

Symmetries and the Complexity of Pure Nash Equilibrium

Felix Brandt^a Felix Fischer^{a,*} Markus Holzer^b

^a*Institut für Informatik, Universität München,
Oettingenstr. 67, 80538 München, Germany*

^b*Institut für Informatik, Technische Universität München
Boltzmannstr. 3, 85748 Garching, Germany*

Abstract

Strategic games may exhibit symmetries in a variety of ways. A characteristic feature, enabling the compact representation of games even when the number of players is unbounded, is that players cannot, or need not, distinguish between the other players. We investigate the computational complexity of pure Nash equilibria in four classes of symmetric games obtained by considering two additional properties: *identical payoff functions* for all players and the ability to *distinguish oneself* from the other players. In contrast to other types of succinctly representable multi-player games, the pure equilibrium problem is tractable in all four classes when only a constant number of actions is available to each player. Identical payoff functions make the difference between TC^0 -completeness and membership in AC^0 , while a growing number of actions renders the equilibrium problem NP-hard for three of the classes and PLS-hard for the most restricted class for which the existence of a pure equilibrium is guaranteed. Our results also extend to larger classes of *threshold symmetric* games where players are unable to determine the exact number of players playing a certain action.

Keywords: game theory, normal-form games, symmetries, pure Nash equilibrium, computational complexity

1 Introduction

In recent years, the computational complexity of game-theoretic solution concepts, both in cooperative and non-cooperative game theory, has come under increas-

* Corresponding author; Telephone: +49 89 2180 9405; Fax: +49 89 2180 9338

Email addresses: brandtf@tcs.ifi.lmu.de (Felix Brandt),
fischerf@tcs.ifi.lmu.de (Felix Fischer), holzer@in.tum.de (Markus Holzer).

ing scrutiny. A major obstacle when considering normal-form games with an unbounded number of players is the exponential size of the explicit representation of the payoffs. More precisely, a general game in normal-form with n players and k actions per player comprises $n \cdot k^n$ numbers. Computational statements over such large objects are somewhat dubious for two reasons [cf. 28]. First, the value of efficient, *i.e.*, polynomial-time, algorithms for problems whose input size is already exponential in a natural parameter (the number of players) is questionable. Secondly, most, if not all, “natural” multi-player games will hardly be given as multi-dimensional payoff matrices but rather in terms of some more intuitive (and compact) representation. A natural and straightforward way to simplify the representation of multi-player games is to somehow formalize similarities between players. As a matter of fact, *symmetric games* have been studied since the early days of game theory [e.g., 34, 15, 25]. The established definition states that a game is symmetric if the payoff functions of all players are *identical* and *symmetric* in the other players’ actions, *i.e.*, it is impossible to distinguish between the other players [35, 24]. When explicitly looking at *multi-player* games, there are other conceivable notions of symmetry. For instance, dropping the requirement of identical payoff functions yields a more general class of multi-player games that still admit a compact representation.

In this paper, we define four classes of succinctly representable symmetric multi-player games and study the computational complexity of finding pure Nash equilibria in games belonging to these classes. It turns out that in all four classes equilibria can be found efficiently if only a constant number of actions is available to each player. Moreover, identical payoff functions for all players further reduce the computational complexity of pure equilibria, an effect that is nullified as soon as there are two different payoff functions. The fact that a player cannot (or does not) distinguish himself from the other players, does not seem to offer any computational advantage. Finally, computing pure equilibria becomes intractable in all four classes of symmetric games when the number of actions grows at least linearly in the number of players.

Unlike Nash equilibria in *mixed* strategies, *i.e.*, probabilistic combinations of actions, pure Nash equilibria are not guaranteed to exist. If they exist, however, pure equilibria have two distinct advantages over mixed ones. For one, requiring randomization in order to reach a stable outcome has been criticized on various grounds. In multi-player games, where action probabilities in equilibrium can be irrational numbers, randomization is particularly questionable. Secondly, pure equilibria as computational objects are usually much smaller in size than mixed ones.

We assume the reader to be familiar with the well-known chain of complexity classes $AC^0 \subset TC^0 \subseteq L \subseteq P \subseteq NP$, and the notions of constant-depth and polynomial-time reducibility [e.g., 7, 27, 21]. AC^0 is the class of problems solvable by uniform constant-depth Boolean circuits with unbounded fan-in, and TC^0 adds so-called threshold gates which output *true* if and only if the number of *true* inputs exceeds a certain threshold. Here, uniformity means that there is an “effi-

cient” algorithm for constructing, for each input length n , the circuit C_n from the circuit family $C = (C_n)_{n \geq 0}$. Different notions of efficiency give rise to different notions of uniformity [31]. We will consider logspace-uniform circuit families, where the mapping $n \mapsto C_n$ is computable in deterministic logarithmic space. In some of our constructions we use sub-circuits described by Chandra et al. [7] as basic building blocks. It is easy to see that all these sub-circuits are logspace-uniform. Finally, L is the class of problems solvable by deterministic Turing machines using only logarithmic space, and P and NP are the classes of problems that can be solved in polynomial time by deterministic and nondeterministic Turing machines, respectively. Furthermore, $\#P$ is the class of counting problems associated with polynomially balanced polynomial-time decidable relations. The class PLS of polynomial local search problems and an appropriate notion of reduction [22] will be introduced as needed.

The remainder of this paper is organized as follows: In the following section, we survey relevant work on symmetric games, succinct representations, and the computational complexity of pure Nash equilibrium. In Section 3, we then formally introduce four different notions of symmetry in strategic games and the solution concept of Nash equilibrium. The main results of this paper, including efficient algorithms as well as hardness results for all four symmetry classes, are given in Section 4. In Section 5, we provide additional results for a more general notion of symmetry. Section 6 concludes the paper and points to some open problems.

2 Related Work

Symmetries in games have been investigated since the earliest days of game theory. Von Neumann [34] and von Neumann and Morgenstern [35] were the first to consider symmetries of *cooperative* games, calling a game in characteristic form symmetric if the value of a coalition depends only on its size. In the context of two-player (non-cooperative) normal-form games, the term symmetric is used to refer to games with a skew-symmetric payoff matrix [e.g., 5, 15], corresponding to strong symmetry in the vocabulary of this paper. Gale et al. [15] provided a (polynomial-time) reduction from arbitrary games to symmetric games which preserves equilibria. Since finding a (possibly mixed) equilibrium in general games has recently been shown PPAD-complete even for games with just two players [10, 8], the same holds for symmetric games as well. The above hardness result has also led to an increased interest in approximate equilibria. In particular, the larger class of *anonymous* games, allowing different payoff functions for different players, has been shown to admit an approximation by a factor depending on the Lipschitz constant of the payoff function and on the square of the number of actions, and a polynomial-time approximation scheme for the case of two actions [11].

To date, most research on symmetries in games has concentrated on games that

require identical payoff functions for all players, called *symmetric* games in this paper. One of the reasons for this may have been the strong focus of the early research in non-cooperative game theory on two-player games, where anonymity as defined in this paper does not impose any restrictions. An early result by Nash [25] shows that there always exists an equilibrium respecting all symmetries of the game, which in symmetric games implies the existence of a symmetric equilibrium, *i.e.*, one where all players play the same (mixed) strategy. Papadimitriou and Roughgarden [28] capitalize on this existence result and show that a Nash equilibrium of a symmetric game with n players and k actions can be computed in P if $k = O(\log n / \log \log n)$. While their tractability results for *correlated equilibrium* [1] do not rely on identical payoff functions and hence apply to anonymous games as well, this is not the case for the results about Nash equilibria. The aforementioned existence of *symmetric* Nash equilibria neither extends to pure equilibria, nor does it hold for anonymous games. For example, Figure 3 on Page 8 shows an anonymous game without a symmetric equilibrium.

Obviously, deciding the existence of a pure Nash equilibrium is easy if the number of candidates for such an equilibrium, *i.e.*, the number of action profiles, is polynomial in the size of the game. This is certainly the case for the explicit representation of a game as a multi-dimensional table of payoffs, but no longer holds if the game is represented succinctly. For example, deciding the existence of a pure equilibrium has been shown to be NP-complete for games in graphical normal form [17, 14] or circuit form [33]. Apart from these generic types, many succinct representations are related to symmetries in that they exploit some form of independence among certain actions or players playing these actions. In congestion games [30], the available actions consist of sets of resources, and the payoff depends on the number of other players that have selected the same resources (*i.e.*, played the same action). Congestion games always have a pure equilibrium [30], and finding one is PLS-complete even for symmetric congestion games and in P in the symmetric network case [13]. For singleton (or simple) congestion games, where only a single resource can be selected, there is a polynomial-time algorithm for finding a social-welfare-maximizing equilibrium [19]. In local-effect games [23], the payoff from an action may also depend on (a function of) the number of agents playing “neighboring” actions. Unlike congestion games and local-effect games, action-graph games [3] can encode arbitrary payoffs. For action-graph games of bounded degree, expected payoffs and the Jacobian of the payoff function can be computed in polynomial time. The latter forms the practical bottleneck step of the algorithm of Govindan and Wilson [18] for finding Nash equilibria, but the algorithm may still take exponentially many steps to converge even for bounded degree. In fact, the pure equilibrium problem is NP-complete for symmetric action-graph games with bounded degree, but becomes tractable if the treewidth is bounded [20]. In general action-graph games, the pure equilibrium problem is NP-complete even if the action-graph is a bounded-degree tree [12]. Finally, Brandt et al. [6] transfer different notions of symmetry to graphical games and obtain mostly negative results concerning the complexity of pure Nash equilibrium.

3 Preliminaries

In this section, we formally define essential game-theoretic concepts, introduce four notions of symmetry in strategic multi-player games, and state several facts concerning these notions.

3.1 Strategic Games

An accepted way to model situations of strategic interaction is by means of a *normal-form game* [e.g., 24].

Definition 1 (normal-form game) A game in normal-form is a tuple $\Gamma = (N, (A_i)_{i \in N}, (p_i)_{i \in N})$ where N is a set of players and for each player $i \in N$, A_i is a nonempty set of actions available to player i , and $p_i : (\times_{i \in N} A_i) \rightarrow \mathbb{R}$ is a function mapping each action profile of the game (i.e., combination of actions) to a real-valued payoff for player i .

A combination of actions $s \in \times_{i \in N} A_i$ is also called a profile of *pure strategies*. This concept can be generalized to (*mixed*) *strategy profiles* $s \in S = \times_{i \in N} S_i$, by letting players randomize over their actions. We have S_i denote the set of probability distributions over player i 's actions, or (*mixed*) *strategies* available to player i . We further write $n = |N|$ for the number of players in a game, s_i for the i th strategy in profile s , and s_{-i} for the vector of all strategies in s but s_i .

3.2 Symmetries in Multi-Player Games

Symmetry as a property of a mathematical object typically refers to its invariance under a certain type of transformation. Symmetries of games usually mean invariance of the payoffs under automorphisms of the set of action profiles induced by some group of permutations of the set of players. Since such an automorphism preserves the number of players that play a particular action, a characteristic feature of symmetries in games is the inability to distinguish between other players. Following Daskalakis and Papadimitriou [11] the most general class of games with this property will be called *anonymous*. Four different classes of games are obtained by considering two additional characteristics: *identical payoff functions* for all players and the ability to *distinguish oneself* from the other players. The games obtained by adding the former property will be called *symmetric*, and presence of the latter will be indicated by the prefix “*self*”. For the obvious reason, we will henceforth talk about games where the set of actions is the same for all players and write $A = A_1 = \dots = A_n$ and $k = |A|$, respectively, to denote this set and its cardinality.




Fig. 1. Inclusion relationships between anonymous, symmetric, self-anonymous, and self-symmetric games

An intuitive way to describe anonymous games is in terms of equivalence classes of the aforementioned automorphism group, using a notion introduced by Parikh [29] in the context of context-free languages. Given a set A of actions, the *commutative image* of an action profile $s \in A^N$ is given by $\#(s) = (\#(a, s))_{a \in A}$ where $\#(a, s) = |\{i \in N : s_i = a\}|$. In other words, $\#(a, s)$ denotes the number of players playing action a in action profile s , and $\#(s)$ is the vector of these numbers for all the different actions. This definition naturally extends to action profiles for subsets of the players.

Definition 2 (symmetries) Let $\Gamma = (N, (A_i)_{i \in N}, (p_i)_{i \in N})$ be a normal-form game and A a set of actions such that $A_i = A$ for all $i \in N$. Γ is called

- anonymous if $p_i(s) = p_i(t)$ for all $i \in N$ and all $s, t \in A^N$ with $s_i = t_i$ and $\#(s_{-i}) = \#(t_{-i})$,
- symmetric if $p_i(s) = p_j(t)$ for all $i, j \in N$ and all $s, t \in A^N$ with $s_i = t_j$ and $\#(s_{-i}) = \#(t_{-j})$,
- self-anonymous if $p_i(s) = p_i(t)$ for all $i \in N$ and all $s, t \in A^N$ with $\#(s) = \#(t)$, and
- self-symmetric if $p_i(s) = p_j(t)$ for all $i, j \in N$ and all $s, t \in A^N$ with $\#(s) = \#(t)$.

When talking about anonymous games, we write $p_i(s_i, x_{-i})$ to denote the payoff of player i under any action profile s with $\#(s_{-i}) = x_{-i}$. For self-anonymous games, $p_i(x)$ is used to denote the payoff of player i under any profile s with $\#(s) = x$. It is easily verified that the class of self-symmetric games equals the intersection of symmetric and self-anonymous games, which in turn are both strictly contained in the class of anonymous games. An illustration of these inclusions is shown in Figure 1. Figure 2 illustrates the different payoff structures for $n = 3$ and $k = 2$. In terms of this characterization, a game is anonymous if the payoff $p_i(s)$ of player $i \in N$ in action profile s depends, besides his own action s_i , only on the number $\#(a, s_{-i})$ of other players playing each of the actions $a \in A$, but not on who plays them. If two players exchange actions, all other players' payoffs remain the same. For two-player games, anonymity does not impose any restrictions (action sets of equal size can simply be achieved by adding dummy actions for one of the players). This may be one of the reasons why anonymity has not received much attention in the past. A game is symmetric if it is anonymous and if the payoff function is the

Γ_1 :	(\cdot, \cdot, \cdot)	(a, g, c)	(a, b, \cdot)	(\cdot, e, f)	Γ_2 :	(a, a, a)	(b, c, b)	(b, b, c)	(e, d, d)
	(\cdot, b, c)	(d, e, \cdot)	(d, \cdot, f)	(\cdot, \cdot, \cdot)		(c, b, b)	(d, d, e)	(d, e, d)	(f, f, f)
Γ_3 :	(\cdot, \cdot, \cdot)	(a, b, c)	(a, b, c)	(d, e, f)	Γ_4 :	(a, a, a)	(b, b, b)	(b, b, b)	(c, c, c)
	(a, b, c)	(d, e, f)	(d, e, f)	(\cdot, \cdot, \cdot)		(b, b, b)	(c, c, c)	(c, c, c)	(d, d, d)

Fig. 2. Relationships between the payoffs of anonymous (Γ_1), symmetric (Γ_2), self-anonymous (Γ_3), and self-symmetric (Γ_4) games for $n = 3$ and $k = 2$. Players 1, 2, and 3 simultaneously choose rows, columns, and tables, respectively, and obtain payoffs according to the vector in the resulting cell. Each lower case letter stands for a payoff value, dots denote arbitrary payoff values. As an example for the separation of the different classes, Γ_1 is *not* symmetric if $a \neq c$ and *not* self-anonymous if $b \neq g$. Γ_2 is *not* self-anonymous if $b \neq c$. Γ_3 is *not* self-symmetric if $a \neq c$.

same for all players. Hence, if two players exchange actions, their payoffs are also exchanged while all other players' payoffs remain the same. Many well-known games like the Prisoner's Dilemma, Rock-Paper-Scissors, or Chicken are examples of symmetric (two-player) games. Multi-player simple congestion games [19] are also symmetric. In a self-anonymous game the payoff of each player depends only on the number $\#(a, s)$ of players playing each of the actions $a \in A$, *including the player himself*. If two players exchange actions, the payoffs of all players remain the same. Matching Pennies is a self-anonymous two-player game, voting with identical weights can be seen as an example for the multi-player case. Finally, in a self-symmetric game the payoff is always the same for all players and stays the same if two players exchange actions. Self-symmetric games thus are a special case of common payoff (or pure coordination) games, in which every action profile with maximum payoff is an equilibrium (since no player can gain by deviating). Other games guaranteed to possess a pure equilibrium, and the complexity of finding an equilibrium in these games, have been investigated by Fabrikant et al. [13].

Interestingly, the ability to distinguish oneself from the other players does not increase the complexity of the pure equilibrium problem when players only have two actions.

Lemma 1 *When there are only two actions available to each player, there exists an AC^0 -reduction from anonymous games to self-anonymous games that preserves pure Nash equilibria and identical payoff functions.*

Proof: Let $\Gamma = (N, (\{a_1, a_2\})_{i \in N}, (p_i)_{i \in N})$ be an anonymous game, and define a new game $\Gamma' = (N, \{a_1, a_2\}^n, (p'_i)_{i \in N})$ such that for all $i \in N$ and for all $x \in \{0, 1, \dots, n-1\}$,

- (1) $p'_i((x, n-x)) > p'_i((x+1, n-x-1))$ if and only if $p_i(a_1, (x, n-x-1)) > p_i(a_2, (x, n-x-1))$,
- (2) $p'_i((x, n-x)) < p'_i((x+1, n-x-1))$ if and only if $p_i(a_1, (x, n-x-1)) < p_i(a_2, (x, n-x-1))$, and

(0, 1, 1)	(0,0,1)	(0,1,0)	(0, 0, 0)
(1,1,1)	(0, 0, 0)	(0, 1, 0)	(1, 0, 1)

Fig. 3. Anonymous game with a unique, non-symmetric Nash equilibrium at the action profile with payoff (1, 1, 1). Players 1, 2, and 3 choose rows, columns, and tables, respectively. Outcomes are denoted as a vector of payoffs for the three players. Action profiles with the same commutative image as the equilibrium are shaded.

$$(3) \quad p'_i((x, n - x)) = p'_i((x + 1, n - x - 1)) \text{ if and only if } p_i(a_1, (x, n - x - 1)) = p_i(a_2, (x, n - x - 1)).$$

Depending on the payoff structure of Γ , it may be necessary to use up to n different payoffs in Γ' , even when Γ contains only two. It is now easily verified that Γ' is self-anonymous in general, and self-symmetric if the original game Γ is symmetric. It should be noted that this construction cannot in general be extended to games where players have more than two actions, because it can lead to cyclic preference relations. For example, the symmetric two-player game Rock-Paper-Scissors cannot be mapped to a corresponding self-symmetric game using the above technique. \square

3.3 Nash Equilibrium

One of the best-known solution concepts for strategic games is Nash equilibrium [25]. In a Nash equilibrium, no player is able to increase his payoff by *unilaterally* changing his strategy.

Definition 3 (Nash equilibrium) *A strategy profile $s \in S$ is called a Nash equilibrium if for each player $i \in N$ and each strategy $s'_i \in S_i$,*

$$p_i(s) \geq p_i((s_{-i}, s'_i)).$$

A Nash equilibrium is called pure if it is a pure strategy profile.

For general games, simply checking the equilibrium condition for each action profile takes time polynomial in the size of their explicit representation. Using a succinct representation for games where the size of the explicit representation grows exponentially in the number of players, which is the case for $k = 2$ already, quickly renders the problem NP-complete [14, 33]. On the other hand, the polynomial size even of the explicit representation for anonymous games with a constant number of actions might suggest that finding pure equilibria is easy by a similar argument as above. This reasoning is flawed, however, since a single entry in the payoff table corresponds to an exponential number of action profiles, and it is very well possible that only a single one of them is an equilibrium while all others are not. The anonymous game given in Figure 3 illustrates this fact.

4 Solving Anonymous and Symmetric Games

In this section, we analyze the computational complexity of pure Nash equilibrium in anonymous and symmetric games with a constant number of actions and a growing number of actions, respectively. We first consider games with a constant number of actions, and show membership of the pure equilibrium problem for complexity classes inside P. We then show that the same problem becomes NP-hard and PLS-hard, respectively, if the number of actions is not bounded.

4.1 A Few Words on Encodings

Since we are interested in games that model real-world situations, we will henceforth restrict ourselves to games that can be represented in space polynomial in their natural parameters, like the number of players or actions [cf. 28]. We will try to characterize games in terms of their natural parameters, while making as few assumptions as possible about any particular encoding. Apart from hardness results in Section 4.3, which will be shown to extend to games with an exponential number of players, we assume that the number of players of a game is polynomial in the size of its representation. We further assume that each player can determine efficiently whether a particular action is a best response for a given action profile of the other players, which obviously is both necessary and sufficient for *playing* a game rationally and efficiently at the same time. Tractability results then hold for any encoding satisfying these properties. Hardness, on the other hand, is established *via* some encoding which allows efficient and rational play.

The most basic way to encode a normal-form game is to explicitly write down a multi-dimensional table listing the payoffs for every single action profile. Certain games can be represented more succinctly because the payoff is the same for action profiles that are equivalent according to some equivalence relation, and needs only be specified once. For anonymous games, this equivalence relation is given by the number of players playing each action. The representation that lists the payoffs for every equivalence class will henceforth be referred to as the *explicit representation* of an anonymous game. There are $\binom{n+k-1}{k-1}$ distributions of n players among k actions. Since these are exactly the equivalence classes of the set of action profiles for $n-1$ players under the commutative image, an anonymous game can be represented using at most $n \cdot k \cdot \binom{n+k-2}{k-1}$ numbers, and is representable using space polynomial in n if and only if k is bounded by a constant. On the other hand, the size of the game becomes super-polynomial in n even for the slightest growth of k . Nevertheless, space polynomial in n may still suffice to encode certain subclasses of symmetric games with a larger number of actions if we use an implicit representation of the payoff functions like a Boolean circuit. It is easy to see that for games with a constant number of actions, any encoding of a game that has size at least linear in the

number of players and satisfies the above assumption of rational and efficient play is equivalent to its explicit representation under polynomial-time reductions.

4.2 Games with a Constant Number of Actions

We begin by investigating games with a constant number of actions. Obviously, solving a game cannot be easier than playing it optimally given that the opponents' actions are known. The most interesting upper bounds for the former problem will thus be obtained when the latter problem is easy. We will therefore assume throughout this section that for any action profile of his opponents, a player can compute the payoff of a particular action in AC^0 , *i.e.*, by evaluating a Boolean circuit with constant depth and bounded fan-in. This particularly holds if the payoff function is given explicitly. It will further be obvious from the proofs that for payoff functions that are harder to compute, the complexity of the pure equilibrium problem exactly matches that of computing the payoff function.

As we have noted earlier, the potential hardness of finding pure equilibria in games with succinct representations stems from the fact that the number of action profiles that are candidates for being an equilibrium is exponential in the size of the representation of the game. While anonymous games certainly satisfy this property, the pure equilibrium problem nevertheless turns out to be tractable. The following theorem concerns games where the number of players is polynomial in the size of the representation.

Theorem 1 *Deciding whether an anonymous or self-anonymous game with a constant number of actions has a pure Nash equilibrium is TC^0 -complete under constant-depth reducibility. Hardness holds even for games with three different payoffs and two different payoff functions.*

Proof: For membership in TC^0 , we propose an algorithm that decides whether there exists a pure Nash equilibrium with a given commutative image. The theorem then follows by observing that the number of different commutative images is polynomial in the number of players if the number of actions is constant.

Let $\Gamma = (N, (A_i)_{i \in N}, (p_i)_{i \in N})$ be an anonymous game, $A = \{a_1, a_2, \dots, a_k\}$ a set of actions such that $A_i = A$ for all $i \in N$. Given the commutative image $x = (x_{a_1}, x_{a_2}, \dots, x_{a_k})$ for some action profile of Γ , call an action $a_\ell \in A$ a *potential best response* for player i in x if $x_{a_\ell} > 0$ and

$$p_i(a_\ell, x_{-\ell}) \geq p_i(a_m, x_{-\ell}) \quad \text{for all } a_m \in A, \quad (1)$$

where $x_{-\ell} = (x_{a_1}, \dots, x_{a_{\ell-1}}, x_{a_{\ell+1}}, \dots, x_{a_k})$.

Fix a particular commutative image $x = (x_{a_1}, x_{a_2}, \dots, x_{a_k})$, and define a bipartite




Fig. 4. Matching problem for the game of Figure 3 (left) and representation of the same problem by a graph with a constant number of vertices (right), as used in the proof of Theorem 1.

graph $G = (V, E)$ such that

$$V = V_1 \cup V_2, \quad V_1 = N, \quad V_2 = \{(a_j, \ell) : a_j \in A, 1 \leq \ell \leq x_j\}, \quad \text{and} \\ E = \{(i, (a_j, \ell)) : a_j \text{ is a potential best response for } i \text{ under } x\}.$$

In other words, the two sides of G correspond to players and actions of Γ , respectively, with multiplicities of the actions chosen according to x . Edges connect each player to his potential best responses. The graph corresponding to the shaded action profiles in Figure 3 is shown on the left of Figure 4. It is now readily appreciated that a pure equilibrium of Γ with commutative image x directly corresponds to a perfect matching of G , and vice versa. Furthermore, by Hall's Theorem, G has a perfect matching if and only if $|\nu(V')| \geq |V'|$ for all $V' \subseteq V_1$, where $\nu(V') = \{v \in V_2 : (u, v) \in E, u \in V'\}$ is the neighborhood of vertex set V' [e.g., 4].

Observe that Hall's condition cannot be verified efficiently in general. We will see, however, that this can indeed be done for G , by considering a new graph obtained from G which possesses only a constant number of vertices. More formally, assume w.l.o.g. that for all $v \in V_1$, $\nu(v) \neq \emptyset$, and define an equivalence relation $\sim \subseteq V \times V$ such that for all $v, v' \in V$, $v \sim v'$ if and only if $\nu(v) = \nu(v')$. By construction of G , and since both the number of actions and the number of possible subsets of actions are constant, the set V/\sim of equivalence classes has constant size, and $V/\sim = (V_1/\sim) \cup (V_2/\sim)$. Each element of V_1/\sim corresponds to the set of players having a particular set of actions as their potential best responses in x . Each element of V_2/\sim corresponds to an action in A . The neighborhood function ν can naturally be extended to equivalence classes by letting for each $U \in V_1/\sim$, $\nu(U) = \{U' \in V_2/\sim : v \in \nu(u) \text{ for some } u \in U, v \in U'\}$. This yields a bipartite graph with vertex set V/\sim , the graph corresponding to the game in Figure 3 is shown on the right of Figure 4. It is now easily verified that G has a perfect matching, and Γ a pure equilibrium, if and only if for every $Y \subseteq V_1/\sim$, $\sum_{X \in Y} |\nu(X)| \geq |Y|$.

We proceed to show that this property can be verified by a threshold circuit with unbounded fan-in, constant depth, and a polynomial number of gates. From the description given below it is easy to see that the constructed circuit is logspace-uniform. Since V_1/\sim has only a constant number of subsets, we can construct a constant depth threshold circuit which uses sub-circuits UNARY-COUNT and UNARY-TO-BINARY as described by Chandra et al. [7] to sum over elements of

p_0	0	...	$\ell + 1$...	$m + 2$	p_1	0	...	$\ell + 1$...	$m + 2$
	...	1	0	1	0		...	1	0	1	0
	2	1	0	1	0	1	2	1	0	1	0

Fig. 5. Payoffs of the game Γ used in the proof of Theorem 1

the equivalence classes, and COMPARISON sub-circuits to verify the inequalities. The former is easily realized with the help of MAJORITY gates. It thus remains to be shown that for any $X \in V_1/\sim$, $|X|$ and $|\nu(X)|$ can be computed in TC^0 . For this, recall that a particular element of V_1/\sim corresponds to the set of players that have a certain set of actions as their set of best responses in x . To compute the number of such players we first construct a circuit of constant depth that uses COMPARISON sub-circuits to check whether Equation 1 is satisfied for a fixed commutative image x , a particular player $i \in N$, and a particular action $a \in A$. To check whether $C \subseteq A$ is the set of best responses for player i under x , we simply combine the outputs of the above circuits for all actions $a \in A$ into a single AND gate, negating the outputs of that for actions $a \notin C$. The desired number of players is then obtained by adding up the outputs of these gates for all players $i \in N$, again using UNARY-COUNT sub-circuits. On the other hand, $|\nu(X)|$ corresponds to the number of players bound to play an action from a certain subset in every action profile with commutative image x , and can easily be obtained by summing over the respective elements of x .

For *hardness*, we reduce the problem of deciding whether exactly ℓ bits of a string of m bits are 1 to deciding the existence of a pure equilibrium in a self-anonymous game. Hardness of the former problem is immediate from that of MAJORITY [e.g., 7]. For a particular m -bit string b , we define a game Γ with $m + 2$ players of two different types 0 and 1 and actions $A = \{0, 1\}$. The i th player of Γ is of type 0 or 1 if the i th bit of b is 0 or 1, respectively. Player $m + 1$ is of type 0, player $m + 2$ is of type 1. The payoffs p_0 and p_1 for the two types are given in Figure 5, the column labeled j specifies the payoff when exactly j players, including the player himself, play action 1. It is easily verified that this is an AC^0 reduction. We claim that Γ possesses a pure equilibrium if and only if exactly ℓ bits of b are 1. We observe the following:

- An action profile s cannot be an equilibrium of Γ if $\#(1, s) \neq \ell + 1$. In this case, the players of one of the two types get a higher payoff at both $\#(1, s) - 1$ and $\#(1, s) + 1$, or at one of these in case $\#(1, s) = 0$ and $\#(1, s) = m + 2$. Since by construction we have at least one player of each type, there always exists a player who can change his action to get a higher payoff.
- If there are $\ell + 1$ players of type 1, the action profile where all players of type 0 play action 0 and all players of type 1 play action 1 is an equilibrium. None of the players of type 0 can gain by changing his action to 1, and none of them can change his action to 0 (because all of them already play 0). A symmetric condition holds for players of type 1.
- In turn, if the number of players of type 1 does not equal $\ell + 1$, an action profile s with $\#(1, s) = \ell + 1$ cannot be an equilibrium. In this case, there must be (i) a

player of type 0 playing action 1 in s , or (ii) a player of type 1 playing 0. This player can change his action to get a higher payoff.

Hence, a pure equilibrium exists if and only if there are $\ell + 1$ players of type 1, *i.e.*, if and only if b has ℓ 1-bits. This completes the reduction. \square

In contrast to anonymous games, if s is a pure equilibrium of a symmetric game, so are all t satisfying $\#(t) = \#(s)$. This is due to the fact that the payoff functions of all players, and thus the situation of all players playing the same action $a \in A$, is identical, as would be the situation of any other player exchanging actions with someone playing a . We exploit this property to show that deciding the existence of a pure equilibrium in symmetric games with a constant number of actions is strictly easier than for anonymous and self-anonymous games.

Theorem 2 *The problem of deciding whether a symmetric game with a constant number of actions has a pure Nash equilibrium is in AC^0 .*

Proof: Like with anonymous games, an action profile s is an equilibrium of a symmetric game if and only if, for all $i \in N$, s_i is a best response to $\#(s_{-i})$, *i.e.*, if

$$p_i(s_i, \#(s_{-i})) \geq p_i(a, \#(s_{-i})) \quad \text{for all } a \in A. \quad (2)$$

For a particular player $i \in N$ and for constant k , checking this inequality requires only a constant number of comparisons and can be done using a circuit of constant depth and polynomial size [*e.g.*, 7]. When it comes to checking Equation 2 for the different players, the observation about action profiles with identical commutative images affords us a considerable computational advantage as compared to, say, anonymous or self-anonymous games. More precisely, we only have to check if Equation 2 is satisfied for a player *playing a certain action*, of which there are at most k . Again, this can be done using a circuit of constant depth and polynomial size if k is a constant.

Finally, to decide whether game Γ has a pure equilibrium, we have to check Equation 2 for the different values of $\#(s)$ for $s \in A^N$. If k is constant, there are only polynomially many of these, so the complete check requires only polynomial size and constant depth. \square

The reasoning in the proof of Theorem 2 also provides a nice illustration of the fact that *every* symmetric game with two actions possesses a pure equilibrium, as recently shown by Cheng et al. [9]. In the case of two actions, p_i depends only on player i 's action (0 or 1) and on the number of other players playing action 1. A pure equilibrium exists if for some m neither the players playing 0 (who see m players playing 1) nor the players playing 1 (who see $m - 1$ other players playing 1) have an incentive to deviate, *i.e.*, $p_i(0, m) \geq p_i(1, m)$ and $p_i(1, m - 1) \geq p_i(0, m - 1)$. For $m = 0$ and $m = n$, one of the conditions is trivially satisfied, because there are no players playing 1 or 0, respectively. It is now straightforward to show that at least one such

m must exist. Alternatively, the existence of pure equilibria in symmetric games with two actions can also be obtained as an immediate consequence of Lemma 1. We can transform every symmetric game with two actions into a self-symmetric game with the same set of equilibria, and every game in the latter class is guaranteed to have at least one pure equilibrium.

As stated earlier, self-symmetric games always possess a pure equilibrium, namely an action profile with maximum payoff for every player. We proceed to show that such an action profile, which has the additional property of maximizing social welfare, *i.e.*, the sum of all players' payoffs, can be found in AC^0 .

Theorem 3 *The problem of finding a social-welfare-maximizing pure Nash equilibrium of a self-symmetric game with a constant number of actions is in AC^0 .*

Proof: Since self-symmetric games belong to the class of common payoff games, any action profile with maximum payoff (for all players) is a social-welfare-maximizing equilibrium (and Pareto dominates any other strategy profile). Finding such an equilibrium is thus equivalent to finding the maximum of $\binom{n+k-2}{k-1}$ integers. The exact number is irrelevant as long as it is polynomial in the size of the input, which is certainly the case if k is bounded by a constant. Chandra et al. [7] have shown that the maximum of m m -bit binary numbers can be computed by an unbounded fan-in, constant-depth Boolean circuit of size polynomial in m . Since m is of course polynomial in the size of the input, the size of this circuit is as well. \square

4.3 Games with a Growing Number of Actions

The proofs we have seen so far could exploit the fact that for constant k the explicit representation of an anonymous game is equivalent, under the appropriate type of reduction, to any kind of payoff function computable in a particular complexity class inside P. This need no longer be the case for unbounded k , because then the size of the explicit representation grows exponentially in n . Such games may of course still admit a polynomial representation, for example if payoff functions are encoded by a Boolean circuit. We will now show that deciding the existence of a pure equilibrium in anonymous, symmetric, and self-anonymous games becomes NP-hard if the number of actions grows in n . For self-symmetric games, which always have a pure equilibrium, the associated search problem will be shown to be PLS-hard. In particular, we show NP-completeness and PLS-completeness, respectively, for games that have a polynomial number of players—like those covered by Theorems 1 and 2—and a number of actions that grows linearly in the number of players. It will be obvious from the proofs that hardness for the respective classes also holds for games with an exponential number of players and logarithmic growth of the number of actions. The corresponding case with a constant number of actions, on the other hand, remains open.

If the number of actions in a game is large enough, they can in principle be used to distinguish the players playing them. We will exploit this fact and prove the following theorems by reductions from satisfiability of a Boolean circuit. For this, recall that circuit satisfiability (CSAT), *i.e.*, deciding whether a Boolean circuit has a satisfying assignment, is NP-complete [e.g., 27]. We provide a reduction from CSAT to the problem of deciding the existence of a pure equilibrium in a special class of games. For a particular circuit C with inputs $M = \{1, 2, \dots, m\}$, we define a game Γ with at least m players and actions $A = \{a_j^0, a_j^1 : j \in M\} \cup \{b\}$. An action profile s of Γ where $\#(a_j^0, s) + \#(a_j^1, s) = 1$ for all $j \in M$, *i.e.*, one where exactly one action of each pair a_j^0, a_j^1 is played, directly corresponds to an assignment c of C , the j th bit of this assignment being 1 if and only if a_j^1 is played. Observe that in this case the auxiliary action b has to be played by exactly $n - m$ players. We can thus distinguish the action profiles of Γ corresponding to a satisfying assignment of C from those corresponding to a non-satisfying assignment and those not corresponding to an assignment at all.

Theorem 4 *Deciding whether a self-anonymous game has a pure Nash equilibrium is NP-complete, even if the number of actions is linear in the number of players and there is only a constant number of different payoffs.*

Proof: Membership in NP is obvious. Since the number of players is polynomial, we can simply guess an action profile and verify that it satisfies the equilibrium condition.

For hardness, we reduce satisfiability of a Boolean circuit C with inputs $M = \{1, 2, \dots, m\}$ to the existence of a pure equilibrium in a game Γ with $n \geq m$ players, actions $A = \{a_j^0, a_j^1 : j \in M\} \cup \{b\}$, and payoff functions p_i as follows:

- If s corresponds to a *satisfying* assignment of C , we let $p_i(s) = 1$ for all $i \in N$.
- Otherwise we let
 - $p_1(s) = 1, p_2(s) = 0$ if $\#(b, s)$ is even,
 - $p_1(s) = 0, p_2(s) = 1$ if $\#(b, s)$ is odd, and
 - $p_i(s) = 1$ for all $i \in N \setminus \{1, 2\}$.

We observe the following:

- In all of the above cases, the payoff of player i only depends on the number of players playing certain actions. If two players exchange actions, the payoff to all players remains the same. Hence, Γ is self-anonymous.
- Clearly, every action profile s corresponding to a satisfying assignment of C is an equilibrium, because in this case all players receive the maximum payoff of 1.
- For an action profile s *not* corresponding to a satisfying assignment of C , either player 1 or player 2 receives a payoff of 0. Furthermore, by choosing his own action to be either b or some other action, this player can determine the parity

of the number of players playing b . Changing the parity strictly increases the player's payoff. This means that s cannot be an equilibrium.

We have hence established a direct correspondence between satisfying assignments of C and pure equilibria of Γ . The transformation from C to Γ essentially works by writing down Boolean circuits that compute p_i . Observing that this can be done in time polynomial in the size of C if $n \leq 2^k$, where $k = |A|$, completes the proof. \square

As the reader may have noticed, the construction used in this proof has players 1 and 2 play *matching pennies* in the case where the Boolean circuit is not satisfied. Not only is this game a well-known example for a game that does not possess a pure equilibrium, it is also self-anonymous on its own. On the other hand, it is readily appreciated that the payoffs in this game do depend on the identities of the players, *i.e.*, the game is not symmetric. We will have to avail of a different construction for the symmetric case.

Theorem 5 *Deciding whether a symmetric game has a pure Nash equilibrium is NP-complete, even if the number of actions is linear in the number of players and there is only a constant number of different payoffs.*

Proof: Membership in NP is again obvious.

For hardness, we provide a reduction from CSAT, mapping a circuit C with inputs $M = \{1, 2, \dots, m\}$ to a game Γ with $n \geq m$ players, actions $A = \{a_j^0, a_j^1 : j \in M\} \cup \{b\}$, and payoff functions p_i as follows:

- If s corresponds to a satisfying assignment of C , we let $p_i(s) = 2$ for all $i \in N$.
- If s does not correspond to a satisfying assignment of C , we distinguish three different cases according to the number $\#(b, s)$ of players playing action b :
 - If $\#(b, s) < n - m$, we let $p_i(s) = 1$ if $s_i = b$, and $p_i(s) = 0$ otherwise.
 - If $\#(b, s) > n - m$, we let $p_i(s) = 0$ if $s_i = b$, and $p_i(s) = 1$ otherwise.
 - Finally, if $\#(b, s) = n - m$, we let
 - $p_i(s) = 2$ if $s_i = a_j^1$ for some $j \in M$, $\#(a_j^0, s) > 0$, and $\#(a_j^1, s) > 0$,
 - $p_i(s) = 1$ if $s_i = a_j^0$ for some $j \in M$, $\#(a_j^0, s) > 0$, and $\#(a_j^1, s) = 0$, and
 - $p_i(s) = 0$ otherwise.

We observe the following:

- For all of the above cases, the payoff of player i only depends on his own action and on the number of players playing certain other actions. If two players exchange actions, their payoffs are also exchanged. Hence, Γ is symmetric.
- Clearly, any action profile corresponding to a satisfying assignment of C is an equilibrium, because in this case all players receive the maximum payoff of 2.
- On the other hand, if s does not correspond to a satisfying assignment, we have one of three different cases, in none of which s is an equilibrium:
 - If $\#(b, s) < n - m$ or $\#(b, s) > n - m + 1$, then there exists a player that receives

- payoff 0 and can change his action to receive a payoff of 1.
- If $\#(a_j^0, s) = 1$ for all $j \in M$, which can only be the case if $\#(b, s) < n - m$, player i can change to some a_m^1 such that $s_i \neq a_m^0$ to increase his payoff from 1 to 2.
- Otherwise, there has to be some player $i \in N$ who gets payoff 0, and, by the pigeonhole principle, some $j \in M$ such that $\#(a_j^0, s_{-i}) = \#(a_j^1, s_{-i}) = 0$. Then, player i can change to a_j^0 to get a higher payoff.

Again, there is a direct correspondence between pure equilibria of Γ and satisfying assignments of C . The transformation from C to Γ essentially works by writing down Boolean circuits that compute p_i . Observing that this can be done in time polynomial in the size of C if $n \leq 2^k$, where $k = |A|$, completes the proof. \square

By each of the previous two theorems and by the inclusion relationships between the different classes of games, we also have the following.

Corollary 1 *Deciding whether an anonymous game has a pure Nash equilibrium is NP-complete, even if the number of actions is linear in the number of players and there is only a constant number of different payoffs.*

Since the proofs of Theorems 4 and 5 work by mapping satisfying assignments of a Boolean circuit to a *certain number* of pure equilibria of a strategic game, we can show that counting the number of pure equilibria in the above classes of games is hard.

Corollary 2 *For anonymous, symmetric, and self-anonymous games, counting the number of pure Nash equilibria is #P-hard, even if the number of actions is linear in the number of players and there is only a constant number of different payoffs.*

Proof: Recall that in the proof of Theorem 4, actions of the game Γ are identified with inputs of the Boolean circuit C . As a direct consequence of anonymity or symmetry, it does not matter which player plays a particular action to assign a value to the corresponding gate. Every satisfying assignment of C thus corresponds to $n!$ equilibria of Γ , so the number of satisfying assignments can be determined by counting the number of pure equilibria, of which there are at most $2^n n!$, and dividing this number by $n!$. Division of two m -bit binary numbers can be done using a circuit with bounded fan-in and depth $O(\log m)$ [2]. For $m = \log(2^n n!) = O(n^2)$, we have depth $O(\log n^2) = O(\log n)$, so the above division can be carried out in NC^1 . We have thus found a reduction of the problem #SAT of counting the number of satisfying assignments of C , which is #P-complete [e.g., 27], to the problem of counting the pure equilibria of Γ . The same line of reasoning applies to the proof of Theorem 5. Analogously to Corollary 1, #P-hardness extends to anonymous games. \square

As we have already outlined above, every self-symmetric game possesses a pure

equilibrium. Theorem 3 states that finding even a social-welfare-maximizing equilibrium is very easy as long as the number of actions is bounded by a constant. If now the number of actions is growing but polynomial in the size of the input, we can start at an arbitrary action profile and check in polynomial time whether some player can change his action to increase the (common) payoff. If this is not the case, we have found an equilibrium. Otherwise, we can repeat the process for the new profile, resulting in a procedure called *best-response dynamics* in game theory. Since the payoff strictly increases in each step, we are guaranteed to find an equilibrium in polynomial time if the number of different payoffs is polynomial. Conversely, we will show that, given a self-symmetric game with a growing number of actions and an exponential number of different payoffs, finding a pure equilibrium is at least as hard as finding a *locally optimal* solution to an NP-hard optimization problem. For this, we formally introduce the class of search problems for which a solution is guaranteed to exist by a local optimality argument.

Definition 4 (local search, PLS) A local search problem is given by (i) a set \mathcal{I} of instances, (ii) a set $\mathcal{F}(x)$ of feasible solutions for each $x \in \mathcal{I}$, (iii) an integer measure $\mu(S, x)$ for each $S \in \mathcal{F}(x)$, and (iv) a set $\mathcal{N}(S, x)$ of neighboring solutions for each $S \in \mathcal{F}(x)$. A solution is *locally optimal* if it does not have a strictly better neighbor, i.e., one with a higher or lower measure depending on the kind of optimization problem.

A local search problem is in the class *PLS* of polynomial local search problems [22] if for every $x \in \mathcal{I}$ there exist polynomial time algorithms for (i) computing an initial feasible solution in $\mathcal{F}(x)$, (ii) computing the measure $\mu(S, x)$ of a solution $S \in \mathcal{F}$, and (iii) determining that S is locally optimal or finding a better solution in $\mathcal{N}(S, x)$.

A problem P in *PLS* is *PLS-reducible* to another problem Q in *PLS* if there exist polynomial time computable functions Φ and Ψ mapping (i) instances x of P to instances $\Phi(x)$ of Q and (ii) solutions S of an instance $\Phi(x)$ of Q to solutions $\Psi(S, x)$ of the corresponding instance x of P such that locally optimal solutions are mapped to locally optimal solutions.

Interestingly, problems in *PLS* have a fully polynomial-time approximation scheme [26].

The proof of the following theorem works along similar lines as those of Theorems 4 and 5 to give a reduction from the *PLS*-complete problem FLIP.

Theorem 6 *The problem of finding a pure Nash equilibrium in a self-symmetric game is PLS-complete, even if the number of actions is linear in the number of players.*

Proof: Neighborhood among action profiles is given by a single player changing his action. Since the number of players and actions is polynomial in the input size,

and since the payoff function is computable in polynomial time, *membership* in PLS is immediate.

For hardness, consider a Boolean circuit C with inputs $M = \{1, 2, \dots, m\}$ and ℓ outputs. Finding an assignment such that the output interpreted as an ℓ -bit binary number is a local maximum under the FLIP neighborhood (*i.e.*, changing a single input bit) is known to be PLS-complete [22, 32]. We provide a PLS reduction to the problem of finding a pure equilibrium in a self-symmetric game by mapping a particular circuit C as described above to a game Γ with $n \geq m$ players, actions $A = \{a_j^0, a_j^1 : j \in M\}$, and a (common) payoff function p as follows:

- If s corresponds to an assignment c of C , we let $p(s) = n + C(c)$, where $C(c)$ denotes the output of C for input c , interpreted as a binary number.
- Otherwise, we let $p(s) = \min(\#(b, s), n - m) + |\{j \in M : \#(a_j^0, s) + \#(a_j^1, s) > 0\}|$. That is, the payoff is at most $n - 1$ and decreases in the minimum number of players that would have to change their action in order to make s correspond to an assignment of C .

We observe the following:

- Obviously, Γ is a common payoff game. Since p is invariant under any permutation of the players in both of the above cases, Γ is self-symmetric.
- If $n \leq 2^k$, where $k = |A|$, a Boolean circuit that computes p can be constructed from C in time polynomial in the size of C . Hence, there exists a polynomial time computable function that maps instances of FLIP to instances of the problem under consideration.
- An action profile a that does not correspond to an assignment of C cannot be an equilibrium of Γ . In this case there always exists $j \in M$ such that a_j^0 and a_j^1 are played by more than one player. At the same time, less than $n - m$ players play b , or no one plays $a_{j'}^0$ or $a_{j'}^1$ for some $j' \in M$. If one of the players playing the former changes to the latter, he gets a higher payoff (actually, all players do).
- There is a direct correspondence between the FLIP neighborhood of C and a single player changing between a_j^0 and a_j^1 for some $j \in M$. Furthermore, changing to an action profile that does not correspond to an assignment of C will get the player strictly less payoff. Thus, there is a direct correspondence between pure equilibria of Γ and local maxima of C under the FLIP neighborhood. Obviously, the assignment corresponding to an action profile can be computed in polynomial time, if such an assignment exists. The conditions of Definition 4 do not require that we map solutions of Γ that are not locally optimal to solutions of C that are not locally optimal. This means that action profiles not corresponding to an assignment can simply be mapped to an arbitrary assignment.

It is easily verified that this satisfies the properties of a PLS reduction. \square

Implicit in the definition of PLS is a *standard algorithm* for finding a locally optimal solution for a given input $x \in \mathcal{I}$: start with an arbitrary feasible solution

$S \in \mathcal{F}(x)$ and repeatedly find a strictly better neighbor until a locally optimal solution $T \in \mathcal{F}(x)$ has been found. The *standard algorithm problem* can be phrased as follows: given x , find the locally optimal solution T output by the standard algorithm on input x . Schäffer and Yannakakis [32] introduce the notion of a *tight* PLS reduction and show that tight reductions preserve both hardness of the standard algorithm problem and exponential worst-case running time of the standard algorithm.

Definition 5 (tight PLS reduction) *Let P, Q be PLS problems. A PLS reduction (Φ, Ψ) from P to Q is called tight if for any instance x of P there exists a set $\mathcal{R} \subseteq \mathcal{F}(\Phi(x))$ with the following properties:*

- (1) \mathcal{R} contains all local optima of $\Phi(x)$.
- (2) For every $p \in \mathcal{F}(x)$, a solution $q \in \mathcal{R}$ satisfying $\Psi(q, x) = p$ can be computed in polynomial time.
- (3) Consider $q_0, q_1, \dots, q_\ell \in \mathcal{F}(\Phi(x))$ such that $q_0, q_\ell \in \mathcal{R}$, $q_i \notin \mathcal{R}$ for all $0 < i < \ell$, $q_{i+1} \in \mathcal{N}(q_i, \Phi(x))$ for all $i < \ell$, and $\mu(q_i) > \mu(q_j)$ if $i > j$. Let $p = \Psi(q_0, x)$, $p' = \Psi(q_\ell, x)$. Then, either $p = p'$ or $p' \in \mathcal{N}(p, x)$.

With some extra work, we can show that the reduction used in the proof of Theorem 6 is tight, and draw additional conclusions about the standard algorithm and the standard algorithm problem.

Corollary 3 *The standard algorithm for finding pure Nash equilibria in self-symmetric games has an exponential worst-case running time. The standard algorithm problem is NP-hard.*

Proof: Johnson et al. [22] have shown that the standard algorithm for FLIP has an exponential worst-case running time, and the standard algorithm problem is NP-hard. By Lemma 3.3 of Schäffer and Yannakakis [32] it thus suffices to show that the reduction in the proof of Theorem 6 is tight. To this end, choose \mathcal{R} to be the set of actions profiles of Γ that correspond to an assignment of C . Obviously, \mathcal{R} contains all optimal solutions, and a payoff profile corresponding to a particular assignment can be computed in polynomial time. The third condition is trivially satisfied because the measure of any solution inside \mathcal{R} is strictly greater than that of any solution outside of \mathcal{R} .

By a slight modification of the proof of Theorem 6, PLS-hardness, exponential worst-case running time of the standard algorithm, and NP-hardness of the standard algorithm problem can also be shown for general (*i.e.*, not necessarily symmetric) common payoff games with $k = 2$. This fact nicely illustrates the influence of symmetry on the hardness of finding, or deciding the existence of, a pure equilibrium.

5 Threshold Symmetries

In order to extend the basic concept of symmetry as the indistinguishability of players, we will now consider games where the players cannot even observe the exact number of players playing a certain action, but only whether this number reaches certain *thresholds*. Let $\Gamma = (N, (A_i)_{i \in N}, (p_i)_{i \in N})$ be a normal-form game and A a set of actions such that $A_i = A$ for all $i \in N$. For $T \subseteq \{1, 2, \dots, n\}$, let $\sim_T \subseteq A^N \times A^N$ be defined as follows: $s \sim_T t$ if for all $a \in A$ and all $x \in T$, $\#(a, s) < x$ if and only if $\#(a, t) < x$. \sim_T naturally extends to action profiles for subsets of N . It is easily verified that for any $T \subseteq \{1, 2, \dots, n\}$, \sim_T is an equivalence relation on the set A^M of action profiles for players $M \subseteq N$. We use \sim_T to generalize Definition 2.

Definition 6 (threshold symmetry) *Let $\Gamma = (N, (A_i)_{i \in N}, (p_i)_{i \in N})$ be a normal-form game, A a set of actions such that $A_i = A$ for all $i \in N$. Let $T \subseteq \{1, 2, \dots, n\}$. Γ is called*

- *T -anonymous if $p_i(s) = p_i(t)$ for all $i \in N$ and all $s, t \in A^N$ with $s_i = t_i$ and $s_{-i} \sim_T t_{-i}$,*
- *T -symmetric if $p_i(s) = p_j(t)$ for all $i, j \in N$ and all $s, t \in A^N$ with $s_i = t_j$ and $s_{-i} \sim_T t_{-j}$,*
- *T -self-anonymous if $p_i(s) = p_i(t)$ for all $i \in N$ and all $s, t \in A^N$ with $s \sim_T t$, and*
- *T -self-symmetric if $p_i(s) = p_j(t)$ for all $i, j \in N$ and all $s, t \in A^N$ with $s \sim_T t$.*

For $T = \{1, 2, \dots, n\}$, these classes are equivalent to those of Definition 2. Moreover, we obtain *Boolean symmetry*, where payoffs only depend on the *support* of an action profile (*i.e.*, the actions that are played by at least one player), for $T = \{1\}$. In general, we call a game *threshold anonymous* (for one of the above classes) if it is T -anonymous for some T (and the corresponding class).

Obviously, the number of payoffs that need to be written down for each player to specify a general T -anonymous game is exactly the number of equivalence classes of \sim_T for action profiles of the other players. A T -anonymous game can be represented using at most $n \cdot k \cdot |A^{n-1} / \sim_T|$ numbers, where X / \sim denotes the quotient set of set X by equivalence relation \sim . For Boolean anonymity, the number of equivalence classes equals the number of k -bit binary numbers where at least one bit is 1, *i.e.*, $2^k - 1$. More generally, there cannot be more than $(|T| + 1)^k$ equivalence classes if $|T|$ is bounded by a constant (since for every action, the number of players playing this action must be between two thresholds), while for $T = \{n\}$ there are as few as $k + 1$. Hence, any T -anonymous game with constant $|T|$ is representable using space polynomial in n if $k = O(\log n)$. It does not matter if the thresholds themselves are constant or not. We are now ready to identify a class of threshold symmetric games for which the pure equilibrium problem is tractable. It should be noted that the proof technique is not limited to this particular class, but in fact applies to the larger class of games for which the kernel of the best response function

has polynomial size.

Theorem 7 *For threshold anonymous games with $k = O(\log n)$ and a constant number of thresholds, deciding the existence of a pure Nash equilibrium is in P.*

Proof: Like in the proof of Theorem 1, we provide an algorithm that checks whether there is an equilibrium in a particular equivalence class $X \in A^N / \sim_T$. Since for $k = O(\log n)$ and $|T| = O(1)$, the cardinality of A^N / \sim_T is polynomial in n , it suffices to show that the algorithm requires only polynomial time for every such set. For a particular element $X \in A^N / \sim_T$, the algorithm is again divided into two phases: (i) computing the set of best responses for each player under X , and (ii) checking whether there is a particular action profile $s \in X$ where each player plays a best response.

In the first phase, and unlike the case $T = \{1, 2, \dots, n\}$ covered by Theorem 1, the action a played by player $i \in N$ may or may not yield a different element of $A^{N \setminus \{i\}} / \sim_T$ against which a should be a best response. Instead of just looking for best responses under elements of T^N , we thus look for best responses under those of U^N , where $U = \{u \leq n : u \in T \text{ or } (u-1) \in T\}$. Since the cardinalities of both U^N and of the set of possible best responses is polynomial if $|T| = O(1)$ and $k = O(\log n)$, the first phase requires only polynomial time.

As for the second phase, we show that it can be reduced to deciding the existence of an integer flow with upper and lower bounds in a directed network with $O(2^k)$ vertices. Since this problem is in P if the number of vertices is polynomial [e.g., 16], observing that 2^k is polynomial in the size of the input if $k = O(\log n)$ completes the proof. Fix $X \in A^N / \sim_T$ and define a directed graph $G = (V, E)$ such that

$$\begin{aligned} V &= \{s, t, t'\} \cup V_1 \cup V_2, \quad V_1 = 2^A, \quad V_2 = A, \quad \text{and} \\ E &= \{s\} \times V_1 \cup \{(A', a) \in V_1 \times V_2 : a \in A'\} \cup V_2 \times \{t\} \cup \{(t, t')\}. \end{aligned}$$

Further define two functions $\ell : E \rightarrow \mathbb{N}$ and $u : E \rightarrow \mathbb{N}$ such that,

$$\begin{aligned} &\text{for all } A' \in V_1, \\ &\quad \ell(s, A') = u(s, A') = |\{i \in N : A' \text{ is the set of best responses of } i \text{ under } X\}|, \\ &\text{for all } A' \in V_1 \text{ and } a \in A', \quad \ell(A', a) = 0 \quad \text{and} \quad u(A', a) = n, \\ &\text{for all } a \in V_2, \quad \ell(a, t) = \min_{x \in X} \#(a, x) \quad \text{and} \quad u(a, t) = \max_{x \in X} \#(a, x), \quad \text{and} \\ &\quad \ell(t, t') = u(t, t') = n. \end{aligned}$$

Figure 6 shows the flow network for the game in Figure 3. Edge capacities have been computed by checking for each player if his action in the respective (shaded) action profile of Figure 3 is a best response. Observe that since this game is not only threshold anonymous but also anonymous, upper and lower bounds are the same for flow leaving vertices in $V_2 = A$.




Fig. 6. Integer flow network used in the proof of Theorem 7, example for the game of Figure 3. Edge e is labeled $(u(e), \ell(e))$.

Obviously every feasible flow from s to t' must have size n . Furthermore, the conditions for flow leaving vertices in V_1 require that there exists an assignment of actions to players such that each player plays a best response, while those for flow leaving edges in V_2 require that the resulting action profile is an element of X . It is thus readily appreciated that a flow from s to t' satisfying lower bounds ℓ and upper bounds u directly corresponds to a Nash equilibrium of Γ , and vice versa. \square

On the other hand, it is rather straightforward to modify the games defined in the proofs of Theorems 4, 5, and 6 to be Boolean if $n = k$. We obtain the following corollary.

Corollary 4 *Deciding the existence of a pure Nash equilibrium is NP-hard for threshold anonymous, threshold symmetric, and threshold self-anonymous games, even if thresholds are Boolean, the number of actions is linear in the number of players, and there is only a constant number of different payoffs. For the same classes, counting the number of pure Nash equilibria is #P-hard.*

For threshold self-symmetric games, finding a pure Nash equilibrium is PLS-hard, even if thresholds are Boolean and the number of actions is linear in the number of players.

Proof: In all constructions, we assume $n = m$ and remove the auxiliary action b . In addition to that, the self-anonymous game used in the proof of Theorem 4 is modified by letting players 1 and 2 play matching pennies on the parity of the number $|\{j \in M : \#(a_j^0, s) > 0\}|$ of 0-actions that are played by at least one player. It is easily verified that the arguments used to show the correspondence between satisfying assignments of the Boolean circuit and pure equilibria of the respective game still go through. Furthermore, the payoff of a particular player in each of these games only depends on whether certain actions are played by at least one player and, potentially, on the player's own action. \square

	$k = O(1)$	$k = O(n)$
anonymous	TC ⁰ -complete	NP-hard
self-anonymous		
symmetric	in AC ⁰	PLS-hard
self-symmetric		

Table 1
Complexity of pure Nash equilibrium in symmetric games

6 Conclusion and Future Work

In this paper, we have introduced four notions of symmetry in strategic multi-player games and investigated the computational complexity of finding pure Nash equilibria. We established that this problem is tractable for games with a constant number of actions, but intractable if the number of actions grows at least linearly in the number of players. It is worth noting that, for games with a constant number of actions, the pure equilibrium problem happens to lie in NC¹ for all symmetry types and is thus open to parallel computation. For games with an exponential number of players in which the number of actions grows sub-logarithmically, the complexity remains open. The main results are summarized in Table 1.

In future work, it would be interesting to extend our tractability results to larger classes of games. For example, games with a certain number of player types, where indistinguishability holds only for players of the same type, can be obtained by restricting Definition 2 to permutations that map players from a certain subset to players of the same set. Given a game in this class, we can construct an anonymous game with the same set of players and an action set that is the Cartesian product of the original set of actions and the set of player types. By assigning a unique minimum payoff to all actions *not* corresponding to the type of the respective player, we can ensure that players only play actions corresponding to their type in every equilibrium of the new game, effectively allowing us to distinguish players of different types in the new game. For games with a constant number of players the size of the new game is polynomial in the size of the original game, and the tractability result of Theorem 1 carries over immediately. A different notion, such that players of the same type have identical payoff functions, does not seem to provide additional structure. As we have already shown, only two different payoff functions suffice to make the pure equilibrium problem TC⁰-hard for a constant number of actions and NP-hard for a growing number of actions. More generally, one might investigate games where payoffs are invariant under particular sets of permutations. For example, von Neumann and Morgenstern [35] regard the number of permutations under which the payoffs of a game are invariant as a measure for the degree of symmetry. The question is in how far the computational complexity of solving a game depends

on the degree of symmetry.

Acknowledgements

This material is based upon work supported by the Deutsche Forschungsgemeinschaft under grants BR 2312/3-1 and BR 2312/3-2. An extended abstract of an earlier version of the paper appeared in the Proceedings of the 24th International Symposium on Theoretical Aspects of Computer Science (STACS). We thank Jan Johannsen for helpful discussions on circuit complexity and local search, and Rob Powers for introducing the first author to the ambiguity of symmetry in games. We further thank the anonymous referees for helpful comments.

References

1. Aumann, R. J., 1974. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics* 1, 67–96.
2. Beame, P., Cook, S. A., Hoover, H. J., 1986. Log depth circuits for division and related problems. *SIAM Journal of Computing* 15 (4), 994–1003.
3. Bhat, N., Leyton-Brown, K., 2004. Computing Nash equilibria of action-graph games. In: *Proceedings of the 20th Annual Conference on Uncertainty in Artificial Intelligence (UAI)*. AUAI Press, pp. 35–42.
4. Bollobás, B., 1998. *Modern Graph Theory*. Springer-Verlag.
5. Borel, E., 1921. La théorie du jeu et les équations intégrales à noyau symétrique. *Comptes Rendus de l'Académie des Sciences* 173, 1304–1308.
6. Brandt, F., Fischer, F., Holzer, M., 2007. Equilibria of graphical games with symmetries. Tech. Rep. TR07-136, *Electronic Colloquium on Computational Complexity (ECCC)*.
7. Chandra, A. K., Stockmeyer, L., Vishkin, U., 1984. Constant depth reducibility. *SIAM Journal on Computing* 13 (2), 423–439.
8. Chen, X., Deng, X., 2006. Settling the complexity of 2-player Nash-equilibrium. In: *Proceedings of the 47th Symposium on Foundations of Computer Science (FOCS)*. IEEE Press, pp. 261–272.
9. Cheng, S.-F., Reeves, D. M., Vorobeychik, Y., Wellman, M. P., 2004. Notes on equilibria in symmetric games. In: *Proceedings of the 6th International Workshop On Game Theoretic And Decision Theoretic Agents (GTDT)*.
10. Daskalakis, C., Goldberg, P., Papadimitriou, C., 2006. The complexity of computing a Nash equilibrium. In: *Proceedings of the 38th Annual ACM Symposium on the Theory of Computing (STOC)*. ACM Press, pp. 71–78.
11. Daskalakis, C., Papadimitriou, C. H., 2007. Computing equilibria in anonymous games. In: *Proceedings of the 48th Symposium on Foundations of Computer Science (FOCS)*. IEEE Computer Society Press, pp. 83–93.

12. Daskalakis, C., Schoenebeck, G., Valiant, G., Valiant, P., 2008. On the complexity of nash equilibria of action-graph games. ArXiv:0802.1604v1.
13. Fabrikant, A., Papadimitriou, C. H., Talwar, K., 2004. The complexity of pure Nash equilibria. In: Proceedings of the 36th Annual ACM Symposium on the Theory of Computing (STOC). ACM Press, pp. 604–612.
14. Fischer, F., Holzer, M., Katzenbeisser, S., 2006. The influence of neighbourhood and choice on the complexity of finding pure Nash equilibria. *Information Processing Letters* 99 (6), 239–245.
15. Gale, D., Kuhn, H. W., Tucker, A. W., 1950. On symmetric games. In: Kuhn, H. W., Tucker, A. W. (Eds.), *Contributions to the Theory of Games*. Vol. 1. Princeton University Press, pp. 81–87.
16. Garey, M. R., Johnson, D. S., 1979. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman.
17. Gottlob, G., Greco, G., Scarcello, F., 2005. Pure Nash equilibria: Hard and easy games. *Journal of Artificial Intelligence Research* 24, 195–220.
18. Govindan, S., Wilson, R., 2003. A global Newton method to compute Nash equilibria. *Journal of Economic Theory* 110 (1), 65–86.
19. Jeong, S., McGrew, R., Nudelman, E., Shoham, Y., Sun, Q., 2005. Fast and compact: A simple class of congestion games. In: Proceedings of the 20th National Conference on Artificial Intelligence (AAAI). AAAI Press, pp. 489–494.
20. Jiang, A. X., Leyton-Brown, K., 2007. Computing pure nash equilibria in symmetric action-graph games. In: Proceedings of the 22nd AAAI Conference on Artificial Intelligence (AAAI). AAAI Press, pp. 79–85.
21. Johnson, D. S., 1990. A catalog of complexity classes. In: van Leeuwen, J. (Ed.), *Handbook of Theoretical Computer Science*. Vol. A. Elsevier, Ch. 2, pp. 67–161.
22. Johnson, D. S., Papadimitriou, C. H., Yannakakis, M., 1988. How easy is local search? *Journal of Computer and System Sciences* 37, 79–100.
23. Leyton-Brown, K., Tennenholtz, M., 2003. Local-effect games. In: Proceedings of the 18th International Joint Conference on Artificial Intelligence (IJCAI). pp. 772–777.
24. Luce, R. D., Raiffa, H., 1957. *Games and Decisions: Introduction and Critical Survey*. John Wiley & Sons Inc.
25. Nash, J. F., 1951. Non-cooperative games. *Annals of Mathematics* 54 (2), 286–295.
26. Orlin, J. B., Punnen, A. P., Schulz, A. S., 2004. Approximate local search in combinatorial optimization. *SIAM Journal of Computing* 33 (5), 1201–1214.
27. Papadimitriou, C. H., 1994. *Computational Complexity*. Addison-Wesley.
28. Papadimitriou, C. H., Roughgarden, T., 2005. Computing equilibria in multi-player games. In: Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA). SIAM, pp. 82–91.
29. Parikh, R., 1966. On context-free languages. *Journal of the ACM* 13 (4), 570–581.
30. Rosenthal, R. W., 1973. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory* 2, 65–67.

31. Ruzzo, W., 1981. On uniform circuit complexity. *Journal of Computer and System Sciences* 22 (3), 365–383.
32. Schäffer, A. A., Yannakakis, M., 1991. Simple local search problems that are hard to solve. *SIAM Journal on Computing* 20 (1), 56–87.
33. Schoenebeck, G., Vadhan, S., 2006. The computational complexity of Nash equilibria in concisely represented games. In: *Proceedings of the 7th ACM Conference on Electronic Commerce (ACM-EC)*. ACM Press, pp. 270–279.
34. von Neumann, J., 1928. Zur Theorie der Gesellschaftspiele. *Mathematische Annalen* 100, 295–320.
35. von Neumann, J., Morgenstern, O., 1947. *The Theory of Games and Economic Behavior*, 2nd Edition. Princeton University Press.