Incentive-Based Efficient Solutions for Public Goods Games

Yair Vaknin
Ben Gurion University
Be'er Sheva, Israel
yairva@post.bgu.ac.il

Amnon Meisels
Ben Gurion University
Be'er Sheva, Israel
am@bgu.ac.il

ABSTRACT

Multi-agent games on networks have nodes that represent agents and edges that represent interactions among agents. An important example in economics and in multi-agent research is the model of public goods games (PGGs). In PGGs agents can either contribute the cost of a single good or free-ride, and agents benefit from copies of the good bought by their neighboring agents. Solutions to games on networks are stable states (i.e., pure Nash equilibria), and in general one is interested in efficient solutions (i.e., of high social welfare). A multi-agent search algorithm for PGGs is proposed. The algorithm uses side payments among the agents during search and a major contribution is the proof that it converges to solutions that are more efficient than the initial strategy profile. An experimental evaluation on randomly generated scale-free networks demonstrates that the proposed algorithm outperforms both best-response dynamics and a former incentive-based PGGs algorithm. The experimental evaluation also explores the behavior of agents' gains with respect to their social capital. This new kind of exploration succeeds in observing patterns over a recently proposed network-based typology of social capital, on solutions of PGGs on networks.

KEYWORDS

Games on networks, Public goods games, Multi-agent search

ACM Reference Format:
Yair Vaknin and Amnon Meisels. 2023. Incentive-Based Efficient Solutions for Public Goods Games. In *Appears at the 5th Games, Agents, and Incentives Workshop (GAIW 2023). Held as part of the Workshops at the 21st International Conference on Autonomous Agents and Multiagent Systems., London, England, May-June 2023*, IFAAMAS, 7 pages.

1 INTRODUCTION

Multi-agent games on networks are commonly represented by the model of *graphical games* [11]. A node in the network represents an agent, and its neighboring nodes represent the agents that interact with the agent in the game. Each agent has a finite domain of strategies and a personal utility function. An important example that has drawn much interest in economics and in multi-agent search is public goods games (PGGs) on networks [3, 4, 6, 22]. In a PGG, agents can either contribute to a public good or free-ride [3, 4]. Agents benefit from copies of the good bought by their neighbors on the network and the benefit for an agent is larger than its cost of the good. A more general version of this game has agents benefit proportionally to the number of their neighboring agents that have bought the good [3].

Solutions to games on networks are stable states (i.e., pure Nash equilibria - PNEs), and in general one is interested in efficient solutions (of high social welfare). The use of side payments among agents as a mean of achieving solutions of higher efficiency in games on networks was first proposed by [7, 10]. A recent extension in this direction is an incentive-based PGGs search algorithm that uses side payments among agents during the iterative search process [13]. Analogous algorithms have been proposed for multi-agent (distributed) optimization problems [18, 19]. The present study proposes a multi agent search algorithm for PGGs on networks that uses side payments among the agents during search. The major result of this paper is that the proposed algorithm converges to a solution with a higher social welfare (SW) than the initial outcome for general PGGs. This is an important improvement over [13] that has no guarantee on the efficiency of its final solution even when it converges.

Performing search on PGGs on networks generates a unique situation in which one can observe the features of solutions experimentally. Experimental studies of multi-agent search on games on networks have the potential of widening the scope of former studies, that considered solution quality [5]. The present study presents an extensive experimental evaluation of solutions (PNEs) for PGGs, and explores the behavior of special classes of agents. One important classification of agents involves their *social capital*. A recent definition of social capital that is based on the local network features of the agents is addressed [9]. The experimental evaluation of the present study solves PGGs on networks and evaluates the effect of the social capital of agents on their personal gains.

Section 2 introduces the three aspects of the paper. Best Shot Public Goods Games (BSPPGs) are introduced first, together with their utility functions and main features. The second aspect is the use of side payments among agents during search and their role as contracts among neighboring agents. The third part of Section 2 presents the basic definitions of social capital, as proposed by [9]. These features are later used to investigate the behavior of classes of agents in solutions for PGGs on networks, in the experimental evaluation.

Section 3 starts by describing a small example that motivates the need for an improved algorithm for proposing side payments among neighboring agents during search. The algorithm that is at the focus of the present study is described and its main guarantees are proven in Section 3.

An extensive experimental evaluation is presented in Section 4, comparing the SW under the proposed algorithm and former ones. The experimental evaluation also studies the recent network-based

Appears at the 5th Games, Agents, and Incentives Workshop (GAIW 2023). Held as part of the Workshops at the 21st International Conference on Autonomous Agents and Multiagent Systems., Abramowitz, Aziz, Ceppi, Dickerson, Hosseini, Lev, Mattei, Tsang, WÄĚs, Zick (Chairs), May-June 2023, London, England.

definition of features that define social capital in [9]. It does so by computing simple correlations among the gain of individual agents in the network and their social capital. The last section lists the conclusions and points at the direction of future research.

2 PRELIMINARIES

2.1 Best Shot Public Goods Games

Many collective decision-making settings feature a tension between the individual interest of the agents and promoting a common good. One model that captures this tension is the public goods game (PGG). Each player in a PGG can decide whether or how much effort to invest for the common good. Everyone in the group (including the individual) profits from all membersâĂŹ efforts, while an investment incurs a cost for the investing agent. Examples of such scenarios include decisions whether or not to report crime in a neighborhood, keep oneâĂŹs yard representable, or purchase a tool that oneâĂŹs friends or neighbors can share ([12]).

A PGG is a networked game of $N = \{1, \ldots, n\}$ agents that reside on a network $G = \{N, E\}$. Each vertex in N represents an agent, and the edges E represent the interaction structure of the game. In the "best shot" public goods game (BSPGG) each player i chooses an action $v_i \in V_i$, where V_i is the set $\{T, F\}$. The choice $v_i = T$ denotes *investing* in some local public good, whereas $v_i = F$ denotes *avoiding* investment. Investment by player i incurs a cost c_i, while the utility of agent i depends also on the investment of its neighbors N_i. A player that can avoid investment and free-ride, relying on the investments of its neighbors, gains the maximal utility. The simplest version of BSPGGs has maximal gain when a single neighbor of an agent buys the good and more general versions have the gain of an agent increasing when more of its neighbors buy the good [3, 4, 13]. We assume that the utility of player i is non-decreasing in the investments of $\{i\} \cup N_i$, and is capped by some $\mathbb{K}$, denoting the maximal number of neighbors that buy the good.

Some common notions are used throughout the paper. v denotes a *strategy profile*, i.e., a vector containing the strategies of all players. v_i is the strategy of player i in v, and v_{-i} represents the strategies of all players excluding player i. Each player i tries to maximize its utility u_i. The strategy that maximizes the utility of a player i given v_{-i} is its *best-response*. A well known iterative procedure is *best-response dynamics*, which is also referred-to as "'best-response algorithm" (or simply "best-response"). Best-response algorithm iterates over all agents in a fixed order and selects the best-response strategy of each agent on its turn.

Denoting by $l_i = |\{j \in N_i : v_j = T\}|$ the number of investing neighbors of player i, the utility of player i under strategy profile v is defined to be

$$u_i(v) = \begin{cases} \min(\mathbb{K}, l_i + 1) - c_i, & \text{if } v_i = T \\ \min(\mathbb{K}, l_i), & \text{if } v_i = F \end{cases} \tag{1}$$

It is very convenient to also define

$$u'_i(v) = \begin{cases} \min(\mathbb{K}, l_i + 1), & \text{if } v_i = T \\ \min(\mathbb{K}, l_i), & \text{if } v_i = F \end{cases}$$

(the utility of agent i when its cost c_i is ignored).

A well known class of general games has the feature that when an agent changes its strategy, its utility change corresponds to the change in some global potential function. This class of games is termed *potential games* [15, 16]. A game is an *ordinal potential game* if there exists a potential function $\Phi : V_1 \times \ldots \times V_n \to \mathbb{R}$ such that $\forall v \in V_1 \times \ldots \times V_n$, $\forall v'_i \in V_i$ it holds that

$$u_i(v'_i, v_{-i}) - u_i(v_i, v_{-i}) > 0 \iff \Phi(v'_i, v_{-i}) - \Phi(v_i, v_{-i}) > 0$$

It has been shown that a BSPGG with the utility function defined by Equation 1 is an ordinal potential game [13]. This in turn implies the convergence of the game to a stable state - a pure Nash Equilibrium (PNE) - under best-response dynamics. A PNE is a state from which no selfish agent would deviate. More formally, a strategy profile v is a PNE if:

$$\forall i \in N, \forall v'_i \in V_i : u_i(v_i, v_{-i}) \geq u_i(v'_i, v_{-i})$$

2.2 Incentive-Based Local Search

Incentive-based search algorithms endow agents with the possibility of sacrificing part of their payoff in order to convince other agents to play a certain strategy. The idea to allow agents to offer payments to neighbors during the search process for a good solution to a BSPGG was proposed in [13]. The side payments offered and accepted among agents are defined by transfer functions. A transfer function $\tau_{i,j}(v)$ expresses the payment that agent i offers player j if the latter plays v_j under strategy profile v. Incentive-based algorithms iterate over all agents in some fixed order, and the use of side payments have been shown to outperform standard local search for distributed optimization problems [18]. The algorithm for BSPGGs on networks proposed by [13] is Algorithm 1. The agents in Algorithm 1 act in a predefined order: each agent in its turn executes procedure **on turn**. The agent whose turn it is to act (that is, select its strategy) is referred to as the *current agent*. The state of the current agent i in strategy profile v is termed $state_i(v)$ throughout the paper; it is either T_l or F_l where T (or F) is the strategy of agent i and $l = |\{j \in N_i : v_j = T\}|$ is the number of its investing neighbors in v. In Algorithm 1 side payments are only offered by neighbors of the current agents in order to convince it to remain in its current strategy. If the current agent's best-response is *investing* (line 1), it updates its strategy accordingly (line 2). If the current agent's best-response is *avoiding* (line 4), the agent notifies its neighbors about its possible strategy deviation (line 5). An agent $j \in N_i$ whose utility may decrease as a result of i's choice of avoiding investment can offer a side-payment in order to make the current agent keep investing (procedure **when received**). Offers of side-payments are *binding contracts* - they are paid upon termination of the algorithm's run, and only if the strategy they refer to is still played upon termination. The current agent sums up all offered payments, and stops investing only if its cost from investing exceeds the sum of all payments (lines 7-8). Similarly to standard best-respone, the algorithm terminates when a pass over all agents does not yield any change in the playersâĂŹ strategies. Algorithm 1 is guaranteed to converge when the contracts are made according to procedure **when received** [13].

Levit et al. [13] exploit the fact that BSPPGs with the utility function defined in Equation 1 are potential games to evaluate Algorithm 1 by comparing it to the best-response algorithm, as both are guaranteed to converge. In an extensive experimental

Algorithm 1 Incentive-Based Search Algorithm

on turn(i)

```
if state_i(v) = F_{l<K} then
    v_i ← T
end if
if state_i(v) = T_{l≥K} then
    send(StrategySelect, i) for all j ∈ N_i
    let τ_{j,i}(v) be the reply of neighbor j
    if c_i > Σ_{j∈N_i} τ_{j,i} then
        v_i ← F
    end if
end if
```

when received(StrategySelect, i)

$$\tau_{j,i}(v) = \begin{cases} 1 & \text{if } state_j(v) = T_{0<l<\mathbb{K}} \\ 1 & \text{if } state_j(v) = F_{0<l<\mathbb{K}} \\ c_j & \text{if } state_j(v) = F_{\mathbb{K}} \\ 0 & \text{otherwise} \end{cases}$$

```
return τ_{j,i}(v)
```

evaluation on randomly generated BSPGGs of a variety of network forms Algorithm 1 was found to converge to solutions of higher SW than best-response [13].

2.3 Social Capital

Social capital plays a central role in many real-life scenarios, from the founding of a company, through the reaching out to get help in a product design, to how workers get hired [9]. It is fundamental to the understanding of welfare of a society and should not be neglected when modeling human interaction.

A complete typology of social capital has been proposed recently by Jackson [9]. That typology is based on the local network topology of agents residing on a social network. Jackson proposes to base the social capital of an agent on its advantages in acquiring information in the social network, its ability to coordinate other player's behavior, and on more features that are listed below. I.e, agent's network topology "represents the consequences of social position in facilitating acquisition of the standard human capital characteristics" [14], and can be thought of as social capital [9].

Jackson counts seven types of social capital:

- **Information capital**: the ability to acquire/spread valuable information to other people through social connections.
- **Brokerage capital**: being in a position to serve as an intermediary between others.
- **Coordination and leadership capital**: being connected to others who do not interact with each other.
- **Bridging capital**: being an exclusive connector between otherwise disparate parties.
- **Favor capital**: the ability to exchange favors with others.
- **Reputation capital**: having others believe that a person or organization is reliable.
- **Community capital**: the ability to sustain cooperative behavior in others.

It is notable that Jackson proposes several **measurable** parameters that capture the social capital of an agent, given its local network topology. The present study investigates three of these parameters:

- **Decay centrality** (defined at [8]): denote N_i^d the players whose distance from player i is exactly d (N_i^1 are the immediate neighbors of i, also denoted N_i). For $0 < p < 1$ and some T, $Dec(i) = \sum_{d=1}^{T} p^d \cdot |N_i^d|$. Implies the **information capital** of player i: p^d of the information of distance d passes on to player i, with a cap of T times the information can be relayed. The experiments on Section 4 assume $p = 0.8$ and $T = 2$.
- **Godfather index**: a count of the number of pairs of a player's neighbors who are not neighbors of each other - $GF(i) = |\{(j,k) : j \in N_i, k \in N_i, j \notin N_k\}|$. Implies the **brokerage and coordination capital** of player i - the potential to serve as a mediator between neighboring players.
- **Support index**: a count of the number of neighbors that a player has who are supported by a common neighbor: $Supp(i) = |\{j \in N_i : \exists k \in N_i, k \in N_j\}|$. Implies the **favor capital** of player i - a neighbor j is more likely to deliver a favor to player i if they have a common connection.

It is natural to expect the power (i.e., social capital) of an agent in a game to determine in many ways its utility from playing on a network. The experiments on Section 4 explore the correlation between these two important entities and find a social capital predictor for the agent's utility in the solutions found by the incentive-based search algorithm that is proposed by the present study.

3 SECURING EFFICIENT EQUILIBRIA

The need for a different incentive-based search algorithm arises from drawbacks of Algorithm 1. Consider the example in Figure 1. Let $\mathbb{K} = 1$ and also let all the costs of investment for all agents be the same ($0 < c = 0.5 = c_1 = c_2 = c_3 = c_4 < 1$). Let the order of agents be their numerical order and the initial strategies of the game be T, F, F, F in that order. A run of Algorithm 1 in the same order of agents behaves as follows: in the first iteration, agent 1 with state $T_{0<\mathbb{K}}$ does not consider a strategy change, because it cannot improve its gain. Agent 2 acts next and stays with strategy F, as again no improvement of its state is possible. Continuing in the same order, agents 3, 4 change their strategy to T because they have no investing neighbor. None of the agents is offered side payments, because neither of them proposes a change from strategy T to F. In the second iteration no agent selects to change its strategy and the algorithm terminates with the outcome (T, F, T, T). In this solution three agents invest and the global SW is therefore $4 - 3 \cdot c$. Note that this is also the final outcome of standard best-response dynamics.

It is important to observe that the example problem has also a more efficient solution, in which the only investing agent is agent 2. The algorithm proposed by the present paper uses side payments in a way that scans all viable states in the neighborhood of the current agent during search. The algorithm is presented next (Algorithm 2).

Similarly to Algorithm 1, the agents execute procedure **on turn** in a fixed order. But now the current agent first asks its neighboring agents for the payments they are willing to offer in order to make it invest (line 1). Unlike Algorithm 1, the current agent does so

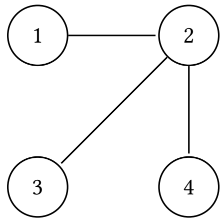

Figure 1: A four agents network example

Algorithm 2 Secure an Efficient Equilibrium

---

on turn(i)

1: **send**(**RequestBid**, i) to all $j \in N_i$
2: let $\tau_{j,i}(v)$ be the reply of neighbor j
3: **if** $c_i \leq (u'_i(T, v_{-i}) - u'_i(F, v_{-i})) + \sum_{j \in N_i} \tau_{j,i}(v)$ **then**
4: $\quad v_i \leftarrow T$
5: **else**
6: $\quad v_i \leftarrow F$
7: **end if**

when received(RequestBid, i**)**

1:

$$\tau_{j,i}(v) = \begin{cases} 1 & \text{if } state_j(v) = B_{0 \leq l < \mathbb{K}-1} \ (B \in \{F, T\}) \\ 1 & \text{if } state_j(v) = B_{l = \mathbb{K}-1} \text{ and } v_i = T \\ c_j & \text{if } state_j(v) = F_{l = \mathbb{K}-1} \text{ and } v_i = F \\ c_j & \text{if } state_j(v) = F_{l = \mathbb{K}} \text{ and } v_i = T \\ 0 & \text{otherwise} \end{cases}$$

2: **return** $\tau_{j,i}(v)$

---

regardless of its current strategy. If an agent's utility from investing (line 3), summed-up with the proposed side-payments, exceeds (or equals) the current agent's cost of investing, then it chooses to invest (line 4). Otherwise (line 5), it avoids investing (line 6). Side-payments act as *binding contracts*, as in Algorithm 1, and are paid upon termination only if the agent under contract keeps-up to the strategy that the contract refers to. The algorithm terminates when a pass through all the agents does not yield any change in the playersâĂŹ strategies.

The computation of the contracts in Algorithm 2 (procedure **when received**) is different than in Algorithm 1. In case agent j has less than $\mathbb{K} - 1$ investing neighbors, it gains 1 if an additional neighbor i chooses to invest. In case agent j has exactly $\mathbb{K} - 1$ investing agents and agent i is one of them, it gains 1 if agent i keeps investing. When agent j's state is $F_{\mathbb{K}-1}$ and $v_i = F$, agent j is only willing to pay c_j in order to make agent i invest. The payment cannot exceed c_j, which is the cost it takes agent j to make the investment itself. When agent j's state is $F_{\mathbb{K}}$ and $v_i = T$, agent j is willing to pay c_j in order to make agent i keep investing. In any other case the payment agent j is willing to offer is 0.

Let us follow the run of Algorithm 2 on the same example. In the first iteration, agent 1 sends a *RequestBid* message to agent 2. The response of agent 2 is $\tau_{2,1}(v) = c_2 = 0.5$. Because it holds that $c_1 = 0.5 <= (u'_1(T, v_{-i}) - u'_1(F, v_{-i})) + \tau_{2,1} = (1 - 0) + 0.5 = 1.5$, agent 1 keeps investing. When agent 2 is the current agent, it send a *RequestBid* message to its neighbors and the responses are $\tau_{1,2} = 0$, $\tau_{3,2} = \tau_{4,2} = 0.5$. Now it holds that $c_2 = 0.5 <= (u'_2(T, v_{-i}) - u'_2(F, v_{-i})) + \tau_{1,2} + \tau_{3,2} + \tau_{4,2} = (1 - 1) + 0 + 0.5 + 0.5 = 1$ and agent 2 chooses to invest. Next are the turns of agents 3 and 4 who choose not to invest, because $c_{i \in \{3,4\}} = 0.5 > (u'_{i \in \{3,4\}}(T, v_{-i}) - u'_{i \in \{3,4\}}(F, v_{-i})) + \tau_{2, i \in \{3,4\}} = (1 - 1) + 0 = 0$. In the next round agent 1 computes that it can avoid investing and still gain maximal utility (specifically, $c_1 = 0.5 > (u'_1(T, v_{-i}) - u'_1(F, v_{-i})) + \tau_{2,1} = (1 - 1) + 0 = 0$). The algorithm completes the second iteration over all agents and terminates because none of the agents changes its strategy. The algorithm terminates with outcome (F, T, F, F). Note that Algorithm 2 converged to the more efficient strategy profile (F, T, F, F) where the SW is $4 - c > 4 - 3 \cdot c$. In fact, to an optimal strategy on this simple example.

It is important to observe that the use of side payments by Algorithm 2 is substantially different than that of Algorithm 1. The new algorithm proposed by the present paper (Algorithm 2) uses side payments **to encourage** free-riding agents to invest.

3.1 Convergence to Efficient Solutions

The general idea of the proof that Algorithm 2 converges to a solution of higher global social welfare is to demonstrate that some function monotonically increases through changes of assignments during the run of the algorithm. It is clear from the definition of the function Φ (Equation 2) that it is bounded and this implies convergence. Furthermore, it will be shown that a run of Algorithm 2 results in an increase of the social welfare. It is most important to note that the following proof holds also for the general version of PGGs where the utility of an agent i is capped by a specific value $\mathbb{K}_i$ that holds for (the individual) agent i and might be different for the other agents. This result is new, because both Algorithm 1 and standard best-response dynamics were only shown to converge, and not necessarily to a better solution. In addition, former studies assumed a global $\mathbb{K}$ that caps the utility of all participating agents in [13].

Let us start by defining the following function

$$\Phi(v) = \sum_{i \in N} \phi_i(v) \tag{2}$$

with

$$\phi_i(v) = \begin{cases} \min(\mathbb{K}_i, l_i + 1) - c_i, & \text{if } v_i = T \\ \mathbb{K}_i - c_i, & \text{if } state_i(v) = F_{\mathbb{K}_i - 1} \\ \min(\mathbb{K}_i, l_i), & \text{if } state_i(v) = F_{l \neq \mathbb{K}_i - 1} \end{cases}$$

Observe the behavior of $\Phi(v)$ during a run of Algorithm 2 on a BSPGGs. To follow a run of Algorithm 2, one only needs to look at *rational strategy changes* - strategy changes that increase the utility of the deviating agent (i.e., the *current agent*), given the side payments it is offered (as defined in procedure **when received**). Note that when agent i performs a strategy change - the value of $\phi_j(v)$ can change only for agent i and for agents in its neighborhood N_i.

During the proof, v denotes the strategy profile from which agent i deviates, and v' denotes the resulting strategy profile. Another notation that is used in the proof is $\phi(B_l)$, which stands for "the value of ϕ for an agent whose state is B_l" (B is either T or F).

Lemma 3.1. *$\Phi(v') \geq \Phi(v)$ for any rational deviation of agent i from F to T.*

Proof. Agent i may change its strategy for one of two reasons - either to increase its own utility (**case 1**), or due to side payments offered to it by its neighbors (**case 2**).

Case 1: $\Phi(v') \geq \Phi(v)$ because the following arguments hold:

- agent i has less then $\mathbb{K}_i$ investing agents and it always holds that $\phi(T_{l<\mathbb{K}_i}) > \phi(F_{l<\mathbb{K}_i})$.
- for each neighboring agent it always holds that $\phi(B_{l+1}) \geq \phi(B_l)$.

Case 2: agent i has at least $\mathbb{K}_i$ investing agents, which implies that $\phi_i(v') - \phi_i(v) = -c_i$. However, there must be a nonempty set of agents $J \subseteq N_i$ who offered agent i side payments that sum up to $\sum_{j\in J} c_j \geq c_i$ to make agent i change to strategy T. For each of these agents it holds that ϕ_j increased by at least c_j (if agent j has exactly $\mathbb{K}_j$ investing agents in the new strategy profile v'), and the overall increase sums up to at least c_i. This implies that the relation $\phi_i(v') + \sum_{j\in J} \phi_j(v') \geq \phi_i(v) + \sum_{j\in J} \phi_j(v)$ still holds. For every other agent $k \in N_i,\ k \notin J$ we have $\phi_k(v') \geq \phi_k(v)$ as in case 1, and $\Phi(v') \geq \Phi(v)$ as required. □

Now we come to the more interesting lemma, that proves the increase of the function even for changes of strategy that have the potential to decrease the value of the function. These changes are from the value T to F. The proof relies on changes of strategy that follow Algorithm 2 and are rational for the current agent. In other words, that do not decrease its gain.

Lemma 3.2. *When agent i rationally deviates from T to F, it holds that $\Phi(v') > \Phi(v)$.*

Proof. For a rational change of strategy of agent i from strategy profile v to take place, two conditions must hold:

(1) no agent $j \in N_i$ is in state $T_{l\leq\mathbb{K}_j-1}$.
(2) no agent $j \in N_i$ is in state $F_{l\leq\mathbb{K}_j-1}$.

To see why these conditions must hold, observe that a neighboring agent in either of the above states offers a side payment of 1 (under procedure **when received** of Algorithm 2) to the current agent i in order to make it not deviate. A rational agent i who is offered such a payment will keep investing.

To complete the proof let us consider agents in N_i for which $\phi_j(v') < \phi_j(v)$. Let us denote this set of agents by J. For this decrease in ϕ_j to hold for agent $j \in J$, the agent must have exactly $\mathbb{K}_j$ investing agents in the strategy profile v. It is easy to see that for these agents $\phi_j(v) - \phi_j(v') = c_j$. This group of agents offer payments that sum up to $\sum_{j\in J} c_j$, but since the current agent i chooses to stop investing one can deduce that $c_i > \sum_{j\in J} c_j$. As was just shown, $\phi_i(v') - \phi_i(v) = c_i$ and $\sum_{j\in J} \phi_j(v) - \sum_{j\in J} \phi_j(v') = \sum_{j\in J} c_j$ which implies that $\Phi(v') - \Phi(v) = \phi_i(v') - \phi_i(v) + \sum_{j\in J} \phi_j(v') - \sum_{j\in J} \phi_j(v) = c_i - \sum_{j\in J} c_j > 0$ as required. □

Now the convergence proof is immediate:

Proposition 3.3. *A run of Algorithm 2 on a BSPGG converges.*

Proof. From Lemma 3.1, a deviation $F \to T$ doesn't decrease $\Phi(v)$. From Lemma 3.2, a deviation $T \to F$ increases $\Phi(v)$ so there can be only finitely many such deviations during a run of Algorithm 2 on a BSPGG. In addition, there can be only a finite number of $F \to T$ deviations before either a $T \to F$ deviation takes place, or the algorithm terminates, □

The fact that Φ is non-decreasing implies a unique property of Algorithm 2 - the outcome it converges to is at least as efficient as its initial outcome. This is not guaranteed by either best-response dynamics or by Algorithm 1 (except for the restricted case of $\mathbb{K} = 1$ [13]) and is a major result of the present study.

Corollary 3.4. *The global utility in outcome v' of a run of Algorithm 2 on a BSPGG is higher than the global utility in the initial strategy profile v.*

Proof. For a general strategy profile s, observe that $\sum_{i\in N} u_i(s) \leq \sum_{i\in N} \phi_i(s)$:

- For an agent i with $state_i(s) \neq F_{\mathbb{K}_i-1}$, it holds that $\phi_i(s) = u_i(s)$.
- For an agent i with $state_i(s) = F_{\mathbb{K}_i-1}$, it holds that $\phi_i(s) = \mathbb{K}_i - c_i > \mathbb{K}_i - 1 = u_i(s)$.

In addition, when the algorithm terminates no agent i can be in state $F_{\mathbb{K}_i-1}$. This is because such an agent can change its strategy to T and increase its utility. Consequently, $\sum_{i\in N} u_i(v') = \sum_{i\in N} \phi_i(v')$ upon termination.

Along with the guarantee $\Phi(v') \geq \Phi(v)$ we have that $\sum_{i\in N} u_i(v) \leq \sum_{i\in N} \phi_i(v) \leq \sum_{i\in N} \phi_i(v') = \sum_{i\in N} u_i(v')$. □

4 EXPERIMENTAL EVALUATION

All experiments were conducted on randomly generated scale-free networks (Barabasi-Albert networks [1]) with 500 agents and various densities. Scale-free networks are commonly taken to resemble real world social networks [2, 21]. The density of a Barabasi-Albert network can be represented by a parameter m - the number of nodes that connect each new node to the already existing network during construction [17]. All results were averaged over 100 runs for each configuration.

The first set of experiments evaluates Algorithm 2 against both best-response dynamics and Algorithm 1.

Algorithm 2 outperforms both algorithms. The parameter $\mathbb{K}$ is first taken to be 5. On low density networks with density 5, the improvement is 44% over best-response and 12% over Algorithm 1. The improvement decreases as the density grows, to 19% over best-response and 1% over Algorithm 1 on dense networks with density 15 (Figure 2). The improvement of Algorithm 2 is more significant when $\mathbb{K} = 15$. It performs 31% better than best-response and 8% better than Algorithm 1 on low density networks with density 5. The improvement peaks to 42% over best-response and 25% over Algorithm 1 when the density is 10, and goes down to 36% and 9% as the density reaches 15 (Figure 3).

All algorithms terminate within few iterations. It takes Algorithm 2 less iterations than best-response to converge, but more than Algorithm 1 (Figures 4, 5).

The second set of experiments explores the correlation between the social capital of agents and the utility they gain in solutions

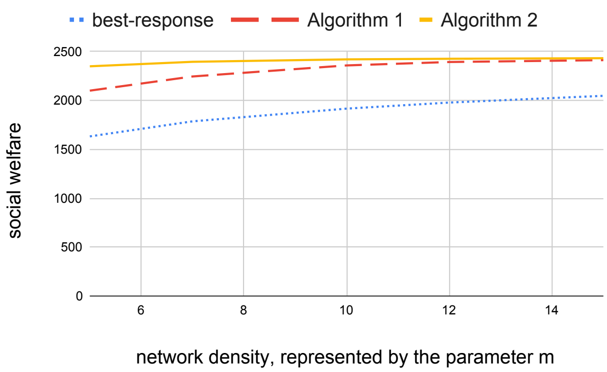

Figure 2: SW for solutions of the 3 Algorithms. $\mathbb{K} = 5$

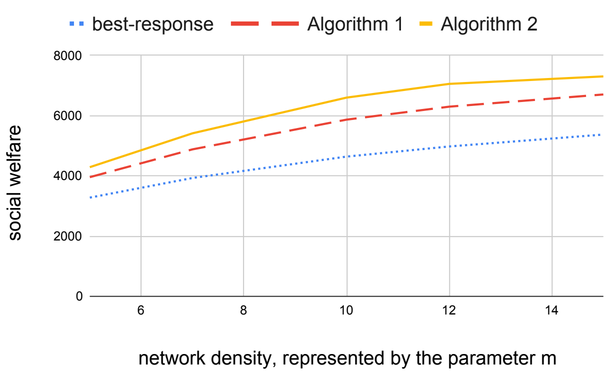

Figure 3: SW for solutions of the 3 Algorithms. $\mathbb{K} = 15$

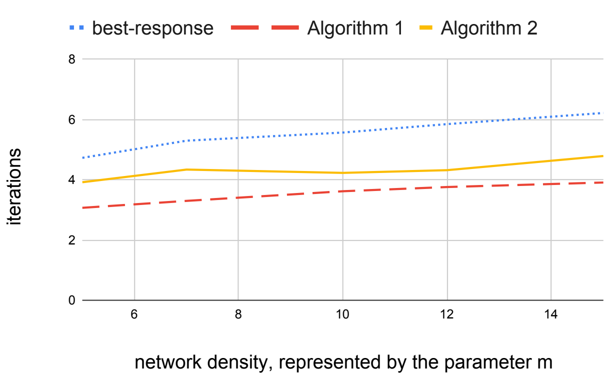

Figure 4: Number of Iterations for the 3 Algorithms. $\mathbb{K} = 5$

to the public goods game. In order to do so, the best-response algorithm was applied to randomly generated problems until convergence (the resulting outcomes of this method are PNEs - the common solution concept for games). Next, the Pearson correlation coefficients between the measurable social capital parameters of

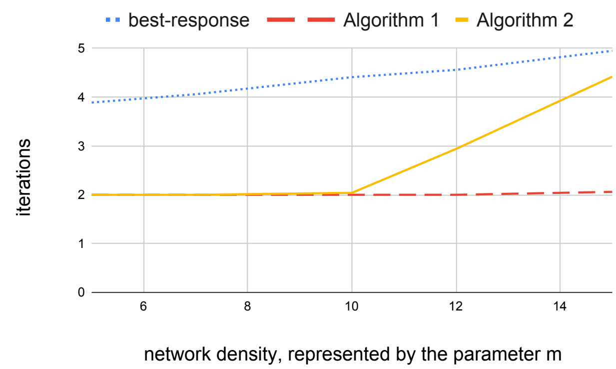

Figure 5: Number of Iterations for the 3 Algorithms. $\mathbb{K} = 15$

agents (as described in Section 2.3) and their utility in all solutions were computed.

The problem of directly measuring the correlation between the proposed measures of social capital and agents' utilities is that the parameters proposed by [9] also correlate with the degree of the agents on the network, and the degree naturally correlates with the utilities of the agents. Take for example an agent i with a high *godfather index* (recall Section 2.3). It has a higher number of pairs of unconnected neighbors than some other agent j with a lower *godfather index*. Therefore, i is also likely to have more neighbors than j in general. In BSPGGs, where the utility is determined by the total number of investing neighbors, agent i is expected to end up with a higher utility than agent j, regardless of its *godfather index*. To try and overcome this bias, the correlation between social capital and utility was computed only for agents with similar number of neighboring agents. Specifically, we explored only the set of agents that have **exactly** $\mathbb{K}$ **neighbors**.

The first experiment was performed on networks with density $m = 5$ (and $\mathbb{K} = 5$). A negative correlation of -0.47, -0.25 was observed between both *decay centrality, support index* and the utilities of agents with exactly $\mathbb{K}$ neighbors. In contrast, the *Godfather index* was found to be positively correlated (0.26) with the utility of the agents. The second experiment was performed on dense networks of $m = 15$ ($\mathbb{K} = 15$). The results were similar in that a negative correlation of -0.64, -0.4 was observed between both *decay centrality, support index* and the utilities of agents with exactly $\mathbb{K}$ neighbors. The *Godfather index* was found to be positively correlated (0.55) with the utility of the agents.

The positive correlation of *godfather index* as well as the negative correlation of both *decay centrality, support index* with the utility seem to imply a single predictor for the utilities of agents in solutions of the best-shot public goods game on networks: **The less connected one's neighbors in a game, the greater its utility is going to be**. This predictor is obviously positively correlated with the *godfather index* and negatively correlated with both *decay centrality, support index*).

The last experiment on the effect of network topology on agents' utility is presented in Figure 6. In this experiment, the average utility of agents with exactly $\mathbb{K} = 5$ neighbors residing on networks

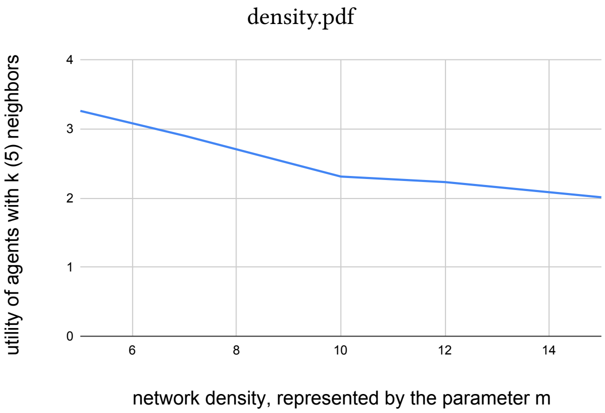

Figure 6: Utility of agents with $\mathbb{K}$ (5) neighbors

of various densities is examined. The utility of this set of agents is relatively high on low density networks with $m = 5$. It is interesting to see that their average utility decreases as the density of the network increases. This result shows that the power of an agent drops as its "relative degree" decreases such that the rest of the agents in the game can make "relations" and deal with more agents.

5 CONCLUSION

An incentive-base search algorithm for finding efficient PNEs in public goods games on networks is proposed. The proposed algorithm is a part of a new approach to public goods games on networks, that uses side payments among the searching agents and guarantees the finding of good solutions for large games. The outcome of the algorithm is at least as efficient as the initial strategy profile. This is a new result for the proposed algorithm, that is not guaranteed by former local search algorithms. The extensive experimental evaluation of the proposed algorithm was performed on large randomly generated scale-free networks of 500 agents and a variety of densities. All runs converged very quickly (within a few iterations) to efficient solutions.

An additional new direction that opens when one actually solves PGGs on networks, is to consider the individual gains of agents. The approach taken by the present study is to investigate the correlation of individual gains with the *social capital* of agents. To this end the recent typology of social capital by Jackson [9] was used. The typology bases social capital features on agents' local network topology. The extensive experimental evaluation explores correlations of the personal gain of agents in solutions obtained by best-response dynamics with several important features of social capital according to the model. The experiments show that local network topology plays an important role in determining the gain of agents in general public goods games on networks.

It is important to observe that the side-payments in the proposed algorithm (as well as in the former study [13]) are not strategic. A selfish neighboring agent might bid strategically and offer a lower payment that could still be sufficient for the current agent to invest. Future study of incentive-based search for public goods games on networks should integrate a truthful bidding mechanism into the search process and eliminate selfish behavior. This was done recently for ADCOPs search [20] and needs to be adapted for the different versions of PGGs.

REFERENCES

[1] Albert-Laszlo Barabasi and Reka Albert. 1999. Emergence of scaling in random networks. *Science* 286, 5439 (1999), 509–512.

[2] Albert-László Barabási and Eric Bonabeau. 2003. Scale-free networks. *Scientific american* 288, 5 (2003), 60–69.

[3] Yann Bramoulle and Rachel Kranton. 2007. Public Goods in Networks. *Journal of Economic Theory* 135 (2007), 478âĂŞ94.

[4] Yann Bramoulle, Rachel Kranton, and Martin DAmours. 2014. Strategic Interaction and Networks. *American Economic Review* 104 (2014), 898âĂŞ930.

[5] Gary Charness, Francesco Feri, Miguel A Meléndez-Jiménez, and Matthias Sutter. 2014. Experimental games on networks: Underpinnings of behavior and equilibrium selection. *Econometrica* 82, 5 (2014), 1615–1670.

[6] Luca DallAsta, Paolo Pin, and Abolfazl Ramezanpour. 2011. Optimal equilibria of the best shot game. *Journal of Public Economic Theory* 13, 6 (2011), 885–901.

[7] Andrea Galeotti, Sanjeev Goyal, Matthew O Jackson, Fernando Vega-Redondo, and Leeat Yariv. 2010. Network games. *The review of economic studies* 77, 1 (2010), 218–244.

[8] Matthew O Jackson. 2010. *Social and economic networks.* Princeton university press.

[9] Matthew O. Jackson. 2020. A typology of social capital and associated network measures. *Soc. Choice Welf.* 54, 2-3 (2020), 311–336.

[10] Matthew O Jackson and Simon Wilkie. 2005. Endogenous games and mechanisms: Side payments among players. *The Review of Economic Studies* 72, 2 (2005), 543–566.

[11] Michael J. Kearns. 2007. Graphical games. In *Vazirani, Vijay V.; Nisan, Noam; Roughgarden, Tim; Tardos, ÃL'va (2007). Algorithmic Game Theory.* Cambridge University Press, 159–178.

[12] David Kempe, Sixie Yu, and Yevgeniy Vorobeychik. 2020. Inducing Equilibria in Networked Public Goods Games through Network Structure Modification. In *Proc. 19th Intern. Conf. Auton. Agents MultiAgent Sys. (AAMAS-2020).* 611âĂŞ619.

[13] Vadim Levit, Zohar Komarovsky, Tal Grinshpoun, and Amnon Meisels. 2018. Incentive-based search for efficient equilibria of the public goods game. *Artif. Intell.* 262 (2018), 142–162.

[14] Glenn C Loury. 1976. *A dynamic theory of racial income differences.* Technical Report. Discussion paper.

[15] D. Monderer and L. S. Shapley. 1996. Potential Games. *Games and Economic Behavior* 14 (1996), 124–143.

[16] Robert W. Rosenthal. 1973. A class of games possessing pure-strategy Nash equilibria. *Int. J. of Game Theory* 2 (1973), 65–67.

[17] Francisco C. Santos and Jorge M. Pacheco. 2005. Scale-free networks provide a unifying framework for the emergence of cooperation. *Physical Review Letters* 95, 9 (2005), 098104.

[18] Yair Vaknin and Amnon Meisels. 2021. Local Maxima in ADCOPs via Side Payments. In *Proc. 20th IEEE/WIC/ACM Intern. Conf. Web Intell. Intellig. Agent Tech. (WI-IAT 21).* Melbourne, Australia.

[19] Yair Vaknin and Amnon Meisels. 2022. Distributed Resource Allocation with Side Payments. In *2022 International Conference on INnovations in Intelligent SysTems and Applications (INISTA).* IEEE, 1–6.

[20] Yair Vaknin and Amnon Meisels. 2022. Search on Asymmetric DCOPs by Strategic Agents. In *Proc. 21th IEEE/WIC/ACM Intern. Conf. Web Intell. Intellig. Agent Tech. (WI-IAT 22).* Niagara Falls, Canada.

[21] Jianjun Wu, Ziyou Gao, Huijun Sun, and Haijun Huang. 2004. Urban transit system as a scale-free network. *Modern Physics Letters B* 18, 19n20 (2004), 1043–1049.

[22] Sixie Yu, Kai Zhou, P. Jeffrey Brantingham, and Yevgeniy Vorobeychik. 2020. Computing Equilibria in Binary Networked Public Goods Games. In *34th Conf. Artif. Intell. (AAAI-2020).* 2310–2317.