbook model with this preference specification, it is well known that the Pareto-efficient quantity of a continuous public good in a perfect information environment is given by $G = \sum_{i=1}^n \beta_i$ and independent of individual endowments; see Cornes and Sandler (1996, Ch. 7).

⁹For technical reasons, we require the action space to be countable. \mathbb{B} defines a continuous set for $\delta \rightarrow 0$.

combinations of actions of the agents of the population, not all of them amounting to T in total.

Definition 1 (Symmetric cost share) Let $\tau := T/n \in \mathbb{B}$ be the symmetric cost share, a player has to contribute in order to accomplish provision of the public good, and define the cost–preference index with $\tau/\beta \in [0, 1]$.

The symmetric cost share τ is contained in the set of feasible actions \mathbb{B} and can be interpreted as a per capita threshold. The ratio τ/β then measures the individual cost of provision against the marginal willingness to pay. It can never exceed unity, as no agent is willing to spend more than her valuation on the provision of the public good.

Let us now focus on individual strategies. From the perspective of an arbitrary player i , there are three relevant states of the society:

Definition 2 (States of the Society)

S1 ('never') $(n-1)b_{-i} < T - \beta$: Even if player i contributes the maximum amount she is willing to pay, the public good will not be provided.

S2 ('pivot') $(n-1)b_{-i} = T - b_i$ for $b_i \leq \beta$: Player i is the pivot individual, which means that i 's participation is essential for the public good being completed.

S3 ('always') $(n-1)b_{-i} > T$: The public good is provided even without the contribution of individual i .

For a uniform behavior of players, $b_i = b$, $\forall i = 1, \dots, n$, the states of the society are given by

$$S1: \quad b < \tau, \quad S2: \quad b = \tau, \quad S3: \quad b > \tau. \quad (4)$$

The analysis now proceeds as follows: We start with a description of the two static games, specify best replies for the different states of the society and derive the corresponding Nash–equilibria for the case of symmetric behavior. We will demonstrate that, similar to Menezes *et al.* (2001), both games are characterized by multiple equilibria. This gives rise to the problem of equilibrium selection which might be solved by learning processes. Additionally, we will show that the contribution and the subscription game differ significantly with respect to uniqueness of best responses and stability of equilibria.

The learning process of section 3 relies on the concept of replicator dynamics of evolutionary game theory. Since the dynamic performance of every single action from the action space is represented by a first–order stochastic difference equation and the game in general is characterized by heterogeneous agents, analytical results are very difficult to obtain. However, to give an intuitive understanding of the key parameters governing

the stochastic imitation dynamics of the learning process, we first discuss *uniform myopic best response learning* as an auxiliary model, which helps to examine the asymptotic properties of the equilibria.

2.2 The contribution game (CG)

The contribution game is characterized by a situation where individual contributions are lost if the aggregate amount collected falls short the threshold T . In short, there is no money-back guarantee in case of provision failure.

The utility function of a typical player i in the contribution game becomes

$$U(x_i, G) = \begin{cases} w - b_i & \text{for } \sum_{j=1}^n b_j < T \\ w - b_i + \beta \ln G & \text{for } \sum_{j=1}^n b_j \geq T \end{cases} \quad (5)$$

Symmetric equilibria in the static contribution game The best responses of agent i for a given state of the society S1 – S3 are unique and can be obtained as follows:

$$b_i^* = \begin{cases} 0 & \text{if S1} \\ T - (n - 1)b & \text{if S2} \\ 0 & \text{if S3} \end{cases} \quad (6)$$

If the society is in state S1, from the perspective of a typical agent i , even contributing the maximum willingness to pay β will not be sufficient to complete the public project. Thus, each unit spent on the public good is wasted and should better be spent on the consumption of the private good, thereby yielding higher utility. The same argument applies for the best response in state S3. In state S2, player i is the pivot agent. Her best response here is to offer the residual between the threshold level and the sum of contributions of the (uniformly acting) other agents. By Definition 2, the residual may fall below but cannot exceed her true valuation β .

Proposition 1 (Nash-equilibria of the static contribution game) *S1 and S2 constitute symmetric Nash-equilibria of the static contribution game according to the best response correspondences*

$$b^* = \begin{cases} 0 & \text{if S1} \\ \tau & \text{if S2} \\ 0 & \text{if S3} \end{cases} \quad (7)$$

S3 does not constitute a Nash-equilibrium.

A state of the society constitutes a Nash-equilibrium, whenever individual best replies lead back to the respective state. As generally expected for the

case of a discrete public good, the static contribution game exhibits multiple equilibria. The best response of contributing nothing in state S1 follows immediately from (7), since any positive amount offered for completion of the public good is wasted. In S2, the society rests in a situation of each agent offering the per capita cost share and the public project is completed. By Definition 1, the best response in state S2 is characterized by $\tau \leq \beta$, with the value depending on the exogenously fixed threshold level $T \geq \beta$. Whenever completion of the public good takes place without an agent's own contribution, her best reply simply is to offer $b = 0$. However, since this argument holds for every player of the game, S3 cannot constitute a symmetric Nash-equilibrium.

Uniform myopic best response dynamics in the contribution game The static game now is extended to a model of myopic best response which can be regarded as a naïve variant of fictitious play. As a simple rule for expectation formation, myopic best response describes the dynamic behavior of agents, who possess only limited information on the actions of other players. Each of the n agents chooses her best reply according to Proposition 1 under the prediction that in period t all other players stick to their previously played strategies. Our analysis will focus on symmetric behavior, as is standard in games of incomplete information with ex ante identical agents (see Menezes *et al.*, 2001). This means all players are assumed to play identical strategies in period t even though multiple best responses to the respective state of the society of the previous period might be available. Under this assumption, the resulting dynamics can be described by a simple Markov process constituted by the transition probabilities between the possible three states of society.

I. Myopic best response equilibria without noise

For the contribution game, best responses are unique, which implies fairly simple uniform myopic best response dynamics:

Proposition 2 *The transition matrix between states S1, S2, and S3 for uniform myopic best response dynamics in the contribution game is given by*

$$P_{CG} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{pmatrix} \quad (8)$$

If the state of the society is S2 in period t , best responses lead back to S2 in $t + 1$. This means that S2 is an absorbing state. If, on the other hand, the society is in S3 in t , best responses will lead to S1 in $t + 1$, while a society in S1 will stay in S1 in $t + 1$. Thus, the uniform myopic

best response dynamics are characterized by two absorbing states, S1 and S2, and one reflecting state S3.

In the case of *uniform* best responses, only one single strategy profile constitutes S2: Each of the n players plays $b_i(t) = \tau$. All differing profiles constitute either S1 or S3. Under the regime of uniform myopic best response dynamics, this means that each profile deviating from $b^* = \tau$ ends up in the absorbing state S1. Hence, S2 is a stationary yet unstable state, while S1 is the only (asymptotically) stable Nash-equilibrium.

The system itself is highly state dependent: If the historically very first state of the society is S2, the society will remain in S2 forever. On the other hand, if the initial state is S1 or S3, the system will converge to S1.

Summarizing the results, two long-run outcomes of uniform myopic best response learning are possible: (a) provision of the public good in S2 or (b) non-provision in S1, the latter being more likely regarding the remarks on state dependency given above.

II. Myopic best response equilibria with noisy imitation in the contribution game

The stability analysis of equilibria can now be extended with a random element. The assumption of uniform behavior is maintained by assuming that the agents' strategies are subject to an aggregate shock. With a small probability ϵ , the agents do not play a best response, but are subject to a common random mistake. The probability of being an outcome of this error is equally distributed among all strategies of the action space. This of course includes the best response strategy, i. e. the possibility of 'erroneously doing the right thing'. The probability ϵ will be called the *mutation probability* and affects the transition probabilities between the three states of the society.

Since we are especially interested in the stability properties of S2, we will treat S1 and S3 as a joint class. The transition probabilities for reaching, leaving or staying in S2 can then be derived as:

$$p_{\leftarrow S2} = \frac{m-1}{m} \epsilon \quad (9)$$

$$p_{S2, S2}^m = 1 - \frac{m-1}{m} \epsilon \quad (10)$$

$$p_{\rightarrow S2} = (p_{S1, S2}^m + p_{S3, S2}^m) = \frac{1}{m} \epsilon. \quad (11)$$

Figure 1 illustrates the corresponding Markov chain. It follows immediately:

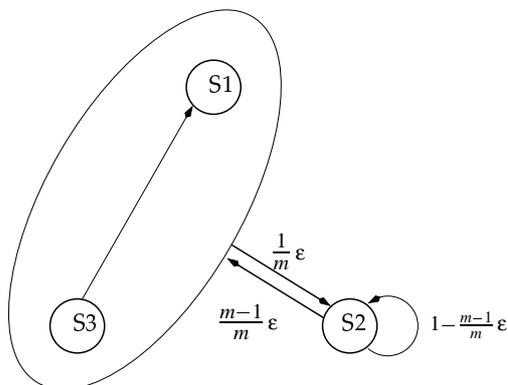


Figure 1: *Transition probabilities of myopic best response with noisy imitation in the contribution game*

Proposition 3 *The pivot state S2 of the contribution game is stochastically unstable under uniform myopic best response dynamics with aggregate noise. S1 is stochastically stable.*

Proof:

$$p_{\leftarrow S2} > p_{\rightarrow S2} \quad \text{for } m > 2$$

S3 is no Nash-equilibrium. □

Hence, S2 is more likely to be left than to be reached and the public project is almost always not completed in the long-run, a result that holds independently of the initial strategy profile.

2.3 The subscription game (SG)

The subscription game is characterized by a money-back guarantee in case of provision failure, that is, if the aggregate contribution is too small and falls short of the threshold T . Then, the endowment w is entirely spent on the consumption of the private good x . As already mentioned above, a money-back guarantee removes the ‘fear incentive’ from individual strategies.

The utility function of a typical player i in the subscription game becomes

$$U(x_i, G) = \begin{cases} w & \text{for } \sum_{j=1}^n b_j < T \\ w - b_i + \beta \ln G & \text{for } \sum_{j=1}^n b_j \geq T \end{cases} \quad (12)$$

Symmetric equilibria in the static subscription game Agent i 's best replies to each of the three states of the society according to Definition 2 can be derived as:

$$b_i^* \begin{cases} \in \mathbb{B} & \text{if } S1 \\ = T - (n - 1)b & \text{if } S2 \\ = 0 & \text{if } S3 \end{cases} \quad (13)$$

The public good is not provided in S1. This result is independent of any arbitrary strategy player i chooses from the action space \mathbb{B} . The payoff is identical for all strategies from the action space, because of the money-back guarantee. The entire budget is spent on the consumption of the private good. In S2, the player is the pivot individual. Her best response is to pay the residual left in order to accomplish the public good. State S3 is characterized by a situation, where the public good is provided even if player i does not contribute. Since any amount exceeding the threshold level is wasted by assumption, the best response is to contribute nothing. A comparison between (6) and (13) shows that the best replies of the subscription and the contribution game only differ with respect to state S1, thereby revealing the strong incentives towards low or zero contributions working in the latter game.

From (13), we can directly derive the best replies in the subscription game for the case of symmetric behavior and determine the Nash-equilibria:

Proposition 4 (Nash-equilibria of the static subscription game) *S1 and S2 constitute symmetric Nash-equilibria of the static subscription game according to the best response correspondences*

$$b^* \begin{cases} \in \mathbb{B} & \text{if } S1 \\ = \tau & \text{if } S2 \\ = 0 & \text{if } S3 \end{cases} \quad (14)$$

S3 does not constitute a Nash-equilibrium.

In S1 the public good is not completed. Due to the money-back guarantee, it is irrelevant what strategy from the action space is chosen. In S2, all agents are pivotal. Each individual equally contributes the minimum per capita amount required to accomplish the public good, that is $b = \tau$, and the threshold level of the public good is accomplished. In S3, the argument given for the contribution game applies.

Uniform myopic best response dynamics in the subscription game We again follow a two-step procedure, by first analyzing the basic, undisturbed dynamics and second, by extending the dynamics with noise in order to identify

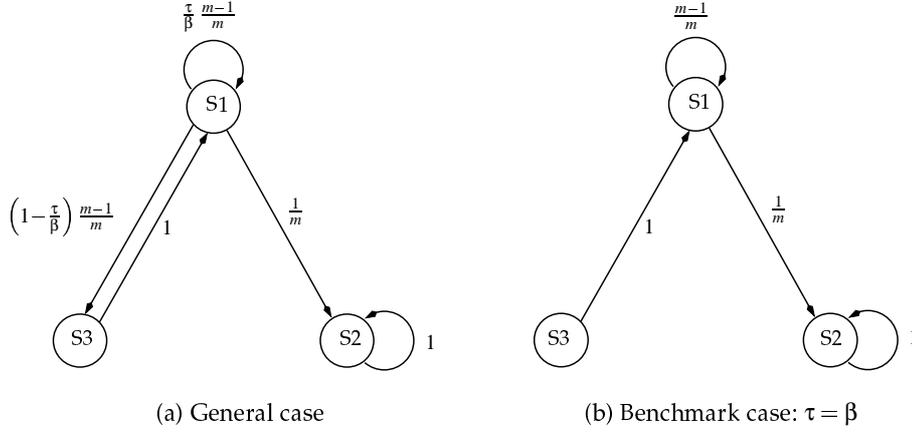


Figure 2: *Transition probabilities of myopic best response in the subscription game*

the influence of small mistakes in the agents' play on the long-run properties of the game.

I. Myopic best response equilibria without noise

Proposition 5 *The transition matrix between states S1, S2, and S3 for uniform myopic best response dynamics in the subscription game is given by*

$$P_{SG} = \begin{pmatrix} \frac{\tau}{\beta} \frac{m-1}{m} & \frac{1}{m} & \left(1 - \frac{\tau}{\beta}\right) \frac{m-1}{m} \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{pmatrix} \quad (15)$$

Proof: see Appendix A .

Figure 2 illustrates the myopic best response dynamics in form of a Markov chain. It is easy to recognize that state S2 is an absorbing state of the process. Once a society has reached S2, it is the unique best response for each agent to repeat contribution of the cost share τ for every iteration of the game. The transition probabilities are given by $p_{S2,S2} = 1$, $p_{S2,S1} = p_{S2,S3} = 0$ and, as can be seen from Figure 2(a), no branch leaves state S2, which is represented by the bottom right node.

The transition matrix (15) also implies that S3 is a reflecting state of the society. According to (14), S3 does not constitute a Nash-equilibrium and the state is left with probability one. From the viewpoint of each single agent, her contribution is not required for the completion of the public project. Consequently, the aggregate contribution will be zero for the next repetition of the game, which simultaneously implies a

probability of of $p_{S3,S2} = 0$ for transiting from S3 towards the ‘pivot’ state S2. Figure 2(a) illustrates that the only branch leaving S3 points towards S1 with probability one.

The corresponding transition probabilities of leaving S1, or staying in this state respectively, turn out to be more sophisticated. Because the public good is not completed in this state with contributions being refunded, by (14), strategies can be chosen arbitrarily from the entire action space. At this point, we only give an intuitive argument for the case of transiting from S1 to the absorbing state S2. Here, in case of symmetric behavior, all agents are required to choose strategy $b = \tau$, while staying in S1 requires a common action of $b < \tau$. If we assume equally distributed chances for each of the m strategies to be played, it follows that $p_{S1,S2} = 1/m$,¹⁰ which is displayed in Figure 2(a) with the respective branch leading from S1 to S2.

Proposition 6 *In the long–run the myopic best response dynamics lock in at the absorbing state S2. In the medium run, the process stays most of the time in the Nash–equilibrium S1 without completion of the public good, which is more likely the larger the number m of strategies contained in the action space \mathbb{B} .*

Proof:

$$\frac{\partial p_{S1,S1}}{\partial m} = \frac{\tau}{\beta} \frac{1}{m^2} > 0.$$

S3 is no Nash–equilibrium. □

Proposition 7 *The higher the cost–preference index τ/β , the less often the public good is provided in the medium run.*

Proof: In the medium run, the transition probability $p_{S1,S1}$ is an inverse measure of how often the public good is provided: The higher $p_{S1,S1}$, the less often aggregate contributions will be high enough to cover the threshold. The change of the transition probability with respect to the cost–preference index is given by

$$\frac{\partial p_{S1,S1}}{\partial \frac{\tau}{\beta}} = \frac{m-1}{m} > 0. \quad \square$$

Given the general results from Proposition 5, we can now focus on the corresponding myopic best response equilibria of a relevant benchmark case: It can be characterized as an all–or–nothing case $\tau = \beta$,

¹⁰This of course means that we rule out focal point arguments. If we did otherwise, the possible candidate for a focal strategy, $b = \tau$, would have a higher probability to be chosen from the action space than other strategies.

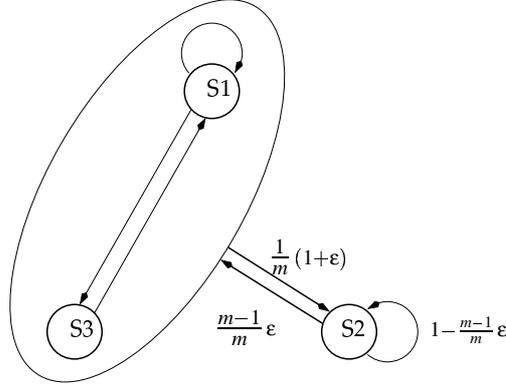


Figure 3: *Transition probabilities of myopic best response with noisy imitation in the subscription game*

which requires truthful revelation of individual preferences. In this scenario, each agent has to contribute her maximum willingness to pay in order to have the public project completed. The probability of transiting from S1 to S3 vanishes in this case, i. e. $p_{S1,S3} = 0$. This means that once S3 is left it can never be revisited again. The probability of staying in S1 adapts correspondingly, $p_{S1,S1} = \frac{m-1}{m}$, while $p_{S1,S2}$, $p_{S2,S2}$, and $p_{S3,S1}$ remain unchanged. This case is displayed in Figure 2(b), which shows the respective Markov chain.

The major insight from this part of the analysis is that uniform myopic best response dynamics are characterized by a Markov process with only one absorbing state. This means that the system will asymptotically converge towards the Nash-equilibrium where every player is pivotal and participates in the completion of the public project, which is provided at exactly the threshold level. Contrary to the contribution game, there is no state-dependency. Since the results on the stability of equilibria in the subscription game only hold for infallible agents, we will now proceed with the analysis of myopic best response equilibria where agents occasionally make mistakes.

II. Myopic best response equilibria and noisy imitation in the subscription game

As before, we extend the myopic best response dynamics with a random element, such that strategies are subject to mutation with probability ϵ . This affects the transition probabilities between the three states of the society. As the main point of interest is whether S2 still is stationary, the focus lies on the relation of the formerly communicating class consisting of S1 and S3 to the pivot state S2. It is important

to note that now S2 ceases to be an absorbing state, because, with the introduction of mutation, there is a positive chance for each state to be left in finite time.

By assuming again a uniform probability $\text{Prob}(b^k) = \frac{1}{m}$ for each of the m strategies $b^k \in \mathbb{B}$ to be played, the transition probabilities for state S2 in the case of noisy imitation can be derived as follows:

$$p_{\leftarrow S2} = \frac{m-1}{m} \varepsilon \quad (16)$$

$$p_{S2, S2}^m = 1 - \frac{m-1}{m} \varepsilon \quad (17)$$

$$p_{\rightarrow S2} = p_{S1, S2} + (p_{S1, S2}^m + p_{S3, S2}^m) = \frac{1}{m} (1 + \varepsilon). \quad (18)$$

Figure 3 illustrates the corresponding Markov chain.

In order to find out, whether the pivot state will be reached more or less often in the long-run than the other states, it is necessary to calculate the stochastic potential of this state.¹¹ Intuitively spoken, the stochastic potential of a state is the difference between the probability of reaching the state and the probability of leaving it. In other words, the stochastic potential determines if a state is easier to be reached than to be left and *vice versa*, and it measures how easy it is to leave the state. The greater the stochastic potential of a state, the more time the process is going to spend in this state in the long-run. The stochastic potential of states S2 in the subscription game is given by

$$SP(S2) = p_{\rightarrow S2} - p_{\leftarrow S2} = \frac{1}{m} [1 - (m-2) \varepsilon]. \quad (19)$$

For the ‘regular’ case, with m being finite and ε being sufficiently small, $SP(S2)$ is positive, such that the pivot state is *stochastically stable*, i. e. the most often visited of all states in the long-run. Nevertheless, in the case of our EA simulations as well as in a world of continuous action spaces, this is not necessarily true: If the action space is continuous, m , the number of actions, goes to infinity. On the other hand, the mutation probability ε should be small, but not too small in order to keep a sensible notion of possible learning errors. In the case of m being infinitely large and ε being strictly larger than zero, the stochastic potential of state S2 is negative:

$$SP(S2) < 0 \quad \text{for } m \rightarrow \infty \quad \text{and } \varepsilon > 0. \quad (20)$$

This means, that the pivot state is easier to be left than to be reached. In the long-run, the process will spend most of its time outside of the pivot state.

¹¹For a definition of the stochastic potential, see e. g. Vega-Redondo (1996, p. 132).

As we know that state S3 also is (mostly almost) reflecting, the only possible state for the process to remain in over time is S1. This, of course, again is bad news. It says that most of the time in the long-run the public good will not be provided, if the agents are likely to make mistakes.

3 Evolutionary Learning in the Discrete Public Good Game

The analysis of the preceding section has shown that both types of discrete public good games exhibit multiple equilibria, one without completion of the public good in S1, the other given by the pivot state S2. While S1 is uniquely characterized by a common action of $b = 0$ in the contribution game, it consists of multiple action profiles in the subscription game, which all satisfy the requirement $\sum b_i < T$ due to the arbitrariness of strategy choice stemming from the money-back guarantee. This gives rise to the problem of equilibrium selection, which can be solved by means of a learning process. We provide a short sketch of the evolutionary mechanism of how knowledge on the quality of strategies spreads out in a population of boundedly rational agents. Contrary to the uniform myopic best responses dynamics discussed in the previous section, here, the evolutionary learning process is capable of describing heterogeneous behavior of agents. The analysis draws from results from evolutionary game theory (Miller and Andreoni, 1991; Vega-Redondo, 1996) and from a companion paper on the evolution of free riding behavior in public good games (Clemens and Riechmann, 2001).

Learning involves that agents replace poorly performing strategies by those performing well. Evolutionary game theory provides us with the concept of replicator dynamics, where learning takes place via imitation of successful strategies. The intertemporal performance of an arbitrary strategy, say $b^k(t) \in \mathbb{B}$ for instance, can be measured by the evolution of the population share $q^k(t)$ of agents, playing this type of strategy in period t . Since time- t -utility of type- k -players, $U^k(t) = w(t) - b^k(t) + \beta \ln G(t)$ also depends on the strategic decisions of all population members via the aggregate contribution level G , the performance of a type- k -strategy can only be calculated in relation to the performance of the entire set of strategies actually played in the population. The measure is relative payoff $\frac{U^k(t)}{\bar{U}(t)}$, with average utility $\bar{U}(t)$ given by

$$\bar{U}(t) = \sum_{k=1}^m q^k(t) U^k(t). \quad (21)$$

The evolution of the population share $q^k(t)$ of agents playing $b^k(t)$ can then be described by an ordinary difference equation

$$q^k(t+1) = q^k(t) \frac{U^k(t)}{\bar{U}(t)}. \quad (22)$$

Equation (22) represents the replicator dynamics of an evolutionary game (Vega-Redondo, 1996, p. 44). The population share of agents playing strategy $b^k(t)$ increases over time if this strategy yields a payoff above average.

In order to describe dynamics of the entire game, the evolution of (22) has to be computed for each of the m strategies $b^k \in \mathbb{B}$. For large actions spaces \mathbb{B} , this method induces extremely high computational costs. Nevertheless, it is shown in Riechmann (2001c) that there is an equivalent method to explicitly computing each of the m equations (22), which is the use of an appropriately constructed Evolutionary Algorithm (EA, Goldberg 1989; Riechmann 2001a,b). This method is applied in this paper.

The core data structure of the algorithm is a set ('population') of real numbers, each number representing a strategy $b_i(t)$ played by each of the n agents in period t . In each period or round of the algorithm, the quality of each agent's strategy is evaluated according to the utility function (5) for the contribution game or equation (12) for the subscription game respectively. The resulting fitness value ('utility') is the criterion deciding on whether the agent repeats using her strategy or whether she will adopt a different one.

The component of the EA determining how many players will play each strategy in the next period is the operator of selection/reproduction. The canonical form of this operator is 'biased roulette wheel selection' (Goldberg, 1989; Goldberg and Deb, 1991). This is a biased probabilistic process generating the following expected population share $q^k(t+1)$ of agents playing strategy b^k in period $t+1$:

$$\mathbb{E} \left(q^k(t+1) \right) = \frac{U^k(t)}{\bar{U}(t)} q^k(t). \quad (23)$$

It is obvious that (23) gives the expectation of the dynamics in (22).¹²

Mutation is the second operator employed in the simulations. It reflects the impact of randomness on strategy choice, which we already discussed in the previous section. In economic applications of EAs, mutation is often interpreted as a metaphor for learning by experiment. Mutation introduces noise into the process of replication, thereby correcting the problem of path

¹²Note that it is only possible to denote the expected number of type- k -players in $t+1$, because otherwise the resulting *absolute* number might not be an integer; see Riechmann (2002).

dependency arising from the process of pure replication by allowing lost strategies to be regained.¹³

The special type of mutation applied in this paper is the one introduced in the survey on ‘Evolution Strategies’ by Bäck *et al.* (1991): After replication has taken place, each agent’s preliminary strategy $\tilde{b}_i(t)$ is slightly changed by addition of a term ε_i , which is drawn from a Gaussian distribution with zero mean and finite σ^2 , the latter denoting the so-called *mutation variance*. The result is agent i ’s final strategy for period t , $b_i(t)$:

$$b_i(t) = \max [\tilde{b}_i(t) + \varepsilon_i(t); 0] \quad \varepsilon_i(t) \sim N(0, \sigma^2). \quad (24)$$

which is non-negative by assumption.

The two operators, selection/reproduction and mutation are repeatedly applied to the population of agents, thus generating processes of population dynamics equivalent to processes of replicator dynamics with noise.

The simulations to be presented in the following section show that the results on the stability of Nash-equilibria already derived for the case of uniform myopic best response extend to the case of heterogeneous behavior described by evolutionary dynamics.

4 Simulations

Questions and Technical Details We were especially interested in the following questions: First, does the learning process support selection between states of the society, especially between S1 and S2? Second, how does the population size affect the learning dynamics? Third, how important is the threshold size, as measured by the cost-preference index, for the long- and medium-run outcomes of the games? What are the most noteworthy differences in learning dynamics between the contribution and the subscription game?

The simulations are based on the evolutionary algorithm as described in the preceding section. In order to derive results for the sensitivity of the learning process with respect to the population size, we performed simulations with $n = 25$ and $n = 100$ agents, the first capturing the notion of a small community, while the latter stands for a comparably more atomistic structure of the economy. We assume three different threshold levels in the range of $\tau/\beta \in \{0.25, 0.5, 1\}$ for both types of games. Especially the last value $\tau/\beta = 1$ requires a large degree of coordination in the population, since each agent is pivotal and has to contribute his maximum willingness to pay in order to get the public good accomplished, and there is no opportunity

¹³If a strategy b^k has ‘died out’ during replication, which means that the respective population share has become zero, there is no way of regaining this strategy via pure imitation.

of ‘cheap riding’. The preference parameter from the utility function and initial endowments were assumed to be identical for all agents with $\beta = 100$ and $w = 200$. The mutation variance is assumed to be a constant value of $\sigma^2 = 0.05$. The strategies of the initial population are randomized, such that $b_i(0)$ is i. i. d. in \mathbb{B} .

The simulation results presented in the following show the development of the population mean of individual contributions and the percentage of successful provisions over time. The plots display averages of 100 simulation runs. Data on the last-period distribution of strategies within the population are based on 2000 simulation runs for a threshold value of $\tau/\beta = 0.5$.

Simulation Results Figures 4 and 5 display the results of the numerical simulations. Figure 4 focuses on the long-run evolution of average contributions b for the two population sizes and the percentage of runs in which the provision of the public good was successful. Figure 5 shows the last-period distribution of strategies in the population. The following observations hold independently of the underlying game:

Observation 1 (Population Size) *The population size affects the learning speed.*

A prominent result from the theoretical analysis of discrete public good games we also derived in section 2 is that efficient equilibria do not depend on group size, since all players are pivotal (Bagnoli and Lipman, 1989). Although our simulations suggest that perhaps the population size might be important, this cannot be ascribed to specific characteristics of the underlying games, but moreover to an important feature of Evolutionary Algorithms, namely that learning is improved with an increase in group size. Smaller populations converge considerably slower than larger ones, because the effects of experimentation are more significant in the latter.¹⁴ This effect is particularly prevalent in our simulations, to the extent that the learning process in the small group of $n = 25$ has not converged, when the simulations were truncated.

Observation 2 (Selection between States of the Society) *In both games, state S1 is far more often visited throughout the learning process than S2.*

Observation 3 (Threshold) *The higher the cost-preference index τ/β the less likely the public good is provided.*

¹⁴For a detailed discussion on the effects of group size on learning dynamics in public good games see Clemens and Riechmann (2001).

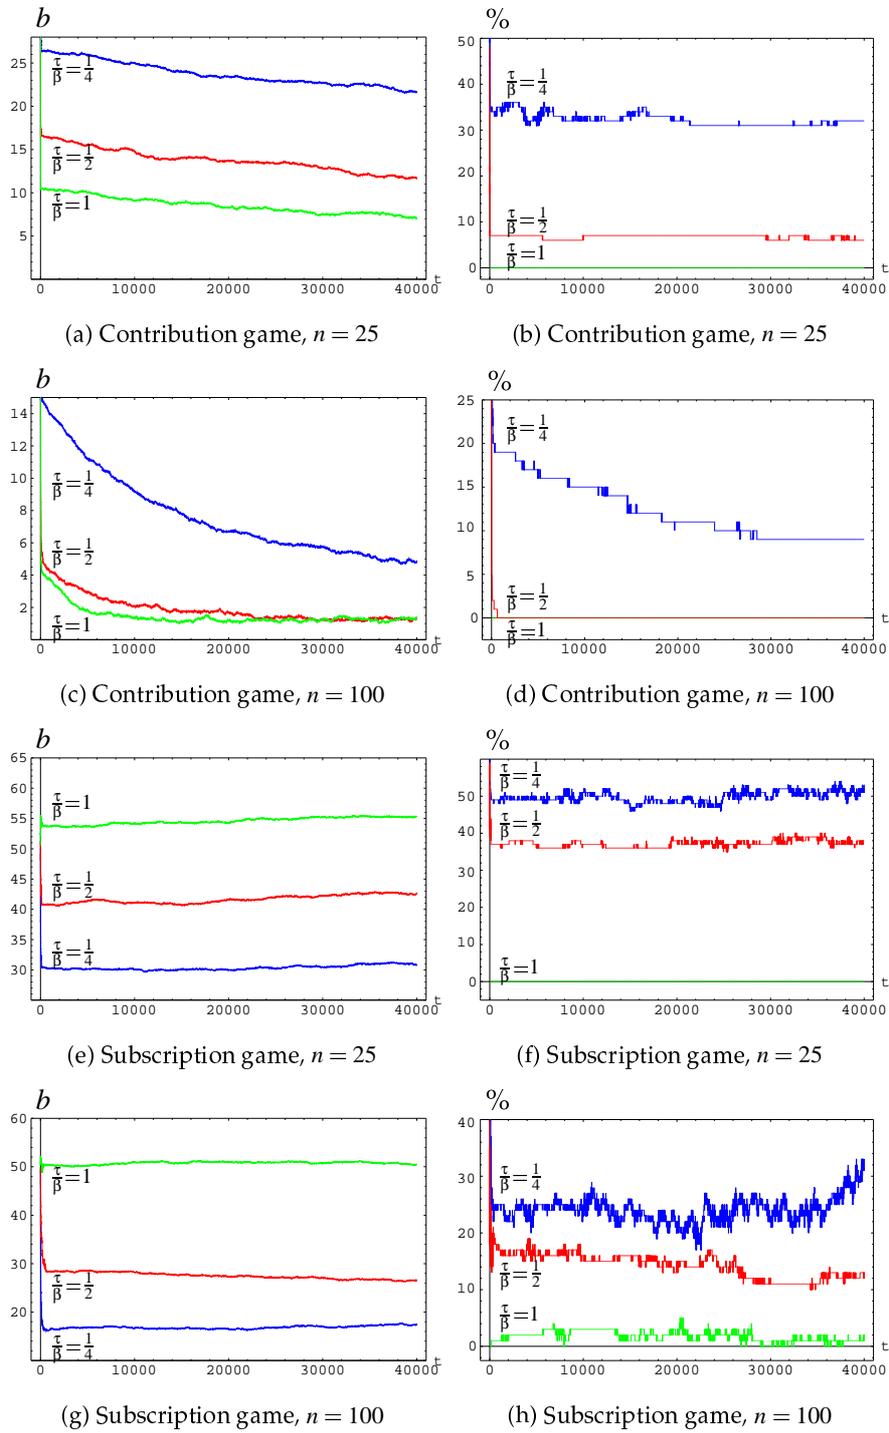


Figure 4: Average Contributions b and Percentage of Successful Provision

Observation 4 (Successful Provision of the Public Good) *Completion of the public good is more successful in the subscription than in the contribution game.*

As can be seen from Subfigures 4(b), 4(d), 4(f) and 4(h) on the right hand side of Figure 4, the public good is not provided during most of the time, which reflects the results from the stability analysis of equilibria in section 2. The results differ with respect to the threshold level and the type of game in consideration.

The highest percentage (50%) of runs with successful provision of the public good can be observed at the lowest threshold level in the subscription game played by the small population; see Figure 4(f). Contrary, the public project is never completed for the highest and hardly ever (7.5%) for the medium threshold in the contribution game, see Figures 4(b) and 4(d). Hence our simulations reflect the statement on the impact of the cost-preference index derived in Proposition 7.

Completion of the public project is more successful in the subscription game compared to the contribution game because of the selective pressure stemming from the refund-rule. Although the pivot state S2 is a Nash-equilibrium in both games, from the viewpoint of a single agent in the contribution game, contributing a positive amount in S1 is immediately punished with a low payoff and induces a switch towards strategies supporting S1 in the long-run. Competition between strategies is weakened due to the money-back guarantee of the subscription game which increases the probability of having enough contributions collected to provide the good.

Observation 5 (Contribution Game) *In the long-run, agents learn not to contribute to the provision of the public good.*

This result stems from the selective pressure already discussed above and is displayed in Figures 4(a) and 4(c), where average contributions converge towards zero independent of the threshold size. If we take a further look into the final distribution of strategies, as displayed by Figures 5(e) and 5(f), we see that strategies are densely concentrated at zero. The presence of non-zero strategies can be ascribed to the effect of mutation. Additionally it can be seen that the threshold level itself is positively related to the speed of convergence.

Laboratory experiments by Isaac *et al.* (1989) even indicate that zero contributions are more frequent in contribution games if compared to a standard continuous public good game. Our simulations support this result and show quite nicely the intensity of selective pressure, which can be seen from a comparison of Figures 5(e) and 5(f) with 5(g) and 5(h).

Observation 6 (Subscription Game I) *In the long-run, agents learn not to voluntarily contribute more than the required cost share τ .*

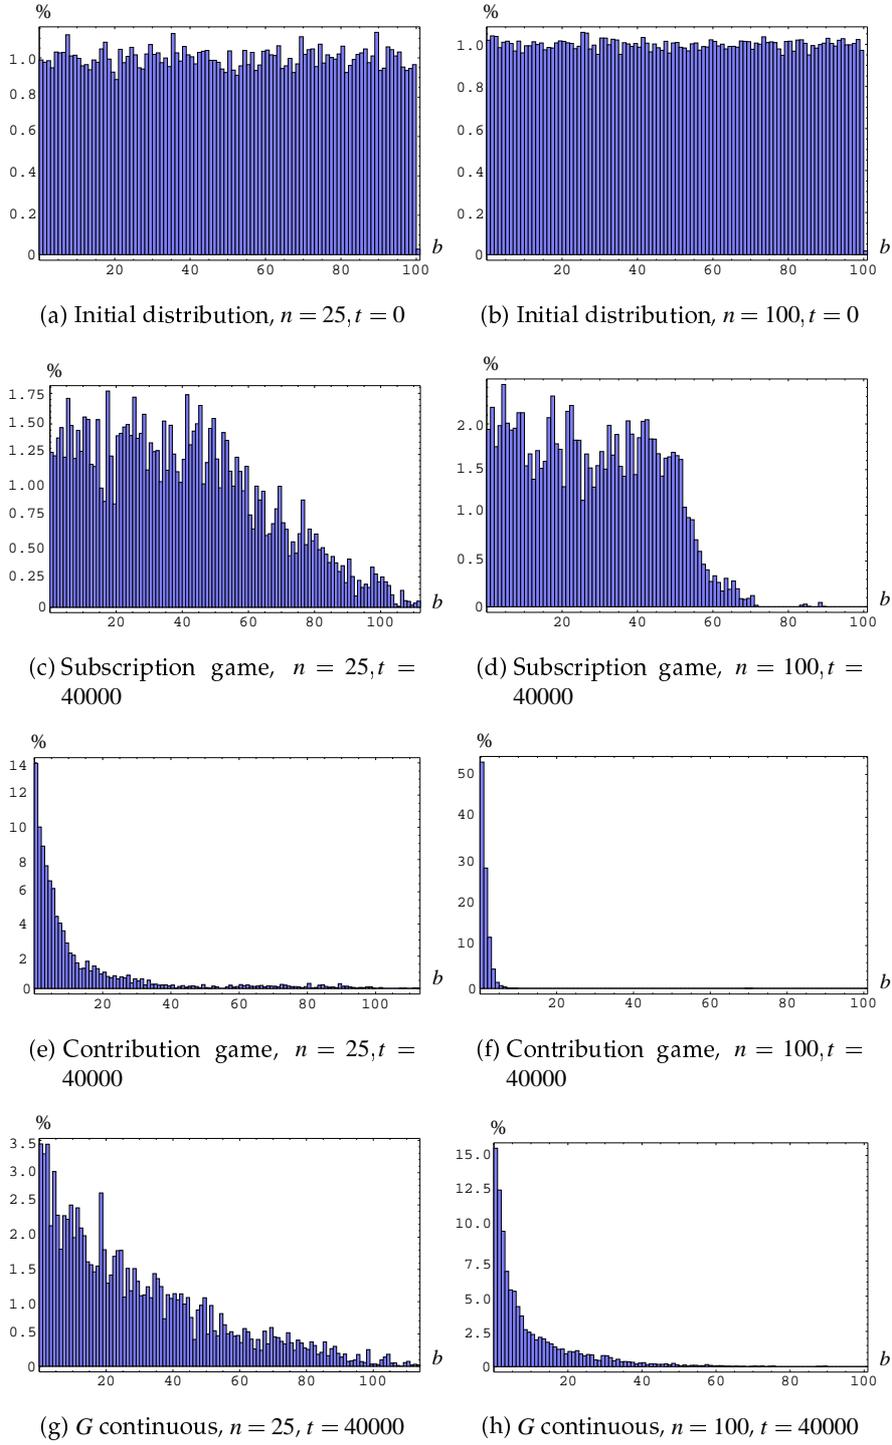


Figure 5: Initial/last period distributions of strategies, 2000 runs, $\tau/\beta = 0.5$

Observation 7 (Subscription Game II) *Individual strategies within the population tend towards being uniformly distributed in the interval $[0, \tau]$.*

These results are depicted in Figures 5(c) and 5(d). They are more obvious in the latter due to a larger extent of convergence. Frequencies of strategies exceeding the per capita threshold $\tau = 50$ decline sharply, while strategies in the interval $[0, \tau]$ do not follow a specific pattern. The reason for this lies in the fact that there is no selective pressure on strategies in the interval $[0, \tau]$. Because of the money-back guarantee all strategies $b < \tau$ yield identical utility which makes any discrimination between them impossible. This is not true for all strategies $b > \tau$, which die out, because all contributions exceeding the threshold level $T = \tau n$ are wasted and there is an incentive for ‘cheap riding’ (Isaac *et al.*, 1989) on the contribution of others.

The complete lack of selective pressure for strategies $b < \tau$ combined with the random elements embedded in the Evolutionary Algorithm causes the convergence of strategies towards a uniform distribution in the long-run.

Observation 8 (Subscription Game III) *In the long run, the average population mean contribution b attains values of above roughly one half of the cost share τ required to complete the public project.*

In our example, a uniform distribution of strategies implies an average individual contribution of $b = \frac{1}{2}\tau$. Although the result, that higher thresholds imply higher average contributions seems counter-intuitive, it can again be explained by the lack of selective pressure in the subscription game. Perhaps it is of some interest for contributions dealing with the question of artificially introducing thresholds in order to make agents participate in public good provision (Morelli and Vesterlund, 2000).

5 Conclusions

In this paper we examined the learning dynamics of boundedly rational agents, who were asked to voluntarily contribute to a discrete public good. The theoretical part of the paper focused on static equilibria and myopic best response dynamics of two different types of the threshold public good game. Whereas in the contribution game individual contributions are lost in case of provision failure, there is a money-back guarantee in the subscription game. Both games exhibit multiple equilibria: an efficient one, where the public good is completed and all agents are pivotal, and inefficient equilibria, where provision is not accomplished. We were able to show that the efficient outcome is an absorbing state of uniform myopic best response learning in the subscription game, while it is an instable equilibrium in the contribution game. With the introduction of randomness in

the course of strategy formation, the analysis could be extended in order to establish results on the stochastic stability of equilibria. The pivot state, i. e. the efficient equilibrium, was shown to be stochastically unstable in the contribution game and demonstrated only to be stochastically stable in the subscription game by imposing additional restrictive conditions regarding the number of strategies available. Moreover it could be shown that in the medium run the society spends more time in states without provision than otherwise. We additionally found that provision failure is positively correlated with the threshold level, a result which is supported by our simulations as well as by laboratory experiments on discrete public good games.

The learning process of heterogeneous agents in our simulations was modeled by an Evolutionary Algorithm, which is closely related to replicator dynamics of evolutionary game theory. In summary, the simulations strongly support our previously derived theoretical results. For high threshold values, the public good was almost never provided in the contribution game, while provision on average was more successful in the subscription game. Nevertheless, even in the subscription game, the public project was not completed in the majority of cases, due to the properties stated on medium-run behavior.

References

- Admati, Anat R. and Perry, Motty (1991), Joint Projects without Commitment, *Review of Economic Studies*, **58**, 259–276.
- Bäck, Thomas, Hoffmeister, Frank, and Schwefel, Hans-Paul (1991), A Survey of Evolution Strategies, in: R. K. Belew and L. B. Booker (eds.), *Proceedings of the 4th International Conference on Genetic Algorithms*, Morgan Kaufmann, San Mateo, CA, pp. 2–9.
- Bagnoli, Mark and Lipman, Barton L. (1989), Provision of Public Goods: Fully Implementing the Core through Private Contributions, *Review of Economic Studies*, **56**, 583–601.
- Bagnoli, Mark and McKee, Michael (1991), Voluntary Contribution Games: Efficient Private Provision of Public Goods, *Economic Inquiry*, **29**, 351–366.
- Bliss, Christopher and Nalebuff, Barry (1984), Dragon-Slaying and Ballroom Dancing: The Private Supply of a Public Good, *Journal of Public Economics*, **25**, 1–12.
- Cadsby, Charles Bram and Maynes, Elizabeth (1999), Voluntary Provision of Threshold Public Goods with Continuous Contributions: Experimental Evidence, *Journal of Public Economics*, **71**, 53–73.
- Cavaliere, Alberto (2000), Coordination and the Provision of Discrete Public Goods by Correlated Equilibria, *Journal of Public Economic Theory*, **3** (3), 235–255.

- Clemens, Christiane and Riechmann, Thomas (2001), *Dynamic Voluntary Contribution to a Public Good: Learning to be a Free Rider*, Diskussionspapier 240, Universität Hannover, Fachbereich Wirtschaftswissenschaften, Hannover.
- Cornes, Richard and Sandler, Todd (1996), *The Theory of Externalities, Public Goods and Club Goods*, Cambridge University Press, Cambridge, 2nd edn.
- Dawes, R. M., Orbell, J. M., Simmons, R. T., and van de Kragt, A. J. C. (1986), Organizing Groups for Collective Action, *American Political Science Review*, December 1986, 1171–1185.
- Fudenberg, Drew and Levine, David K. (1998), *The Theory of Learning in Games*, MIT Press, Cambridge, MA.
- Goldberg, David E. (1989), *Genetic Algorithms in Search, Optimization, and Machine Learning*, Addison-Wesley, Reading.
- Goldberg, David E. and Deb, Kalyanmoy (1991), A Comparative Analysis of Selection Schemes used in Genetic Algorithms, in: Gregory J. E. Rawlins (ed.), *Foundations of Genetic Algorithms*, Morgan Kaufmann, San Mateo, CA, pp. 69–93.
- Gradstein, Mark (1992), Time Dynamics and Incomplete Information in the Private Provision of Public Goods, *Journal of Political Economy*, **100** (5), 581–597.
- (1994), Efficient Provision of a Discrete Public Good, *International Economic Review*, **35**, 877–897.
- Isaac, R. Mark, Schmidtz, D., and Walker, James M. (1989), The Assurance Problem in a Laboratory Market, *Public Choice*, **62**, 217–236.
- van de Kragt, A. J. C., Orbell, J. M., and Dawes, R. M. (1983), The Minimal Contributing Set as a Solution to Public Goods Problems, *American Political Science Review*, March 1983, 112–122.
- Ledyard, John (1995), Public Goods: A Survey of Experimental Research, in: John H. Kagel and Alvin E. Roth (eds.), *The Handbook of Experimental Economics*, Princeton University Press, Princeton, pp. 111–194.
- Marimon, Ramon (1993), Adaptive Learning, Evolutionary Dynamics and Equilibrium Selection in Games, *European Economic Review*, **37**, 603–611.
- Marx, Leslie M. and Matthews, Steven A. (2000), Dynamic Voluntary Contribution to a Public Project, *Review of Economic Studies*, **67**, 327–358.
- Menezes, Flavio M., Monteiro, Paulo K., and Temimi, Akram (2001), Private Provision of Discrete Public with Incomplete Information, *Journal of Mathematical Economics*, **35**, 493–514.
- Miller, John H. and Andreoni, James (1991), Can Evolutionary Dynamics Explain Free Ridings in Experiments?, *Economics Letters*, **36**, 9–15.
- Morelli, Massimo and Vesterlund, Lise (2000), Provision Point Mechanism and Over Provision of Public Goods, Working Paper 00-14, Ohio State

University.

- Nitzan, Shmuel and Romano, R. (1990), Private Provision of a Discrete Public Good with Uncertain Cost, *Journal of Public Economics*, **42**, 357–370.
- Palfrey, T. R. and Rosenthal, H. (1984), Participation and the Provision of Discrete Public Goods: A Strategic Analysis, *Journal of Public Economics*, **24**, 171–193.
- Riechmann, Thomas (2001a), Genetic Algorithm Learning and Evolutionary Games, *Journal of Economic Dynamics and Control*, **25**, 1019–1037.
- (2001b), *Learning in Economics. Analysis and Application of Genetic Algorithms*, Physica-Verlag, Heidelberg.
- (2001c), Two Notes on Replication in Evolutionary Modelling, Diskussionspapier 239, Universität Hannover, Fachbereich Wirtschaftswissenschaften, Hannover.
- (2002), Cournot or Walras? Agent based Learning, Rationality, and Long Run Results in Oligopoly Games, Diskussionspapier 261, Universität Hannover, Fachbereich Wirtschaftswissenschaften, Hannover.
- Samuelson, Larry (1997), *Evolutionary Games and Equilibrium Selection*, MIT Press, Cambridge, MA.
- Suleiman, R. and Rapoport, A. (1992), Provision of Step-Level Public Goods with Continuous Contribution, *Journal of Behavioral Decision Making*, **5**, 133–153.
- Vega-Redondo, Fernando (1995), Public Subscription and Private Contributions, *Social Choice and Welfare*, **12**, 193–212.
- (1996), *Evolution, Games, Economic Behavior*, Oxford University Press, Oxford.

A Myopic best response in the subscription game

Proof of Proposition 5

- (i) *State S1*: every possible action from the action space \mathbb{B} is a best reply to S1 due to the money-back guarantee. We assume identical behavior of all agents and that each action $b^k \in \mathbb{B}$ is equally likely to be played, i. e. has probability $\text{Prob}(b^k) = \frac{1}{m} \forall b^k \in \mathbb{B}$.

According to Proposition 4, a transition from S1 to S2 is only accomplished if agents simultaneously play strategy $b^* = \tau$. Consequently, the probability is $p_{S1, S2} = 1/m$, while non-uniform behavior of agents implies $\tilde{p}_{S1, S2} = (\frac{1}{m})^n$.

For a transition from S1 to S1, the common action of agents must be $b < \tau$. The determination of $p_{S1, S1}$ involves to count the number of strategies less than the cost share:

Let $\mathbb{T} \subset \mathbb{B}, \mathbb{T} := \{b^k \in \mathbb{B} | b^k \leq \tau\}$ be the set of actions less or equal τ . Let furthermore $\#(\mathbb{T})$ denote the number of elements of set \mathbb{T} . The number of actions smaller than τ can be determined as $\#(\mathbb{T}) - 1$. The probability of staying in

state S1 is given by $p_{S1,S1} = \frac{\#(\mathbb{T})-1}{\#(\mathbb{B})} = \frac{\#(\mathbb{T})-1}{m}$. The remaining task is to count the elements in \mathbb{T} .

The distance between two neighboring elements in \mathbb{B} is given by $\delta := |b^k - b^{k-1}| \forall k \in \{1, 2, \dots, m\}$. Between the m elements of \mathbb{B} there are exactly $m - 1$ distances, from which follows

$$\delta = \frac{\beta}{m-1}$$

\mathbb{T} is a subset of \mathbb{B} . The distance between elements in \mathbb{T} is equal to the distance δ between elements in \mathbb{B} . There are exactly $\#(\mathbb{T}) - 1$ distances between $b^1 = 0$ and τ . It follows that

$$\tau = (\#(\mathbb{T}) - 1) \delta. \quad (\text{A.1})$$

Consequently, $\#(\mathbb{T})$ can be obtained as

$$\#(\mathbb{T}) = \frac{\tau}{\beta} (m - 1) + 1.$$

Finally, the probability of transiting from S1 to S1 can be derived as

$$p_{S1,S1} = \frac{\#(\mathbb{T}) - 1}{m} = \frac{\tau}{\beta} \frac{m - 1}{m} \quad (\text{A.2})$$

The probability of transiting from S1 to S3 can be derived residually

$$\begin{aligned} p_{S1,S3} &= 1 - p_{S1,S1} - p_{S1,S2} \\ &= 1 - \frac{\tau}{\beta} \frac{m - 1}{m} - \frac{1}{m} = \left(1 - \frac{\tau}{\beta}\right) \frac{m - 1}{m}. \end{aligned} \quad (\text{A.3})$$

- (ii) *State S2*: obvious from Proposition 4. The contribution of each agent is necessary for the completion of the public good, which provides the incentive not to deviate from the equilibrium strategy $b^* = \tau$

$$p_{S2,S2} = 1 \implies p_{S2,S1} = p_{S2,S3} = 0 \quad (\text{A.4})$$

- (iii) *State S3*: obvious from Proposition 4

$$p_{S3,S1} = 1 \implies p_{S3,S2} = p_{S3,S3} = 0 \quad (\text{A.5})$$

If the society is located in S3, the unique best response for each agent i is to signal $b_i = 0$, since completion of the public project (in theory) does not require player i 's contribution. As this is true for every player participating in the game, the aggregate contribution will be zero when the game is repeated in the next period. From this argument it becomes evident that it is not possible to switch from S3 to S2 (the 'pivot' state) in one move, that is $p_{S3,S2} = 0$. Consequently, the probability of transiting from S3 to S1 is given by $p_{S3,S1} = 1$. \square