

Public Goods Games in Directed Networks

CHRISTOS PAPADIMITRIOU and BINGHUI PENG, Columbia University

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.

CCS Concepts: • **Theory of computation** → **Algorithmic game theory; Exact and approximate computation of equilibria; Network games.**

Additional Key Words and Phrases: public goods game, network game, equilibrium computation.

ACM Reference Format:

Christos Papadimitriou and Binghui Peng. 2021. Public Goods Games in Directed Networks. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC '21), July 18–23, 2021, Budapest, Hungary*. ACM, New York, NY, USA, 18 pages. <https://doi.org/10.1145/3465456.3467616>

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-rivalry¹). Scientific knowledge [42], 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 [3] for a classical framework for the public goods problem within which a unique Nash equilibrium exists.

Networks are perfect arenas for public goods games [7]. 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 [45] 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

¹Non-rivalry was called *collective consumption* by Paul Samuelson, who initiated the study of the subject [39].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

EC '21, July 18–23, 2021, Budapest, Hungary

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8554-1/21/07...\$15.00

<https://doi.org/10.1145/3465456.3467616>

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 [45]), pure Nash equilibria exist — typically corresponding to independent or dominating sets of the graph — and are easy to find.

Undirected graphs have limitations as models 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 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; while social networks are often directed: Twitter, Instagram, Flickr, and others. There has been some work on public goods in directed networks, e.g. [32], 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 F_2 . The main part of the proof entails showing that all other functions make the equilibrium problem NP-hard (Theorem 3.4).

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 [11, 38]. 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 *max* case of the problem (the utility is the maximum of the neighbors minus the good’s cost) is one where it can be seen through a simple averaging argument that a mixed equilibrium exists; we show that this too is PPAD-complete to find, this time by a reduction from mixed Nash equilibria in two-player win-lose games [1, 16].

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* [36]? 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 [19] and [43], but several substantial adaptations and innovations are needed.

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

- Sweeping intractability results for the equilibrium problem of public goods games in directed networks, including a novel PPAD-completeness proof through threshold games, an intriguing analysis of polynomial special cases for the pure equilibrium problem culminating in a precise P/NP-complete dichotomy, and even a very different PPAD-completeness proof for divisible goods.
- The formulation of public goods games in 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 (Section 3). In the discussion section we point out that the complexity of such classes of equilibrium problems is open even for undirected networks.
- An approximation algorithm when the treewidth is $O(\log n / \log \log n)$, through the development of new and enhanced techniques for approximating equilibrium problems in graphical games with small treewidth.

1.1 Related work

Bramoullé and Kranton [7] 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 [2, 5, 8, 17, 23, 24, 30, 32, 41, 45] for follow-up works. Bramoullé et. al. [8] 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 [2] differentiates private provision from public provision, and again characterizes the existence and uniqueness of a Nash equilibrium through the lowest eigenvalue of the graph. Public goods games were first generalized to directed graphs in [32], who provide sufficient conditions for pure Nash equilibria to exist. The only complexity result regarding such public goods games we are aware of is [45]: 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. We refer interested readers to the surveys [6, 26, 28] for a general coverage of this area.

Our work uses certain ideas from *graphical games* [29]. It is NP-hard to find a pure Nash equilibrium [27] and PPAD-hard to compute, even approximately, a mixed Nash equilibrium [11, 18, 38] of a general graphical game with maximum degree 3. However, the problem is tractable in several settings [19, 20, 43]. Daskalakis and Papadimitriou [19] 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 Leeuwen [43] provided an algorithm that computes a pure Nash equilibrium in $\text{poly}(s^w, |M|)$, 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 [35] to capture one particular genre of total search functions, encompassing the notion of equilibrium. The PPAD-completeness of Nash equilibria was established in [11, 18] and extended recently in [37, 38]. Over the past decades, a broad range of problems have been proved to be PPAD-hard, including equilibrium computation [1, 12, 16, 21], market equilibrium [9, 10, 14, 15, 44], equilibrium in auction [13, 25], fair allocation [33], min-max optimization [22] and problems in financial networks [40].

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., $N(i) = \{i\} \cup \{j \mid (j, i) \in E\}$. 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 that 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 $p : \mathbb{R}_+ \rightarrow \mathbb{R}_+$.

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 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. This can be achieved by assuming that all values of X are rational (not a significant loss of generality), while p is also rational but with a large denominator (technically, larger than the square of the largest denominator used in the values of X). 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,

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

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

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

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

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

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

$$\forall i \in V, \forall s_i \in \text{Supp}(\Delta_i), s'_i \in S_i : \mathbb{E}_{s_{-i} \sim \Delta_{-i}} [u_i(s_i, s_{-i})] \geq \mathbb{E}_{s_{-i} \sim \Delta_{-i}} [u_i(s'_i, s_{-i})] - \epsilon. \quad (2)$$

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 Figure 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.




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

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 $s = (s_1, \dots, s_n)$, 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 Figure. 2 and Figure. 3).




Fig. 2. Variable gadget.




Fig. 3. Clause gadget.

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.

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. \square

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 $X(1) \leq p$; 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 $k \geq 0, X(k + 1) \geq X(k) + p$; 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 $k \geq 0, X(k + 1) < X(k) + p$ 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 $i \in [n]$

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

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 :

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

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

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 Figure 2 and Figure 3), but we reduce from several different NP-hard problems. First, observe that when X is not flat, steep, or alternating, there must be a $k \geq 0$ such that $X(k + 1) > X(k) + p$,

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

$X(k+2) < X(k+1) + p$. 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 $k > 0$. 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 the previous paragraph.

We only need to be cautious about one exception: $X(1) > p, \dots, X(k+1) > X(k) + p, X(k+2) < X(k+1) + p, X(k+3) > X(k+2) + p, X(k+4) < X(k+3) + p$, 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. It is easy to check that 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. \square

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 $y - \epsilon \leq x \leq y + \epsilon$. 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.*

We start with a high level overview of the proof. We reduce from the ϵ -GCIRCUIT problem (see Definition 4.5), which is shown to be PPAD-hard for sufficiently small constant $\epsilon > 0$ by Rubinstein [38]. 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. The threshold game is semi-anonymous, and thus we can only modify the local structure of the graph in order to construct all 9 types of gates $\{G_{\xi}, G_{\times\xi}, G_{=}, G_{+}, G_{-}, G_{<}, G_{\wedge}, G_{\vee}, G_{\neg}\}$. We start by constructing an elementary gadget $G_{\frac{1}{2}-}$ (see Figure ??), and use it as a building block to gradually construct most of the gates. We restrict the players' equilibrium strategies to be in $[0, \frac{1}{2} + \epsilon]$ in the *arithmetic* (non-logic) gates. The logic gates ($G_{\wedge}, G_{\vee}, G_{\neg}$) are special: in order to prove PPA-hardness for constant ϵ , we need them to be error-resilient, in that they do not amplify errors of the input. We achieve this by restricting players' equilibrium strategy to be $\{0, 1\}$ (instead of $\{0, \frac{1}{2}\}$) when doing logic operations, and construct *transformer gadgets* $G_{\frac{1}{2}, 1}, G_{1, \frac{1}{2}}$ to map between the domains $\{0, \frac{1}{2}\}$ and $\{0, 1\}$.

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 $\mathbf{x} = (x_1, \dots, x_n) \in [0, 1]^n$ is an equilibrium if it satisfies

$$x_i = \begin{cases} 0 & \sum_{j \in N_i} x_j > t \\ 1 & \sum_{j \in N_i} x_j < t \\ \text{arbitrary} & \sum_{j \in N_i} x_j = t \end{cases} . \quad (3)$$

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

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 $\mathbf{x} = (x_1, \dots, x_n) \in [0, 1]^n$ of a threshold game $\mathcal{G}(V, E, t)$ satisfies

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

We next establish the equivalence between threshold games and public good games. The proof can be found in full version of this paper [34].

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

Next, we give the definition of generalized circuits.

Definition 4.5. (*Generalized circuit [11]*) 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; $v_1, v_2 \in V \cup \{\text{nil}\}$ are the input nodes, $\alpha \in \mathbb{R} \cup \{\text{nil}\}$ 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 $T, T' \in \mathcal{T}$, $T = (G, v_1, v_2, v, \alpha)$ 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 $G_{\xi}, G_{\times\xi}$ 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 $\mathbf{x} : V \rightarrow [0, 1]$ ϵ -approximately satisfies \mathcal{S} , if it satisfies the constraints shown in table 1.

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

Table 1

To complete the proof we must reduce ϵ -GCIRCUIT to computing an ϵ -approximate equilibrium of the threshold game.

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

The proof of this result can be found in the full version. It is based on an elementary game gadget $G_{\frac{1}{2}-}(|v_1, v_2|v)$, where $v_1, v_2 \in V \cup \{nil\}$ are input players, $v \in V$ is the output player. The gadget consists of a directed cycle, where v_1, v_2 have directed edges to the auxiliary player v_a , v_a has a directed edge to the auxiliary player v_b , v_b points to the output player v and there is a directed edge from v to v_a . The output player v could have many outgoing edges, but it only has one incoming edge from the internal node v_b . The intention is that, at equilibrium, $\mathbf{x}[v] = \max\{\frac{1}{2} - \mathbf{x}[v_1] - \mathbf{x}[v_2], 0\}$. The core of the proof, which we omit here, entails the construction and deployment of gates $G_{=}$, G_{+} , G_{-} , whose design is based on the elementary gadget described above.

Combining Theorem 4.7 and 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, Bramoulle et. al. [7] prove that there is always a pure Nash equilibrium³. 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 public goods games, for $\epsilon = 1/\text{poly}(n)$.*

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 [1, 16] that finding an ϵ -Nash of two-player win-lose game is PPAD-hard for $\epsilon = 1/\text{poly}(n)$. 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 $\text{poly}(n) \cdot \epsilon$ -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 [31], 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

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

³They 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 [7]: 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 $\mathbf{e} = \{1, \dots, 1\}^T$ denotes the n -dimension all 1 vector. The proof is by contradiction. Suppose $|x| < \frac{1}{4n}$; it then follows that

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

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

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

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 $x > \frac{1}{4n}, y > \frac{1}{4n}$. 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 $4n\epsilon$ -Nash of the win-lose game (R, C) . \square

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 $\mathbf{s} = (s_1, \dots, s_n)$ of the public goods game, we can find a symmetric $3n\epsilon$ -Nash of game (A, B) .*

PROOF. Let $\mathbf{s} = (s_1, \dots, s_n)$ be an ϵ -approximate pure Nash equilibrium of the public goods game, then for any agent $i \in [n]$, we have $\sum_{j \in N(i)} s_j \geq 1 - \epsilon$. 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) $\sum_{j \in N_j} s_i \leq 1$, then the best response is $b_i(s_{-i}) = 1 - \sum_{j \in N} s_i$ 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

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

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

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

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

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

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

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 $\epsilon = 1/\text{poly}(n)$. 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 Figure 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 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. \square

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 decomposition⁴ 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 $\text{twd}(G)$, 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 $\text{twd}(G)$ 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 [19], we get rid of the exponential dependence on d . Alas, we make no assumptions on the sparsity of the graph.

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

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 $\text{poly}(n) \cdot \min\{2d/\epsilon, 16 \log(n)/\epsilon\}^{O(w)}$. In particular, when the treewidth is $O\left(\frac{\log n}{\log \log n}\right)$, we can find an ϵ -Nash equilibrium in $\text{poly}(n) \cdot \left(\frac{1}{\epsilon}\right)^{O(w)}$ 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 [19]. To achieve this, we need to circumvent several difficulties, explained below. Like the proof in [19], 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 [19] works on the *primal* graph (see [19] for the definition), whose treewidth can be $w \cdot d$, yielding an exponential dependence on d . Instead, our algorithm directly works on the the original graph through *dynamic programming*, with no exponential dependence on d . Our dynamic programming method bares some similarities with the approach in [43]. 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.

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 $S^\delta = [\delta]$, where $\delta = \max\{\epsilon/2d, \epsilon/16 \log n\}$. We first show that there exists an ϵ -approximate *pure* Nash equilibrium in 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 $\mathbf{x} = (x_1, \dots, x_n)$ 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

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

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 $(\tilde{x}_1, \dots, \tilde{x}_n)$ is an ϵ -approximate equilibrium with positive probability, therefore proving its existence.

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

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

and

$$\Pr\left(\sum_{j \in N_i} \tilde{x}_j - \sum_{j \in N_i} x_j \geq \epsilon\right) = \Pr\left(\sum_{j \in N_i} (\tilde{x}_j - t_j\delta) - \sum_{j \in N_i} (\mathbb{E}[\tilde{x}_j] - t_j\delta) \geq \epsilon\right) \leq \exp\left(-\frac{\epsilon}{3\delta}\right) \leq n^{-2}. \quad (6)$$

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

$$\Pr \left(\left| \sum_{j \in N_i} \tilde{x}_j - \sum_{j \in N_i} x_j \right| \geq \epsilon \right) = \Pr \left(\left| \sum_{j \in N_i} (\tilde{x}_j - t_j \delta) - \sum_{j \in N_i} (\mathbb{E}[\tilde{x}_j] - t_j \delta) \right| \geq \epsilon \right) \leq 2 \exp \left(-\frac{\epsilon^2}{2\delta (\sum_{j \in N_i} \mathbb{E}[\tilde{x}_j] - \sum_{j \in N_i} t_j \delta)} \right) \leq 2 \exp \left(-\frac{\epsilon}{4\delta} \right) \leq n^{-2}. \quad (7)$$

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

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

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

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 Figure 4 for an illustration).

- (1) Leaf node.
- (2) Forget node. Such a node i has only one child i' , and $X_{i'} = X_i \setminus \{v\}$ 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 exact 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 [4].




Fig. 4. 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 $T_i : [\delta]^{|X_i|} \times [\delta]^{|X_i|} \rightarrow \{0, 1\}$ for each node X_i , and we note that the size of the table is bounded by $\delta^{-O(w)}$. Ideally, we would set an 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 $v \in V_i \setminus X_i$, the vertex v satisfies the equilibrium condition,

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

We remark that vertices in X_i do not need to satisfy the equilibrium condition, and we only record the summation of their neighbors in $V_i \setminus X_i$. 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 $X_{i'} = X_i \setminus \{v\}$ 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 $X_i = X_{i_1} \cup X_{i_2}$. Then for any $s, c \in [\delta]^{|X_i|}$, we set $T_i(s, c) = 1$ iff there exists $c_1, c_2 \in [\delta]^{|X_{i_1}|}, [\delta]^{|X_{i_2}|}$, 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 $(s, c) \in [\delta]^{|X_r|} \times [\delta]^{|X_r|}$, 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 $G(V, E)$, we can find an ϵ -approximate pure Nash equilibrium in time $\text{poly}(n) \cdot \min\{2d/\epsilon, 16 \log(n/\epsilon)\}^{O(w)}$ time. In particular, when the treewidth is $O(\log n / \log \log n)$, we can find an ϵ -approximate pure Nash equilibrium in $\text{poly}(n) \cdot (\frac{1}{\epsilon})^{O(w)}$ time.*

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

7 DISCUSSION

We explore 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.

Are public good games in directed networks real? It can be argued⁵ that some of the directed graphs we evoked in the introduction (towns that are downwind or upriver from one another, or

⁵Many thanks to the reviewers for bringing this point up.

the relationship “B is on A’s way to work”) are *transitive*, and it is easy 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 asymmetric and non-transitive, and so are infection networks in much of epidemic modeling. In addition, we hope that our techniques and techniques may be a public good of some value to the community.

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 $\{1, \dots, 7\}$ and the edges go from i to $i + 1, i + 2, 1 + 4 \bmod 7$. But this does not immediately give us an NP-hardness proof, nor does it generalize to other composition functions with p steps.

We believe that the general form of the utility function of public goods games in networks (Section 3) has not been articulated in the past, and it does lead to an interesting complexity classification problem. We note that the equilibrium situation in this generality is open in the *undirected* case. Take for example the case where $p < 1$ and X is the threshold function $X(i) = 1$ if $i \geq 2$ and zero otherwise; we believe that this case is NP-complete. The complexity dichotomy problem in undirected networks with general utility functions is a very interesting open problem raised by this work.

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.

ACKNOWLEDGMENTS

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

REFERENCES

- [1] Tim Abbott, Daniel Kane, and Paul Valiant. 2005. On the complexity of two-player win-lose games. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS'05)*. IEEE, 113–122.
- [2] Nizar Allouch. 2015. On the private provision of public goods on networks. *Journal of Economic Theory* 157 (2015), 527–552.
- [3] Theodore Bergstrom, Lawrence Blume, and Hal Varian. 1986. On the private provision of public goods. *Journal of Public Economics* 29, 1 (1986), 25–49.

- [4] Hans L Bodlaender and Arie MCA Koster. 2008. Combinatorial optimization on graphs of bounded treewidth. *Comput. J.* 51, 3 (2008), 255–269.
- [5] Leonardo Boncinelli and Paolo Pin. 2012. Stochastic stability in best shot network games. *Games and Economic Behavior* 75, 2 (2012), 538–554.
- [6] Yann Bramoullé and Rachel Kranton. [n.d.]. Games Played on Networks. In *The Oxford Handbook of the Economics of Networks*.
- [7] Yann Bramoullé, Rachel Kranton, et al. 2007. Public goods in networks. *Journal of Economic Theory* 135, 1 (2007), 478–494.
- [8] Yann Bramoullé, Rachel Kranton, and Martin D’amours. 2014. Strategic interaction and networks. *American Economic Review* 104, 3 (2014), 898–930.
- [9] Wei Chen, Pinyan Lu, Xiaorui Sun, Bo Tang, Yajun Wang, and Zeyuan Allen Zhu. 2011. Optimal pricing in social networks with incomplete information. In *International Workshop on Internet and Network Economics*. Springer, 49–60.
- [10] Xi Chen, Decheng Dai, Ye Du, and Shang-Hua Teng. 2009. 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, 273–282.
- [11] Xi Chen, Xiaotie Deng, and Shang-Hua Teng. 2009. Settling the complexity of computing two-player Nash equilibria. *Journal of the ACM (JACM)* 56, 3 (2009), 1–57.
- [12] Xi Chen, David Durfee, and Anthei Orfanou. 2015. On the complexity of Nash equilibria in anonymous games. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*. 381–390.
- [13] Xi Chen, Christian Kroer, and Rachitesh Kumar. 2021. The Complexity of Pricing for Second-Price Auctions. In *Proceedings of the 22nd ACM Conference on Electronic Commerce*.
- [14] Xi Chen, Dimitris Paparas, and Mihalis Yannakakis. 2013. The complexity of non-monotone markets. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*. 181–190.
- [15] Xi Chen and Shang-Hua Teng. 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, 647–656.
- [16] Xi Chen, Shang-Hua Teng, and Paul Valiant. 2007. The approximation complexity of win-lose games. In *SODA*, Vol. 7. 159–168.
- [17] Luca Dall’Asta, Paolo Pin, and Abolfazl Ramezani. 2011. Optimal equilibria of the best shot game. *Journal of Public Economic Theory* 13, 6 (2011), 885–901.
- [18] Constantinos Daskalakis, Paul W Goldberg, and Christos H Papadimitriou. 2009. The complexity of computing a Nash equilibrium. *SIAM J. Comput.* 39, 1 (2009), 195–259.
- [19] Constantinos Daskalakis and Christos H Papadimitriou. 2006. Computing pure Nash equilibria in graphical games via Markov random fields. In *Proceedings of the 7th ACM conference on Electronic commerce*. 91–99.
- [20] Constantinos Daskalakis and Christos H Papadimitriou. 2015. Approximate Nash equilibria in anonymous games. *Journal of Economic Theory* 156 (2015), 207–245.
- [21] Constantinos Daskalakis, Grant Schoenebeck, Gregory Valiant, and Paul Valiant. 2009. On the complexity of Nash equilibria of action-graph games. In *Proceedings of the twentieth annual ACM-SIAM symposium on Discrete algorithms*. SIAM, 710–719.
- [22] Constantinos Daskalakis, Stratis Skoulakis, and Manolis Zampetakis. 2021. The complexity of constrained min-max optimization. In *Proceedings of the 53th Annual ACM SIGACT Symposium on Theory of Computing*.
- [23] Matthew Elliott and Benjamin Golub. 2019. A network approach to public goods. *Journal of Political Economy* 127, 2 (2019), 730–776.
- [24] Michal Feldman, David Kempe, Brendan Lucier, and Renato Paes Leme. 2013. Pricing public goods for private sale. In *Proceedings of the fourteenth ACM conference on Electronic commerce*. 417–434.
- [25] Aris Filos-Ratsikas, Yiannis Giannakopoulos, Alexandros Hollender, Philip Lazos, and Diogo Poças. 2021. On the Complexity of Equilibrium Computation in First-Price Auctions. In *Proceedings of the 22nd ACM Conference on Electronic Commerce*.
- [26] 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.
- [27] Georg Gottlob, Gianluigi Greco, and Francesco Scarcello. 2005. Pure Nash equilibria: Hard and easy games. *Journal of Artificial Intelligence Research* 24 (2005), 357–406.
- [28] Matthew O Jackson and Yves Zenou. 2015. Games on networks. In *Handbook of game theory with economic applications*. Vol. 4. Elsevier, 95–163.
- [29] Michael Kearns, Michael L Littman, and Satinder Singh. 2001. Graphical models for game theory. In *the 17th Conference in Uncertainty in Artificial Intelligence, UAI (2001)*.
- [30] David Kempe, Sixie Yu, and Yevgeniy Vorobeychik. 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. 611–619.

- [31] Carlton E Lemke and Joseph T Howson, Jr. 1964. Equilibrium points of bimatrix games. *Journal of the Society for industrial and Applied Mathematics* 12, 2 (1964), 413–423.
- [32] Dunia López-Pintado. 2013. Public goods in directed networks. *Economics Letters* 121, 2 (2013), 160–162.
- [33] Abraham Othman, Christos Papadimitriou, and Aviad Rubinstein. 2016. The complexity of fairness through equilibrium. *ACM Transactions on Economics and Computation (TEAC)* 4, 4 (2016), 1–19.
- [34] Christos Papadimitriou and Binghui Peng. 2021. Public Goods Games in Directed Networks. In *Proceedings of the 22nd ACM Conference on Electronic Commerce*. Full version can be found at <https://arxiv.org/abs/2106.00718>.
- [35] Christos H Papadimitriou. 1994. On the complexity of the parity argument and other inefficient proofs of existence. *Journal of Computer and system Sciences* 48, 3 (1994), 498–532.
- [36] Neil Robertson and Paul D. Seymour. 1986. Graph minors. II. Algorithmic aspects of tree-width. *Journal of algorithms* 7, 3 (1986), 309–322.
- [37] Aviad Rubinstein. 2016. Settling the complexity of computing approximate two-player Nash equilibria. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 258–265.
- [38] Aviad Rubinstein. 2018. Inapproximability of Nash equilibrium. *SIAM J. Comput.* 47, 3 (2018), 917–959.
- [39] Paul A Samuelson. 1954. The pure theory of public expenditure. *The review of economics and statistics* (1954), 387–389.
- [40] Steffen Schuldenzucker, Sven Seuken, and Stefano Battiston. 2017. Finding Clearing Payments in Financial Networks with Credit Default Swaps is PPAD-complete. In *8th Innovations in Theoretical Computer Science Conference (ITCS 2017)*.
- [41] Kijung Shin, Euiwoong Lee, Dhivya Eswaran, and Ariel D Procaccia. 2017. Why you should charge your friends for borrowing your stuff. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*. 395–401.
- [42] Joseph E Stiglitz. 1999. Knowledge as a global public good. *Global public goods* 1, 9 (1999), 308–326.
- [43] Antonis Thomas and Jan van Leeuwen. 2015. Pure Nash equilibria in graphical games and treewidth. *Algorithmica* 71, 3 (2015), 581–604.
- [44] Vijay V Vazirani and Mihalis Yannakakis. 2011. Market equilibrium under separable, piecewise-linear, concave utilities. *Journal of the ACM (JACM)* 58, 3 (2011), 1–25.
- [45] Sixie Yu, Kai Zhou, P Jeffrey Brantingham, and Yevgeniy Vorobeychik. 2020. Computing Equilibria in Binary Networked Public Goods Games.. In *AAAI*. 2310–2317.