Games and Economic Behavior 139 (2023) 161–179

Contents lists available at ScienceDirect

Games and Economic Behavior

journal homepage: www.elsevier.com/locate/geb

Public goods games in directed networks

Christos Papadimitriou, Binghui Peng *

Department of Computer Science, Columbia University, United States of America

ARTICLE INFO

Article history:
Received 8 September 2021
Available online 1 March 2023

Keywords:
Public goods
Graphical games
Equilibrium computation
Complexity theory
Algorithms

ABSTRACT

Public goods games in undirected networks are generally known to have pure Nash equilibria, which are easy to find. In contrast, we prove that, in directed networks, a broad range of public goods games have intractable equilibrium problems: The existence of pure Nash equilibria is NP-hard to decide, and mixed Nash equilibria are PPAD-hard to find. We define general utility public goods games, and prove a complexity dichotomy result for finding pure equilibria, and a PPAD-completeness proof for mixed Nash equilibria. Even in the divisible goods variant of the problem, where existence is easy to prove, finding the equilibrium is PPAD-complete. Finally, when the treewidth of the directed network is appropriately bounded, we prove that polynomial-time algorithms are possible.

1. Introduction

A public good is a resource which, once produced, is available to all (non-excludability), and can be enjoyed collectively by many agents (non-rivalry[1]). Scientific knowledge (Stiglitz, 1999), open-source software, vaccination for an infectious disease, volunteer work, information resources, and clean environment are fine examples of public goods. Since public goods can be produced at a cost and contribute to the utility of others, they enable a variety of strategic behaviors such as free-riding. Game theoretic formulations of public goods have been extensively studied by economists — see Bergstrom et al. (1986) for a classical framework for the public goods problem within which a unique Nash equilibrium exists.

Networks are perfect arenas for public goods games (Bramoullé et al., 2007). Networks model the fact that a particular public good, such as a piece of software or protection due to the immunization of an individual, may not be accessible by all, but only by the neighbors of the node where it is produced. A node's utility then is an nondecreasing function of the goods in the neighborhood, minus the cost of the goods produced by the node. Almost all of the literature deals with the homogeneous case, where all nodes have the same two strategies (produce the common goods at a cost, or not) and the same utility function (see Yu et al. (2020) for an exception); in fact, the nondecreasing functions max and sum are typically considered. In this paper, we assume that all nodes have the same utility (even though they have different circumstances due to network connectivity, and hence the game is not symmetric unless the graph is), and we consider very general utility functions. There is extensive work on public goods in undirected networks (see the related work subsection), and the rough consensus seems to be that, in just about all variants of the problem (again, with the exception of Yu et al., 2020), pure Nash equilibria exist — typically corresponding to independent or dominating sets of the graph — and are easy to find.

Undirected graphs cannot model accurately all possible kinds of utility transfer. The ability to enjoy the public goods produced by others is not necessarily symmetric — for example, clean air in a neighboring city is of no use if that city is

* Corresponding author.
E-mail addresses: christos@columbia.edu (C. Papadimitriou), bp2601@columbia.edu (B. Peng).

[1] Non-rivalry was called *collective consumption* by Paul Samuelson, who initiated the study of the subject (Samuelson, 1954).

https://doi.org/10.1016/j.geb.2023.02.002
0899-8256/

downwind; if my house is on a co-worker's way to work, then the good of carpooling to work produced by my co-worker benefits me, but not vice-versa. For another example, in peer-to-peer networks, agents can selfishly leave a torrent or altruistically participate in its provision, and this creates a non-reciprocal beneficial relationship. Importantly, social networks are often directed, e.g. Twitter, Instagram, Flickr, beneficial resource (e.g., information, ideas) usually flows along the direction of social network. Therefore, it is of interest to explore public goods games on directed networks. There has been some work on public goods in directed networks, e.g. López-Pintado (2013), where sufficient conditions for the existence of equilibria are developed. The general impression one gets from the literature is that the matter of equilibria in the directed case is more subtle.

This paper is a comprehensive exploration of the complexity of the equilibrium problem in public goods games on directed graphs.

The simplest and most widely studied variant of the problem is the indivisible case with the max utility: the decision a node faces is whether or not to produce the good; and a node does not need to produce the good if one or more of its predecessors have it. It turns out to be quite intricate. It is easy to see that a pure equilibrium may not exist (consider a directed odd cycle), and it turns out that it is NP-complete to decide if a pure equilibrium does exist. We give a simple reduction to that effect (Theorem 3.1).

We then generalize this NP-completeness result to a full *complexity dichotomy* of nondecreasing utility functions. We identify three families of utility functions that can be solved in polynomial time: The *flat functions*, *the steep functions*, and *the alternating functions.* The first two have trivial equilibria where all nodes abstain or all nodes produce the good, respectively. In the case of alternating functions, finding a pure Nash equilibrium is shown to be equivalent to solving a system of equations in $\mathbb{F}_2$. The main part of the proof entails showing that all other functions make the equilibrium problem NP-hard (Theorem 3.4).

It is interesting to note that this generalized problem, with an arbitrary nondecreasing utility, has not been considered in the case of undirected graphs. As we point out in Section 7, this problem is quite nontrivial, in the sense that it is NP-complete for some utility functions, while of course for others it is polynomial. Determining the precise dichotomy in the undirected case seems a very challenging open problem that is left open here.

Since pure equilibria in these games are fraught with non-existence and NP-completeness, can we find in polynomial time a mixed Nash equilibrium (guaranteed to exist by Nash's theorem)? We prove (Theorem 4.1) that this problem is PPAD-complete, even in the simplest case of the max utility; this is perhaps the most technically demanding proof in this paper. We reduce from the generalized circuit problem, proved to be PPAD-hard in Rubinstein (2018); Chen et al. (2009b). The reduction requires several new ideas, including the definition of a new kind of intermediate game — in addition to several that already exist in this literature — which we call *the threshold game*, and we believe is of interest in its own right. Finally, when the goods are divisible, the *sum* case of the problem (the utility is the summation of the neighbors minus the good's cost) is also PPAD-complete, this time by a reduction from mixed Nash equilibria in two-player win-lose games (Chen et al., 2007; Abbott et al., 2005).

All of our complexity results hold for sparse networks, with indegrees and outdegrees at most three. But how about networks that are tree-like *in the sense of graph minors* (Robertson and Seymour, 1986)? We show that, when the (underlying undirected) network has *bounded treewidth*, essentially all versions of the Nash equilibrium problem of network public goods games can be solved, or at least approximated arbitrarily close, in polynomial time. Our algorithm and techniques are inspired by Daskalakis and Papadimitriou (2006) and Thomas and van Leeuwen (2015), but several substantial adaptations and innovations are needed.

Our contributions. In summary, our main contributions are these:

- The formulation of public goods games in directed networks with a general objective function — beyond the two functions treated in the literature, max and sum — leading to a surprisingly rich and diverse family of problems and a precise P/NP-complete dichotomy (Section 3).
- Sweeping intractability results for the equilibrium problem of public goods games in directed networks, including a PPAD-completeness proof through threshold games (Section 4), an intriguing analysis of polynomial special cases for the pure equilibrium problem culminating in a precise P/NP-complete dichotomy (Section 3), and even a very different PPAD-completeness proof for divisible goods (Section 5).
- An approximation algorithm when the treewidth is $O(\frac{\log n}{\log\log n})$, through the development of new and enhanced techniques for approximating equilibrium problems in graphical games with small treewidth (Section 6).

1.1. Related work

Bramoullé and Kranton (2007) initiated the study of public goods in a network. They consider a type of pure Nash equilibrium called *specialized equilibrium*, and prove that such equilibria are *stable* under small perturbations, *universal* (always exist), and in fact *computable* by a natural distributed algorithm, since they correspond to maximal independent sets of the graph; see Dall Asta et al. (2011); Boncinelli and Pin (2012); López-Pintado (2013); Feldman et al. (2013); Bramoullé et al. (2014); Allouch (2015); Shin et al. (2017); Elliott and Golub (2019); Yu et al. (2020); Kempe et al. (2020); Bervoets and Faure (2019); Acemoglu et al. (2015) for follow-up works. Bramoullé et al. (2014) extended the theory to imperfectly substitutable public goods, and proved the existence of a unique Nash equilibrium, assuming that the graph's lowest eigenvalue is sufficiently small. Allouch (2015) differentiates private provision from public provision, and again characterizes the existence and uniqueness of a Nash equilibrium through the lowest eigenvalue of the graph. Bervoets and Faure (2019) examine the stability of Nash equilibrium and provide necessary and sufficient conditions for Nash equilibria to be asymptotically stable via the best-response dynamic. Public goods games were first generalized to directed graphs in López-Pintado (2013), who provide sufficient conditions for pure Nash equilibria to exist. The only complexity result regarding such public goods games we are aware of is Yu et al. (2020): finding a pure Nash equilibrium of a discrete version of the public goods game, albeit in the far more general case of *heterogeneous* agents, is NP-hard. A concurrent work of Bayer et al. (2021) studies the best response dynamic of public goods in directed graph and provides sufficient conditions for convergence. We refer interested readers to the surveys (Jackson and Zenou, 2015; Galeotti et al., 2010; Bramoullé and Kranton, 2015) for a general coverage of this area.

Our work uses certain ideas from *graphical games* (Kearns et al., 2001). It is NP-hard to find a pure Nash equilibrium (Gottlob et al., 2005) and PPAD-hard to compute, even approximately, a mixed Nash equilibrium (Daskalakis et al., 2009a; Rubinstein, 2018; Chen et al., 2009b) of a general graphical game with maximum degree 3. However, the problem is tractable in several settings (Daskalakis and Papadimitriou, 2006; Thomas and van Leeuwen, 2015; Daskalakis and Papadimitriou, 2015). Daskalakis and Papadimitriou (2006) developed a polynomial-time approximation scheme (PTAS) for computing an ϵ-approximate Nash equilibrium when the game has bounded strategy size, the network has bounded neighborhood size and $O(\log n)$ treewidth. Thomas and van Leeuwen (2015) provided an algorithm that computes a pure Nash equilibrium in poly(s^w, $|M|$), where s is the strategy size, w the *treewidth* of the graph and $|M|$ the size of the payoff matrix. We use similar ideas in our main algorithmic result for computing Nash equilibria in public goods problems for networks of bounded treewidth, but we have to address the problem that, in the present case, the parameter $|M|$ of this algorithm is exponential.

The PPAD complexity class was introduced by Papadimitriou (1994) to capture one particular genre of total search functions, encompassing the notion of equilibrium. The PPAD-completeness of Nash equilibria was established in Daskalakis et al. (2009a); Chen et al. (2009b) and extended recently in Rubinstein (2018, 2016). Over the past decades, a broad range of problems have been proved to be PPAD-hard, including equilibrium computation (Daskalakis et al., 2009b; Abbott et al., 2005; Chen et al., 2007, 2015; Mehta, 2014), market equilibrium (Chen and Teng, 2009; Chen et al., 2009a; Vazirani and Yannakakis, 2011; Chen et al., 2011, 2013; Chaudhury et al., 2022; Chen et al., 2022), equilibrium in auction (Chen et al., 2021; Filos-Ratsikas et al., 2021; Chen and Peng, 2023), fair allocation (Othman et al., 2016; Chaudhury et al., 2021, 2020), min-max optimization (Daskalakis et al., 2021) and problems in financial networks (Schuldenzucker et al., 2017).

Subsequent work. Deligkas et al. (2022) improve the inapproximability constant of threshold game to 1/6 and show it to be tight, it thus gives improved inapproximability constant for public goods game.

2. Model

A *public goods game* is a game with n players, defined through a directed graph $G(V, E)$ without loops, where $V = \{1, \dots, n\}$ is the set of players. We use $N(i)$ to denote the *neighborhood* of i, namely incoming neighbors of agent i, i.e., $N(i) = \{i\} \cup \{j | (j, i) \in E\}$. We assume common game theoretic terms and notation, such as strategy, strategy profile, pure Nash equilibrium and (mixed) Nash equilibrium. If $\mathbf{s} = (s_1, \dots, s_n)$ is a strategy profile, we use s_{-i} to denote actions adopted by all agents *except* i.

As is almost always done with public goods games, we assume that *all players have the same strategy space and the same utility function*. In the indivisible good (discrete) case, the strategy space of all players is $S = \{0, 1\}$, while in the divisible (continuous) case $S = [0, \infty)$. To define the *utility function* U_i of a player i, we start with defining the *price* or *cost* p of producing the good s_i, common to all players. In the indivisible case, it is a single real $p(s_i) = p > 0$. In the divisible case it is a function $p : \mathbb{R}_+ \to \mathbb{R}_+$.

Once p has been fixed, the common utility function of agent i for the strategy profile $\mathbf{s}$ is $U_i(\mathbf{s}) = X_i(\mathbf{s}) - p(s_i)$, where $X_i : S^{|N(i)|} \to \mathbb{R}$ is a *symmetric social composition function of the strategies played by the players in* $N(i)$. Since players may have different indegrees, and thus different sizes of neighborhood, we assume for uniformity that the common social composition function X is a *symmetric* function from S^n to the reals, where the strategies of players not in $N(i)$ are all set to zero — a value that does not affect X. The composition functions studied by the vast majority of the literature is the *max* (or *best shot*, or *or*) function in the indivisible case, picking the maximum of the neighborhood's $0-1$ choices, while in the divisible case the composition functions *max* and *sum* is used.

In indivisible good games with *max* composition, in the literature it is always assumed that $p \neq 1$, because otherwise $p = 1$ creates ties between contributing and free-riding. For more general indivisible good games and social composition functions X, we shall also avoid ties between contributing and free-riding – it forms a measure 1 set among all possible choice of p. This completes the definition of the common general utility function U, and thus of the game.

We are interested in the standard concepts of pure and mixed Nash equilibrium. A strategy profile $\mathbf{s} = (s_1, \dots, s_n)$ of the public good game is a *(pure) Nash equilibrium*, if no agent can derive better utility by changing their own strategy,

$$u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i}) \quad \forall i \in V, \forall s_i, s_i' \in S_i.$$

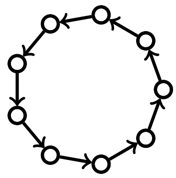

Fig. 1. Odd cycles have no pure Nash equilibrium.

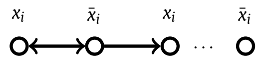

Fig. 2. Variable gadget.

In a *mixed Nash equilibrium* $\mathbf{\Delta} = (\Delta_1, \ldots, \Delta_n)$, each agent i plays a distribution Δ_i over its strategy set S_i, and satisfies

$$\mathop{\mathbb{E}}_{s_i \sim \Delta_i, s_{-i} \sim \Delta_{-i}} [u_i(s_i, s_{-i})] \geq \mathop{\mathbb{E}}_{s_{-i} \sim \Delta_{-i}} \left[u_i(s_i', s_{-i})\right] \quad \forall i \in V, \forall s_i' \in S_i. \tag{1}$$

Define $\mathrm{Supp}(\Delta_i)$ to be the support of the distribution Δ_i, i.e., $\mathrm{Supp}(\Delta_i) = \{s_i | s_i \in S_i, \Delta_i(s_i) > 0\}$. Then the definition in (1) is equivalent to

$$\forall i \in V, \forall s_i \in \mathrm{Supp}(\Delta_i), s_i' \in S_i : \mathop{\mathbb{E}}_{s_{-i} \sim \Delta_{-i}} [u_i(s_i, s_{-i})] \geq \mathop{\mathbb{E}}_{s_{-i} \sim \Delta_{-i}} \left[u_i(s_i', s_{-i})\right].$$

An ϵ*-approximately well supported Nash equilibrium* (ϵ-Nash) is then defined as

$$\forall i \in V, \forall s_i \in \mathrm{Supp}(\Delta_i), s_i' \in S_i : \mathop{\mathbb{E}}_{s_{-i} \sim \Delta_{-i}} [u_i(s_i, s_{-i})] \geq \mathop{\mathbb{E}}_{s_{-i} \sim \Delta_{-i}} \left[u_i(s_i', s_{-i})\right] - \epsilon. \tag{2}$$

When dealing with pure Nash equilibria, it is often easier to use the notion of a decision function. Given a composition function $X : \mathbb{N} \to \mathbb{R}$, define $f : \mathbb{N} \to \{0, 1\}$ as

$$f(t) = \begin{cases} 1 & X(t+1) - X(t) > p \\ 0 & X(t+1) - X(t) < p. \end{cases}$$

Recall that we assume there are no ties in the social composition function so f is well-defined. The decision function completely characterizes the best response in a pure NE. In particular, for each agent, suppose there are t neighbors who choose to produce the good; then the agent's best response is $f(t)$.

3. Pure Nash equilibria: a dichotomy

In this section we characterize the complexity of finding pure equilibria, focusing first on the best shot (max, or) function. In contrast to undirected networks, where every *maximal* independent set corresponds to a pure Nash equilibrium, pure Nash equilibria may not exist in directed graphs (see Fig. 1). We show in this section that determining whether a pure Nash equilibrium exists is NP-complete, and then generalize this to a sweeping complexity dichotomy result, characterizing precisely — modulo the P≠NP conjecture — the kinds of utility functions that have tractable Nash equilibrium problems.

Theorem 3.1. *Deciding whether a pure Nash equilibrium exists in an indivisible public good game with the max social composition function is NP-complete.*

Proof. In an equilibrium profile $\mathbf{s} = (s_1, \ldots, s_n)$, for each agent i, we have $s_i = 1$ if $\sum_{j \in N_i} s_j = 0$, and $s_i = 0$ otherwise. In another words, an agent would purchase the good if, and only if, none of its predecessors possesses the good. The reduction is from 3SAT, and employs the following two gadgets (see Fig. 2 and Fig. 3).

Variable gadget. For each variable x_i, we construct a directed path with $2k_i$ nodes, where k_i is the number of times x_i appears in the 3SAT instance. The path is directed, with the exception that there is a bi-directional edge between the first two nodes. The bi-directional edge forces the choice (exactly one of the first two nodes has the good), and the rest of the path propagates it (either all odd nodes have the good and all even nodes do not, or the other way around).

Clause gadget. The clause gadget consists of two parts. The left part is an OR gadget, in that node 4 must equal the disjunction of nodes 0, 1, and 2. To see this, suppose that none of these three nodes has the good; then node 3 must have it, and so node 4 does not. And if one or more of nodes 0, 1, 2 has the good, then 3 does not have the good, and thus 4 must have it.

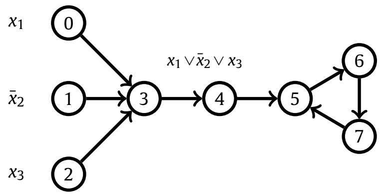

Fig. 3. Clause gadget.

The right part forces the clause to be true — that is, in any equilibrium profile, agent 4 must play strategy 1 and buys the good. This is because if 4 does not provide the good, then 5, 6, 7 is an isolated odd circle, which cannot exist in a pure Nash equilibrium. On the other hand, players 4 and 6 buying the good, and players 5 and 7 not buying it, is a pure Nash equilibrium of the four rightmost nodes. In summary, the clause gadget ensures that at least one of the nodes 0, 1, 2 buys the good.

Putting things together, given any 3SAT instance we can construct a public good game by composing variable gadgets and clause gadgets in the obvious way, so that the pure Nash equilibria of the public good game are in one to one correspondence with the satisfiable solutions of the 3SAT instance, concluding the proof. □

We want to generalize this result to *any* social composition function X, and so we start with the question: *For which composition functions is the pure Nash equilibrium problem polynomial-time solvable?* Consider a symmetric, non-decreasing function $X : \{0,1\}^n \mapsto \mathbb{R}_+$ without loss of generality with $X(0^n) = 0$. Because of symmetry, we can treat X as a function from $\mathbb{N}$ to $\mathbb{R}_+$, since its value depends on $\sum_{i=1}^n s_i$; we shall use the same symbol for this form of X,[2] and recall that $X(0) = 0$. Since X is monotone, it can be also thought as a sequence of nonnegative *steps*. Call X *flat* if $X(1) \le p$; that is, the first step of X does not provide sufficient incentive to produce the good. Obviously, all flat functions have the all-zero pure Nash equilibrium, and so the problem is trivial. Call now X *steep* if for all $k \ge 0$, $X(k+1) \ge X(k) + p$; that is, all steps are at least p. Then all nodes have an incentive to produce the good no matter what anybody else is doing, and so the all-ones solution is a pure Nash equilibrium, and again the problem is trivial. We have shown:

Lemma 3.2. *The pure Nash equilibrium problem is in P if the utility function is flat. Ditto for steep functions.*

Are there any other tractable cases? It turns out, that there is one more: Call X *alternating* if for all $k \ge 0$, $X(k+1) < X(k) + p$ if k is odd, and $X(k+1) > X(k) + p$ if k is even.

Lemma 3.3. *The pure Nash equilibrium problem is in P if the utility function is alternating.*

Proof. Let $s = (s_1, \dots, s_n)$ be the equilibrium profile with $s_i \in \{0,1\}$. Based on the definition of alternating utility function, we have that for any $i \in [n]$

$$s_i = \begin{cases} 1 & \sum_{(j,i)\in E} s_i = 0 (\bmod 2) \\ 0 & \sum_{(j,i)\in E} s_i = 1 (\bmod 2). \end{cases}$$

That is, a player chooses to produce when there is an even number of neighboring players who produce the good. Hence, the equilibrium problem reduces to the solution of a linear system of equations in $\mathbb{F}_2$ with one $0-1$ variable s_i per player, with one equation for each player i:

$$s_i + \sum_{(j,i)\in E} s_j = 1 (\bmod 2).$$

This can be solved in polynomial time with Gaussian elimination, say. □

We next establish that, unless P = NP, *these are the only tractable cases:*

Theorem 3.4. *If the utility function does not belong in these three classes: (1) flat; (2) steep; or (3) alternating, then the pure Nash equilibrium problem is NP-complete.*

Proof. We use the variable gadgets and the clause gadgets in the proof of Theorem 3.1 (see Fig. 2 and Fig. 3), but we reduce from several different NP-hard problems. First, observe that when X is not flat, steep, or alternating, there must be

[2] That is, we assume that X has values for all integers, not limited to the size of the network; this is obviously a harmless convention.

a $k \geq 0$ such that $X(k+1) > X(k) + p$, $X(k+2) < X(k+1) + p$. We start by assuming that $k = 0$, that is, $X(1) > p$ and $X(2) < X(1) + p$. We divide the proof into four cases.

Case 1. Suppose $X(3) < X(2) + p$, $X(4) < X(3) + p$, then it is easy to check that the construction of Theorem 3.1 works. Indeed, a function satisfying $X(1) > p, X(2) < X(1)+p, X(3) < X(2)+p, X(4) < X(3)+p$ is, for the purposes of the network constructed in the proof of the previous theorem, equivalent to the max function.

Case 2. Suppose $X(3) < X(2) + p$, $X(4) > X(3) + p$. We can still use the network constructed in the proof of Theorem 3.1. The difference is that we reduce from Not-all-equal SAT, since the clause gadget is satisfied if and only if one or two literals are true.

Case 3. Suppose $X(3) > X(2) + p$, $X(4) > X(3) + p$. Again, we consider the network constructed in the proof of Theorem 3.1. We can check that the clause gadget is satisfiable if and only if exactly one of the literal is true. The NP-hardness then comes from One-in-3SAT.

Case 4. Suppose $X(3) > X(2) + p$, $X(4) < X(3) + p$. Since we assume X is not alternating, there exists $t \geq 1$ satisfying $X(2t+1) > X(2t) + p$, $X(2t+2) < X(2t+1) + p$, and $X(2t+3), X(2t+4)$ does *not* obey $X(2t+3) > X(2t+2) + p$, $X(2t+4) < X(2t+3) + p$. We create $2t$ new players who have no incoming edges and directed edges to all other nodes. These $2t$ players will provide the good at equilibrium, and thus the remaining players start the game with $2t$ copies of the good already. The game for the original player is then changed to the function $\tilde{X}(j) = X(j - 2t)$, and NP-completeness follows from cases (1–3).

Finally, suppose that $k > 0$. Whenever (1) $X(k+3) < X(k+2) + p, X(k+4) < X(k+3) + p$; or (2) $X(k+3) < X(k+2) + p, X(k+4) > X(k+3) + p$; or (3) $X(k+3) > X(k+2) + p, X(k+4) > X(k+3) + p$, we add k new players who have no incoming edges, and directed edges to all other nodes. At equilibrium, these nodes will provide the good, and so the remaining players will start the game with k copies of the good already provided. Therefore, the game for the remaining players will be as if $X(j)$ was changed to $X(j-k)$, that is to say, to a function covered by Case (1-3).

There is one case left, i.e., for some $k > 0$: $X(1) > p, \ldots, X(k+1) > X(k) + p, X(k+2) < X(k+1) + p, X(k+3) > X(k+2) + p, X(k+4) < X(k+3) + p$, we can not reduce to Case 4 since X could be alternating after $X(k)$. We still create k new players and direct them to all other agents, *except* for node 3 in every clause gadget, for which we only connect $k-1$ players to it. Our argument in Theorem 3.1 works, with one modification: the clause gadget is satisfiable if and only if exactly two of the literals are true. This, again, is NP-complete, as we can reduce from One-in-3SAT. □

4. PPAD-hardness of mixed Nash equilibria

We next examine mixed Nash equilibria of indivisible public goods games. In a mixed Nash equilibrium, agents randomize over the two actions and choose to buy the public good with some probability. We denote by s_i the probability that agent i purchases the good. Also, by $x = y \pm \epsilon$ we mean that $y - \epsilon \leq x \leq y + \epsilon$. Throughout this section, whenever $\pm$ is used, we always cap the min at 0 and max at 1. For ease of presentation, we assume $U = 1$ and $p < 1$ throughout the proof. The following result is the main technical contribution of this paper.

Theorem 4.1. *There exists some constant $\epsilon > 0$, such that it is PPAD-hard to find an ϵ-Nash of the indivisible public goods game.*

Our reduction consists of two steps. We first introduce an intermediate game, called the *threshold game* (see Definition 4.2), where each individual's strategy depends solely on the *summation* of its neighbors' strategies. The threshold game exhibits rich algorithmic and complexity structure, which we believe could be of independent interest. We show a correspondence between equilibrium profiles of threshold games and those of public goods games (Lemma 4.4); hence it suffices to demonstrate PPAD-hardness of finding an ϵ-approximate equilibrium of threshold games. It is performed via a reduction from the ϵ-GCIRCUIT problem, which is shown to be PPAD-hard for sufficiently small constant $\epsilon > 0$ by Rubinstein (2018).

4.1. Equivalence between public goods games and threshold games

We first introduce the *threshold game*.

Definition 4.2 *(Threshold game). A threshold game $\mathcal{G}(V, E, t)$ is defined on a directed graph $G = (V, E)$, with a threshold t ($0 < t < 1$). The vertices of the graph represent players with strategy space $[0, 1]$. A strategy profile $\mathbf{x} = (x_1, \ldots, x_n) \in [0, 1]^n$ is an equilibrium if it satisfies*

$$x_i = \begin{cases} 0 & \sum_{j \in N_i} x_j > t \\ 1 & \sum_{j \in N_i} x_j < t \\ \textit{arbitrary} & \sum_{j \in N_i} x_j = t \end{cases}. \tag{3}$$

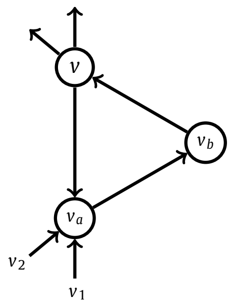

Fig. 4. Elementary gadget.

Note that x_i can be an arbitrary number in $[0, 1]$ if $\sum_{j\in N_i} x_i = t$.

We define the ϵ-approximate equilibrium in a threshold game as follows:

Definition 4.3 *(ϵ-approximate equilibrium of threshold game). Let $\epsilon > 0$ be a constant satisfying $\epsilon < t < 1 - \epsilon$. An ϵ-approximate equilibrium $\mathbf{x} = (x_1, \ldots, x_n) \in [0, 1]^n$ of a threshold game $\mathcal{G}(V, E, t)$ satisfies*

$$x_i = \begin{cases} 0 \pm \epsilon & \sum_{j\in N_i} x_j > t + \epsilon \\ 1 \pm \epsilon & \sum_{j\in N_i} x_i < t - \epsilon \\ \textit{arbitrary} & \sum_{j\in N_i} x_i \in [t - \epsilon, t + \epsilon] \end{cases} . \tag{4}$$

We establish the equivalence between threshold games and public good games. The proof can be found in Appendix.

Lemma 4.4. *There is a polynomial time reduction between the threshold game and the public good game. Specifically, (1) given any threshold game $\mathcal{G}(V, E, t)$ with $0 < t < 1$, we can construct a public good game and map any ϵ-Nash of the public goods game to an 8ϵ-approximate equilibrium of threshold game $\mathcal{G}(V, E, t)$, for $\epsilon < \min\{0.1, \frac{t}{8}, \frac{1-t}{8}\}$; (2) given any public good game with $U = 1, 0 < p < 1$, we can construct a threshold game $\mathcal{G}(V, E, t)$ and map any ϵ-approximate equilibrium of threshold game to an $c_p\epsilon$-Nash of public goods game, where $c_p = -4p \log p$ is a constant depending only on p.*

4.2. Reducing generalized circuits to threshold games

We give the definition of generalized circuits.

Definition 4.5 *(Generalized circuit (Chen et al., 2009b)). A generalized circuit is a tuple $(V, \mathcal{T})$, where V is a set of nodes and $\mathcal{T}$ is a collection of gates. Every gate $T \in \mathcal{T}$ is a 5-tuple $T = (G, v_1, v_2, v, \alpha)$, where $G \in \{G_\xi, G_{\times\xi}, G_=, G_+, G_-, G_<, G_\wedge, G_\vee, G_\neg\}$ is the type of the gate; $v_1, v_2 \in V \cup \{nil\}$ are the input nodes, $\alpha \in \mathbb{R} \cup \{nil\}$ is a real parameter and v is the output node.*

The collection $\mathcal{T}$ of gates must satisfy the following important property. For every two gates $T, T' \in \mathcal{T}$ $(T \neq T')$, $T = (G, v_1, v_2, v, \alpha)$ and $T' = (G', v_1', v_2', v', \alpha')$, we must have $v \neq v'$.

The ϵ-GCIRCUIT is the problem of finding an ϵ-approximate assignment for the generalized circuit. Notice that we replace G_ξ, $G_{\times\xi}$ with $G_{\frac{1}{2}}$, $G_{\times\frac{1}{2}}$ for ease of proof.

Definition 4.6. *Given a generalized circuit $\mathcal{S} = (V, \mathcal{T})$, we say an assignment $\mathbf{x} : V \to [0, 1]$ ϵ-approximately satisfies $\mathcal{S}$, if it satisfies the constraints shown in Table 1.*

We then prove

Theorem 4.7. *It is PPAD-hard to find an ϵ-approximate equilibrium of the threshold game, for some constant $\epsilon > 0$.*

The key step is to construct an elementary gadget $G_{\frac{1}{2}-}$ (see Fig. 4), and use it as a building block to gradually construct most of the gates. We fix the threshold $t = \frac{1}{2}$ in the rest of the proof. Furthermore, we restrict the equilibrium strategy in $[0, \frac{1}{2} + \epsilon] \cup \{1\}$, since for any ϵ-approximate equilibrium $\mathbf{x} = (x_1, \ldots, x_n)$, we could set

$$\tilde{x}_i = \begin{cases} x_i & x_i \le \frac{1}{2} + \epsilon \\ 1 & \text{otherwise} \end{cases}$$

Table 1
Gadgets.

Gate	Constraint
$G_{\frac{1}{2}}(v)$	$\mathbf{x}[v] = \frac{1}{2} \pm \epsilon$
$G_{\times\frac{1}{2}}(\mid v_1 \mid v)$	$\mathbf{x}[v] = \frac{1}{2} \cdot \mathbf{x}[v_1] \pm \epsilon$
$G_{=}(\mid v_1 \mid v)$	$\mathbf{x}[v] = \min\{\mathbf{x}[v_1], \frac{1}{2}\} \pm \epsilon$
$G_{+}(\mid v_1, v_2 \mid v)$	$\mathbf{x}[v] = \min\{\mathbf{x}[v_1] + \mathbf{x}[v_2], \frac{1}{2}\} \pm \epsilon$
$G_{-}(\mid v_1, v_2 \mid v)$	$\mathbf{x}[v] = \max\{\mathbf{x}[v_1] - \mathbf{x}[v_2], 0\} \pm \epsilon$
$G_{<}(\mid v_1, v_2 \mid v)$	$\mathbf{x}[v] = \begin{cases} \frac{1}{2} \pm \epsilon & \mathbf{x}[v_1] < \mathbf{x}[v_2] - \epsilon \\ 0 \pm \epsilon & \mathbf{x}[v_1] > \mathbf{x}[v_2] + \epsilon \end{cases}$
$G_{\wedge}(\mid v_1, v_2 \mid v)$	$\mathbf{x}[v] = \begin{cases} \frac{1}{2} \pm \epsilon & \mathbf{x}[v_1] = \frac{1}{2} \pm \epsilon \wedge \mathbf{x}[v_2] = \frac{1}{2} \pm \epsilon \\ 0 \pm \epsilon & \mathbf{x}[v_1] = 0 \pm \epsilon \vee \mathbf{x}[v_2] = 0 \pm \epsilon \end{cases}$
$G_{\vee}(\mid v_1, v_2 \mid v)$	$\mathbf{x}[v] = \begin{cases} 1 \pm \epsilon & \mathbf{x}[v_1] = \frac{1}{2} \pm \epsilon \vee \mathbf{x}[v_2] = \frac{1}{2} \pm \epsilon \\ 0 \pm \epsilon & \mathbf{x}[v_1] = 0 \pm \epsilon \wedge \mathbf{x}[v_2] = 0 \pm \epsilon \end{cases}$
$G_{\neg}(\mid v_1 \mid v)$	$\mathbf{x}[v] = \begin{cases} \frac{1}{2} \pm \epsilon & \mathbf{x}[v_1] = 0 \pm \epsilon \\ 0 \pm \epsilon & \mathbf{x}[v_1] = \frac{1}{2} \pm \epsilon \end{cases}$

and we can easily verify that $\tilde{\mathbf{x}} = [\tilde{x}_1, \ldots, \tilde{x}_n]$ is still an ϵ-approximate equilibrium. We use the strategies of players in the threshold game to represent an ϵ-approximate assignment to ϵ-GCIRCUIT, and build all 9 types of gates in $\{G_{\xi}, G_{\times\xi}, G_{=}, G_{+}, G_{-}, G_{<}, G_{\wedge}, G_{\vee}, G_{\neg}\}$. We start from constructing an elementary game gadget $G_{\frac{1}{2}-}(\mid v_1, v_2 \mid v)$ (see Fig. 4), where $v_1, v_2 \in V \cup \{nil\}$ are input players, $v \in V$ is the output player. The output player v could have many out-coming edges, but it only has one in coming edge from the internal node v_b. The elementary gadget $G_{\frac{1}{2}-}(\mid v_1, v_2 \mid v)$ serves as a building block for later constructions and proves useful throughout our proof. Ideally, it poses the constraint that $\mathbf{x}[v] = \max\{\frac{1}{2} - \mathbf{x}[v_1] - \mathbf{x}[v_2], 0\}$ in an equilibrium.

Lemma 4.8. *Consider the game gadget $G_{\frac{1}{2}-}(\mid v_1, v_2 \mid v)$ constructed in Fig. 4. In any ϵ-approximate equilibrium of the threshold game $\mathcal{G}(V, E, \frac{1}{2})$, we have $\mathbf{x}[v] = \max\{\frac{1}{2} - \mathbf{x}[v_1] - \mathbf{x}[v_2], 0\} \pm \epsilon$. In particular, if we set $v_2 = nil$, then $\mathbf{x}[v] = \max\{\frac{1}{2} - \mathbf{x}[v_1], 0\} \pm \epsilon$; if we set $v_1, v_2 = nil$, then $\mathbf{x}[v] = \frac{1}{2} \pm \epsilon$.*

Proof. Consider the in-coming neighbors of player v_a, $N_a = \{v, v_1, v_2\}$. In an ϵ-approximate equilibrium, if $\mathbf{x}[v] + \mathbf{x}[v_1] + \mathbf{x}[v_2] > \frac{1}{2} + \epsilon$, we have $\mathbf{x}[v_a] = 0 \pm \epsilon$ and $\mathbf{x}[v_b] = 1 \pm \epsilon$. It then follows that $\mathbf{x}[v] = 0 \pm \epsilon$. This implies $\mathbf{x}[v_1] + \mathbf{x}[v_2] > \frac{1}{2}$, and thus $\max\{\frac{1}{2} - \mathbf{x}[v_1] - \mathbf{x}[v_2], 0\} \pm \epsilon = 0 \pm \epsilon$. This satisfies the equilibrium condition. If $\mathbf{x}[v] + \mathbf{x}[v_1] + \mathbf{x}[v_2] < \frac{1}{2} - \epsilon$, then we have $\mathbf{x}[v_a] = 1 \pm \epsilon$ and $\mathbf{x}[v_b] = 0 \pm \epsilon$, this again implies $\mathbf{x}[v] = 1 \pm \epsilon$. This contradicts with the fact that $\mathbf{x}[v] + \mathbf{x}[v_1] + \mathbf{x}[v_2] < \frac{1}{2} - \epsilon$. In summary, we have $\mathbf{x}[v] = \max\{\frac{1}{2} - \mathbf{x}[v_1] - \mathbf{x}[v_2], 0\} \pm \epsilon$. □

We next construct $G_{=}$ and G_{+}. We assume that the inputs of these gates belong to $[0, \frac{1}{2} + \epsilon]$, *except* for the COPY gate. It is not a loss of generality since: (1) if the input node v_1 is also the output node of another gate, then its value is guaranteed to be in $[0, \frac{1}{2} + \epsilon]$ by our construction below; (2) otherwise, we can always apply a COPY gate (see below) to restrict its value in $[0, \frac{1}{2} + \epsilon]$.

(1) COPY $G_{=}(\mid v_1 \mid v)$. Concatenating $G_{\frac{1}{2}-}(\mid v_1 \mid v_2)$ with $G_{\frac{1}{2}-}(\mid v_2 \mid v)$, then we have $\mathbf{x}[v_2] = \max\{\frac{1}{2} - \mathbf{x}[v_1], 0\} \pm \epsilon$, and $\mathbf{x}[v] = \max\{\frac{1}{2} - \mathbf{x}[v_2], 0\} \pm \epsilon = \min\{\mathbf{x}[v_1], \frac{1}{2}\} \pm 2\epsilon$.
(2) ADD $G_{+}(\mid v_1, v_2 \mid v)$. Concatenating $G_{\frac{1}{2}-}(\mid v_1, v_2 \mid v_3)$ with $G_{\frac{1}{2}-}(\mid v_3 \mid v)$, then we have $\mathbf{x}[v_3] = \max\{\frac{1}{2} - \mathbf{x}[v_1] - \mathbf{x}[v_2], 0\} \pm \epsilon$ and $\mathbf{x}[v] = \max\{\frac{1}{2} - \mathbf{x}[v_3], 0\} \pm \epsilon = \min\{\mathbf{x}[v_1] + \mathbf{x}[v_2], \frac{1}{2}\} \pm 2\epsilon$.

Thanks to a concurrent work of Filos-Ratsikas et al. (2021), this already suffices to prove the PPAD-hardness of threshold game. In particular, we use

Lemma 4.9 *(Proposition 5.3 of Filos-Ratsikas et al. (2021)). There exists a constant $\epsilon > 0$ such that the problem of ϵ-GCIRCUIT with gate-types $G_+(\,|\,v_1, v_2\,|\,v)$ and $G_{\frac{1}{2}-}(\,|\,v_1, v_2\,|\,v)$ is PPAD-complete.*

Combining Lemma 4.8 and Lemma 4.9, we have proved Theorem 4.7. Together with Lemma 4.4, the proof of Theorem 4.1 is complete.

Theorem 4.1 only covers the max utility function. We *conjecture* that it holds for all utility functions except for the polynomial cases discussed in Theorem 3.4. There are several interesting challenges in extending the proof to this direction.

5. Divisible goods

For divisible public goods games in directed graphs, we study the three most studied utility functions, and completely characterize the equilibrium problem:

- For the summation utility (the utility of each node is the sum of the amounts of goods provided by its predecessors), a pure Nash equilibrium always exists, but it is PPAD-complete to find one.
- For the best-shot function (max), a pure Nash equilibrium may not exist and it is PPAD-hard to find a mixed Nash equilibrium; the reasons are quite similar to the indivisible case.
- Finally, for the weakest-link function (min), it turns out that there are always multiple trivial pure Nash equilibria, as no player has the incentive to supply any amount of the good.

5.1. Summation

When the utility function is the summation, Bramoullé et al. (2007) prove that there is always a pure Nash equilibrium.[3] Here, we prove it is PPAD-hard to find one, when the network is directed. In fact, we prove a slightly stronger result: call a strategy profile $\mathbf{s} = (s_1, \dots, s_n)$ an *ϵ-approximate pure Nash equilibrium*, if $s_i = b_i(s_{-i}) \pm \epsilon$, where $b_i(\cdot)$ is the best response of agent i.

Theorem 5.1. *It is PPAD-hard to find an ϵ-approximate pure Nash equilibrium of divisible public goods games with summation utility function, for $\epsilon = 1/\operatorname{poly}(n)$.*

We reduce from the mixed Nash equilibrium problem in two-player win-lose games. A two-player game (R, C) is win-lose if $R, C \in \{0, 1\}^{n\times n}$. It is known (Chen et al., 2007; Abbott et al., 2005) that finding an ϵ-Nash of two-player win-lose game is PPAD-hard for $\epsilon = 1/\operatorname{poly}(n)$. It is an interesting question whether one can improve the hardness of approximation to constant $\epsilon > 0$.

Given an instance (R, C) of a two player win-lose game, we construct a divisible public goods game on a directed network, such that we can map any ϵ-approximate pure Nash equilibrium of the public goods game to a $\operatorname{poly}(n)\cdot\epsilon$-Nash of two-player win-lose game (R, C). In order to do so, we first symmetrize the win-lose game (Lemma 5.2), and then reduce it to the public goods game (Lemma 5.3).

For convenience, we assume that $R, C \in \{-1, 0\}^{n\times n}$ and that there is no weakly-dominated strategy for both row and column players. Moreover, we assume every column (row) of $R(C)$ contains at least one 0 entry — otherwise, there is a trivial pure Nash equilibrium. Define a symmetric game (A, B) as follows:

$$A = \begin{pmatrix} -\mathbf{1} & R \\ C^T & -\mathbf{1} \end{pmatrix} \text{ and } B = \begin{pmatrix} -\mathbf{1} & C \\ R^T & -\mathbf{1} \end{pmatrix},$$

where $-\mathbf{1}$ denotes an $n \times n$ all -1 matrix. We notice that $A, B \in \{-1, 0\}^{2n\times 2n}$ and $A = B^T$. The above symmetrization is standard in the literature (Lemke and Howson, 1964), and it is known that for any symmetric Nash equilibrium (x, y) of (A, B), $(x/|x|, y/|y|)$ is a Nash equilibrium for (R, C). The following lemma states that approximation is preserved:

Lemma 5.2. *Suppose (x, y) is a symmetric ϵ-Nash equilibrium of game (A, B). Then $(\tilde{x}, \tilde{y}) = (x/|x|, y/|y|)$ is a $4n\epsilon$-Nash of the win-lose game (R, C).*

Proof. We first prove that $|x|, |y| \geq \frac{1}{4n}$. Notice that

$$A\begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} -\mathbf{1} & R \\ C^T & -\mathbf{1} \end{pmatrix}\begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} -|x|\mathbf{e} + Ry \\ -|y|\mathbf{e} + C^T x \end{pmatrix},$$

[3] They actually prove it for undirected networks, but their proof generalizes easily to the directed case, as the best response function for each agent i is still continuous in s_{-i}, and existence of a pure Nash equilibrium follows from Brouwer's fix point theorem. In fact, when the valuation function is strictly concave, it can be shown that there are *only* pure Nash equilibria (Bramoullé et al., 2007): It is always better to replace the mixed strategy with its mean value, but also, replacing by the mean value does not create pure Neq.

where $\mathbf{e} = \{1, \ldots, 1\}^T$ denotes the n-dimension all 1 vector. The proof is by contradiction. Suppose $|x| < \frac{1}{4n}$; it then follows that

$$\max_{i\in[n]} -|y| + (C^T x)_i \le -|y| < -1 + \frac{1}{4n},$$

where the second step follows from $|y| = 1 - |x| > 1 - \frac{1}{4n}$ by our assumption

$$\begin{aligned}\max_{i\in[n]} -|x| + (Ry)_i &> -\frac{1}{4n} + \max_{i\in[n]}(Ry)_i \ge -\frac{1}{4n} + \frac{1}{n}\sum_i (Ry)_i = -\frac{1}{4n} + \frac{1}{n}\sum_{j\in[n]} y_j \sum_{i\in[n]} R_{ij} \\ &\ge -\frac{1}{4n} + \frac{1}{n}\sum_{j\in[n]} y_j \cdot (1-n) \ge -\frac{1}{4n} + \frac{1}{n}\cdot(1-n) = -1 + \frac{3}{4n}.\end{aligned}$$

The first step follows from $|x| < \frac{1}{4n}$. We replace max with average in the second step. The fourth step comes from the fact that there exists at least one 0 entry for each column of the payoff matrix R. Thus we have

$$\max_{i\in[n]} -|x| + (Ry)_i - \max_{i\in[n]} -|y| + (C^T x)_i > \frac{1}{2n} > \epsilon,$$

which contradicts with the fact that $|y| > 1 - \frac{1}{4n}$. Therefore, we have $|x| > \frac{1}{4n}$ and $|y| > \frac{1}{4n}$. Consequently, for any $i, j \in [n]$, $x_i > 0$, we have $(-|x| + (Ry)_i) - (-|x| + (Ry)_j) = (Ry)_i - (Ry)_j > \epsilon$. Hence, we have $(R\tilde{y})_i - (R\tilde{y})_j > 4n\epsilon$ for any $i, j \in [n]$ and $\tilde{x}_i > 0$. The same holds for the column player, confirming that $(\tilde{x}, \tilde{y}) = (x/|x|, y/|y|)$ is a 4nϵ-Nash of the win-lose game (R, C). □

Define $E = -A^T - I$ and $D = -A^T = E + I$; note that $E \in \{0, 1\}^n$ and the diagonal entries E_{ii} are zero. Now we claim:

Lemma 5.3. *Let E be the adjacency matrix of the directed network of a public goods game (with divisible goods game and summation utility of the players). Then from any ϵ-pure Nash equilibrium $\mathbf{s} = (s_1, \ldots, s_n)$ of the public goods game, we can find a symmetric $3n\epsilon$-Nash of game (A, B).*

Proof. Let $\mathbf{s} = (s_1, \ldots, s_n)$ be an ϵ-approximate pure Nash equilibrium of the public goods game, then for any agent $i \in [n]$, we have $\sum_{j\in N(i)} s_j \ge 1 - \epsilon$. Otherwise, agent i would increase its effort. Moreover, we claim that $\sum_{j\in N(i)} s_j > 1 + \epsilon$ implies $s_i = 0 \pm \epsilon$. This holds because (1) if $\sum_{j\in N_i} s_i > 1$, then the best response of agent i is $b_i(s_{-i}) = 0$, it then follows $x_i = 0 \pm \epsilon$; (2) $\sum_{j\in N_i} s_i \le 1$, then the best response is $b_i(s_{-i}) = 1 - \sum_{j\in N_i} s_i$ and $s_i - b_i(s_{-i}) = \sum_{i\in N(i)} s_i - 1 > \epsilon$, which contradicts with the equilibrium condition. In summary, for all $i \in [n]$, we have

$$\begin{aligned}&(D^T\mathbf{s})_i \ge 1 - \epsilon \\ &s_i = 0 \pm \epsilon \text{ or } (D^T s)_i = 1 \pm \epsilon.\end{aligned}$$

Denote $\mathbf{s}' = \max\{\mathbf{s} - \epsilon, 0\}$, then we have

$$\begin{aligned}&(D^T\mathbf{s}')_i \ge 1 - (n+1)\epsilon \\ &\mathbf{s}'_i = 0 \text{ or } (D^T\mathbf{s}')_i = 1 \pm n\epsilon.\end{aligned}$$

Since $A = -D^T$, we have

$$\begin{aligned}&(A\mathbf{s}')_i \le -1 + (n+1)\epsilon \\ &\mathbf{s}'_i = 0 \text{ or } (A\mathbf{s}')_i = -1 \pm n\epsilon.\end{aligned}$$

Since $|\mathbf{s}'| \ge \sum_{j\in N(1)} s'_j \ge 1 - (n+1)\epsilon$, we conclude that $\mathbf{s}'/|\mathbf{s}'|$ is a symmetric $3n\epsilon$-Nash of the game (A, B) □

Combining Lemma 5.3 and Lemma 5.2, we conclude that it is PPAD-hard to find an ϵ-approximate pure Nash equilibrium of public goods game, for $\epsilon = 1/\operatorname{poly}(n)$. This concludes the proof of Theorem 5.1.

5.2. Best-shot rule

When the utility function is the best-shot rule (i.e., the utility of a node is the maximum of the provisions by its predecessors), there is a simple proof that there is no pure Nash equilibrium. First we prove that, in any pure Nash equilibrium, an agent plays either 0 or 1 (not any number between (0, 1)). Then the result follows from the example shown in Fig. 1.

For mixed Nash equilibria, we have the following theorem, shown through a simple reduction from the indivisible case.

Theorem 5.4. *When the utility function is the best-shot rule (max), it is PPAD-hard to find a mixed Nash equilibrium of the divisible public goods game.*

Proof. We set the valuation function to be $U(s) = \max\{1, s\}$. We assume in an equilibrium profile, the player prefers a mixed combination over action $0, 1$ to a strategy $s \in (0, 1)$ if they have the same utility guarantee. We then prove that, in a mixed Nash equilibrium, an agent will only play a mixed strategy over actions 0 and 1. To see this, first, in an equilibrium profile, no player chooses to play s with $s > 1$ in the support of its mixed strategy, since it can decrease it to 1, which reduces the cost and does not affect the utility. Next, if a player chooses to play $s \in (0, 1)$ in the support of its mixed strategy, then we claim it is always better to replace s with a convex combination of 0 and 1, i.e. chooses 0 with probability $1 - s$ and chooses 1 with probability s. We divide into two cases. (1) If the max production of neighbors is $s' \geq s$. Then the utility for later profile gets larger while the cost remains the same. (2) If the max production of neighbors is $s' < s$, then both the cost and utility remains the same.

Assuming all agent play mixed strategies over actions 0 and 1 in the equilibrium profile, it is not hard to modify the proof of Theorem 4.1 to show that it is PPAD-hard to find a mixed Nash equilibrium. We conclude the proof here. □

6. The bounded treewidth algorithm

When the treewidth of the underlying graph is bounded by $O\left(\frac{\log n}{\log\log n}\right)$, we develop a PTAS for computing an ϵ-Nash of the (indivisible) public goods game. We first recall the definition of *tree decomposition*.

Definition 6.1 *(Tree decomposition). A tree decomposition*[4] *of a graph $G(V, E)$ is a tree T, with nodes $X_1, \dots X_{|T|}$. Each node X_i is a subset of V, and it satisfies:*

1. *The union of X_i equals V.*
2. *For each edge $(u, v) \in E$, there exists a node X_i that contains both vertices u and v.*
3. *For any vertex $u \in V$, the set of tree nodes that contain the vertex u forms a connected sub-tree of T.*

The width *of a tree decomposition is defined as* $\max_{1\leq i\leq |T|} |X_i| - 1$ *and the* treewidth *of a graph G, denoted as* $\mathrm{twd}(G)$*, is the minimum width among all tree decompositions of the graph G.*

We will call the vertices of T *nodes*, and those of G *vertices*. The treewidth $\mathrm{twd}(G)$ will be abbreviated by w, while d is the maximum degree of G. Our main result is shown below. Comparing with the general result of Daskalakis and Papadimitriou (2006), we get rid of the exponential dependence on d. Alas, we make no assumptions on the sparsity of the graph.

Theorem 6.2. *Given an indivisible public goods game defined on a network $G(V, E)$, we can find an ϵ-Nash equilibrium in time* $\mathrm{poly}(n) \cdot \min\{2d/\epsilon, 16\log(n)/\epsilon\}^{O(w)}$*, where w is the treewidth of the graph. In particular, when the treewidth is $w = O\left(\frac{\log n}{\log\log n}\right)$, we can find an ϵ-Nash equilibrium in* $\mathrm{poly}(n) \cdot \left(\frac{1}{\epsilon}\right)^{O(w)}$ *time.*

First, a few notes about the proof. The time complexity of our algorithm depends minimally on d, while d is in the exponent of the algorithm in Daskalakis and Papadimitriou (2006). To achieve this, we need to circumvent several difficulties, explained below. Like the proof in Daskalakis and Papadimitriou (2006), we first need to show the existence of an approximate Nash equilibrium with probabilities that are multiples of a small real $\delta > 0$. Simply applying the total variation bound gives $\delta = O(\frac{\epsilon}{d})$, which is not coarse enough. In Lemma 6.3, we use a probabilistic argument showing the existence of an approximate Nash equilibrium after discretizing the strategy space. In particular, we randomly round a Nash equilibrium, for $\delta = O(\frac{\epsilon}{\log n})$ and utilize the concentration property. Another difficulty is that the algorithm (Daskalakis and Papadimitriou, 2006) works on the *primal* graph (see Daskalakis and Papadimitriou (2006) for the definition), whose treewidth can be $w \cdot d$, yielding an exponential dependence on d. Instead, our algorithm directly works on the original graph through *dynamic programming*, with no exponential dependence on d. Our dynamic programming method bares some similarities with the approach in Thomas and van Leeuwen (2015). However, we must modify significantly that algorithm, whose running time has a polynomial dependency on the size of the payoff matrix, which in our case be exponential. Finally, we note algorithm is general enough to handle any composition function that is additive.

Now, to prove the theorem, by Lemma 4.4, it suffices to show how to compute an ϵ-approximate equilibrium of a threshold game $\mathcal{G}(V, E, t)$. Again, we assume $t = 1/2$ for simplicity. We discretize the strategy space of each player to

[4] In defining treewidth, we ignore directions of the edges.

$S^\delta = [\delta]$, where $\delta = \max\{\epsilon/2d, \epsilon/16\log n\}$. We first show that there exists an ϵ-approximate *pure* Nash equilibrium in the discretized strategy space.

Lemma 6.3. *For any threshold game $\mathcal{G}(V, E, \frac{1}{2})$, there exists an ϵ-approximate equilibrium when we restrict the strategy space to be $[\delta]^n$, where $\delta = \epsilon/16\log n$.*

Proof. Suppose $\mathbf{x} = (x_1, \ldots, x_n)$ is an equilibrium profile of the threshold game $\mathcal{G}(V, E, \frac{1}{2})$. For any $i \in [n]$, suppose $x_i \in [t_i\delta, (t_i+1)\delta]$, then we randomly round x_i to $\tilde{x}_i \in \{t_i\delta, (t_i+1)\delta\}$, and we have

$$\tilde{x}_i = \begin{cases} (t_i+1)\,\delta & \text{with prob. } \frac{x_i}{\delta} - t_i \\ t_i\delta & \text{with prob. } 1 - \frac{x_i}{\delta} + t_i. \end{cases}$$

We remark that $\mathbb{E}[\tilde{x}_i] = x_i$ and $(\tilde{x}_i - t_i\delta)$ is a binary random variable that takes value in $\{0, \delta\}$, with mean $(x_i - t_i\delta)$. The rest of the proof establishes that $(\tilde{x}_1, \ldots, \tilde{x}_n)$ is an ϵ-approximate equilibrium with positive probability, therefore proving its existence.

By the multiplicative Chernoff bound, for any $i \in [n]$, if $\sum_{j\in N_i}(x_i - t_i\delta) = \sum_{j\in N_i}(\mathbb{E}[\tilde{x}_i] - t_i\delta) < \epsilon$, then we have

$$\Pr\left(\sum_{j\in N_i}\tilde{x}_j - \sum_{j\in N_i}x_j \le -\epsilon\right) = 0 \tag{5}$$

and

$$\begin{aligned}\Pr\left(\sum_{j\in N_i}\tilde{x}_j - \sum_{j\in N_i}x_j \ge \epsilon\right) &= \Pr\left(\sum_{j\in N_i}\left(\tilde{x}_j - t_j\delta\right) - \sum_{j\in N_i}\left(\mathbb{E}[\tilde{x}_j] - t_j\delta\right) \ge \epsilon\right) \\ &\le \exp\left(-\frac{\epsilon}{3\delta}\right) \le n^{-2}.\end{aligned} \tag{6}$$

If $\sum_{j\in N_i}(x_i - t_i\delta) = \sum_{j\in N_i}(\mathbb{E}[\tilde{x}_i] - t_i\delta) \in [\epsilon, 2]$, then we have

$$\begin{aligned}\Pr\left(\left|\sum_{j\in N_i}\tilde{x}_j - \sum_{j\in N_i}x_j\right| \ge \epsilon\right) &= \Pr\left(\left|\sum_{j\in N_i}\left(\tilde{x}_j - t_j\delta\right) - \sum_{j\in N_i}\left(\mathbb{E}[\tilde{x}_j] - t_j\delta\right)\right| \ge \epsilon\right) \\ &\le 2\exp\left(-\frac{\epsilon^2}{2\delta\left(\sum_{j\in N_i}\mathbb{E}[\tilde{x}_j] - \sum_{j\in N_i}t_i\delta\right)}\right) \\ &\le 2\exp\left(-\frac{\epsilon}{4\delta}\right) \le n^{-2}.\end{aligned} \tag{7}$$

If $\sum_{j\in N_i}(x_i - t_i\delta) = \sum_{j\in N_i}(\mathbb{E}[\tilde{x}_i] - t_i\delta) > 2$, it then follows that $x_i = 0$ and we have

$$\Pr\left(\sum_{j\in N_i}\tilde{x}_j < 1\right) \le \Pr\left(\sum_{j\in N_i}\left(\tilde{x}_j - t_j\delta\right) < 1\right) \le \exp\left(-\frac{1}{4\delta}\right) \le n^{-2}. \tag{8}$$

Combining Eq. (5) (6) (7) (8) and using an union bound, we conclude that $(\tilde{x}_1, \ldots, \tilde{x}_n)$ satisfies equilibrium condition with probability at least $(1 - 1/n)$, completing the proof. □

We next provide an algorithm that finds an ϵ-approximate equilibrium based on dynamic programming. A *nice tree decomposition* is a tree decomposition T that only contains the following four types of nodes (see Fig. 5 for an illustration).

1. Leaf node.
2. Forget node. Such a node i has only one child i', and $X_{i'} = X_i \backslash \{v\}$ for some vertex $v \in V$.
3. Introduce node. Such a node i has only one child node i', and $X_{i'} = X_i \cup \{v\}$ for some vertex $v \in V$.
4. Join node. Such a node i has exactly two children nodes i_1, i_2, and $X_i = X_{i_1} = X_{i_2}$.

Any tree decomposition can be converted into a nice tree decomposition, of size at most $w \cdot |V|$, in linear time without enlarging the width (Bodlaender and Koster, 2008).

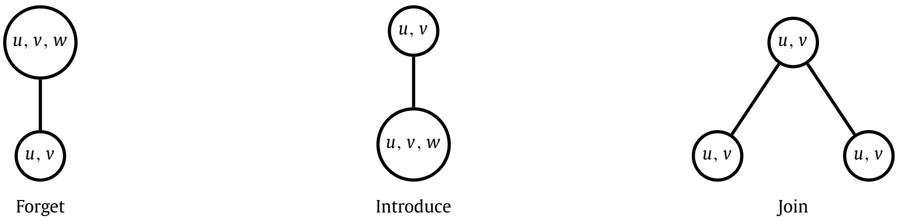

Fig. 5. An illustration for three types of nodes of a nice tree decomposition.

Now, we have

Lemma 6.4. *Given a threshold game $\mathcal{G}(V, E, \frac{1}{2})$ with the strategy space $[\delta]^n$, and a nice tree decomposition T of the graph $G(V, E)$, we can compute an ϵ-approximate equilibrium in $\delta^{-O(w)}$ time.*

Proof. Given a nice tree decomposition T, we compute an ϵ-approximate equilibrium via a bottom-up approach. For any node $X_i \in T$, we use V_i to denote all vertices contained in X_i and its sub-tree. We compute a table $T_i : [\delta]^{|X_i|} \times [\delta]^{|X_i|} \to \{0, 1\}$ for each node X_i, and we note that the size of the table is bounded by $\delta^{-O(w)}$. Intuitively, the first set of $[\delta]^{|X_i|}$ entries enumerate the strategy of vertex in X_i and the second set of $[\delta]^{|X_i|}$ entries enumerate the total supply from their neighbors in $V_i \backslash X_i$. Concretely, a table entry $T_i(s_1, \dots, s_{|X_i|}, c_1, \dots, c_{|X_i|}) = 1$, iff there exists a strategy profile $(p_1, \dots, p_{|V_i|})$ of vertex set V_i, such that

(i) for any vertex $v \in V_i \backslash X_i$, the vertex v satisfies the equilibrium condition,
(ii) for any vertex $v \in X_i$, $p_v = s_v$ and $\sum_{j \in N_v \cap (V_i \backslash X_i)} p_v = c_v$, i.e., the summation of vertex v's neighbor in $V_i \backslash X_i$ is c_v.

We note that vertices in X_i do not need to satisfy the equilibrium condition, and we only record the summation of their neighbors in $V_i \backslash X_i$.

Next, we show how to do update the table in a bottom-up manner.

(1) Leaf. For any leaf $X_i \in T$ and $s, c \in [\delta]^{|X_i|}$, we set $T_i(s, c) = 1$ if and only if $c = (0, \dots, 0)$.
(2) Forget. Suppose $X_{i'} = X_i \cup \{v\}$ is the parent node, $s, c \in [\delta]^{|X_i|}$ and $s_v, c_v \in [\delta]$, we set $T_{i'}(s, s_v, c, c_v) = 1$ if $T_i(s, c) = 1$ *and* $c_v = 0$; we set $T_{i'}(s, s_v, c, c_v) = 0$ otherwise.
(3) Introduce. Suppose $X_{i'} = X_i \backslash \{v\}$ is the parent node and $s, c \in [\delta]^{|X_{i'}|}$, we set $T_{i'}(s, c) = 1$ iff there exists $(s, s_v, c, c_v) \in [\delta]^{2|X_i|}$, such that $T_i(s, s_v, c, c_v) = 1$ *and* the vertex v satisfies the equilibrium condition, i.e., (i) if $c_v + \sum_{j \in N_v \cap X_i} s_j > \frac{1}{2}$, then $s_v = 0 \pm \epsilon$; (ii) $c_v + \sum_{j \in N_v \cap X_i} s_j < \frac{1}{2}$, then $s_v = 1 \pm \epsilon$.
(4) Join. Suppose node X_i has two children, X_{i_1}, X_{i_2}, and $X_i = X_{i_1} = X_{i_2}$. Then for any $s, c \in [\delta]^{|X_i|}$, we set $T_i(s, c) = 1$ iff there exists $c_1, c_2 \in [\delta]^{|X_i|}$, such that $T_{i_1}(s, c_1) = 1$, $T_{i_2}(s, c_2) = 1$, *and* for any $j \in X_i$, $c[j] = \min\{c_1[j] + c_2[j], 1\}$.

After we reach the root r and complete the table T_r, we verify equilibrium conditions for all vertices $v \in X_r$. To be more specific, if there exists a configuration $(s, c) \in [\delta]^{|X_i|} \times [\delta]^{|X_i|}$, such that $T_r(s, c) = 1$ *and* all vertices v in V_r satisfy the equilibrium, i.e., (i) if $c_v + \sum_{j \in N_v \cap V_r} s_j > \frac{1}{2} + \epsilon$ then $s_v = 0 \pm \epsilon$; (ii) if $c_v + \sum_{j \in N_v \cap V_r} s_j < \frac{1}{2} - \epsilon$ then $s_v = 1 \pm \epsilon$; we then confirm that there exists an ϵ-approximate equilibrium. We can find one by either fixing the strategy of all vertices $v \in V_r$ to be s_v, and recursively computing equilibrium profiles in the sub-tree; or we can associate a satisfiable assignment (if there exists one) for each entry during the dynamic programming process. We output that there is no ϵ-approximate equilibrium profile otherwise. $\square$

Combining Lemma 6.4 and Lemma 6.3, we conclude the proof of Theorem 6.2.

We can show a similar result for divisible public good games with the summation rule. Again, when the treewidth of the underlying graph is bounded by $O(\log n / \log\log n)$, there is a PTAS for finding an ϵ-approximate pure Nash equilibrium of the public goods game:

Theorem 6.5. *Given a divisible public goods game with summation utility defined on a directed network, we can find an ϵ-approximate pure Nash equilibrium in time* $\mathrm{poly}(n) \cdot \min\{2d/\epsilon, 16\log(n)/\epsilon\}^{O(w)}$ *time. In particular, when the treewidth is $O(\log n / \log\log n)$, we can find an ϵ-approximate pure Nash equilibrium in* $\mathrm{poly}(n) \cdot (\frac{1}{\epsilon})^{O(w)}$ *time.*

The proof is similar to Theorem 6.2, and is omitted.

7. Pure Nash equilibria: the general undirected case

We started our treatment of public goods games by pointing out that the well-behaved problem with the max composition function in undirected graphs becomes NP-complete in the directed case. But then we went on to consider arbitrary composition functions, and proved the dichotomy result for directed graphs (Theorem 3.4). It is natural to ask, how hard is the classification problem of general composition functions in undirected graphs? We know it is easy when the composition function is the max, but are there hard functions?

We next show that indeed there are:

Theorem 7.1. *It is NP-hard to find a pure Nash equilibrium in a public goods game on an* undirected *network with a general composition function.*

We have so far dealt with homogeneous public goods games in which all agents have the same decision function. In this proof we shall also consider heterogeneous public goods games, where agents could have different decision functions.

Proof. We first show that the heterogeneous problem can be reduced to the homogeneous problem, and then prove the NP-hardness of the heterogeneous problem.

Given a heterogeneous game defined on graph $G=(V,E)$, $|V|=n$ in which $f_i:\mathbb{N}\to\{0,1\}$ denotes the decision function of agent i ($i\in[n]$). We construct a homogeneous instance defined on graph $G'=(V',E')$. We first specify the common decision function:

$$f(t)=\begin{cases}1 & t\in\{0,1\}\\ 0 & t\in\{2,\dots,n-1\}\\ f_i(j) & t=ni+j, \qquad i\in[n],\ j\in\{0,\dots,n-1\}\\ 0 & t\geq n(n+1)\end{cases}$$

The node set is $V'=V_1'\cup V_2'$ with $|V_1'|=n$ and $|V_2'|=n^2(n+1)/2$. Intuitively, we use nodes in V_1' to simulate the agents of V, while nodes in V_2' are auxiliary agents who help differentiate the decision functions of the nodes in V_1'. Concretely, let $V_1'=\{v_1,\dots,v_n\}$, and there is an edge between v_i and v_j if and only if $(i,j)\in E$. At the same time, the node v_i is connected to ni different nodes in V_2'; there are no edge connections between nodes in V_2'. We observe that in a pure Nash equilibrium, agents in V_2' always choose to produce, since they only have one neighbor, and $f(0)=f(1)=1$. As for each v_i, it has ni neighbor nodes in V_2' who produce the good, and therefore, its decision function reduce to $f_i(t-ni)$. The reduction is complete.

We next reduce from the Exact-3-Cover problem to the heterogeneous problem. In the Exact-3-Cover problem we are given a finite ground set $X=\{x_1,\dots,x_{3q}\}$ and a collection $\mathcal{C}=\{C_1,\dots,C_m\}$ of 3-element subsets of X; the goal is to find a collection of subsets $C_1,\dots,C_q$ that (exactly) cover the ground set. It is known that the Exact-3-Cover problem is NP-hard (Garey and Johnson, 1979).

Let the node set $V=V_1\cup V_2\cup V_3$ with $V_1=\{v_{1,1},\dots,v_{1,3q}\}$, $V_2=\{v_{2,1},\dots,v_{2,m}\}$ and $V_3=\{v_{3,1},v_{3,2}\}$. The nodes $v_{3,1}$ and $v_{3,2}$ are connected to all nodes in V_1 and V_2, and they also connect to each other. The connection between V_1 and V_2 is determined by the Exact-3-Cover instance, i.e., node $v_{1,i}$ is connected to $v_{2,j}$ iff $x_i\in C_j$. Nodes in V_2 are connected to each other. We next specify the decision function of each node.

- Nodes in V_1 have the same decision function,

$$f_1(t)=\begin{cases}1 & t=0\\ 0 & t\geq 1\end{cases}$$

i.e., the agent chooses to purchase the good iff none of its neighbors buy it.
- Nodes in V_2 have the same decision function.

$$f_2(t)=\begin{cases}1 & t\leq q-1\\ 0 & t\geq q\end{cases}$$

i.e., the agent chooses to purchase the good iff no more than $(q-1)$ of its neighbors buy the good.
- Nodes in V_3 have the following decision function.

$$f_{3,1}(t)=\begin{cases}0 & t\in\{0,\dots,q\}\\ 0 & t=q+2i-1 \quad i\in\mathbb{N}_+\\ 1 & t=q+2i \quad i\in\mathbb{N}_+\end{cases} \quad\text{and}\quad f_{3,2}(t)=\begin{cases}0 & t\in\{0,\dots,q\}\\ 1 & t=q+2i-1 \quad i\in\mathbb{N}_+\\ 0 & t=q+2i \quad i\in\mathbb{N}_+.\end{cases}$$

To show correctness, first suppose there is a solution to the original Exact-3-Cover instance, say the solution is $C_{i_1},\dots,C_{i_q}$. Then it is easy to verify that making agent v_{2,i_j} ($j\in[q]$) purchase the good forms a pure NE.

Next, suppose there is no exact cover, then we prove there is no pure Nash equilibrium. First, we claim that in any pure Nash equilibrium, agents $v_{3,1}, v_{3,2}$ do not produce the good and there are at most q agents produce the good in $V_1 \cup V_2$. On the contrary, suppose there are at least $s > q$ agents in V_1, V_2 buy the good.

1. Suppose $s - q$ is odd. If $v_{3,1}$ purchases the good, then $v_{3,2}$ would free-ride, and this contradicts the equilibrium condition of $v_{3,1}$. Similarly, if $v_{3,1}$ chooses to free-ride, then $v_{3,2}$ should buy the good and this contradicts the equilibrium condition of $v_{3,1}$
2. Suppose $s - q$ is even. If $v_{3,1}$ produces the good, then $v_{3,2}$ also buys the good, which contradicts with the fact that $v_{3,1}$ buys. On the other had, if $v_{3,1}$ does not produce the good, then so does $v_{3,2}$, the equilibrium condition of $v_{3,1}$ is violated.

We conclude that no more than q agents in $V_1 \cup V_2$ produce the good. To see both $v_{3,1}$ and $v_{3,2}$ should choose to free-ride, suppose there less than $(q-1)$ agents in $V_1 \cup V_2$ produce the good, then they clearly choose to free-ride. If there are exactly q agents that produce the good and at least one of $v_{3,1}, v_{3,2}$ also produce the good, then according to the decision function both of them should produce and this, again, contradicts the equilibrium condition of $v_{3,1}$.

Next, suppose there are a_1 agents in V_1 purchase the good and a_2 agents in V_2 purchases, and $a_1 + a_2 \le q$. The equilibrium condition of V_1 implies that each agent of V_1 and its neighbor must have at least one piece of the good. Since each agent in V_2 is connected to 3 agents in V_1, this indicates $a_1 + 3a_2 \ge 3q$. This implies $a_1 = 0$, $a_2 = q$, and each agents in V_1 is connected to one of these q agents, hence forming a solution of the Exact-3-Cover problem. This concludes the proof. □

Hence in the undirected case of common goods problem there are easy utilities/decision functions — e.g. the max utility, but also the utilities shown in Section 4 to be easy even for directed networks. Also, we now know that there are hard ones — the kind of decision function that is created in the reduction from the heterogeneous to the homogeneous case. Finding the exact characterization of easy cases seems a formidable problem which we leave open here.

We conclude with a complexity result in a different direction: Consider the *second to max* utility in which all nodes need two units of the good. This is a steep function, and hence it has the trivial, all-zero Nash equilibrium. But notice that this equilibrium is not Pareto optimal in general — in contrast, the max utility guarantees Pareto optimality of the equilibrium.

Theorem 7.2. *It is NP-hard to find a Pareto optimum pure Nash equilibrium in a public goods game in an* undirected *network with the second-to-max utility.*

Proof. An *exact doubly dominating set* of a graph $G = (V, E)$ is defined as a subset of node $V' \subseteq V$ such that every node has in its neighborhood (which includes itself) exactly two nodes in V'. It is NP-complete to determine whether an undirected graph has such a doubly dominating set (Chellali et al., 2005). Meanwhile, it is easy to verify that an exact doubly dominating set (if there exists one) is a Pareto optimum pure Nash equilibrium of the public goods game. Hence finding a Pareto optimum pure NE is NP-hard with second-to-max utility. □

Finally, we propose a (challenging) open question on completely characterizing the computational complexity of pure NE equilibria under general utility function in undirected network

Open Problem 7.3. *Can one provide a P/NP dichotomy characterization on the computational complexity of pure NE for undirected graph with general utility function, in the same spirits of Theorem 3.4?*

8. Discussion

We have explored the complexity of equilibria in public goods games played on directed graphs. One striking conclusion is the ubiquity of PPAD-completeness in this domain. For a number of quite different reasons, very different variants of the problem are shown to share the same fate — and a rather sophisticated fate at that. This is in stark contrast with the corresponding public goods games in undirected graphs, where the consensus is that equilibria are rather boring (but see the discussion below of some intriguing problems in undirected networks raised by this work). Note that graphical games are already intractable when they are symmetric — but, of course, this is because the local normal form games in each neighborhood can simulate any asymmetry.

Does the equilibrium problem for public good games on directed networks come up in the real world? It can be argued that some of the directed graphs we evoked for motivation in the introduction (towns that are downwind or upriver from one another, or the relationship "B is on A's way to work") are *transitive*, and it is not hard to see that public good games on such directed graphs have trivial equilibrium problems. On the other hand, many social networks with sharing features are indeed asymmetric and non-transitive, and so are infection networks in much of epidemic modeling. Another example is peer-to-peer content sharing networks such as BitTorrent, if one assumes that all nodes have already publicly committed to being either contributors or free riders; the decision is whether a node will download the content from the source (produce the good), or will obtain it from a contributing neighbor. In fact, this latter scenario gives rise to a two-stage game, which may be interested in its own right.

For the indivisible case, we found that there are three special cases of utilities that admit polynomial time solution: Flat utilities, steep utilities, plus a third polynomial case, alternating utilities, which is quite unexpected and intriguing (its algorithm relies on the solution of a system of equations in $\mathbb{F}_2$). We show that these are the only tractable cases. But there is an interesting *variant* of this problem which is quite mysterious: Suppose that we allow the utility function to be such that *certain steps of X have height exactly* p, and therefore nodes can be indifferent between buying the good and free-riding. We suspect that this variant is subject to the same dichotomy, but it seems much harder to prove. Consider for example the function $X(1) = 1 > p, X(k) = 1 + p$ for all $k > 1$. Then it is easy to see that, in this case, odd cycles *do* have an equilibrium, with all players producing the good: the p step makes them indifferent to doing so. This deprives us of a valuable gadget. It turns out that there is a 7-node, 21-edge gadget with no equilibrium for this case: the node set is $\{1, \ldots, 7\}$ and the edges go from i to $i + 1, i + 2, i + 4 \bmod 7$. But this does not immediately give us an NP-hardness proof, nor does it generalize to other composition functions with p steps.

The divisible good games under the summation utility are something of a mystery when it comes to *mixed* equilibria. As with other games with uncountable strategy spaces, it is not easy to characterize mixed Nash equilibria in a tangible, useful way. We believe that positive results may be possible here: Could it be that there are always mixed Nash equilibria with small support, and in fact they are easy to find? There are reasons for hope for a truly positive result in this case.

The intractability of simple Nash equilibrium problems in common goods games in directed networks is an indication that asymmetry in social systems — a notion intuitively coterminous with unfairness — may consistently lead to instability. Can the intractability proofs help identify the features of the directed networks, and of the agents and their utilities, which are at the root of such instability? This could lead to principles for better design of social networks, or beneficial interventions therein.

Finally, we believe that the open problems pointed out in the Section 7, namely complexity dichotomy results for finding Nash equilibria and Pareto-optimal solutions in common good games in undirected graphs under general utility functions, are very interesting and quite challenging.

Declaration of competing interest

None.

Data availability

No data was used for the research described in the article.

Acknowledgment

The authors would like to thank Xi Chen and anonymous EC/GEB reviewers for their very helpful feedback. This research was supported by National Science Foundation grants CCF-1763970 AF and CCF-1910700 AF, and a grant from Softbank.

Appendix A. Omitted proof from Section 4

We aim to prove

Lemma 4.4. *There is a polynomial time reduction between the threshold game and the public good game. Specifically, (1) given any threshold game $\mathcal{G}(V, E, t)$ with $0 < t < 1$, we can construct a public good game and map any ϵ-Nash of the public goods game to an 8ϵ-approximate equilibrium of threshold game $\mathcal{G}(V, E, t)$, for $\epsilon < \min\{0.1, \frac{t}{8}, \frac{1-t}{8}\}$; (2) given any public good game with $U = 1, 0 < p < 1$, we can construct a threshold game $\mathcal{G}(V, E, t)$ and map any ϵ-approximate equilibrium of threshold game to an $c_p\epsilon$-Nash of public goods game, where $c_p = -4p \log p$ is a constant depending only on p.*

Proof. We first reduce the threshold game to the public good game. Given an instance of the threshold game $\mathcal{G}(V, E, t)$, we construct a public good game as follows. We keep the network $G(V, E)$ unchanged and set the value of the good to be $U = 1$ and the price to be $p = e^{-t} \in (0, 1)$. For any ϵ-Nash $\mathbf{s} = (s_1, \cdots, s_n)$ of the public good game, we construct an ϵ-approximate equilibrium $\mathbf{x} = (x_1, \cdots, x_n)$ of $\mathcal{G}(V, E, t)$ as

$$x_i = \min\{-\log(1 - s_j), 1\} \in [0, 1], \forall i.$$

Consider any agent i in the public good game, its utility is specified as

$$U(s_i, s_{-i}) = \begin{cases} 1 - p & s_i = 1 \\ 1 - \prod_{j \in N_i}(1 - s_j) & s_i = 0, \end{cases}$$

thus we have

$$U(1, s_{-i}) - U(0, s_{-i}) = \prod_{j \in N_i} (1 - s_j) - p.$$

We divide into three cases.

Case 1. $\prod_{j \in N_i} (1 - s_j) - p > \epsilon$. This implies $s_i = 1$ and $x_i = \min\{-\log(1 - s_i), 1\} = 1$. Now we have

$$\prod_{j \in N_i} (1 - s_j) - p > \epsilon \Rightarrow \prod_{j \in N_i} (1 - s_j) > p + \epsilon \Rightarrow \log \prod_{j \in N_i} (1 - s_j) > \log(p + \epsilon)$$
$$\Rightarrow \sum_{j \in N_i} -\log(1 - s_j) < -\log(p + \epsilon) < -\log p = t.$$

Since $\max_{j \in N_i}\{-\log(1 - s_j)\} \le \sum_{j \in N_i} -\log(1 - s_j) < t < 1$, we have $\sum_{j \in N_i} x_j = \sum_{j \in N_i} -\log(1 - s_j) < t$, this satisfies the equilibrium condition of the threshold game.

Case 2. $\prod_{j \in N_i} (1 - s_j) - p < -\epsilon$. This implies $s_i = 0$ and $x_i = 0$. Similar to the first case, we have

$$\prod_{j \in N_i} (1 - s_j) - p < -\epsilon \Rightarrow \prod_{j \in N_i} (1 - s_j) < p - \epsilon \Rightarrow \log \prod_{j \in N_i} (1 - s_j) < \log(p - \epsilon)$$
$$\Rightarrow \sum_{j \in N_i} -\log(1 - s_j) > -\log(p - \epsilon) > -\log p = t.$$

Since $t < 1$ and $-\log(1 - s_j) > 0$ for $\forall j \in N_i$, we conclude that $\sum_{j \in N_i} x_j = \sum_{j \in N_i} \min\{-\log(1 - s_j), 1\} > t$, this satisfies the equilibrium condition of the threshold game.

Case 3. $\prod_{j \in N_i} (1 - s_j) - p \in [-\epsilon, \epsilon]$. This time s_i can be any number in $[0, 1]$, so does x_i. We need to verify that $\sum_{j \in N_i} x_i \in [t - 8\epsilon, t + 8\epsilon]$. We have

$$\prod_{j \in N_i} (1 - s_j) - p \in [-\epsilon, \epsilon] \Rightarrow \prod_{j \in N_i} (1 - s_j) \in [p - \epsilon, p + \epsilon]$$
$$\Rightarrow \sum_{j \in N_i} -\log(1 - s_j) \in [-\log(p + \epsilon), -\log(p - \epsilon)].$$

When $\epsilon < \min\{0.1, \frac{t}{8}, \frac{1-t}{8}\}$, we can prove that $[-\log(p + \epsilon), -\log(p - \epsilon)] \in [t - 8\epsilon, t + 8\epsilon]$. We defer the calculation to Lemma A.1. Now we have $\sum_{j \in N_j} x_j = \sum_{j \in N_i} \min\{-\log(1 - s_j), 1\} = \sum_{j \in N_i} -\log(1 - s_j) \in [t - 8\epsilon, t + 8\epsilon]$, which satisfies the equilibrium condition.

We next show there is a polynomial time reduction from public good games to threshold games. Similar as above, given an instance of public good game defined on $G(V, E)$, $U = 1$, $0 < p < 1$, we construct a threshold game on the same network (V, E), with $t = \frac{1}{2}$. Given an ϵ-approximate equilibrium $\mathbf{x} = (x_1, \dots, x_n)$ of the threshold game, we recover an $-4p\log(p)\epsilon$-Nash $\mathbf{s} = (s_1, \dots, s_n)$ of the public good game as follows,

$$s_i = \begin{cases} 1 - p^{2x_i} & x_i \le \frac{1}{2} + \epsilon \\ 1 & \text{otherwise.} \end{cases}$$

For any agent i, if $\sum_{j \in N_i} x_j > \frac{1}{2} + \epsilon$, then $x_i = 0$ and $s_i = 0$ by definition. It then follows that $U(1, s_{-i}) - U(0, s_{-i}) = \prod_{j \in N_i} (1 - s_j) - p \le p^{\sum_{j \in N_i} 2x_j} - p \le p^{1+2\epsilon} - p < 0$. Hence, it satisfies the equilibrium condition. If $\sum_{j \in N_i} x_j < \frac{1}{2} - \epsilon$, then $x_i = 1$ and $s_i = 1$. Meanwhile, we have $U(1, s_{-i}) - U(0, s_{-i}) = p^{\sum_{j \in N_i} 2x_j} - p > p^{1-2\epsilon} - p > 0$. Finally, if $\sum_{j \in N_i} x_j \in [\frac{1}{2} - \epsilon, \frac{1}{2} + \epsilon]$, we have $U(1, s_{-i}) - U(0, s_{-i}) = \prod_{j \in N_i} (1 - s_j) - p = p^{\sum_{j \in N_i} 2x_j} - p \in [p(p^{2\epsilon} - 1), p(p^{-2\epsilon} - 1)] \in [2p\log(p)\epsilon, -4p\log(p)\epsilon]$. Here we use the facts that $\lambda \le e^{\lambda} - 1 \le 2\lambda$ for $\lambda < 1$. Therefore, we have verified that $\mathbf{s} = (s_1, \dots, s_n)$ is an $-4p\log(p)\epsilon$-Nash of the public good game. Hence, setting $c_p = -4p\log p$, we conclude the proof. □

Lemma A.1. *For any* $0 < t < 1$ *and* $0 < \epsilon < \min\{0.1, \frac{t}{8}, \frac{1-t}{8}\}$, *we have*

1. $-\log(e^{-t} - \epsilon) < t + 8\epsilon$,
2. $-\log(e^{-t} + \epsilon) > t - 8\epsilon$.

Proof. We have

$$-\log(e^{-t} - \epsilon) < t + 8\epsilon \Leftrightarrow \log(e^{-t} - \epsilon) > -(t + 8\epsilon) \Leftrightarrow e^{-t} - \epsilon > e^{-(t+8\epsilon)}$$
$$\Leftrightarrow e^{-t}\left(1 - e^{-8\epsilon}\right) > \epsilon \Leftarrow 1 - e^{-8\epsilon} > 3\epsilon.$$

By simple calculations, we can show $1 - e^{-8\epsilon} > 3\epsilon$ for $\epsilon < 0.1$. On the other side, we have

$$-\log(e^{-t} + \epsilon) > t - 8\epsilon \Leftrightarrow \log(e^{-t} + \epsilon) < -(t - 8\epsilon) \Leftrightarrow e^{-t} + \epsilon < e^{-t} \cdot e^{8\epsilon} \Leftrightarrow e^{-t}(e^{8\epsilon} - 1) \geq \epsilon$$

This follows from the fact that $e^{-t}(e^{8\epsilon} - 1) \geq \frac{1}{3}(e^{8\epsilon} - 1) \geq \frac{8}{3}\epsilon > \epsilon$. $\square$

References

Abbott, T., Kane, D., Valiant, P., 2005. On the complexity of two-player win-lose games. In: 46th Annual IEEE Symposium on Foundations of Computer Science (FOCS'05). IEEE, pp. 113–122.
Acemoglu, D., García-Jimeno, C., Robinson, J.A., 2015. State capacity and economic development: a network approach. Am. Econ. Rev. 105 (8), 2364–2409.
Allouch, N., 2015. On the private provision of public goods on networks. J. Econ. Theory 157, 527–552.
Bayer, P., Kozics, G., Szőke, N.G., 2021. Best-response dynamics in directed network games. ArXiv preprint. arXiv:2101.03863.
Bergstrom, T., Blume, L., Varian, H., 1986. On the private provision of public goods. J. Public Econ. 29 (1), 25–49.
Bervoets, S., Faure, M., 2019. Stability in games with continua of equilibria. J. Econ. Theory 179, 131–162.
Bodlaender, H.L., Koster, A.M., 2008. Combinatorial optimization on graphs of bounded treewidth. Comput. J. 51 (3), 255–269.
Boncinelli, L., Pin, P., 2012. Stochastic stability in best shot network games. Games Econ. Behav. 75 (2), 538–554.
Bramoullé, Y., Kranton, R., 2015. Games played on networks. In: The Oxford Handbook of the Economics of Networks.
Bramoullé, Y., Kranton, R., et al., 2007. Public goods in networks. J. Econ. Theory 135 (1), 478–494.
Bramoullé, Y., Kranton, R., D'amours, M., 2014. Strategic interaction and networks. Am. Econ. Rev. 104 (3), 898–930.
Chaudhury, B.R., Garg, J., McGlaughlin, P., Mehta, R., 2020. Dividing bads is harder than dividing goods: on the complexity of fair and efficient division of chores. ArXiv preprint. arXiv:2008.00285.
Chaudhury, B.R., Garg, J., McGlaughlin, P., Mehta, R., 2021. Competitive allocation of a mixed manna. In: Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA). SIAM, pp. 1405–1424.
Chaudhury, B.R., Garg, J., McGlaughlin, P., Mehta, R., 2022. Competitive equilibrium with chores: combinatorial algorithm and hardness. In: Proceedings of the 23rd ACM Conference on Electronic Commerce.
Chellali, M., Khelladi, A., Maffray, F., 2005. Exact double domination in graphs. Discuss. Math., Graph Theory 25 (3), 291–302.
Chen, T., Chen, X., Peng, B., Yannakakis, M., 2022. Computational hardness of the Hylland-Zeckhauser scheme. In: Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA). SIAM, pp. 2253–2268.
Chen, W., Lu, P., Sun, X., Tang, B., Wang, Y., Zhu, Z.A., 2011. Optimal pricing in social networks with incomplete information. In: International Workshop on Internet and Network Economics. Springer, pp. 49–60.
Chen, X., Peng, B., 2023. Complexity of equilibria in first-price auctions under general tie-breaking rules. In: Proceedings of the 55th Annual ACM SIGACT Symposium on Theory of Computing.
Chen, X., Teng, S.-H., 2009. Spending is not easier than trading: on the computational equivalence of Fisher and arrow-Debreu equilibria. In: International Symposium on Algorithms and Computation. Springer, pp. 647–656.
Chen, X., Teng, S.-H., Valiant, P., 2007. The approximation complexity of win-lose games. In: SODA, vol. 7, pp. 159–168.
Chen, X., Dai, D., Du, Y., Teng, S.-H., 2009a. Settling the complexity of arrow-Debreu equilibria in markets with additively separable utilities. In: 2009 50th Annual IEEE Symposium on Foundations of Computer Science. IEEE, pp. 273–282.
Chen, X., Deng, X., Teng, S.-H., 2009b. Settling the complexity of computing two-player Nash equilibria. J. ACM 56 (3), 1–57.
Chen, X., Paparas, D., Yannakakis, M., 2013. The complexity of non-monotone markets. In: Proceedings of the Forty-Fifth Annual ACM Symposium on Theory of Computing, pp. 181–190.
Chen, X., Durfee, D., Orfanou, A., 2015. On the complexity of Nash equilibria in anonymous games. In: Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing, pp. 381–390.
Chen, X., Kroer, C., Kumar, R., 2021. The complexity of pacing for second-price auctions. In: Proceedings of the 22nd ACM Conference on Electronic Commerce.
Dall Asta, L., Pin, P., Ramezanpour, A., 2011. Optimal equilibria of the best shot game. J. Public Econ. Theory 13 (6), 885–901.
Daskalakis, C., Papadimitriou, C.H., 2006. Computing pure Nash equilibria in graphical games via Markov random fields. In: Proceedings of the 7th ACM Conference on Electronic Commerce, pp. 91–99.
Daskalakis, C., Papadimitriou, C.H., 2015. Approximate Nash equilibria in anonymous games. J. Econ. Theory 156, 207–245.
Daskalakis, C., Goldberg, P.W., Papadimitriou, C.H., 2009a. The complexity of computing a Nash equilibrium. SIAM J. Comput. 39 (1), 195–259.
Daskalakis, C., Schoenebeck, G., Valiant, G., Valiant, P., 2009b. On the complexity of Nash equilibria of action-graph games. In: Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms. SIAM, pp. 710–719.
Daskalakis, C., Skoulakis, S., Zampetakis, M., 2021. The complexity of constrained min-max optimization. In: Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing.
Deligkas, A., Fearnley, J., Hollender, A., Melissourgos, T., 2022. Pure-circuit: strong inapproximability for PPAD. In: 53rd Annual IEEE Symposium on Foundations of Computer Science (FOCS'22). IEEE.
Elliott, M., Golub, B., 2019. A network approach to public goods. J. Polit. Econ. 127 (2), 730–776.
Feldman, M., Kempe, D., Lucier, B., Paes Leme, R., 2013. Pricing public goods for private sale. In: Proceedings of the Fourteenth ACM Conference on Electronic Commerce, pp. 417–434.
Filos-Ratsikas, A., Giannakopoulos, Y., Hollender, A., Lazos, P., Poças, D., 2021. On the complexity of equilibrium computation in first-price auctions. In: Proceedings of the 22nd ACM Conference on Electronic Commerce.
Galeotti, A., Goyal, S., Jackson, M.O., Vega-Redondo, F., Yariv, L., 2010. Network games. Rev. Econ. Stud. 77 (1), 218–244.
Garey, M.R., Johnson, D.S., 1979. Computers and Intractability, vol. 174. Freeman, San Francisco.
Gottlob, G., Greco, G., Scarcello, F., 2005. Pure Nash equilibria: hard and easy games. J. Artif. Intell. Res. 24, 357–406.
Jackson, M.O., Zenou, Y., 2015. Games on Networks. Handbook of Game Theory with Economic Applications, vol. 4. Elsevier, pp. 95–163.
Kearns, M., Littman, M.L., Singh, S., 2001. Graphical models for game theory. In: The 17th Conference in Uncertainty in Artificial Intelligence, UAI.
Kempe, D., Yu, S., Vorobeychik, Y., 2020. Inducing equilibria in networked public goods games through network structure modification. In: Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems, pp. 611–619.
Lemke, C.E., Howson Jr, J.T., 1964. Equilibrium points of bimatrix games. J. Soc. Ind. Appl. Math. 12 (2), 413–423.
López-Pintado, D., 2013. Public goods in directed networks. Econ. Lett. 121 (2), 160–162.
Mehta, R., 2014. Constant rank bimatrix games are ppad-hard. In: Proceedings of the Forty-Sixth Annual ACM Symposium on Theory of Computing, pp. 545–554.
Othman, A., Papadimitriou, C., Rubinstein, A., 2016. The complexity of fairness through equilibrium. ACM Trans. Econ. Comput. 4 (4), 1–19.

Papadimitriou, C.H., 1994. On the complexity of the parity argument and other inefficient proofs of existence. J. Comput. Syst. Sci. 48 (3), 498–532.
Robertson, N., Seymour, P.D., 1986. Graph minors. II. Algorithmic aspects of tree-width. J. Algorithms 7 (3), 309–322.
Rubinstein, A., 2016. Settling the complexity of computing approximate two-player Nash equilibria. In: 2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS). IEEE, pp. 258–265.
Rubinstein, A., 2018. Inapproximability of Nash equilibrium. SIAM J. Comput. 47 (3), 917–959.
Samuelson, P.A., 1954. The pure theory of public expenditure. Rev. Econ. Stat., 387–389.
Schuldenzucker, S., Seuken, S., Battiston, S., 2017. Finding clearing payments in financial networks with credit default swaps is ppad-complete. In: 8th Innovations in Theoretical Computer Science Conference (ITCS 2017).
Shin, K., Lee, E., Eswaran, D., Procaccia, A.D., 2017. Why you should charge your friends for borrowing your stuff. In: Proceedings of the 26th International Joint Conference on Artificial Intelligence, pp. 395–401.
Stiglitz, J.E., 1999. Knowledge as a global public good. Glob. Public Goods 1 (9), 308–326.
Thomas, A., van Leeuwen, J., 2015. Pure Nash equilibria in graphical games and treewidth. Algorithmica 71 (3), 581–604.
Vazirani, V.V., Yannakakis, M., 2011. Market equilibrium under separable, piecewise-linear, concave utilities. J. ACM 58 (3), 1–25.
Yu, S., Zhou, K., Brantingham, P.J., Vorobeychik, Y., 2020. Computing equilibria in binary networked public goods games. In: AAAI, pp. 2310–2317.