Efficient Computation and Best-Response Dynamics in Anonymous Two-Action Games with Linear Utilities

Michail Fasoulakis[1], Evangelos Markakis[2,3], Ioannis Panageas[4], Christodoulos Santorinaios[2,3], and Jingming Yan[4]

[1]Royal Holloway, University of London, Egham, UK
[2]Athens University of Economics and Business, Athens, Greece
[3]Archimedes/Athena RC, Athens, Greece
[4]University of California, Irvine, USA

Abstract

The computational complexity of anonymous games has been a central theme in algorithmic game theory, with foundational contributions by Daskalakis and Papadimitriou [2015], Goldberg and Turchetta [2017], and Cheng et al. [2017] establishing that the landscape of equilibrium computation admits a PTAS. In this paper, we investigate the class of n-player anonymous games with linear utilities and two actions. While finding an equilibrium in anonymous games is PPAD-hard for arbitrary utility functions—even with a constant number of actions [Chen et al., 2015]—we show that this linear structure allows for significant algorithmic improvements. Exploiting the fact that payoffs depend solely on the first moment of the players' strategy distribution, we provide the first algorithm for computing exact Nash equilibria that runs in time polynomial in the number of players n. We achieve this by establishing a novel combinatorial characterization of the Nash equilibria. Finally, we establish a novel connection between Nash equilibria in these games and non-convex non-concave min-max optimization. Despite the technical challenges posed by this structural property, we design a sequential best-response dynamic that provably converges to an ϵ-Nash equilibrium in $\mathcal{O}(1/\epsilon)$ steps.

1 Introduction

In many real-world scenarios—from voting systems and auctions to internet routing and social dynamics—the identity of individual players is often irrelevant; what matters is the aggregate behavior of the population. This observation motivates the study of *Anonymous Games*, a fundamental class of succinctly represented games where a player's utility depends solely on their own action and the aggregate distribution of strategies by the other players, rather than their specific individual choices.

Anonymous games offer a compelling framework for modeling large-scale interactions where agents are indistinguishable from the perspective of their payoff functions. The literature on these games is extensive, reflecting their applicability to economics and operations research [Blonski, 1999, Milchtaich, 1996, Blonski, 2005, Kalai, 2005]. This wide applicability has led to a rich line of work in algorithmic game theory focusing on the computational complexity of Nash equilibria.

On the positive side, the works of Daskalakis and Papadimitriou [2007, 2008, 2009, 2015] provided a Polynomial Time Approximation Scheme (PTAS) with running time for anonymous games with a constant number k of pure strategies. Further notable works and improvements include Goldberg and Turchetta [2017], Cheng et al. [2017] (and references therein). For the specific case of $k = 2$, Daskalakis [2008] provided an efficient PTAS, while Babichenko [2013] demonstrated that best-response dynamics converge to approximate Nash equilibria provided the population size is

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

sufficiently large. On the negative side, it has been shown by Chen et al. [2015] that computing an approximate Nash equilibrium is PPAD-hard for exponentially small approximation accuracy, even for games with $k = 7$ strategies. Very recently, Liu [2026] strengthens the hardness results by establishing PPAD-completeness for computing Nash equilibria even with $k = 2$ strategies per player.

The computational intractability of Nash equilibria in general anonymous games motivates the search for restricted utility structures that allow for efficient computation. This work concentrates on *Linear Anonymous Games*, a subclass where the utility of each player is a linear function on the aggregate distribution of strategies played by others. We address the following fundamental question:

What is the computational complexity of finding a Nash equilibrium in Linear Anonymous Games with a constant number of strategies? Moreover, is it possible to design game dynamics that provably converge to a Nash equilibrium?

Our contributions. We design a polynomial time algorithm for computing a Nash equilibrium in n-player two-action anonymous games with linear utilities. An informal statement of our main contribution is given below:

Theorem 1 (Informal; Formal Version in Theorem 3)**.** *There exists a poly(n)-time algorithm for computing an (exact) Nash equilibrium in linear anonymous games with n players and two actions. The result is based on a characterization of the structure of Nash equilibria (Lemma 4).*

Linearity is a natural assumption in many economic models, approximating scenarios where marginal effects remain constant—such as simple market clearing mechanisms or basic peer effects. However, it is crucial to recognize that linearity in payoffs does not, by itself, guarantee computational tractability. A pertinent comparison can be drawn to Congestion Games [Rosenthal, 1973]. Congestion games possess a potential function, ensuring the existence of pure Nash equilibria and placing the task of finding one in the complexity class PLS (Polynomial Local Search) [Johnson et al., 1988, Fabrikant et al., 2004]. Strikingly, even when the latency functions on resources are restricted to be linear, computing a pure Nash equilibrium in general potential/congestion games remains PLS-complete [Ackermann et al., 2008]. More recently, Anagnostides et al. [2026b] establish CLS-completeness for computing (mixed) Nash equilibria in linear congestion games. These results suggest that "simple" linear structures can still encode computationally intractable optimization landscapes.

We further highlight that restricting each player to two actions does not, by itself, guarantee computational tractability. In particular, Rubinstein [2018] establishes that computing an approximate Nash equilibrium in polymatrix games is PPAD-complete, even when every player has only two actions and her expected payoff for each action is linear in the mixed strategies. A closely related example is provided by two-action threshold games [Papadimitriou and Peng, 2021], where each player's best response depends on whether the aggregate strategy of her neighbors exceeds a given threshold. Despite this simple structure, computing an approximate Nash equilibrium remains PPAD-hard. The key distinction is that these games allow player-specific interactions, whereas anonymity requires each player's utility to depend only on the aggregate strategies of all other players. Theorem 1 demonstrates that this additional structure, together with two actions and linear utilities, enables polynomial-time computation of exact Nash equilibria.

We complement our results on computing equilibria by analyzing discrete-time sequential Best response (SBR) dynamics and prove convergence to an ϵ-approximate Nash equilibrium in $O\left(1/\epsilon\right)$ iterations. Best-response dynamics – arguably one of the most well-studied game dynamics – is an iterative process where agents sequentially update their actions to the optimal strategy given the current joint action profile of the other players. This process is intrinsically guaranteed to converge to a pure-strategy Nash equilibrium in potential games [Monderer and Shapley, 1996, Rosenthal, 1973]. Proving convergence of SBR in our context was particularly challenging, as the setting can be cast as a min-max problem over a nonconvex-nonconcave landscape. For such landscapes, no natural dynamics are known to converge to a Nash equilibrium (the only result we are aware of uses a Lemke-Howson type algorithm [Daskalakis et al., 2023]). Furthermore, recent results [Bernasconi et al., 2026, Anagnostides et al., 2025, Bernasconi and Castiglioni, 2026, Anagnostides et al., 2027, Bernasconi et al., 2027] indicate that general computation in these settings is unlikely to be achievable in polynomial time. To circumvent these computational hurdles, we exploit the structural characterization of Nash equilibria in our setting to design SBR dynamics that guarantee efficient convergence. An informal statement is given below:

Theorem 2 (Informal; Formal Version in Theorem 4)**.** *Suppose that players employ SBR dynamics (defined in Algorithm 4). After at most $O(1/\epsilon)$ number of iterations, SBR dynamics reaches an ϵ approximate Nash equilibrium.*

Technical Overview. As we study games with 2 actions, we introduce a mixing threshold (τ_i) for each player i, as the specific aggregate mass on strategy 1 that makes player i indifferent between her two choices. Based on how their preferences change relative to this threshold, players are categorized as *maximizers* (incentives aligned with the population) or *minimizers* (incentives opposed to the population). We further develop a "self-inclusion" function ϕ that maps the total mass on a strategy to a player's best-response probability. By analyzing the properties of ϕ, we establish a structural theorem (Lemma 4) showing that an equilibrium exists where at most one minimizer uses a mixed strategy. This characterization reduces the search for a Nash equilibrium to a combinatorial search over $O(n^3)$ possible strategy supports, leading to an $O(n^4)$ algorithm.

The second main contribution of the paper is the design of SBR dynamics that converge to a Nash equilibrium (Theorem 4). Because the min-max formulation is inherently nonconvex-nonconcave, standard optimization techniques are insufficient to establish convergence. This limitation is explicitly demonstrated in Example 1. We circumvent this challenge by reformulating the equilibrium search as a fixed point problem over the self-inclusion functions. Specifically, we evaluate a candidate aggregate mass S by passing it through every player's self-inclusion function $\phi_i(\cdot)$ to determine their individual best-responses. A Nash equilibrium is achieved only when the resulting aggregate mass aligns perfectly with the proposed value S. By leveraging the structural characterization (Lemma 4), we collapse the search space from exponential size to $O(n)$. This allows the proposed SBR dynamics to reach an ϵ-approximate equilibrium in at most $O(1/\epsilon)$ iterations.

2 Preliminaries

Notation. We use boldface letters (i.e., $\boldsymbol{x}, \boldsymbol{y}$) to denote vectors, and we let x_i represent the ith coordinate of vector $\boldsymbol{x}$. We define $\|\cdot\|_2$, $\|\cdot\|_1$, and $\|\cdot\|_\infty$ to be the Euclidean norm, the ℓ_1 norm, and ℓ_∞ norm respectively. Unless specified otherwise, we abbreviate the Euclidean norm as $\|\cdot\|$. We use the standard asymptotic notation $O(\cdot), \Theta(\cdot)$, and $\Omega(\cdot)$ to suppress absolute constants.

Anonymous Games. Our work focuses on the class of anonymous games with two actions, henceforth referred to as actions or strategies 0 and 1. In these games, the payoff of a player depends on her own action and the distribution of actions taken by the other players; that is, the payoff is independent of the specific identities of the players selecting each strategy. Formally, an anonymous game is defined as a tuple $G = (\mathcal{N}, \mathcal{S}, \mathcal{U})$, where $\mathcal{N} = [n] = \{1, \ldots, n\}$ denotes the set of players, $\mathcal{S}$ denotes the set of common actions among the players strategies – thus $\mathcal{S} = \{0, 1\}$ in our case – and $\mathcal{U} = \{u_i\}_{i \in \mathcal{N}}$ is the set of the utility functions of the players.

Given a game G, a pure strategy profile $\boldsymbol{s}$ is a binary vector of size n, specifying the action chosen by each player. For any player i, we denote by $\boldsymbol{s}_{-i}$ the strategy profile of all other players except player i. Because the game is anonymous, the payoff of player i depends solely on the total number of other players selecting action 1. We denote this aggregate count by $c_{-i}(\boldsymbol{s}_{-i})$, which is formally defined as $c_{-i}(\boldsymbol{s}_{-i}) = \sum_{j \in \mathcal{N} \setminus i} s_j$.

With the above notation, we can express the utility of a player i in a profile $\boldsymbol{s}$ solely as a function dependent on her binary action s and the count $c_{-i} \in [0, n-1]$. Assuming the standard normalization of utilities in $[0, 1]$, we have that $u_i : \{0, 1\} \times [0, n-1] \to [0, 1]$. Thus, the game is *succinctly representable* (in the sense of Papadimitriou and Roughgarden [2008]) by $O(n^2)$ numbers.

Linear Utilities. In this paper, we work on the subclass of anonymous games with linear utilities. For the two-action case, the utility function of player i takes the form

$$u_i(s, \boldsymbol{s}_{-i}) = a_{i,s} \cdot c_{-i}(\boldsymbol{s}_{-i}) + b_{i,s},$$

where $a_{i,s}$ and $b_{i,s}$ are rational coefficients for each player $i \in \mathcal{N}$ and pure strategy $s \in \{0, 1\}$.

The players can also randomize and use mixed strategies, which are represented as Bernoulli probability distributions. We use p_i to denote the probability that player i selects strategy 1 and the vector $\boldsymbol{p} = (p_1, \ldots, p_n) \in [0, 1]^n$ to denote a mixed strategy profile. For any player i, the mixed strategy

profile of the other players is given by $\boldsymbol{p}_{-i}$. By the linearity of expectation, the expected payoff of player i, when she selects pure strategy s and the remaining players play according to $\boldsymbol{p}_{-i}$, is:

$$\mathbb{E}\left[u_i(s, \boldsymbol{p}_{-i})\right] = a_{i,s} \cdot c_{-i}(\boldsymbol{p}_{-i}) + b_{i,s}.$$

Here, the expectation is taken over the distribution of the aggregate strategies of the other players. We naturally extend the definition of $c_{-i}(\cdot)$ to handle mixed strategy profiles, i.e., $c_{-i}(\boldsymbol{p}_{-i}) = \sum_{j \in \mathcal{N} \setminus i} p_j$. Finally, if player i employs a mixed strategy p_i, her expected payoff becomes:

$$\mathbb{E}\left[u_i(p_i, \boldsymbol{p}_{-i})\right] = p_i \cdot \mathbb{E}\left[u_i(1, \boldsymbol{p}_{-i})\right] + (1 - p_i) \cdot \mathbb{E}\left[u_i(0, \boldsymbol{p}_{-i})\right].$$

Note that the above expression depends entirely on p_i and $c_{-i}(\boldsymbol{p}_{-i}) = \sum_j p_j - p_i$. Hence, each player's utility depends only on her action and the first moment (expectation) of the cumulative distribution.

Nash Equilibrium. A strategy profile $\boldsymbol{p}^*$ constitutes an ϵ-approximate Nash Equilibrium when no player can increase her expected utility by more than ϵ through a unilateral deviation. Formally,

$$\text{for all } i \in [n] \text{ and } s \in \{0, 1\}: \quad \mathbb{E}\left[u_i(p_i^*, \boldsymbol{p}_{-i}^*]\right) \geq \mathbb{E}\left[u_i(s, \boldsymbol{p}_{-i}^*)\right] - \epsilon.$$

Here, the expectations are taken over the mixed strategies of the players as defined previously. When $\epsilon = 0$, the profile $\boldsymbol{p}^*$ is referred to as an (exact) Nash Equilibrium.

3 A polynomial time algorithm for two actions via characterization of equilibria

We start our technical exposition by focusing on the computation of equilibria in two-action anonymous games. Since we will frequently need to compare the relative utilities of these two actions, we introduce the following definition.

Definition 1. *Denote the difference in expected payoffs between the two pure actions for player i as:*

$$\Delta u_i(\boldsymbol{p}_{-i}) \triangleq \mathbb{E}\left[u_i(1, c_{-i}(\boldsymbol{p}_{-i})\right] - \mathbb{E}\left[u_i(0, c_{-i}(\boldsymbol{p}_{-i})\right] = a_i \cdot c_{-i}(\boldsymbol{p}_{-i}) - b_i,$$

where $a_i = a_{i,1} - a_{i,0}$ and $b_i = b_{i,0} - b_{i,1}$.

Notably, the sign of Δu_i governs the behavior of player i, regarding her best response to $\boldsymbol{p}_{-i}$. This motivates the following definition (the specific case where $a_i = 0$ is deferred to a later discussion).

Definition 2 (Mixing Threshold)**.** *For each player i, we define the unique root of Δu_i as the mixing threshold*

$$\tau_i \triangleq \frac{b_i}{a_i}.$$

From Definition 1, player i is indifferent between actions 0 and 1 when the expected aggregate count of the other players choosing action 1 is exactly equal to τ_i. Under this condition, she may play a mixed strategy. In any other case where $c_{-i}(\boldsymbol{p}_{-i}) \neq \tau_i$, she has a unique pure best response. Notice that whether this pure best response is strategy 0 or 1 depends entirely on the sign of a_i. This leads to the following insightful definition, where we classify players based on the slope of their utility difference function $\Delta u_i(\boldsymbol{p}_{-i})$.

Definition 3 (Player types)**.** *We partition the set of players into two disjoint groups $\mathcal{N} = \mathcal{N}_{\min} \cup \mathcal{N}_{\max}$, defined as follows:*

- *minimizers: if $a_i < 0$ then player $i \in \mathcal{N}_{\min}$,*
- *maximizers: if $a_j > 0$ then player $j \in \mathcal{N}_{\max}$.*

We provide a formal justification for this naming convention in the subsequent section. For now, we offer the following intuition: suppose player i observes an increase in the expected aggregate count $c_{-i}(\boldsymbol{p}_{-1})$. If player i is a maximizer, this increase incentivizes her to shift more probability mass toward action 1, thereby further amplifying the expected total count of action 1. Conversely, if player i is a minimizer, she will respond to the increase by reducing her probability of playing action 1, effectively counteracting the total count.

Players with a dominant strategy. Thus far, we have excluded the case where $a_i = 0$, explicitly in Definition 3 and implicitly in Definition 2. Observe that if $a_i = 0$ for some player i, her expected utility difference depends solely on the constant $-b_i$. As a result, she either has a dominant strategy (based on the sign of b_i), or she is a degenerate player (if $b_i = 0$). In either scenario, we can fix the action of player i to a pure strategy (either her dominant best response or an arbitrary pure action if $b_i = 0$) and remove her from the analysis of the game. Furthermore, a player whose mixing threshold falls outside the feasible range of aggregate counts (i.e., $\tau_i \leq 0$ or $\tau_i \geq n-1$) also has a dominant strategy and can be removed from the game. Therefore, for the remainder of the paper, we may assume that the thresholds are well-defined and for all players it holds that $\tau_i \in (0, n-1)$.

Since we will often need to refer to the number of players of each group separately in the following sections, we slightly change the notation: we use $n_y = |\mathcal{N}_{\max}|$ to refer to the number of maximizers, and respectively $n_x = |\mathcal{N}_{\min}|$ for the minimizers and $n = n_x + n_y$. Moreover, we use $\boldsymbol{y}$ for the strategy profile of the maximizers and $\boldsymbol{x}$ for the minimizers. Analogously, we define $\boldsymbol{y}_{-i}, \boldsymbol{x}_{-i}$, and we set $y_{-i} = \sum_{j \in \mathcal{N}_{\max}} y_j - y_i$ and $x_{-i} = \sum_{j \in \mathcal{N}_{\min}} x_j - x_i$. Last but not least, if a threshold τ belongs to a maximizer y_i, we denote it by τ_i^y and if it belongs to a minimizer x_i, we denote it by τ_i^x.

Next, we analyze anonymous games with a single player type. The detailed analysis is included in Section A in the Appendix. This analysis provides an important foundation for characterizing the structure of Nash equilibria in general two-action anonymous games in the next subsection. For games with only maximizers, we make the following observation:

Observation 1 (Maximizers-only instance)**.** *A linear anonymous game with two actions and only maximizers (as defined in Definition 3) admits (at least) two pure Nash equilibria: the all-1 and the all-0 strategy profiles.*

Based on the above observation, for each maximizer, we define a best-response function as a function of the total mass assigned to action 1. We refer to this mapping as *self-inclusion best response*. Specifically, given a possibly mixed strategy profile, let S be the total expected count of players choosing strategy 1. Then,

$$\phi_i^{\max}(S) \triangleq \begin{cases} 1 & S > \tau_i^y + 1 \\ S - \tau_i^y & S \in [\tau_i^y, \tau_i^y + 1] \\ 0 & S < \tau_i^y \end{cases}. \qquad \text{(Maximizer Self-inclusion BR)}$$

We further define the *aggregate self-inclusion best response*: $\phi^{\max}(S) \triangleq \sum_{i \in \mathcal{N}_{\max}} \phi_i^{\max}(S)$. Similarly, for games with only minimizers, we define the following self-inclusion best-response correspondence:

$$\phi_i^{\min}(S) \triangleq \begin{cases} 0 & S > \tau_i^x + 1 \\ \{0, S - \tau_i^x, 1\} & S \in [\tau_i^x, \tau_i^x + 1] \\ 1 & S < \tau_i^x \end{cases}. \qquad \text{(Minimizer Self-inclusion BR)}$$

We note that the main difference between the above best-response correspondence and (Maximizer Self-inclusion BR) is that the former is not single-valued. A detailed discussion is provided in Section A in the Appendix.

3.1 Characterization of equilibria in anonymous games with two actions

In this subsection, we aim to show that there exists always an equilibrium such that at most one minimizer i plays a mixed strategy $x_i \in (0, 1)$. All minimizers with larger thresholds play 1 and those with smaller thresholds play 0. In the sequel, we assume that the minimizers' thresholds are sorted such that $\tau_1^x \geq \tau_2^x \geq \cdots \geq \tau_{n_x}^x$ and, for the ease of presentation, we define the boundary cases $\tau_0^x = n + m$ and $\tau_{n_x+1}^x = 0$. Next, we define the function $H(S) = S - \phi^{\max}(S)$. We make the following observation:

Observation 2. *There exists a Nash equilibrium with total mass S^* if and only if for any minimizer i there exists some $x_i \in \phi_i^{\min}(S^*)$ such that $H(S^*) = \sum_{i \in \mathcal{N}_{\min}} x_i$.*

The straightforward proof is deferred to Section B. Building on this observation, we can characterize the conditions under which there are either no minimizers or exactly one minimizer who plays a mixed

strategy. In particular, we establish the following sufficient condition under which no minimizer plays a mixed strategy.

Lemma 1 (Zero mixing minimizers criterion)**.** *Assume that there exists a k with $S \in [\tau^x_{k+1}, \tau^x_k + 1]$ such that $H(S) = k$. Then, there is an equilibrium where exactly k minimizers play 1 and the remaining $n_x - k$ play 0.*

The proof follows directly from the definition of (Minimizer Self-inclusion BR) and is provided in Section B. Moreover, in the case where exactly one minimizer plays a mixed strategy, we establish the following technical lemmas:

Lemma 2 (Mixing Equilibrium Condition)**.** *Define $C_k = \tau^x_k - (k-1)$. Suppose there exists a k such that $C_k \in (\phi^{\max}(\tau^x_k), \phi^{\max}(\tau^x_k + 1))$. It holds that an equilibrium exists with $x_k \in (0, 1)$ being the sole mixer among the minimizers. All players with larger thresholds play 1 and all with smaller thresholds play 0.*

Lemma 3 (Mixing Failure)**.** *Assume that Lemma 2 does not hold for some index k, i.e. $C_k \notin (\phi^{\max}(\tau^x_k), \phi^{\max}(\tau^x_k + 1))$. Then, equivalently,*

$$H(\tau^x_k) \leq k - 1, \text{ or } H(\tau^x_k + 1) \geq k.$$

The proof leverages the fact that, for a minimizer x_k to play a mixed strategy, it must hold that $x_k = S - \tau^x_k$. We refer the reader to the Appendix for further details. We are now in a position to state the following technical lemma:

Lemma 4 (Characterization of Nash equilibrium for the minimizers)**.** *There exists an equilibrium where at most one minimizer plays a mixed strategy.*

The proof is based on a combinatorial argument showing that either the condition stated in Lemma 1 or the condition characterized in Lemma 2 holds. The detailed proof is deferred to Section B.

3.2 The algorithm

Given the characterization, designing the algorithm is fairly straightforward. We will essentially brute force the identification of the players into three groups. We use the notation I_1, I_0, I_{mix} for the set of players that play action 1 with probability 1, 0 or mixed respectively. So the last ingredient we need is how to solve the problem given the support. Algorithm 2 (which is deferred to the Appendix together with Algorithm 1) handles this matter. We must address that Algorithm 2 does not differentiate between the minimizers and the maximizers when players from both groups mix, hence we use τ_i instead of τ^x_i, τ^y_i.

Recall that for any player i who mixes, regardless of her type, her mixing probability must equal $S - \tau_i$. In Algorithm 2, we exploit this property to obtain a closed-form expression for the mixing probabilities of the players in I_{mix}. If the solution to this system of linear equations results in a value outside the interval $[0, 1]$, we infer that the configuration is infeasible. Moreover, we have the following lemma:

Lemma 5. *Given a feasible support of a Nash equilibrium of a linear anonymous game with n players and two actions, Algorithm 2 constructs the profile in $O(n)$ time. If the support is infeasible, the algorithm identifies this in the same time.*

Now we are able to provide the pseudocode of the algorithm (Algorithm 3) that computes Nash equilibria in anonymous games with two actions. For ease of presentation, we do not describe in detail the pre-processing step that identifies the type of a player and computes her threshold.

In brief, Algorithm 3 operates as follows: it iterates over all possible pairs (l, r) representing the leftmost and rightmost mixing maximizers, respectively. For each interval $[l, r]$, the algorithm identifies the identity of the mixing minimizer k (if one exists). These agents are assigned to the set I_{mix}, while the remaining agents are partitioned into I_0 and I_1 based on our structural characterization. Finally, each configuration $(I_0, I_1, I_{\text{mix}})$ is validated or discarded via Algorithm 2. The following theorem provides a formal guarantee for Algorithm 3:

Theorem 3. *Algorithm 3 computes a Nash equilibrium in linear anonymous games with n players and two actions in $O(n^4)$ time.*

Algorithm 3: An $O(n^4)$ algorithm for anonymous two-action games with linear utilities

Pre-process player types and sort the thresholds.
If every player is a maximizer **then return** $(1, \ldots, 1)$.
If every player is a minimizer **then return** Algorithm 1.
for $l \in [1, n_y + 1]$ **do**
  for $r \in [l-1, n_y]$ **do**
    $K_0 = \{i \in \mathcal{N}_{\max} : i < l\}$.
    $K_1 = \{i \in \mathcal{N}_{\max} : i > r\}$.
    $I_{\text{mix}} = \{i \in \mathcal{N}_{\max} : i \in [l, r]\}$.
    for $k \in \mathcal{N}_{\min}$ **do**
      $I_0 = K_0 \cup \{i \in \mathcal{N}_{\min} : i > k\}$.
      $I_1 = K_1 \cup \{i \in \mathcal{N}_{\min} : i < k\}$.
      if $(\boldsymbol{x}, \boldsymbol{y}) =$ Algorithm 2$(I_0, I_1 \cup \{k\}, I_{\text{mix}})$ is **not** "Infeasible" **then return** $(\boldsymbol{x}, \boldsymbol{y})$.
      if $(\boldsymbol{x}, \boldsymbol{y}) =$ Algorithm 2$(I_0 \cup \{k\}, I_1, I_{\text{mix}})$ is **not** "Infeasible" **then return** $(\boldsymbol{x}, \boldsymbol{y})$.
      if $(\boldsymbol{x}, \boldsymbol{y}) =$ Algorithm 2$(I_0, I_1, I_{\text{mix}} \cup \{k\})$ is **not** "Infeasible" **then return** $(\boldsymbol{x}, \boldsymbol{y})$.

Remark 1 (Rational Nash Equilibrium and FPTAS)**.** *Theorem 3 addresses the exact computation of Nash equilibria. Analogous to two-player normal-form games, linear anonymous games admit a rational Nash equilibrium. This property arises because the expected utility of each player i is bilinear with respect to the mixed strategy vector $\mathbf{p}$; by fixing the set of players using pure strategies, those using mixed strategies must satisfy a system of linear constraints.*

Furthermore, because the first moment of the Poisson binomial distribution characterizing the agents' strategies uniquely determines the utilities, the approach of Daskalakis and Papadimitriou [2015] yields a FPTAS. The latter is true even if one assumes anonymous games with polynomial utilities of constant degree d; in that case the FPTAS will run in time poly $\left(n, \frac{1}{\epsilon^d}\right)$.

Moreover, by using prefix sums of the players' thresholds, we can further improve the running time of Theorem 3, as stated in the following corollary.

Corollary 1. *An exact Nash equilibrium in a two-action anonymous game with n players and linear utilities can be computed in $O(n^3)$ arithmetic operations.*

4 Connection to min-max optimization

In this section we highlight an interesting connection between two-action anonymous games with linear utilities and min-max optimization. For any joint strategy profile $\boldsymbol{p}$, consider the function:

$$\Phi(\boldsymbol{p}) = \frac{1}{2} \sum_i \sum_{j \neq i} p_i p_j - \sum_i \frac{b_i}{a_i} p_i,$$

where p_i denotes the probability that the ith player plays action 1. Evaluating the change in utility when player x_i switches from action 1 to 0, we obtain

$$\Phi(1, \boldsymbol{p}_{-i}) - \Phi(0, \boldsymbol{p}_{-i}) = \Big(\sum_{j \neq i} p_j - \sum_i \frac{b_i}{a_i} \Big) + \sum_{j \neq i} \frac{b_i}{a_i} = \sum_{j \neq i} p_j - \frac{b_i}{a_i} = \frac{1}{a_i} \Delta u_i(\boldsymbol{p}_{-i}),$$

where $\Delta u_i(\boldsymbol{p}_{-i})$ is defined in Definition 1. Since p_i captures the probability of playing action 1, we conclude that players with $a_i > 0$ wish to maximize $\Phi(\boldsymbol{p})$, whereas players with $a_i < 0$ wish to minimize it. This formulation justifies the naming convention we introduced in the previous section.

Following from the notations we introduce in Section 3, we write the joint strategy profile as $\boldsymbol{p} = (\boldsymbol{x}, \boldsymbol{y})$, where $\boldsymbol{x}$ represents the strategy profile of the minimizers and $\boldsymbol{y}$ represents the strategy profile of the maximizers. This yields the following min-max optimization task

$$\min_{\boldsymbol{x}} \max_{\boldsymbol{y}} \Phi(\boldsymbol{x}, \boldsymbol{y}) = \frac{1}{2} \sum_i \sum_{j \neq i} x_i x_j + \sum_i \sum_j x_i y_j + \frac{1}{2} \sum_i \sum_{j \neq i} y_i y_j - \sum_i \tau_i^x x_i - \sum_j \tau_j^y y_j, \quad (1)$$

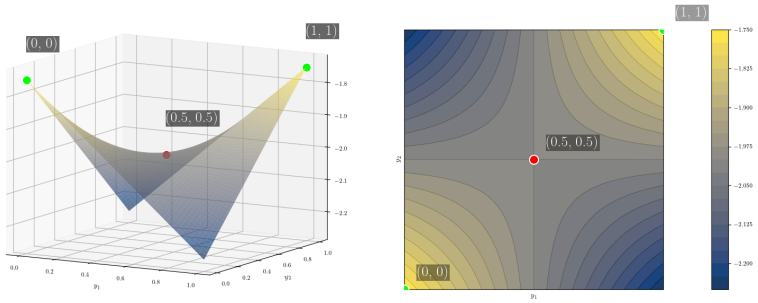

Figure 1: Value plots for $\Phi(\boldsymbol{x}, \boldsymbol{y})$ in Example 1 where $\boldsymbol{x} = (1/2, 1/2)$. There are two maximizers $\boldsymbol{y} = (0, 0)$ and $\boldsymbol{y}' = (1, 1)$. These maximizers induce opposite gradients for $\boldsymbol{x}$: $\nabla_{\boldsymbol{x}}\Phi(\boldsymbol{x}, \boldsymbol{y}) < 0$ and $\nabla_{\boldsymbol{x}}\Phi(\boldsymbol{x}, \boldsymbol{y}') > 0$. As a result, $(1/2, 1/2)$ is a first-order stationary point of the max function $\Phi^\star(\boldsymbol{x})$.

where τ_i^x and τ_j^y are defined in Definition 2. For the ith minimizer x_i, notice that the gradient is

$$\nabla_{x_i}\Phi(\boldsymbol{x}, \boldsymbol{y}) = \sum_{j \neq i} x_j + \sum_{k=1}^{n_y} y_k - \tau_i^x = \frac{1}{a_i}(u_{x_i}(1, \boldsymbol{x}_{-i}, \boldsymbol{y}) - u_{x_i}(0, \boldsymbol{x}_{-i}, \boldsymbol{y})).$$

Thus, a positive utility difference $u_{x_i}(1, \boldsymbol{x}_{-i}, \boldsymbol{y}) - u_{x_i}(0, \boldsymbol{x}_{-i}, \boldsymbol{y})$ yields a negative gradient $\nabla_{x_i}\Phi(\boldsymbol{x}, \boldsymbol{y}) < 0$ (since $a_i < 0$ for x_i), which in turn drives x_i towards 1. A similar argument holds for the maximizer y_j, with the signs reversed.

It is straightforward to verify that $\Phi(\boldsymbol{x}, \boldsymbol{y})$ is jointly nonconvex-nonconcave. In general, nonconvex-nonconcave optimization is PPAD-hard [Daskalakis et al., 2021, Bernasconi and Castiglioni, 2026]. Nevertheless, existing positive results leverage the hidden structures of the objective to bypass the computational hardness and establish convergence guarantees [Kalogiannis et al., 2024].

One initial approach is to find a Nash equilibrium through seeking a first-order stationary point for the min-max optimization problem formulated in (1). Standard min-max optimization techniques typically establish convergence by firstly identifying a first-order stationary point of the max function, $\Phi^\star(\boldsymbol{x}) = \max_{\boldsymbol{y} \in [0,1]^{n_y}} \Phi(\boldsymbol{x}, \boldsymbol{y})$ [Jin et al., 2020], and then extending it to a Nash equilibrium via Danskin's theorem [Anagnostides et al., 2026a]. However, the following example demonstrates that even if we successfully find a first-order stationary point of $\Phi^\star(\boldsymbol{x})$, it generally cannot be extended to a Nash equilibrium in two-action anonymous games.

Example 1. *Let G be a four-player anonymous game containing two minimizers x_1, x_2 with thresholds $\tau_1^x = \tau_2^x = 2$ and two maximizers y_1, y_2 with $\tau_1^y = \tau_2^y = \frac{3}{2}$.*

In a brief summary, at $\boldsymbol{x} = (1/2, 1/2)$, there are two maximizers $\boldsymbol{y} = (0, 0)$ and $\boldsymbol{y}' = (1, 1)$ which induce opposite gradients for $\boldsymbol{x}$. Thus, $\boldsymbol{x} = (1/2, 1/2)$ is a stationary point of the max function $\Phi^\star(\boldsymbol{x})$. However, $\boldsymbol{x}$ cannot be extended to a Nash equilibrium as there is no equilibrium profile $(\boldsymbol{x}^*, \boldsymbol{y}^*)$ with $\boldsymbol{x}^* = \boldsymbol{x}$. The detailed analysis is deferred to Section C in the Appendix. The preceding example serves as a strong indicator that, despite the elegant connection to min-max optimization, the inherently nonconvex-nonconcave nature of the formulation presents significant computational hurdles. As demonstrated, even successfully identifying an exact first-order stationary point of the max function $\Phi^\star(\boldsymbol{x})$ does not guarantee the recovery of a Nash equilibrium.

Consequently, treating the anonymous game purely as a black-box min-max optimization task is insufficient for establishing positive algorithmic results. In order to design dynamics that converge to a Nash equilibrium, it is necessary to incorporate additional structural characterizations of the equilibrium to constrain the search space.

5 Sequential best-response dynamics

In this section we propose sequential best-response (SBR) dynamics for finding approximate Nash equilibria in two-action anonymous games with linear utilities. Recall that, given the sum S, the

best-responses correspondence of a minimizer x_i is denoted as $\phi_i^{\min}(S)$. Moreover, the best-response function for a maximizer y_i is given by $\phi_i^{\max}(S)$.

Notice that in any Nash equilibrium $(\boldsymbol{x}^*, \boldsymbol{y}^*)$, if we let $S^* = \sum_{i=1}^{n_x} x_i^* + \sum_{j=1}^{n_y} y_j^*$ denote the sum of the probabilities, it holds that $x_i^* \in \phi_i^{\min}(S^*)$ and $y_j^* = \phi_j^{\max}(S^*)$. Consequently, the task of finding a Nash equilibrium can be reduced to finding the sum S^* and the corresponding $x_i^* \in \phi_i^{\min}(S^*)$ such that $\sum_{i=1}^{n_x} x_i^* = S^* - \sum_{j=1}^{n_y} \phi_j^{\max}(S^*)$. The technical challenge, however, is that for any given sum S, the best-response definition of the minimizers allows each x_i to take up to three possible values, which renders the search space exponential.

We can bypass this technical difficulty by leveraging a key observation from Lemma 4: there exists a Nash equilibrium where at most one minimizer x_i^* plays a mixed strategy. For any sum S, let $k = \lfloor S - \sum_{i=1}^{n_y} \phi_i^{\max}(S) \rfloor$. Assuming the indices are sorted in decreasing order of the thresholds τ_i^x, we can uniquely determine the strategy profile for the minimizers. Specifically, the first k minimizers play the pure strategy 1, the $(k+1)^{th}$ minimizer plays a mixed strategy with probability $S - \sum_{i=1}^{n_y} \phi_i^{\max}(S) - k$, and the remaining minimizers play 0.

We define the following best-response functions for the minimizers and maximizers respectively, which will be utilized in Algorithm 4. Let $L = \sum_{j \neq i} x_j + \sum_{k=1}^{n_y} y_k$. For the minimizers and the maximizers, we define

$$BR_i^{\min}(S, \boldsymbol{x}, \boldsymbol{y}) \triangleq \begin{cases} 1 & S > L + 1 \\ S - \sum_{j \neq i} x_j - \sum_{k=1}^{n_y} y_k & S \in [L, L+1] \\ 0 & S < L \end{cases},$$

$$BR_i^{\max}(S, \boldsymbol{x}, \boldsymbol{y}) = \begin{cases} 1 & S > \tau_i^y + 1 \\ S - \tau_i^y & S \in [\tau_i^y, \tau_i^y + 1] \\ 0 & S < \tau_i^y \end{cases}.$$

Algorithm 4: SBR dynamics

Set $S_0 = 0$ and **sort** τ_i^x, τ_j^y in decreasing order.
for $t = 1 \dots T$ **do**
  initialize $\boldsymbol{x}_t = \boldsymbol{0}, \boldsymbol{y}_t = \boldsymbol{0}$.
  $S_t = S_{t-1} + n/T$.
  for $i = 1 \dots n_y$ **do** $y_{t,i} = BR_i^{\max}(S_t, \boldsymbol{x}_t, \boldsymbol{y}_t)$.
  for $j = 1 \dots n_x$ **do** $x_{t,j} = BR_j^{\min}(S_t, \boldsymbol{x}_t, \boldsymbol{y}_t)$.

We are now ready to present the proposed dynamics to an approximate Nash equilibria. We note that the sequential update routine in Algorithm 4 ensures that at most one minimizer x_i adopts a mixed strategy, a formal proof is provided in Lemma D.3. Our main result follows.

Theorem 4. *Suppose $\boldsymbol{x}$ and $\boldsymbol{y}$ follow the SBR dynamics defined in Algorithm 4. By setting $T \geq 2n^{7/2}/\epsilon$, there exists an iteration $t \in [T]$ such that $(\boldsymbol{x}_t, \boldsymbol{y}_t)$ is an ϵ-approximate Nash equilibrium.*

6 Experiments

We complement our theoretical results with an experimental evaluation of the dynamics we analyzed in Section 5. We have implemented and tested Algorithm 4 on games of various sizes. To evaluate the performance of our algorithm, we have used as a metric the accuracy of the approximation, i.e., ϵ. The main take-away message of our experiments is that the empirical performance of Algorithm 4 is consistent with the theoretical $O(1/T)$ best-iterate rate, as dictated by Theorem 4.

Our evaluation was mostly with randomly generated examples. Note that in order to specify a game, it suffices to determine the thresholds of the players. Therefore we generated thresholds uniformly at random, for different sizes of the game population. For all the experiments with random games, we used up to $T = 10^5$ iterations and terminated the algorithm once the accuracy was below 10^{-9}.

In Figure 2, we examine the performance for a game with 120 minimizers and 80 maximizers, as well as for a game with 400 minimizers and 600 maximizers. In both subfigures of Figure 2, we can clearly see the $O(1/T)$ best-iterate rate. Additional experiments can also be found in Appendix E.

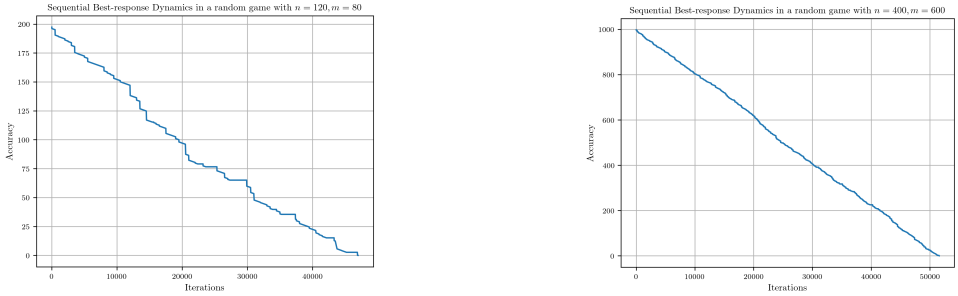

Figure 2: Sequential Best-response Dynamics in large random games.

7 Conclusion

In this work, we developed a polynomial-time algorithm for computing exact Nash equilibria in two-action anonymous games with linear utilities. Additionally, we proposed sequential best-response dynamics that converge to approximate Nash equilibria in $O(1/\epsilon)$ iterations. One important question that remains open is whether efficient algorithms can be designed to find exact Nash equilibrium in linear anonymous games with $k > 2$ strategies, a problem we conjecture to be PPAD-hard.

Acknowledgments and Disclosure of Funding

This work has been partially supported by the project MIS 5154714 of the National Recovery and Resilience Plan "Greece 2.0" funded by the European Union under the NextGenerationEU Program. It has also been supported by Input Output (iog.io) through their funding of the Athena Blockchain Technology Lab. Ioannis Panageas and Jingming Yan were supported by National Science Foundation grant CCF-2454115.

Competing interests disclosure. The authors declare no competing interests.

References

Heiner Ackermann, Heiko Röglin, and Berthold Vöcking. On the impact of combinatorial structure on congestion games. *Journal of the ACM (JACM)*, 55(6):1–22, 2008.

Ioannis Anagnostides, Ioannis Panageas, Tuomas Sandholm, and Jingming Yan. The complexity of symmetric equilibria in min-max optimization and team zero-sum games. *Advances in Neural Information Processing Systems (NeurIPS)*, 2025.

Ioannis Anagnostides, Fivos Kalogiannis, Ioannis Panageas, Emmanouil-Vasileios Vlatakis-Gkaragkounis, and Stephen McAleer. Algorithms and complexity for computing nash equilibria in adversarial team games. *Games and Economic Behavior*, 2026a.

Ioannis Anagnostides, Ioannis Panageas, and Jingming Yan. The complexity of nash equilibrium in network congestion and coordination games, 2026b. URL https://arxiv.org/abs/2609.33225.

Ioannis Anagnostides, Ioannis Panageas, Tuomas Sandholm, and Jingming Yan. The computational complexity of team zero-sum games. In *Symposium on Discrete Algorithms (SODA)*, 2027.

Yakov Babichenko. Best-reply dynamics in large binary-choice anonymous games. *Games and Economic Behavior*, 81:130–144, 2013.

Martino Bernasconi and Matteo Castiglioni. The complexity of min-max optimization with product constraints. In *Symposium on Theory of Computing (STOC)*, 2026.

Martino Bernasconi, Matteo Castiglioni, Andrea Celli, and Gabriele Farina. Improved hardness results for min-max optimization with coupled constraints. In *Conference on Economics and Computation (EC)*, 2026.

Martino Bernasconi, Matteo Castiglioni, Andrea Celli, and Alexandros Hollender. The complexity of min–max optimization for quadratic polynomials. In *Symposium on Discrete Algorithms (SODA)*, 2027.

Matthias Blonski. Anonymous games with binary actions. *Games and Economic Behavior*, 28(2): 171–180, 1999.

Matthias Blonski. The women of Cairo: Equilibria in large anonymous games. *Journal of Mathematical Economics*, 41(3):253–264, 2005.

Xi Chen, David Durfee, and Anthi Orfanou. On the complexity of Nash equilibria in anonymous games. In *Symposium on Theory of Computing (STOC)*, pages 381–390, 2015.

Yu Cheng, Ilias Diakonikolas, and Alistair Stewart. Playing anonymous games using simple strategies. In *Symposium on Discrete Algorithms (SODA)*, pages 616–631, 2017.

Constantinos Daskalakis. An efficient ptas for two-strategy anonymous games. In *International Workshop on Internet and Network Economics*, pages 186–197. Springer, 2008.

Constantinos Daskalakis and Christos Papadimitriou. Computing equilibria in anonymous games. In *Symposium on Foundations of Computer Science (FOCS)*, pages 83–93. IEEE, 2007.

Constantinos Daskalakis and Christos H Papadimitriou. Discretized multinomial distributions and Nash equilibria in anonymous games. In *Symposium on Foundations of Computer Science (FOCS)*, pages 25–34. IEEE, 2008.

Constantinos Daskalakis and Christos H Papadimitriou. On oblivious PTAS's for Nash equilibrium. In *Symposium on Theory of Computing (STOC)*, pages 75–84, 2009.

Constantinos Daskalakis and Christos H. Papadimitriou. Approximate Nash equilibria in anonymous games. *J. Econ. Theory*, 156:207–245, 2015.

Constantinos Daskalakis, Stratis Skoulakis, and Manolis Zampetakis. The complexity of constrained min-max optimization. In *Symposium on Theory of Computing (STOC)*, pages 1466–1478, 2021.

Constantinos Daskalakis, Noah Golowich, Stratis Skoulakis, and Emmanouil Zampetakis. Stay-on-the-ridge: Guaranteed convergence to local minimax equilibrium in nonconvex-nonconcave games. In *Conference on Learning Theory (COLT)*, pages 5146–5198. PMLR, 2023.

Alex Fabrikant, Christos Papadimitriou, and Kunal Talwar. The complexity of pure Nash equilibria. In *Symposium on Theory of Computing (STOC)*, pages 604–612, 2004.

Paul W. Goldberg and Stefano Turchetta. Query complexity of approximate equilibria in anonymous games. *J. Comput. Syst. Sci.*, 90:80–98, 2017.

Chi Jin, Praneeth Netrapalli, and Michael Jordan. What is local optimality in nonconvex-nonconcave minimax optimization? In *International Conference on Machine Learning (ICML)*, pages 4880–4889. PMLR, 2020.

David S Johnson, Christos H Papadimitriou, and Mihalis Yannakakis. How easy is local search? *Journal of Computer and System Sciences*, 37(1):79–100, 1988.

Ehud Kalai. Partially-specified large games. In *International workshop on internet and network economics*, pages 3–13. Springer, 2005.

Fivos Kalogiannis, Jingming Yan, and Ioannis Panageas. Learning equilibria in adversarial team markov games: A nonconvex-hidden-concave min-max optimization problem. *Advances in Neural Information Processing Systems (NeurIPS)*, 37:92832–92890, 2024.

Zhengyang Liu. The complexity of exact and approximate nash equilibria in anonymous games, 2026. URL https://arxiv.org/abs/2610.16848.

Igal Milchtaich. Congestion games with player-specific payoff functions. *Games and economic behavior*, 13(1):111–124, 1996.

Dov Monderer and Lloyd S Shapley. Potential games. *Games and economic behavior*, 14(1):124–143, 1996.

Christos Papadimitriou and Binghui Peng. Public goods games in directed networks. In *Conference on Economics and Computation (EC)*, pages 745–762, 2021.

Christos H. Papadimitriou and Tim Roughgarden. Computing correlated equilibria in multi-player games. *J. ACM*, 55(3):14:1–14:29, 2008.

Robert W Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2(1):65–67, 1973.

Aviad Rubinstein. Inapproximability of Nash equilibrium. *SIAM Journal on Computing*, 47(3): 917–959, 2018. doi: 10.1137/15M1039274.

A Instances with only one type of players

We start this section by analyzing the simpler instances where only one group of players is present. For the anonymous game where only maximizers are present, we make the following observation:

Observation 1 (Maximizers-only instance)**.** *A linear anonymous game with two actions and only maximizers (as defined in Definition 3) admits (at least) two pure Nash equilibria: the all-1 and the all-0 strategy profiles.*

Proof. If every player plays strategy 1 then $y_{-i} = n - 1 > \tau_i^y$ for each player y_i. Thus each maximizer is best responding and, therefore, the profile is a Nash equilibrium. Similarly, if every player plays strategy 0 then $y_{-i} = 0 < \tau_i^y$ with the same conclusion. □

Recall that in the case where we only have maximizers, the incentives of the players naturally align. This alignment holds generally: the more mass a player assigns to strategy 1, the more mass the other maximizers are inclined to assign to strategy 1 as well. This monotonicity allows us to define a best-response function as a function of the total mass assigned to action 1. We refer to this mapping as *self-inclusion best response*, since it encodes the best response probability for a maximizer y_i to select strategy 1, as a function of the entire profile rather than solely the strategies of the other players. Specifically, given a possibly mixed strategy profile, let S be the total expected count of players choosing strategy 1. Then,

$$\phi_i^{\max}(S) \triangleq \begin{cases} 1, & S > \tau_i^y + 1 \\ S - \tau_i^y, & S \in [\tau_i^y, \tau_i^y + 1] \\ 0, & S < \tau_i^y \end{cases}.$$

The above formula states that the maximizer y_i should play a pure strategy as long as $S \notin [\tau_i^y, \tau_i^y + 1]$, and otherwise should play a mixed strategy which assigns probability $S - \tau_i^y$ to strategy 1. We focus on the intermediate case, as the other two cases follow analogously. In the intermediate regime, player i is deploying a mixed strategy, hence the payoffs she receives when playing the two pure strategies are equivalent. Let y_i be the probability that she places on strategy 1. We start from the relation $y_{-i} + y_i = S$, which is derived by the linearity of expectation. Substituting player y_i's indifference condition $y_{-i} = \tau_i^y$ yields $y_i = S - \tau_{y_i}$. Thus, our construction is sound. The following observation is immediate:

Observation 3. *The functions $\phi_i^{\max}$ are continuous.*

Notice that $\phi^{\max}(0) = 0$ and $\phi^{\max}(n) = n$ and these values for S correspond precisely to the two Nash equilibria we identified previously in Observation 1. Hence, the two equilibria are fixed points of this function. Observe that this holds in general; every Nash equilibrium of the maximizers-only instance is a fixed point of $\phi^{\max}$. This motivates our next definition.

$$\psi^{\max}(S) \triangleq S - \phi^{\max}(S). \qquad \text{(Maximizers Fixed Point Function)}$$

Observation 4. *There exists a Nash equilibrium of a maximizer-only instance with total mass S^* if and only if $\psi^{\max}(S^*) = 0$.*

To complete our previous analysis, our next goal is to understand the structure of the equilibria in games with minimizers only. Ideally, we would like a similar equilibrium-encoding function. We begin by defining the self-inclusion best-response correspondence, which is no longer single-valued:

$$\phi_i^{\min}(S) \triangleq \begin{cases} 0, & S > \tau_i^x + 1, \\ \{0, S - \tau_i^x, 1\}, & S \in [\tau_i^x, \tau_i^x + 1], \\ 1, & S < \tau_i^x. \end{cases}$$

This construction is similar to Maximizer Self-inclusion BR. The main difference, apart from the switch of the pure strategies, is that a minimizer has 3 possible best responses when $S \in [\tau_i^x, \tau_i^x + 1]$. The best response $S - \tau_i^x$ is justified as in the case of maximizers. The pure strategy 0 is a best response since it leads to $x_{-i} = S > \tau_i^x$ and analogously for strategy 1. Next, we propose an algorithm that computes a pure Nash equilibrium in two-action anonymous games with only the minimizers.

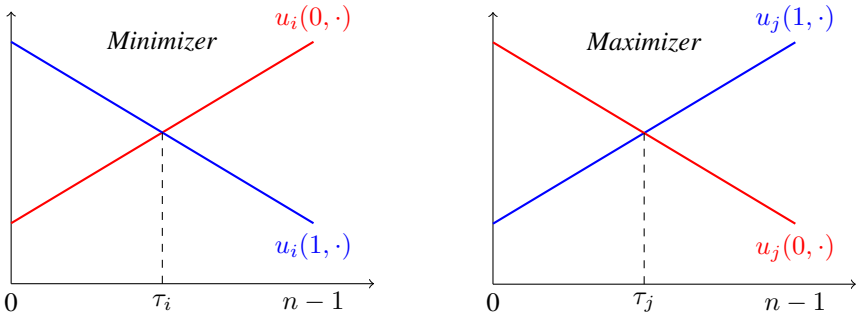

Figure 3: The utility function of players i and j. The threshold τ_i marks the point where the player i changes her best response, with minimizers playing strategy 1 left of τ_i and strategy 0 afterwards and maximizers doing the opposite.

Algorithm 1: Computation of Nash equilibrium for a minimizer-only instance

Initialize with the all-1 profile $\mathbf{s} = (1, \dots, 1)$.
Sort the players in increasing order of τ_i^x.
$i = 1$.
while $i \leq n$ *and* $\tau_i^x < n - i$ **do**
  Set maximizer x_i to play zero.
  $i = i + 1$.
return $\mathbf{s}$

Lemma A.1 (Minimizer-only instance)**.** *Algorithm 1 outputs a pure Nash equilibrium in a linear anonymous game with two actions and n_x minimizers in $O(n_x \log n_x)$ time.*

Proof. First, note that Algorithm 1 produces all the assignments of the form $(0, \dots, 0, 1, \dots, 1)$ so our goal is to show that at least one equilibrium of this form exists (due to our assumption that $\tau_i^x \in (0, n_x - 1)$ for all $i \in \mathcal{N}_{\min}$, it is easy to see that the all-1 and all-0 profiles cannot be equilibria). Next, note that if two players i, j play the same strategy, then they observe the same mass k on action 1. Wlog assume that $\tau_i^x \leq \tau_j^x$. If player x_i's best response is 1, i.e. $k \leq \tau_i^x$, then $k \leq \tau_j^x$ and action 1 is a best response for player x_j as well. On the other hand, if player x_j's best response is action 0, i.e. $k \geq \tau_j^x$ then $k \geq \tau_i^x$ and x_i's best response is also 0. Thus, it suffices to make sure that the players where the strategy changes from 1 to 0 best respond. In other words, there exists $k \in \{1, \dots, n_x - 1\}$ such that $\tau^x_{n_x - k} \leq k$ and $\tau^x_{n_x - k + 1} \geq k - 1$. Guaranteeing the existence of such a pair of inequalities is straightforward: for $k = 1$ the second one translates to $\tau^x_{n_x} \geq 0$ which holds. If $\tau^x_{n_x - 1} \leq 1$ we conclude. Otherwise, i.e. $\tau^x_{n_x - 1} > 1$, the second inequality holds for $k = 2$. Repeating the process gives us a solution since $\tau_1^x < n_x - 1$, which is the first condition for $k = n_x - 1$. Finally, note that the complexity of the algorithm is dominated by the sorting time. □

B Omitted proofs from Section 3 and Algorithm 2

We begin this section by providing the proof for Observation 2.

Proof of Observation 2. By Observation 4, we have that $\phi^{\max}(S^*) = \sum_{i \in \mathcal{N}_{\max}} \phi_i^{\max}(S^*) = \sum_{i \in \mathcal{N}_{\max}} y_i$. Substituting this into H gives us $H(S^*) = S^* - \sum_{i \in \mathcal{N}_{\max}} y_i = \sum_{i \in \mathcal{N}_{\min}} x_i$. The other direction i straightforward from the definition of the self-inclusion BR function. □

We continue with the proof of Lemma 1.

Proof of Lemma 1. For $i \leq k$ we have that $\tau_i^x \geq \tau_k^x$, thus $S \leq \tau_k^x + 1 \leq \tau_i^x + 1$. This implies $1 \in \phi_i^{\min}(S)$. For $i > k$, we have that $\tau_i^x \leq \tau_{k+1}^x$, thus $S \geq \tau_{k+1}^x \geq \tau_i^x$. This implies $0 \in \phi_i^{\min}(S)$. By picking $x_i = 1$ in the first case and $x_i = 0$ in the second and summing we obtain $\sum_{i \in \mathcal{N}_{\min}} x_i = k = H(S)$ satisfying the equilibrium condition of Observation 2. □

Next, we provide proofs of Lemma 2 and Lemma 3.

Proof of Lemma 2. For player x_k to be the sole mixer, we require $x_i = 1$ for $i < k$ and $x_i = 0$ for $i > k$. The total mass at equilibrium becomes $S = (k-1) + x_k + \phi^{\max}(S)$. For $x_k \in (0,1)$ it must be that $x_k = S - \tau_k^x$, which implies $S \in (\tau_k^x, \tau_k^x + 1)$. Substituting x_i yields:

$$S = (k-1) + S - \tau_k^x + \phi^{\max}(S) \implies \phi^{\max}(S) = \tau_k^x - k + 1 = C_k.$$

Since $\phi^{\max}(S)$ is continuous and non-decreasing and we assumed that $C_k \in (\phi^{\max}(\tau_k^x), \phi^{\max}(\tau_k^x + 1))$, the existence of such an S is guaranteed by the Intermediate Value Theorem. □

Proof of Lemma 3. We have two cases based on whether C_k lies left or right of the given interval.

Case 1: $C_k \leq \phi^{\max}(\tau_k^x)$. Then, by the definition of C_k we get that $\tau_k^x - (k-1) \leq \phi^{\max}(\tau_k^x)$ or, equivalently, $\tau_k^x - \phi^{\max}(\tau_k^x) \leq k - 1$. Using now the definition of H gives us $H(\tau_k^x) \leq k - 1$.

Case 2: $C_k \geq \phi^{\max}(\tau_k^x + 1)$. Again by the definition of C_k we get that $\tau_k^x + 1 - \phi^{\max}(\tau_k^x + 1) \geq k$ and after rearranging the terms and using the definition of H we obtain $H(\tau_k^x + 1) \geq k$. □

We continue with the proof of Lemma 4, which states that there exists a Nash equilibrium where at most one minimizer plays a mixed strategy.

Proof of Lemma 4. If there is some k such that mixing condition (Lemma 2) holds then we have found a profile where exactly one minimizer mixes and the theorem holds. If no such k exists we want to show that there exists an equilibrium where all the mimimizers play a pure strategy, i.e. Lemma 1 holds. Assume for contradiction, that no such profile exists. By the negation of Lemma 1, for every $k \in \{0, \ldots, n_x\}$ there is no $S \in I_k = [\tau_{k+1}^x, \tau_k^x + 1]$ such that $H(S) = k$. By the Intermediate Value Theorem, $H(S)$ must be strictly greater than k or strictly less than k on the entire interval I_k.

We are going to use induction to show that $H(S) > k$ for $S \in I_k$.

1. For $k = 0$ on $I_0 = [\tau_1^x, n_x + n_y + 1]$, we have that $H(n_x) = n_x$ (due to the assumption that all the thresholds are in $(0, n_x + n_y)$, and we get that $\phi^{\max}(n_x + n_y) = n_y$). Thus $H(S) > 0$ for all $S \in I_0$, implying $H(\tau_1^x) > 0$.
2. For $k = 1$ on $I_1 = [\tau_2^x, \tau_1^x + 1]$, Lemma 3 requires $H(\tau_1^x) \leq 0$ or $H(\tau_1^x + 1) \geq 1$. Since $H(\tau_1^x) > 0$, the second condition must hold. Since no pure equilibrium exists for $k = 1$, $H(S) \neq 1$ on I_1. Since $H(\tau_1^x + 1) \geq 1$ and H is continuous, we must have $H(S) > 1$ on I_1. This implies $H(\tau_2^x) > 1$.
3. By induction, assume $H(\tau_k^x) > k - 1$. For the interval $I_k = [\tau_{k+1}^x, \tau_k^x + 1]$, Lemma 3 requires $H(\tau_k^x) \leq k-1$ or $H(\tau_k^x + 1) \geq k$. The induction hypothesis forces $H(\tau_k^x + 1) \geq k$. Absence of pure equilibria forces $H(S) > k$ for all $S \in I_k$, implying $H(\tau_{k+1}^x) > k$.

For $k = n_x$ this implies $H(S) > n_x$ for all $S \in [0, \tau_{n_x}^x + 1]$. However, at $S = 0, H(S) = 0 - \phi^{\max}(0) \leq 0$. This contradicts the continuity of $H(S)$. Thus, at least one pure strategy equilibrium must exist. □

We continue by stating a lemma that characterizes the strategy profile of the maximizers at equilibria.

Lemma B.2 (Characterization of Nash equilibrium for the maximizers)**.** *In every Nash equilibrium the maximizers that play a mixed strategy have consecutive thresholds.*

Proof of Lemma B.2. If at most one maximizers mix the lemma follows immediately. Assume for contradiction that there exists a pair of maximizers i, j with $\tau_i^y < \tau_j^y$ that play a mixed strategy and a maximizer k with $\tau_k^x \in (\tau_i^y, \tau_j^y)$ that plays a pure strategy. Moreover, let S be the total mass on action 1. Then we know that $S = \tau_i^y + y_i = \tau_j^y + y_j$. If player k plays 0 then she observes $S = \tau_j^y + y_j > \tau_k^y$ and thus she should have been playing 1. Similarly, if she plays 1 then she observes $S - 1 = \tau_i^y + y_i - 1 < \tau_k$ and she should have been playing 0. In both cases, we reach a contradiction and therefore player k mixes. This completes the proof. □

Algorithm 2: Computation of a Nash equilibrium given the supports

Input: I_0, I_1, I_{mix}.
$T = \sum_{i \in I_{\text{mix}}} \tau_i$.
if $|I_{\text{mix}}| = 1$ and $|I_1| \neq T$ **then return** "Infeasible".
else if $|I_{\text{mix}}| = 1$ and $|I_1| = T$ **then** compute the feasibility range $[L, U]$ from the thresholds for players in I_0 and I_1, choose $S \in [L, U]$ and set $p_{mix} = S - |I_1|$
else
  for $i \in I_{\text{mix}}$ **do**
    if $\frac{T-|I_1|}{|I_{\text{mix}}|-1} - \tau_i \notin [0, 1]$ **then return** "Infeasible".
    Player i plays 1 with probability $\frac{T-|I_1|}{|I_{\text{mix}}|-1} - \tau_i$
  $S = |I_1| + \sum_{i \in I_{mix}} p_i$
for $i \in I_0$: **if** $a_i(S - \tau_i) > 0$ **then return** "Infeasible".
for $i \in I_1$: **if** $a_i(S - \tau_i - 1) < 0$ **then return** "Infeasible".
for $i \in I_1$: Player i plays 1.
for $i \in I_0$: Player i plays 0.
return the mixed strategy profile

We proceed with the proof of Lemma 5.

Proof of Lemma 5. Let S be the total mass on action 1. Observe that $S = |I_1| + \sum_{i \in I_{\text{mix}}} p_i$, where we use p_i as a shorthand for x_i for the minimizers or y_i for the maximizers. For the mixers, we have that $p_i = S - \tau_i$. Substituting this gives

$$S = |I_1| + \sum_{i \in I_{\text{mix}}} z_i = |I_1| + \sum_{i \in I_{\text{mix}}} (S - \tau_i) = |I_1| + |I_{\text{mix}}| \cdot S - \sum_{i \in I_{\text{mix}}} \tau_i \implies (|I_{\text{mix}}| - 1) \cdot S = T - |I_1|,$$

where $T = \sum_{i \in I_{\text{mix}}} \tau_i$. We have three cases.

(a) $|I_{\text{mix}}| = 1$ and $|I_1| \neq T$. The profile does not correspond to a Nash equilibrium.

(b) $|I_{\text{mix}}| = 1$ and $|I_1| = T = \tau_j$, where j is the only agent that mixes. In this case, the system is underdetermined, and we have $S = |I_1| + p_j$ for some $p_j \in [0, 1]$. However, not every choice of p_j necessarily corresponds to a Nash equilibrium, since the players in I_0 and I_1 must also be best responding. Their best-response conditions impose linear inequalities on S, which, together with $S \in [|I_1|, |I_1| + 1]$, define a feasible interval $[L, U]$. If $L > U$, the support is infeasible. Otherwise, choosing any $S \in [L, U]$ and setting $p_j = S - |I_1|$ yields a Nash equilibrium.

(c) $|I_{\text{mix}}| \neq 1$. Then $t = \frac{T-|I_1|}{|I_{\text{mix}}|-1}$ and $z_i = \frac{T-|I_1|}{|I_{\text{mix}}|-1} - \tau_i$ for $i \in I_{\text{mix}}$. Then the support corresponds to a Nash equilibrium profile if and only if $z_i \in [0, 1]$ for all i.

For a given support, the computation of T and a sweep of the players require linear time, and we conclude the proof. □

Next, we provide the proof for Theorem 3.

Theorem 3. *Algorithm 3 computes a Nash equilibrium in linear anonymous games with n players and two actions in $O(n^4)$ time.*

Proof. The correctness of the algorithm follows directly from the fact that it brute forces every profile that matches the characterization of a Nash equilibrium as in Lemmas B.2 and 4. Specifically, l is the position of the leftmost maximizer that mixes, with $l = n_y + 1$ corresponding to the case where every maximizer plays 0. Then, r is the position of the rightmost maximizer, with $r = 0$ corresponding to the case where every maximizer plays 1. The interval $[l, r]$ corresponds to the maximizers with consecutive thresholds that mix (Lemma B.2). Next, the inner loop iterates over eack minimizer k, assigns those left and right of player k based on the characterization (Lemma 4) and then checks all

three possibilities for k. Regarding the complexity, computing the type and threshold of a player takes constant time, thus the preprocessing step is dominated by the sorting time of the thresholds. Also, Algorithm 1 runs in $O(n \log n)$ time. Proceeding to the main body of the algorithm, we have three nested loops, each costing $O(n)$. Therefore the total number of considered configurations is $O(n^3)$. Each configuration requires a linear sweep of the players to compute I_0, I_1, I_{mix} and 3 calls of Algorithm 2 that cost $O(n)$, as established in Lemma 5. Hence, the total time is $O(n^4)$. □

Finally, we give the proof of Corollary 1.

Corollary 1. *An exact Nash equilibrium in a two-action anonymous game with n players and linear utilities can be computed in $O(n^3)$ arithmetic operations.*

Proof. Algorithm 3 enumerates $O(n^3)$ candidate supports, each verified in $O(n)$ time. We show that the feasibility of each support can instead be checked in constant time after preprocessing with prefix sums.

Suppose the mixing maximizers form a consecutive block $[l, r]$, with thresholds sorted in decreasing order. Their best-response conditions reduce to

$$\tau_l^y \le S \le \tau_r^y + 1.$$

Similarly, the pure maximizers only require

$$S \le \tau_{l-1}^y, \qquad S \ge \tau_{r+1}^y + 1.$$

For a candidate mixing minimizer k, the pure minimizers' best-response conditions then reduce to

$$S \le \tau_{k-1}^x + 1, \qquad S \ge \tau_{k+1}^x.$$

Therefore, the strategy of player k introduces at most two additional inequalities. By precomputing prefix sums of the maximizers' thresholds, the aggregate threshold $\sum \tau_i$ of the mixing players can be obtained in constant time. The aggregate mass S can then be determined directly when the number of mixers is different from one following the analysis for Algorithm 2. In the single-mixer case, feasibility reduces to checking a constant number of interval constraints.

Thus, each candidate support can be verified in $O(1)$ time. Once a feasible support is identified, the equilibrium reconstruction takes $O(n)$ time. The overall running time is therefore $O(n^3)$. □

C Omitted analysis from Section 4

Example 1. *Let G be a four-player anonymous game containing two minimizers x_1, x_2 with thresholds $\tau_1^x = \tau_2^x = 2$ and two maximizers y_1, y_2 with $\tau_1^y = \tau_2^y = \frac{3}{2}$.*

First, observe that any Nash equilibrium of the game G must take one of the following forms

- $x_1^* = x_2^* = \frac{1}{3}, y_1^* = y_2^* = \frac{5}{6}$;
- $x_1^* = 0, x_2^* \in [\frac{1}{2}, 1], y_1^* = y_2^* = 1$;
- $x_1^* \in [\frac{1}{2}, 1], x_2^* = 0, y_1^* = y_2^* = 1$;
- $x_1^* = 0, x_2^* = 1, y_1^* = y_2^* = \frac{1}{2}$;
- $x_1^* = 1, x_2^* = 0, y_1^* = y_2^* = \frac{1}{2}$.

Meanwhile, let $\boldsymbol{x} = \left(\frac{1}{2}, \frac{1}{2}\right)$ and consider the strategy profile $\boldsymbol{y}^\star$ that maximizes the function $\Phi(\boldsymbol{x}, \boldsymbol{y})$; that is, $\boldsymbol{y}^\star \in \operatorname{argmax}_{\boldsymbol{y} \in [0,1]^2} \Phi(\boldsymbol{x}, \boldsymbol{y})$. One can verify that at $\boldsymbol{x} = (\frac{1}{2}, \frac{1}{2})$, exactly two profiles maximize $\Phi(\boldsymbol{x}, \boldsymbol{y})$, namely, $\boldsymbol{y}^\star \in \{(0,0), (1,1)\}$. These maximizers yield a maximum value of $\Phi^\star(\boldsymbol{x}) = -\frac{7}{4}$. If we set $\boldsymbol{y} = (0,0)$, it holds that

$$\nabla_{x_1} \Phi(\boldsymbol{x}, \boldsymbol{y}) = \nabla_{x_2} \Phi(\boldsymbol{x}, \boldsymbol{y}) = x_2 + y_1 + y_2 - \tau_1^x = -\frac{3}{2}.$$

On the other hand, if we set $\boldsymbol{y}' = (1,1)$, we have

$$\nabla_{x_1} \Phi(\boldsymbol{x}, \boldsymbol{y}') = \nabla_{x_2} \Phi(\boldsymbol{x}, \boldsymbol{y}') = x_2 + y_1 + y_2 - \tau_1^x = \frac{1}{2}.$$

Since $0 \in \operatorname{conv}(-\frac{3}{2}, \frac{1}{2})$, from Danskin's theorem, we conclude that $\boldsymbol{x} = (\frac{1}{2}, \frac{1}{2})$ is a first-order stationary point of the max function $\Phi^\star(\boldsymbol{x})$. However, this point cannot be extended to a Nash equilibrium, as there is no Nash equilibrium $(\boldsymbol{x}^*, \boldsymbol{y}^*)$ where $\boldsymbol{x}^* = \boldsymbol{x}$.

D Omitted proofs from Section 5

First, we establish a lemma demonstrating that under the SBR dynamics, at most one minimizer plays a mixed strategy.

Lemma D.3. *For any iteration $t \in [T]$, after the sequential updates on $\boldsymbol{x}_t$, at most one minimizer $x_{t,i}$ plays a mixed strategy.*

Proof. We prove by contradiction. Assume the indices are sorted such that the thresholds $\tau_{x,i}$ are in decreasing order. Suppose that after the sequential updates, there exist two minimizers $x_{t,i}$ and $x_{t,j}$ such that $x_{t,i} \in (0,1)$ and $x_{t,j} \in (0,1)$. Without loss of generality, let us assume that $i < j$. Recall that $x_{t,l}$ is initialized to 0 for all $l \in [n_x]$. Consider the ith iteration, in which $x_{t,i}$ is updated. Since $x_{t,i} \in (0,1)$, the definition of $BR_i^{\min}$ implies that

$$[0,1] \ni S - \sum_{l \neq i} x_l - \sum_{k=1}^{n_y} y_k = S - \sum_{l<i} x_l - \sum_{k=1}^{n_y} y_k.$$

Consequently, for any $j > i$, it holds that

$$\begin{aligned} S - \sum_{l \neq j} x_l - \sum_{k=1}^{n_y} y_k &= S - \sum_{l<j} x_l - \sum_{k=1}^{n_y} y_k \\ &\leq S - \sum_{l=0}^{i} x_l - \sum_{k=1}^{n_y} y_k \\ &= 0. \end{aligned}$$

This contradicts the assumption that $x_{t,j}$ plays a mixed strategy. The proof is complete. □

We proceed to prove a technical lemma regarding the continuity of the utility function.

Lemma D.4. *Let the utility function be defined as in (1). Then $\Phi(\boldsymbol{x}, \boldsymbol{y})$ is $n^{3/2}$-Lipschitz continuous with respect to $\boldsymbol{x}$ and $\boldsymbol{y}$.*

Proof. Notice that

$$|\nabla_{x_i} \Phi(\boldsymbol{x}, \boldsymbol{y})| = \left| \sum_{j \neq i} x_j + \sum_{k=1}^{n_y} y_k - \tau_{x,i} \right| \leq n,$$

where the last inequality holds due to the assumption that $\tau_{x,i} \in [0, n]$. Since $\nabla_{\boldsymbol{x}} \Phi(\boldsymbol{x}, \boldsymbol{y}) \in \mathbb{R}^{n_x}$, we have

$$\|\nabla_{\boldsymbol{x}} \Phi(\boldsymbol{x}, \boldsymbol{y})\| \leq n\sqrt{n_x} \leq n^{3/2}.$$

Thus it is the case that

$$|\Phi(\boldsymbol{x}, \boldsymbol{y}) - \Phi(\boldsymbol{x}', \boldsymbol{y})| \leq n^{3/2} \|\boldsymbol{x} - \boldsymbol{x}'\|.$$

A similar argument holds for $\boldsymbol{y}$, the proof is complete. □

Finally, we restate the statement of Theorem 4 and show the proof.

Theorem 4. *Suppose $\boldsymbol{x}$ and $\boldsymbol{y}$ follow the SBR dynamics defined in Algorithm 4. By setting $T \geq 2n^{7/2}/\epsilon$, there exists an iteration $t \in [T]$ such that $(\boldsymbol{x}_t, \boldsymbol{y}_t)$ is an ϵ-approximate Nash equilibrium.*

Proof. By Lemma 4, there exists a Nash equilibrium $(\boldsymbol{x}^*, \boldsymbol{y}^*)$ such that at most one minimizer x_i^* plays a mixed strategy, while all players with larger thresholds play 1 and all players with smaller thresholds play 0. Moreover, by setting $S^* = \sum_{i=1}^{n_x} x_i^* + \sum_{j=1}^{n_y} y_j^*$, it holds that $x_i^* \in \phi_i^{\min}(S^*)$ and $y_i^* = \phi_i^{\max}(S^*) = BR_i^{\max}(S, \boldsymbol{x}^*, \boldsymbol{y}^*)$.

Based on the update rule in Algorithm 4, we conclude that there exists an iteration $t \in [T]$ such that $|S_t - S^*| \leq \frac{n}{2T}$. Since $BR_i^{\max}(S, \boldsymbol{x}, \boldsymbol{y})$ is 1-Lipschitz continuous with respect to S, it is the case that $|y_{t,i} - y_i^*| \leq \frac{n}{2T}$ for all $i \in [n_y]$. Thus it holds that

$$\begin{aligned}\left|\sum_{i=1}^{n_x} x_{t,i} - \sum_{i=1}^{n_x} x_i^*\right| &= \left|S_t - \sum_{j=1}^{n_y} y_{t,j} - S^* + \sum_{j=1}^{n_y} y_j^*\right| \\ &\leq |S_t - S^*| + \left|\sum_{j=1}^{n_y} y_{t,j} - \sum_{j=1}^{n_y} y_j^*\right| \\ &\leq (n_y + 1)\frac{n}{2T}.\end{aligned}$$

Assume the indices of the minimizers are sorted such that $\tau_i^x \geq \tau_j^x$ for all $i > j$. Let x_i^* be the last player in $\boldsymbol{x}^*$ with a positive probability. By Lemma 4, we have $x_j^* = 1$ for all $j < i$ and $x_j^* = 0$ for all $j > i$. Furthermore, due to Lemma D.3, there exists an index $k \in [n_x]$ such that $x_{t,j} = 1$ for all $j < k$ and $x_{t,j} = 0$ for all $j > k$. By setting $T \geq n^2 \geq n(n_y + 1)$, we guarantee that $|\sum_{i=1}^{n_x} x_{t,i} - \sum_{i=1}^{n_x} x_i^*| \leq 1$. Consequently, $\boldsymbol{x}_t$ and $\boldsymbol{x}^*$ differ in at most two coordinates. Therefore, it holds that

$$\|\boldsymbol{x}_t - \boldsymbol{x}^*\| \leq \|\boldsymbol{x}_t - \boldsymbol{x}^*\|_1 \leq \frac{n^2}{T},$$

where the last inequality holds because each coordinate changes by at most $(n_y + 1)\frac{n}{2T}$.

Finally, by Lemma D.4, the utility function is $n^{3/2}$-Lipschitz continuous with respect to strategy profile $\boldsymbol{x}$ and $\boldsymbol{y}$. We conclude that $(\boldsymbol{x}_t, \boldsymbol{y}_t)$ is a $\frac{2n^{7/2}}{T}$-approximate Nash equilibrium. Setting $T \geq \frac{2n^{7/2}}{\epsilon}$ concludes the proof. □

E Additional Discussion on Experiments

Algorithm 4 was implemented in Python 3.9.6. All the experiments were executed on a Macbook Air M4. The code and the randomly generated data referred to within Section 6 can be found in the supplementary material.

We complement here the exposition within Section 6, by showing the performance of Algorithm 4 on two additional games. First, we start with the 4-player game presented already as Example 1 in Section 4. We ran our dynamics for $T = 10^4$ iterations, after which the accuracy already falls below 10^{-12}, attaining convergence to the mixed Nash equilibrium $((0, 1/2), (1, 1))$. The best iterate accuracy can be seen in the left subfigure of Figure 4. In the right subfigure, we see the evolution of the players' probability of using the first strategy, and their convergence to the equilibrium $((0, 1/2), (1, 1))$.

For completeness, we also include a small randomly generated game to demonstrate that the algorithm's performance is similar to the case of much larger random games, as already discussed in Section 6. The thresholds are presented below, in Table 1. For ease of presentation, we have sorted the minimizers' thresholds in descending order and maximizers' thresholds in ascending order.

Minimizers	Maximizers
6.1671	1.1650
5.0401	3.0451
4.5624	5.6486
1.5751	8.0493
1.1680	8.1943

Table 1: The randomly generated thresholds for the game of Figure 5.

In the left subfigure of Figure 5 we observe once again that the rate of convergence is of the form $O(\frac{1}{T})$. Here, the piecewise nature of the function is more pronounced though.

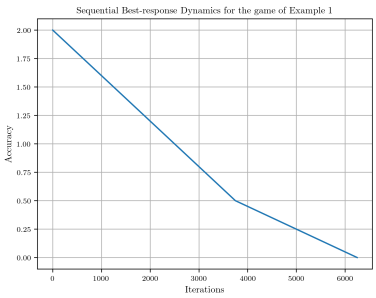

(a) The accuracy of the best iterate. Consistently with our theory, the convergence rate is $O(\frac{1}{T})$.

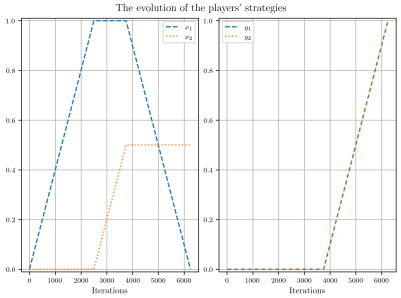

(b) The evolution of the players' strategies. We see that our dynamics converge to $((0, 1/2), (1, 1))$.

Figure 4: Sequential Best-response Dynamics on the game of Example 1, with 2 minimizers ($\tau_1 = \tau_2 = 2$) and 2 maximizers ($\tau_3 = \tau_4 = 3/2$).

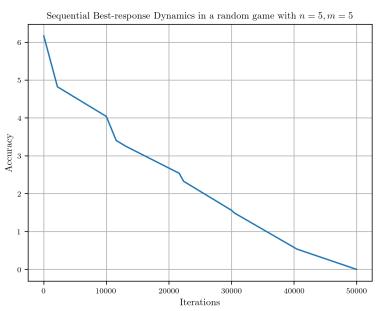

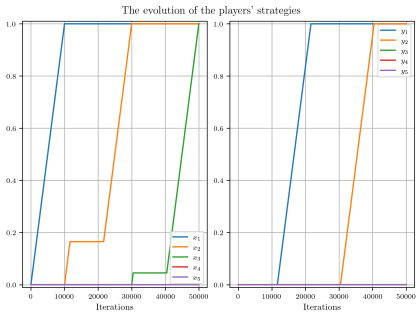

Figure 5: Sequential Best-response Dynamics in the game described in Table 1.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: All claims made are established within Sections 3 to 5 and all missing proofs can be found in the Appendix.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: We have discussed that our work is limited both by the number of actions and the utility functions of the players. In particular, our algorithm works for anonymous games with 2 actions per player, and with the type of utility functions defined in Section 2.

   Guidelines:

- The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [Yes]

   Justification: All proofs for our theorems and lemmata are provided in Sections 3 to 5 and in the Appendix.

   Guidelines:

   - The answer [N/A] means that the paper does not include theoretical results.
   - All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
   - All assumptions should be clearly stated or referenced in the statement of any theorems.
   - The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
   - Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
   - Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: The experimental results of the paper are about the proposed dynamics, the description of which is given in the form of pseudocode. Moreover, the experimental parameters are specified and the random data generated are included partly in the main body of the paper and partly in the Appendix and the Supplementary Material.

   Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [Yes]

   Justification: We provide our source code and the randomly generated data in the zip file.

   Guidelines:

   - The answer [N/A] means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).

- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [Yes]

   Justification: Our experiments involve the implementation of a simple algorithm (Algorithm 4 in Section 5). The only relevant parameter is the number of iterations T, and we explicitly state in our experimental evaluation the values we used for T.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [No]

   Justification: We have not included error bars as all the figures that we have included in Section 6 concern the performance of our algorithm on particular games.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
   - The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
   - The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
   - The assumptions made should be given (e.g., Normally distributed errors).
   - It should be clear whether the error bar is the standard deviation or the standard error of the mean.
   - It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
   - For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
   - If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: We mention in Section 6 that all our experiments were run on a single personal laptop. The running time of each experiment ranged between a few seconds and a few minutes.

   Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

   Answer: [Yes]

   Justification: Our research does not raise any ethical issues.

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [N/A]

    Justification: Our research is focused on theoretical foundations and therefore is not tied to any deployment or immediate application that could have direct societal impact.

    Guidelines:

    - The answer [N/A] means that there is no societal impact of the work performed.
    - If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
    - Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
    - The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
    - The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
    - If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

Answer: [N/A]

Justification: The paper poses no such risks.

Guidelines:

- The answer [N/A] means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [N/A]

Justification: The paper does not use existing assets.

Guidelines:

- The answer [N/A] means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [N/A]

Justification: The paper does not release new assets.

Guidelines:

- The answer [N/A] means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [N/A]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [N/A]

Justification: The paper does not involve research with human subjects.

Guidelines:

- The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

Answer: [N/A]

Justification: The core method developed in this research does not involve LLMs as any important, original, or non-standard components.

Guidelines:

- The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.