The Computational Intractability of Not Worst Responding

Mete Şeref Ahunbay[†]
Paul W. Goldberg[†]
Edwin Lock[‡]
Panayotis Mertikopoulos[*]
Bary S. R. Pradelski[†*]
Bassel Tarbush[†]

[*]CNRS
[‡]King's College London
[†]University of Oxford

July 11, 2026

Finding, counting, or determining the existence of Nash equilibria, where players must play optimally given each others' actions, are known to be computational intractable problems. We ask whether weakening optimality to the requirement that each player merely avoid worst responses—arguably the weakest meaningful rationality criterion—yields tractable solution concepts. We show that it does not: any solution concept with this minimal guarantee is "as intractable" as pure Nash equilibrium. In general games, determining the existence of *no-worst-response* action profiles is `NP-complete`, finding one is `NP-hard`, and counting them is `#P-complete`. In potential games, where existence is guaranteed, the search problem is `PLS-complete`. Computational intractability therefore stems not only from the requirement of optimality, but also from the requirement of a minimal rationality guarantee *for each player*. Moreover, relaxing the latter requirement gives rise to a tractability trade-off between the strength of individual rationality guarantees and the fraction of players satisfying them.

Keywords: best responses, not-worst responses, minimal rationality, satisficing equilibrium, pure Nash equilibrium, normal form game, NP-completeness, PLS-completeness
Contact: `mete.ahunbay@cs.ox.ac.uk`, `paul.goldberg@cs.ox.ac.uk`, `edwin.lock@kcl.ac.uk`, `panayotis.mertikopoulos@imag.fr`, `bary.pradelski@cnrs.fr`, `bassel.tarbush@economics.ox.ac.uk`
Acknowledgements: PG was supported by UKRI/EPSRC grants EP/X040461/1 and EP/X038548/1. MŞF, PM, and BSRP were supported by PEPR project FOUNDRY, ANR23-PEIA-0003.

1 Introduction

It is well-known that determining the existence, finding, and counting Nash equilibria in games is computationally intractable (cf., Gottlob, Greco, and Scarcello, 2003; Fabrikant, Papadimitriou, and Talwar, 2004; Chen, Deng, and Teng, 2009; Daskalakis, Goldberg, and Papadimitriou, 2009). Nash equilibrium requires each player to select an optimal action in response to the actions of the other players. We ask whether weakening optimality to merely requiring each player to avoid *worst* responses—arguably the weakest meaningful rationality criterion—yields tractable solution concepts. It does not. We study succinctly represented games with pure actions and find:

> *Any solution concept that guarantees that no player plays a worst response to the actions of the other players is computationally intractable.*

In particular, determining the existence of such action profiles, finding one, or counting how many exist is computationally equivalent to the same problems for pure Nash equilibrium. Even for potential games, finding such action profiles is intractable.

Avoiding worst responses is a weak individual rationality requirement. For example, in a congestion game, a player may simply seek to avoid her worst route rather than optimize—yet it is generally intractable to find an action profile in which no player worst-responds. Likewise, even if players were to jointly seek an action profile that all can regard as acceptable—that is, an action profile in which no one plays a worst response—identifying such an agreement is computationally intractable. Taken together, our results suggest that the source of intractability does not stem merely from optimization (as required by Nash equilibrium) but also from requiring a minimal rationality guarantee for each player.

Computational complexity has been argued to provide a minimal plausibility requirement for economic solution concepts: predictions that rely on solving fundamentally hard problems may be descriptively or normatively untenable (Gilboa and Zemel, 1989; Camara, 2022). In game theory, this perspective has put in question the plausibility of (pure) Nash equilibrium, as the concept has been shown to be computationally intractable in a variety of game classes (see discussion in Section 1.2 below). What we show is that any solution concept that requires all players to not worst-respond fails this fundamental plausibility requirement.

1.1 Our contribution

We provide a complexity-theoretic characterization of minimal individual rationality guarantees in succinctly represented finite games played with pure strategies. [1] We show that relaxing optimality alone does not restore computational tractability: computational hardness persists whenever minimal guarantees apply "universally"; that is, to all players simultaneously. We then determine the extent to which universality must be relaxed in order to recover computational tractability.

Concretely, our results establish the following: in succinctly represented games with an arbitrarily fixed number of actions per player, deciding whether there exists an action profile in which no player plays a worst response is computationally as hard as deciding the existence of a pure Nash equilibrium (Theorem 1, NP-complete). The same equivalence holds for finding such an action profile (Corollary 1, NP-hard) and for counting the number of such action profiles (Corollary 2, #P-complete). The NP hardness proofs rely on a reduction from CircuitSAT$^?$. Note that, when each player has two actions, our results follow from known hardness results for pure Nash equilibrium. Our analysis reveals that even with many actions, finding action profiles in which no player plays a worst response is computationally hard.

We then turn our attention to potential games, a class of games with many economic and technical applications. Although potential games are guaranteed to admit pure Nash equilibria (Rosenthal, 1973; Monderer and Shapley, 1996), we show that finding an action profile in which no player plays a worst response is computationally equivalent to finding a pure Nash equilibrium (Theorem 2, PLS-complete). A new reduction is required to prove this result because the one from CircuitSAT$^?$ does not preserve potentialness. A key ingredient to our second reduction is the use of circuit gadgets whose existence we prove using the probabilistic method and the Lovász Local Lemma.

Having established that action profiles with universal guarantees; that is, guarantees for all players, are computationally hard for a variety of problems, we turn to non-universal guarantees. When guarantees are required only for a fraction of players, we identify a sharp threshold separating computational tractability from hardness. Specifically, the problem of finding an action profile in which a fraction at least α of players choose an action that is in their top-β fraction of actions is computationally tractable when $\alpha < \beta$ (Theorem 3, RP), and computationally hard when $\alpha > \beta$ (Theorem 3, NP-hard).[2] In other words, there is a threshold along which there is a tractability

[1] More precisely, we study circuit games. These are games in which the payoff functions are represented by Boolean circuits. See Section 2.2 for details.

[2] We note that the precise formulation of Theorem 3 leaves a small wedge between the two cases.

trade-off between individual rationality guarantees and the fraction of players satisfying them.

From a technical perspective, our results rely on reductions between decision, search, and counting problems in games, and on standard complexity classes such as NP, PLS, and RP. We provide a brief overview of these notions in Section 2.4. In particular, our reductions show that avoiding worst responses is computationally equivalent to pure Nash equilibrium across multiple complexity-theoretic dimensions.

1.2 Literature

Our results relate to several strands of the literature on computational complexity in game theory, bounded rationality, and approximate equilibrium concepts.

Computational complexity of equilibrium. Throughout the computational game theory literature, complexity results are stated for games given in a *succinct representation*, typically by polynomial-size Boolean circuits that compute players' payoffs. This convention rules out trivial hardness arising from exponentially large payoff tables and allows computational hardness to be interpreted as a genuine feasibility constraint rather than an artefact of input size.

The computational study of equilibrium in games originates in the seminal work of Gilboa and Zemel (1989), who showed that certain equilibrium-related problems in finite games can be computationally intractable. Subsequent work identified PPAD as the relevant complexity class for equilibrium computation (Papadimitriou, 2001) and showed that finding mixed Nash equilibria is PPAD-hard even in small games (Chen, Deng, and Teng, 2009; Daskalakis, Goldberg, and Papadimitriou, 2009).

For pure strategies, deciding whether a pure Nash equilibrium exists in a general succinctly represented finite game is NP-complete and the associated search problem is NP-hard (Gottlob, Greco, and Scarcello, 2003). Even in the class of potential games, where pure Nash equilibria are guaranteed to exist (Rosenthal, 1973; Monderer and Shapley, 1996), finding a pure Nash equilibrium is PLS-complete (Fabrikant, Papadimitriou, and Talwar, 2004), and best-response dynamics require exponential convergence time in the worst case (Schäffer and Yannakakis, 1991; Anshelevich, Dasgupta, Kleinberg, Tardos, Wexler, and Roughgarden, 2008). Our results show that action profiles in which no player worst-responds (in pure strategies) are computationally equivalent to pure Nash equilibrium along the same complexity dimensions.

Related work studies the approximate notion of ϵ-equilibrium (Radner, 1980). For mixed strategies, computing an ϵ-equilibrium remains intractable for sufficiently small

ϵ (cf., e.g., Daskalakis, Mehta, and Papadimitriou, 2006; Rubinstein, 2015; Schoenebeck and Vadhan, 2012). For pure strategies, Schoenebeck and Vadhan (2012) establish several hardness results for pure ϵ-equilibria. In this literature, ϵ must be interpreted relative to the scale of utilities, so the restrictiveness of the concept varies across games, players, and action profiles. In contrast, in line with the ordinal nature of pure Nash equilibrium, our notion of not worst-responding provides a universal guarantee across games, action profiles, and players.

Finally, Babichenko, Papadimitriou, and Rubinstein (2016) study approximate equilibria in which only a fraction of players may have large regret and show that this relaxation does not generally restore tractability. Our results on partial guarantees relate to this line of work while focusing on pure actions. Taken together, this literature shows that computational hardness is a pervasive feature of equilibrium concepts, affecting mixed, pure, and approximate formulations alike.

Complexity in economics. Complexity considerations have a long tradition in economic analysis. For example, computation models have been studied in the context of repeated games (Neyman, 1985; Rubinstein, 1986; Kalai and Stanford, 1988), choice (Salant, 2011), and learning (Wilson, 2014). Many articles in this strand of the literature rely on finite automata. Camara (2022) broke new ground by studying choice in a decision-theoretic setting through the lens of computational complexity theory. We adopt this approach to study games.

Bounded rationality. Bounded rationality, or behaviour that is not consistent with optimisation over all possible actions, has been documented in a wide range of settings (cf. Artinger, Gigerenzer, and Jacobs, 2022; de Clippel and Rozen, 2024, for recent surveys). In games, a variety of solution concepts aim to capture boundedly rational behaviour, including ϵ-equilibrium (Radner, 1980), and $\mathbf{k}$-satisficing equilibrium (Pradelski and Tarbush, 2025), where $\mathbf{k} = (k_i)_i$ is a vector of positive integers, one for each player i. The latter relaxes pure Nash equilibrium by requiring each player i to play one of her top k_i actions in response to the actions of the other players. Our results show that if, for each player i, k_i is strictly smaller than the number of i's actions, then the problem of determining the existence of, finding, or counting such action profiles is generally intractable.

2 Preliminaries

We introduce our notation and recall standard notions from game theory and computational complexity theory.

2.1 Games

A *game* is a tuple

$$\left(N, \{A_i\}_{i \in N}, \{u_i\}_{i \in N} \right)$$

consisting of a set of players $N = \{1, \ldots, n\}$, a set of actions $A_i = \{0, \ldots, m_i - 1\}$ for each player $i \in N$, as well as a payoff function $u_i : A \to \mathbb{R}$ for each $i \in N$, where $A := \times_{i \in N} A_i$ denotes the set of all *action profiles*. In a standard abuse of notation, we write $(a_i, a_{-i}) \in A$ for the action profile in which player i plays action $a_i \in A_i$ and $a_{-i} \in \times_{j \in N \setminus \{i\}} A_j$ denotes the actions played by the players in $N \setminus \{i\}$.

We say that an action $a_i \in A_i$ is a *best-response* to a_{-i} for player i if $u_i(a_i, a_{-i}) \geq u_i(a'_i, a_{-i})$ for all $a'_i \in A_i$, and it is a *worst-response* if $u_i(a_i, a_{-i}) < u_i(a'_i, a_{-i})$ for all $a'_i \in A_i \setminus \{a_i\}$. For any $\beta \in [0, 1]$ and any action profile a, we say that $a_i \in A_i$ is a *top β action* for player i in response to a_{-i} if the number of actions a'_i for which $u_i(a'_i, a_{-i}) > u_i(a_i, a_{-i})$ is less than $\beta \cdot m_i$. In other words, a_i is a top β action for player i if it is one of the best $\lceil \beta \cdot m_i \rceil$ actions of player i in response to a_{-i}.

An action profile $a \in A$ is a *pure Nash equilibrium* if, for each player $i \in N$, a_i is a best-response to a_{-i}. An action profile $a \in A$ is a *no-worst-response action profile* if, for each $i \in N$, a_i is not a worst-response to a_{-i}. Finally, for any game $(N, \{A_i\}_{i \in N}, \{u_i\}_{i \in N})$, action profile $a \in A$, and any $\alpha, \beta \in [0, 1]$, we say that a *fraction at least α of players play a top β action* at a if there is a set $S \subseteq N$ with $|S| \geq \lceil \alpha \cdot n \rceil$ such that, for each player $i \in S$, a_i is a top β action in response to a_{-i}.

2.2 Boolean circuits and circuit games

A *Boolean circuit* C represents a multivariate function that takes n binary inputs, $(x_1, \ldots, x_n)$, with 1 representing *true* and 0 representing *false*. The circuit has one or more outputs, which are computed through a sequence of three types of elementary logical gates: AND ($\wedge$), OR ($\vee$), and NOT ($\neg$). The fan-in of each gate is at most 2, so it accepts no more than two input wires. The circuit forms a directed acyclic graph: the n input bits feed into gates, whose outputs feed into subsequent gates, until the final gates produces the circuit's outputs. See Figure 1 for an example with one output gate. The *size* of the circuit C is the number of its gates. We write $\tau \in \{0, 1\}^n$ for a truth assignment to the inputs $x_1, \ldots, x_n$, and $C(\tau)$ for the output of circuit C given input τ.

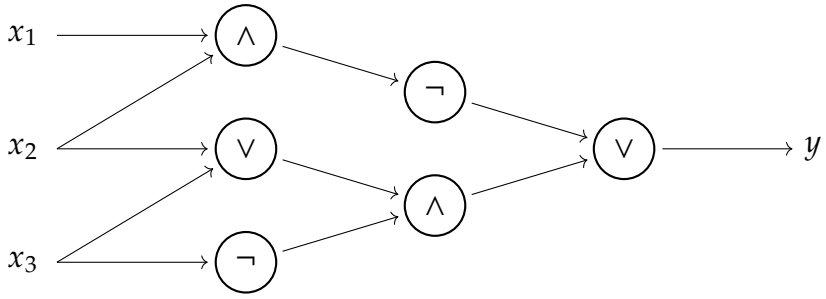


Figure 1: A Boolean circuit of size 6 and depth 3 with three inputs (x_1, x_2, x_3) and one output. The gates are $\wedge$ (AND), $\vee$ (OR), and $\neg$ (NOT). Truth assignment $\tau = (1, 1, 1)$ evaluates to 0 (false), while assignment $\tau' = (1, 1, 0)$ evaluates to 1 (true).

When C is a circuit with one output, then we say that truth assignment τ *satisfies* C if $C(\tau) = 1$.

Circuits will also provide the computational model for our games. Specifically, we work within the formalism of Schoenebeck and Vadhan (2012) and assume throughout that our games are circuit games, that is, the utility functions u_i in our games are represented by Boolean circuits have multiple outputs that encode two integers whose quotient represent rational payoffs. We let $\mathbf{G}$ denote the set of all such games, and $\mathbf{A}(\mathbf{G})$ be the set of all action profiles of all games in $\mathbf{G}$.

Definition 1 (Feigenbaum, Koller, and Shor 1995; Schoenebeck and Vadhan 2012)**.** A *circuit game* is a game $(N, \{A_i\}_{i\in N}, \{u_i\}_{i\in N})$ in which the number of actions m_i of each agent are represented by binary strings, and every payoff function u_i is encoded by a circuit C_i with k_i output gates, $\sum_{i\in N} s_i$ inputs such that $A_i \subseteq \{0,1\}^{s_i}$ for each player i and $C_i(a) = u_i(a)$ for all $a \in A$. We assume that the payoffs of each u_i are the quotient of two integers encoded (by binary strings) by the output of C_i, and thus the payoffs are rational numbers.

2.3 Computational problems

For our computational problems, we provide a high-level description of the central concepts (without, for example, concerning ourselves with encodings into languages). We identify a *computational problem* with a relation $f \subseteq \texttt{In} \times \texttt{Out}$ whose inputs consist of elements of a set `In` and whose outputs are elements of a set `Out`. The 'circuit satisfiability' problem, for example, is the relation $\texttt{CircuitSAT}^? \subseteq \Sigma \times \{0,1\}$ that maps from the set of all Boolean circuits Σ to $\{0,1\}$, where $(C, 1) \in \texttt{CircuitSAT}^?$ if and only if circuit C has a satisfying truth assignment.

In the context of games, we will be primarily interested in three classes of computational problems. The first are decision problems (denoted ?), the second are counting problems (denoted #), and the third are search problems (denoted Q).

Definition 2 (Decision problems: $\texttt{PNE}^?$, $\texttt{NWR}^?$, and $\langle\alpha,\beta\rangle^?$)**.**
The problem of deciding whether there exists a pure Nash equilibrium is the relation $\texttt{PNE}^? \subseteq \mathbf{G} \times \{0,1\}$ where $(g,1) \in \texttt{PNE}^?$ if and only if $g \in \mathbf{G}$ admits a pure Nash equilibrium. Similarly, $\texttt{NWR}^? \subseteq \mathbf{G}\times\{0,1\}$ denotes the problem of deciding whether there exists a no-worst-response action profile. Finally, for any $\alpha,\beta \in [0,1]$, $\langle\alpha,\beta\rangle^? \subseteq \mathbf{G}\times\{0,1\}$ denotes the problem of deciding whether there exists an action profile in which a fraction at least α of players play a top β action.

Definition 3 (Search problems: $\texttt{PNE}^{\mathrm{Q}}$, $\texttt{NWR}^{\mathrm{Q}}$, and $\langle\alpha,\beta\rangle^{\mathrm{Q}}$)**.**
The problem of finding a pure Nash equilibrium is the relation $\texttt{PNE}^{\mathrm{Q}} \subseteq \mathbf{G} \times \mathbf{A}(\mathbf{G})$ defined by $(g,a) \in \texttt{PNE}^{\mathrm{Q}}$ if and only if the action profile a is a pure Nash equilibrium of $g \in \mathbf{G}$. An algorithm solves $\texttt{PNE}^{\mathrm{Q}}$ if, on input $g \in \mathbf{G}$, it outputs some a such that $(g,a) \in \texttt{PNE}^{\mathrm{Q}}$ or it returns "infeasible" if the game g does not have a pure Nash equilibrium. The problem $\texttt{NWR}^{\mathrm{Q}}$ is defined analogously. Finally, $\langle\alpha,\beta\rangle^{\mathrm{Q}}$ is the search problem of finding an action profile in which a fraction at least α of players play a top β action.

Definition 4 (Counting problems: $\texttt{PNE}^\#$, $\texttt{NWR}^\#$, and $\langle\alpha,\beta\rangle^\#$)**.**
The problem of counting the number of pure Nash equilibria is the relation $\texttt{PNE}^\# \subseteq \mathbf{G} \times \{0,1,2,\dots\}$ where $(g,k) \in \texttt{PNE}^\#$ if and only if $g \in \mathbf{G}$ admits exactly k pure Nash equilibria. Similarly, $\texttt{NWR}^\# \subseteq \mathbf{G} \times \{0,1,2,\dots\}$ denotes the problem of counting the number of no-worst-response action profiles, and $\langle\alpha,\beta\rangle^\# \subseteq \mathbf{G} \times \{0,1,2,\dots\}$ denotes the problem of counting the number of action profiles in which a fraction at least α of players play a top β action.

Definition 5 (Parametrised problems)**.** For each of our computational problems on games defined above, f, and each positive integer m, we define the parametrised problem m-f to be the problem f within the class of games with m actions per player. For example, m-$\texttt{NWR}^{\mathrm{Q}}$ denotes the problem of finding a no-worst-response action profile in games with m actions per player.

2.4 Computational complexity

We say that a computational problem is *solvable* if there is an algorithm that computes its relation. For example, we say that $\texttt{PNE}^?$ is solvable if there is an algorithm that, given any $g \in \mathbf{G}$, correctly determines whether g has a pure Nash equilibrium.

For any computational problems $f \subseteq \texttt{In} \times \texttt{Out}$ and $f' \subseteq \texttt{In}' \times \texttt{Out}'$, a *polynomial-time reduction from f to f'* is a pair $(\texttt{push}, \texttt{pull})$ of polynomial-time computable functions $\texttt{push} : \texttt{In} \to \texttt{In}'$ and $\texttt{pull} : \texttt{In} \times \texttt{Out}' \to \texttt{Out}$ such that for all $x \in \texttt{In}$ and $y' \in \texttt{Out}'$,

$$(\texttt{push}(x), y') \in f' \text{ implies } (x, \texttt{pull}(x, y')) \in f,$$

and the existence of some y such that $(x, y) \in f$ implies that there is some y' for which $(\texttt{push}(x), y') \in f'$; i.e., solvable instances of f are reduced to solvable instances of f'. In cases where the `pull` function is independent of the original problem instance, we simply take it that $\texttt{pull} : \texttt{Out}' \to \texttt{Out}$.

The reduction turns any instance of problem f into an instance of problem f' and the answer of problem f' is turned into an answer of problem f. This means that if one can solve f', then one can use the reduction $(\texttt{push}, \texttt{pull})$ to also solve f; hence (provided `push`, `pull` are easy to compute) f' is at least as "hard" as f. See the sketch in Figure 2.

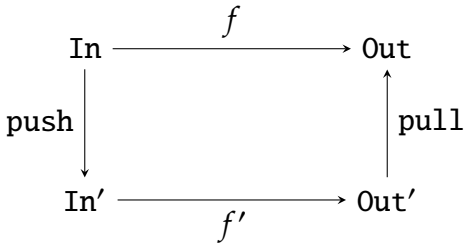


Figure 2: Reduction from f to f'.

2.4.1 Complexity of decision problems

The class `NP` is the class of all decision problems whose solutions, when guessed, can be verified in polynomial time. As an example, we can show that $\texttt{PNE}^? \in \texttt{NP}$; consider any $g \in \mathbf{G}$, and guess an arbitrary action profile a. We can check whether any unilateral deviations are utility improving for any player in polynomial time. If not, return true; otherwise, return false.

The class `P` is the class of all decision problems that are solvable by a deterministic polynomial-time algorithm. The $\texttt{P} \neq \texttt{NP}$ conjecture asserts that the class of problems in `P` is strictly smaller than the class `NP`, that is, there exist problems whose solutions can be efficiently verified but not efficiently found.

A problem f' is `NP-hard` if, for every problem $f \in \texttt{NP}$, there is a polynomial-time reduction from f to f'. A decision problem f' is `NP-complete` if $f' \in \texttt{NP}$ and is `NP-hard`. Moreover, polynomial-time reduction s are transitive. To show that a problem f' is

NP-complete, it therefore suffices to show that the problem is in NP and that there is a polynomial-time reduction from an NP-complete problem f to f'. The decision problem CircuitSAT$^?$, deciding whether a given circuit C is satisfiable, is NP-complete (Karp, 1972).

Finally, the class RP is the class of all decision problems for which there is a polynomial-time randomized algorithm (a false-biased Monte Carlo algorithm) that returns 0 on any 0 instance, and that falsely returns 0 on a 1 instance with probability at most $\frac{1}{2}$.

2.4.2 Complexity of search problems

We will also be interested in the complexity class PLS (Polynomial Local Search). A problem belongs to PLS if it can be formulated as searching for a local optimum in a landscape where:

(i) *Solutions have polynomial length.* Every feasible solution can be encoded in a string whose length is polynomial in the instance size.

(ii) *Feasible solutions are easy to generate.* There is a polynomial-time algorithm that, given an instance, produces some valid starting solution.

(iii) *There is an objective function.* Each feasible solution s is assigned a value $v(s) \in \mathbb{Z}$.

(iv) *Solutions are easy to verify and evaluate.* Given a candidate solution, one can check in polynomial time whether it is feasible and compute $v(s)$.

(v) *There is a neighbourhood structure.* There is a relation Γ on feasible solutions; we say s' is a *neighbour* of s if $(s, s') \in \Gamma$. A solution s is a *local optimum* if $v(s) \geq v(s')$ for all neighbours s' of s.

(vi) *Neighbours are easy to enumerate and compare.* For any solution, there is a polynomial-time algorithm that either certifies the solution is a local optimum (no neighbour improves the objective) or returns a strictly better neighbouring solution.

Because any improving step increases the objective and the objective is bounded, the process must terminate at a local optimum (thus implying the existence of local optima). However, the number of improving steps may be exponential, so membership in PLS does not guarantee an efficient algorithm; only that the *structure* of local search applies.

A problem is PLS-hard if every problem in PLS can be reduced to it in polynomial time. A PLS-reduction from problem f to problem f' consists of a polynomial-time

reduction (push, pull) with additional structure. The push function maps instances of f to instances of f', and the second function pull maps an instance I of f and a solution to an instance of f' back to a solution of the original instance, such that: for any instance I of f, and any local optimum s^* of push(I), the solution pull(I, s^*) is a locally optimal solution of I. A problem f' is PLS-complete if $f' \in$ PLS and is PLS-hard. Moreover, PLS-reductions are transitive. To show that a problem f' is PLS-complete, it therefore suffices to show that the problem is in PLS and that there is a reduction from a PLS-complete problem f to f'. It is known that finding a pure Nash equilibrium in (succinctly represented) potential games is PLS-complete (Fabrikant, Papadimitriou, and Talwar, 2004).

Establishing PLS-completeness for a problem is strong evidence that no polynomial-time algorithm finds a local optimum in general, because such an algorithm would solve all local-search problems efficiently. Canonical PLS-complete problems include finding a pure Nash equilibrium in potential games (Fabrikant, Papadimitriou, and Talwar, 2004). In this setting, when players myopically best-respond, convergence of a local search is guaranteed but may take an exponential number of steps in the number of players.

2.4.3 Complexity of counting problems

The complexity class #P is the analogue of NP for counting problems: it consists of all counting problems whose solutions, when guessed, can be verified in polynomial time.

A counting problem f' is #P-hard if for every problem $f \in$ #P, there is a polynomial-time computable reduction from f to f' that preserves counts. A problem f' is #P-complete if $f' \in$ #P and is #P-hard. Moreover, #P-reductions are transitive. To show that a problem f' is #P-complete, it therefore suffices to show that the problem is in #P and that there is a reduction from a #P-complete problem f to f'. The counting version of CircuitSAT$^?$ is the relation CircuitSAT$^\#$ $\subseteq \Sigma \times \{0, 1, 2, \dots\}$ that maps from the set of all Boolean circuits Σ to the non-negative integers, where $(C, k) \in$ CircuitSAT$^\#$ if and only if C admits k satisfying truth assignments. The CircuitSAT$^\#$ problem is known to be #P-complete.

3 Decision and counting problems in general games

We begin our investigation of the intractability of not worst-responding by studying the problem of determining whether a given circuit game has a no-worst-response action

profile. Here, and throughout, we show hardness for the computational problems in which the number of actions for all players are fixed to some integer $m \geq 2$. This also implies hardness for the non-parametric version of the problems, in which the number of actions is part of the input.

Theorem 1. *For every integer $m \geq 2$, the decision problem m-$\texttt{NWR}^?$ of determining the existence of a no-worst-response action profile, restricted to games with m actions per player, is* $\texttt{NP-complete}$.

It is easy to check that $\texttt{NWR}^? \in \texttt{NP}$, as we can guess a no-worst-response action profile and verify in polynomial time that no player is worst-responding. To show that $\texttt{NWR}^?$ is $\texttt{NP-hard}$, we define a reduction from $\texttt{CircuitSAT}^?$ and then show that it is polynomial. Note that our proof implies that no-worst-response action profiles are not guaranteed to exist (this was previously pointed out by Pradelski and Tarbush, 2025).

Reduction 1 (Reduction from $\texttt{CircuitSAT}^?$ to $\texttt{NWR}^?$)**.** Let C be a Boolean circuit with inputs $x_1, \ldots, x_n$, and fix some $m \geq 2$. We construct a game g from C with nm players (i, ℓ), where $i \in \{1, \ldots, n\}$ and $\ell \in \{1, \ldots, m\}$. Each player has actions $A_{i\ell} = \{0, \ldots, m-1\}$.

To construct the payoff function of each player, we first define a function $f : A \to \{0, 1\}$ as follows: let $f(a) = 1$ if, for each $i \in \{1, \ldots, n\}$, the players $(i, 1), \ldots, (i, m)$ play the same action $a_{i\ell} \in \{0, 1\}$ and $C(\tau) = 1$ for τ defined by $\tau_i = a_{i\ell}$. Otherwise, let $f(a) = 0$.

This allows us to define payoffs as

$$u_{i\ell}(a) = \underbrace{(m+1)f(a)}_{\text{global term}} + \underbrace{\left[\left(\ell + \sum_{\ell'=1}^{m} a_{i\ell'}\right) \bmod m\right]}_{\text{local term}}, \tag{1}$$

which completes the description of the reduction. △

Intuitively, the global term is a potential which reflects whether the actions map to a satisfying truth assignment of the circuit C. Meanwhile, for each group $(i, \ell)_{\ell \in \{1,\ldots,m\}}$, the local term is the utility function of a harmonic game (Candogan, Menache, Ozdaglar, and Parrilo, 2011). This leads to the following straightforward observation.

Observation 1. Each player's payoff from action profile a is at least $m + 1$ if $f(a) = 1$ and at most m if $f(a) = 0$. This implies the following for any action profile a.

(i) If $f(a) = 1$, then each player gets a payoff of at least $m + 1$ from the global term. By unilaterally deviating to, say, action $a'_{i\ell}$, the player forces $f(a'_{i\ell}, a_{-i\ell}) = 0$, so her payoff strictly reduces.

(ii) If $f(a) = 0$, then for each group $(i, \ell)_{\ell \in \{1,\dots,m\}}$, one player has a payoff of 0, as the global term in (1) is 0 and the local term as a function of ℓ is a bijection from $\{1, \dots, m\}$ to $\{0, \dots, m-1\}$. Any such player with payoff 0 can profitably deviate by playing any other action in $A_{i\ell} \setminus \{a_{i\ell}\}$, so she is strictly worst-responding at a.

Proof of Theorem 1. Let C be a Boolean circuit with inputs $(x_1, \dots, x_n)$, and let g be the game as constructed above. Suppose circuit C is unsatisfiable. By construction of f, we have $f(a) = 0$ for any action profile, and so Observation 1 tells us that every action profile has some players who worst-respond.

Now suppose that C is satisfiable, and let τ^* be a satisfying truth assignment for C. Define the action profile a given by $a_{i\ell} = \tau_i^*$ for each $i \in \{1, \dots, n\}$ and $\ell \in \{1, \dots, m\}$. It is immediate by construction that $f(a) = 1$, so by Observation 1 we see that each player gets a payoff of at least $m + 1$. Moreover, unilaterally deviating strictly reduces a player's payoff, so a is a pure Nash equilibrium (and so also a no-worst-response action profile).

In summary, $g \in \texttt{NWR}^?$ if and only if $C \in \texttt{CircuitSAT}^?$. Finally, it is straightforward to construct Boolean circuits to compute each $u_{i\ell}$ in polynomial time, so we arrive at a polynomial-time reduction. □

Since $\texttt{NWR}^{\mathsf{Q}}$ cannot be easier than $\texttt{NWR}^?$, we immediately obtain the following corollary.

Corollary 1. *For every integer $m \geq 2$, the search problem m-$\texttt{NWR}^{\mathsf{Q}}$ of finding a no-worst-response action profile, restricted to games with m actions per player, is* NP-hard.

The game g constructed from circuit C has k no-worst-response action profiles (in particular, k pure Nash equilibria) exactly when there are k satisfying assignments for C. Moreover, it is easy to check that $\texttt{NWR}^\# \in \#\texttt{P}$. And, since $\texttt{CircuitSAT}^\#$ is #P-complete, we see the same for $\texttt{NWR}^\#$.

Corollary 2. *For every integer $m \geq 2$, the counting problem m-$\texttt{NWR}^\#$ of determining the number of no-worst-response action profiles, restricted to games with m actions per player, is* #P-complete.

Proof of Corollary 2. In the reduction from circuit C to game g, we see that all action profiles a for g have a worst-responding player if $f(a) = 0$. This is the case in particular if C is not satisfiable. Moreover, when C is satisfiable, a is a no-worst-response action profile if and only if $f(a) = 1$. It is straightforward that these correspond one-to-one to satisfying truth assignments for C, so $\texttt{NWR}^\#$ is #P-hard. □

4 Complexity of not-worst-responding in potential games

We now turn our attention to potential games.

Definition 6. A game $(N, \{A_i\}_{i\in N}, \{u_i\}_{i\in N}) \in \mathbf{G}$ is a *potential game* if there exists a function $\phi : A \to \mathbb{R}$ such that for each $a \in A$, each $i \in N$, and each $a'_i \in A_i$, we have $u_i(a'_i, a_{-i}) - u_i(a) = \phi(a'_i, a_{-i}) - \phi(a)$.

It is well-known that potential games are guaranteed to have a pure Nash equilibrium (Rosenthal, 1973; Monderer and Shapley, 1996). Moreover, the search problem of finding such a Nash equilibrium is known to be PLS-complete (Fabrikant, Papadimitriou, and Talwar, 2004). Our main result in this section will be that finding no-worst-response action profiles in potential games has the same complexity. To this end, we first remark that finding a Nash equilibrium in potential games is PLS-complete even if each player has two actions. This allows us to reduce from this restricted problem to NWR$^{\mathsf{Q}}$.

Fact 1 (following from Pardalos and Jha 1992)**.** The search problem 2-PNE$^{\mathsf{Q}}$ of finding a pure Nash equilibrium, restricted to potential games with two actions per player, is PLS-complete.

For completeness, we discuss how this follows. First, it is immediate that the problem lies in PLS. To show hardness, we reduce from the local binary quadratic program problem. Recall that a binary quadratic program is specified by an $n \times n$ matrix Q and a length n vector q, both integral. The problem asks to find an optimal solution of

$$\max_{x\in\{0,1\}^n} q^\top x + \frac{1}{2} x^\top Q x. \tag{2}$$

The local variant, in turn, asks to find a 0-1 vector x such that no change in a single coordinate improves the value of the solution. Pardalos and Jha (1992) show that this search problem is PLS-complete. This local optimisation variant admits an obvious representation as a potential game with n players where each player $i \in N = \{1, 2, \ldots, n\}$ has action set $A_i = \{0, 1\}$, and payoff functions are given by

$$u_i(a) = \sum_{j=1}^{n} \left[q_j a_j + \sum_{k=1}^{n} \frac{1}{2} Q_{jk} a_j a_k \right].$$

Then the set of pure Nash equilibria of the game is in bijection with the set of local optima of (2).

We are now ready to develop our main result in this section.

Theorem 2. *For every integer $m \geq 2$, the search problem m-NWR$^{\mathsf{Q}}$ of finding a no-worst-response action profile, restricted to potential games with m actions per player, is* PLS-complete.

Theorem 2 follows from a series of lemmas. Lemma 1 below establishes PLS membership. The remaining arguments and results in this section establish PLS-hardness.

Lemma 1. For every integer $m \geq 2$, the search problem m-NWR$^{\mathsf{Q}}$ of finding a no-worst-response action profile, restricted to potential games with m actions per player, is in PLS.

Proof of Lemma 1. Suppose we have a potential game with an ordered player set $N = \{1, 2, ..., n\}$ and m actions per player. Firstly, note that the number of actions for each player is constant and in our model it is straightforward to generate an initial action profile a^*. Secondly, we see that we may evaluate the potential ϕ up to a constant term. Indeed, by assumption, we have access to a valid starting solution $a^* \in A$. Let a' be any other action profile. We shall iteratively let $a^k \in A$ be defined, for $0 \leq k \leq n$ and for $a_i^k = a_i'$ if $i \leq k$ and $a_i^k = a_i^*$ otherwise. Now, by the potential game property,

$$\phi(a') - \phi(a^*) = \sum_{k=1}^{n} \phi(a^k) - \phi(a^{k-1}) = \sum_{k \in N} u_k(a^k) - u_k(a^{k-1}).$$

Therefore, the potential is computable at any action profile a', up to a constant which renders $\phi(a^*) = 0$. Since all utilities are rational numbers whose numerators and denominators have bit representation of length bounded in the size of the circuit, two facts follow. First we notice that, denoting by $|C| = \max_i |C_i|$ the maximum size of players' circuits, for any $k = 1, 2, ..., n$, the bit representation of $u_k(a^k) - u_k(a^{k-1})$ requires at most $4|C|+1$ bits; with $2|C|+1$ bits for the numerator and $2|C|$ bits for the denominator. This is because we have $u_i(a^k) - u_i(a^{k-1}) = a_k/b_k + c_k/d_k = (a_k d_k + b_k c_k)/(b_k d_k)$, where (a_k, b_k, c_k, d_k) have bit representations of size $\leq |C|$. As a consequence, $\phi(a') - \phi(a_*)$ requires $O(n|C|)$ bits in its expression as the ratio of two integers in binary. Second, for any action profile a', the potential function can take value at most $n \cdot 2^{|C|}$. This is because, in the above we have $u_i(a^k) - u_i(a^{k-1}) \leq 2^{|C|}$ for each k due to the fact that the utilities are represented as the ratio of two numbers in binary representation, output by a circuit of size $\leq |C|$.

We shall therefore fix an auxiliary potential for the game, letting $\hat{\phi}(a) = \phi(a) - \phi(a^*)$ if some player i is worst-responding, and $\hat{\phi}(a) = n \cdot 2^{|C|} + 1$ otherwise. At each action profile a, the potential may be evaluated in polynomial time, by evaluating $u_i(a_i', a_{-i})$ for each player i and each action $a_i' \in A_i$ and checking for the worst-response property, which requires $1 + (m-1)n$ utility evaluations and $(m-1)n$ pairwise comparisons of

the values of the utilities. Checking for a local improvement in $\hat{\phi}$ is similarly possible in polynomial time. □

We now turn to showing PLS-hardness. We note that the reduction from $\texttt{CircuitSAT}^?$ we invoked in the proof of Theorem 1 is insufficient to infer Theorem 2; in particular, the local term of (1) is cyclic, so the reduced game is not a potential game. Thus, to prove hardness, our first step is to develop a PLS-reduction ($\texttt{push}, \texttt{pull}$) from $\texttt{PNE}^{\mathsf{Q}}$ to $\texttt{NWR}^{\mathsf{Q}}$ that preserves the potential property of the game.

Reduction 2 (Reduction from $\texttt{PNE}^{\mathsf{Q}}$ to $\texttt{NWR}^{\mathsf{Q}}$)**.** Fix integers $\hat{m} \geq 2$ and $q = 6\hat{m}2^{\hat{m}}$. Suppose we are given a game $g = (N, \{A_i\}_{i\in N}, \{u_i\}_{i\in N}) \in \mathbf{G}$ with 2 actions per player. [3]

From this, we now construct a game $\hat{g} = (\hat{N}, \{\hat{A}_i\}_{i\in\hat{N}}, \{\hat{u}_i\}_{i\in\hat{N}}) \in \mathbf{G}$ with $\hat{m}$ actions per player as follows. The set of players

$$\hat{N} \coloneqq \{(i, \ell) : i \in N \text{ and } \ell \in \{1, \ldots, q\}\}$$

consists of n groups $[i] \coloneqq \{(i, \ell) : \ell \in \{1, \ldots, q\}\}$, one for each $i \in N$. Each player (i, ℓ) has action set $\hat{A}_{i\ell} \coloneqq \{0, \ldots, \hat{m} - 1\}$. For any group $[i]$, define $a_{[i]}$ as the restriction of action profile a to the group's players, and let $\hat{A}_{[i]} = \{0, \ldots, \hat{m}-1\}^q$ be the group's action space. To define the payoff functions, we first introduce a *circuit gadget* $\rho_i : \hat{A}_{[i]} \to A_i$ for each $i \in N$. This gadget aggregates the actions $\hat{a}_{[i]}$ of a group in the game $\hat{g}$ into a single action $\rho_i(\hat{a}_{[i]})$ of player i in the game g. By applying the gadget to all groups' actions, we map an action profile $\hat{a}$ in the new game $\hat{g}$ back to an action profile $\rho(\hat{a})$ in the original game g by defining $\rho : \hat{A} \to A$ as

$$\rho(\hat{a}) = (\rho_i(\hat{a}_{[i]}))_{i\in N}.$$

We now define the payoff of each player (i, ℓ) in game $\hat{g}$ as

$$\hat{u}_{i\ell}(\hat{a}) = u_i(\rho(\hat{a})). \tag{3}$$

Crucially, we require that each gadget ρ_i satisfies the following property ($\star$): at each action profile in $\hat{A}_{[i]}$ among players in the group $[i]$, and for each action a_i not currently taken by player i in g, there is some player in the group for whom any action change must change the action of player i to a_i. Formally:

($\star$) For each $i \in N$, each $\hat{a}_{[i]} = (\hat{a}_{i\ell}, \hat{a}_{-i\ell}) \in \hat{A}_{[i]}$, and each $a_i \in A_i \setminus \{\rho_i(\hat{a}_{[i]})\}$, there is

[3] We note here that the reduction, and the subsequent proofs, extend immediately to any fixed number m of actions for the players in the game g, but for our purposes $m = 2$ suffices.

some $\ell \in \{1, \ldots, q_i\}$ such that $\rho_i(\hat{a}'_{i\ell}, \hat{a}_{-i\ell}) = a_i$ holds for each $\hat{a}'_{i\ell} \neq \hat{a}_{i\ell}$.

Setting $m = 2$ in Lemma 2 below establishes the existence of circuit gadgets satisfying ($\star$). △

Lemma 2. Fix integers $m, \hat{m} \geq 2$, and $q \geq 16m^{\hat{m}-1}\hat{m} \ln m$. There exists a function $f : \{1, \ldots, \hat{m}\}^q \to \{1, \ldots, m\}$ satisfying the following: for any $x \in \{1, \ldots, \hat{m}\}^q$ and any $k \in \{1, \ldots, m\} \setminus \{f(x)\}$, there is some $\ell \in \{1, \ldots, q\}$ such that $f(x'_\ell, x_{-\ell}) = k$ for each $x'_\ell \neq x_\ell$.

Proof of Lemma 2. We start by arguing that the problem is identical to the existence of a certain grid colouring. For any vertex x in a grid $G := \{1, \ldots, \hat{m}\}^q$ and any $\ell \in \{1, \ldots, q\}$, we refer to

$$L_\ell(x) := \{(z, x_{-\ell}) : z \in \{1, \ldots, \hat{m}\}\}$$

as the *line* going through x in the ℓth direction. For each vertex x in the grid, we assign a *colour* $f(x) \in \{1, \ldots, m\}$. We'll say that a vertex $x \in G$ is *bad* if there is some $k \neq f(x)$ such that, for every direction ℓ, there is some vertex $y \in L_\ell(x) \setminus \{x\}$ for which $f(y) \neq k$. Proving the lemma is therefore equivalent to showing that there is some value of q (which we determine below) and a function f that assigns colours from $\{1, \ldots, m\}$ to vertices in the grid G such that no vertex is bad.

We now show that such a colouring exists via the probabilistic method, using the Lovász Local Lemma (LLL); see Mitzenmacher and Upfal (2005). For each vertex in the grid G, independently assign a colour uniformly at random from $\{1, \ldots, m\}$. For each $x \in G$, let B_x denote the event that x is bad. For the events $\{B_x\}_{x \in G}$, the LLL states that if (i) $\mathbb{P}[B_x] \leq p$ for each $x \in G$ and (ii) each B_x is independent of all the other events except for at most d of them, then

$$e(d + 1)p \leq 1$$

implies that there is a non-zero probability that none of the vertices are bad. We next determine values for p and d (in terms of q) in our problem. It is then sufficient to choose a value of q that is large enough such that the LLL condition holds.

(i) Fix a vertex $x \in G$ and a colour $k \in \{1, \ldots, m\}$. Fixing some direction $\ell \in \{1, \ldots, q\}$, the probability that every $y \in L_\ell(x) \setminus \{x\}$ has colour k is $\lambda := (\frac{1}{m})^{\hat{m}-1}$. The probability that every direction has some vertex whose colour is not k is therefore $(1 - \lambda)^q$. By a union bound over all colours $k \neq f(x)$, we obtain that

$$\mathbb{P}[B_x] \leq (m - 1)(1 - \lambda)^q \leq me^{-\lambda q} =: p.$$

(ii) The event B_x depends only on the colours of vertices lying on a line through x. Therefore, the events B_x and $B_{x'}$ are independent whenever x and x' differ in at least three entries (because then there is no vertex y that lies on a line through both x and x'). It follows that the event B_x is mutually independent of all but at most d other events, where

$$d \leq 1 + q(\hat{m} - 1) + \binom{q}{2}(\hat{m} - 1)^2 \leq q^2\hat{m}^2 - 1.$$

The term 1 counts x itself, the term $q(\hat{m} - 1)$ counts all vertices at a Hamming distance of 1 from x, and the term $\binom{q}{2}(\hat{m} - 1)^2$ counts all vertices at a Hamming distance of 2 from x.

Using the upper bounds from arguments (i) and (ii) above, the LLL condition becomes $eq^2\hat{m}^2me^{-\lambda q} \leq 1$. Re-arranging gives us

$$\frac{q^2}{e^{\lambda q}} \leq \frac{1}{e\hat{m}^2m}.$$

Since $x^2 \leq e^x$ for all $x > 0$, we can set $x = \frac{\lambda}{2}q$ to conclude that $\frac{q^2}{e^{\lambda q}} \leq \frac{4}{\lambda^2}e^{-\frac{\lambda}{2}q}$. For the LLL condition to hold, it is therefore sufficient to have

$$\frac{4}{\lambda^2}e^{-\frac{\lambda}{2}q} \leq \frac{1}{e\hat{m}^2m},$$

which re-arranges to

$$q \geq \frac{2}{\lambda}\left(\ln(4e\hat{m}^2m) - 2\ln\lambda\right) = 2m^{\hat{m}-1}\left(1 + \ln(4) + 2\ln(\hat{m}) + (2\hat{m} - 1)\ln(m)\right).$$

Upper bounding each element of the sum in parentheses by $2\hat{m}\ln(m)$ gives a crude upper bound of $(1 + \ln(4) + 2\ln(\hat{m}) + (2\hat{m} - 1)\ln(m)) \leq 8\hat{m}\ln(m)$. For the LLL condition to hold, it is therefore sufficient that

$$q \geq 16m^{\hat{m}-1}\hat{m}\ln(m),$$

and this completes the proof. □

We now prove that Reduction 2 is indeed a PLS-reduction, in a series of lemmas.

Lemma 3. If game g is a potential game, then $\hat{g}$ is a potential game.

Proof of Lemma 3. Suppose game g is a potential game, and let ϕ be its potential function. Fix a player (i, ℓ) and an action profile $\hat{a} = (\hat{a}_{i\ell}, \hat{a}_{-i\ell})$ in the game $\hat{g}$. Now, let

$\hat{a}' = (\hat{a}'_{i\ell}, \hat{a}_{-i\ell})$ be an action profile in which the player unilaterally deviates from $\hat{a}$. We also define $a = \rho(\hat{a})$ and $a' = \rho(\hat{a}')$. By construction of ρ, we see that $a' = (\rho_i(\hat{a}'_{[i]}), a_{-i})$. Applying the payoff definition (3), we see that

$$\begin{aligned} u_{i\ell}(\hat{a}'_{i\ell}, \hat{a}_{-i\ell}) - u_{i\ell}(\hat{a}) &= u_i(\rho(\hat{a}'_{i\ell}, \hat{a}_{-i\ell})) - u_i(\rho(\hat{a})) \\ &= u_i(a'_i, a_{-i}) - u_i(a) \\ &= \phi(a'_i, a_{-i}) - \phi(a) \\ &= \phi(\rho(\hat{a}'_{i\ell}, \hat{a}_{-i\ell})) - \phi(\rho(\hat{a})). \end{aligned}$$

So $\hat{\phi}$ defined by $\hat{\phi}(\hat{a}) = \phi(\rho(\hat{a}))$ for any action profile $\hat{a}$ of game $\hat{g}$ is a potential function for game $\hat{g}$, proving that $\hat{g}$ is a potential game. □

Every potential game has a pure strategy Nash equilibrium, which is necessarily a local maximum of the potential function (Monderer and Shapley, 1996). Since any pure Nash equilibrium is necessarily a no-worst-response action profile, we infer the existence of the latter:

Corollary 3. *Game $\hat{g}$ admits a no-worst-response action profile $\hat{a}$.*

More interestingly, our reduction ensures that no-worst-response action profiles stand in correspondence with the pure Nash equilibria of the original game.

Lemma 4. Fix some action profile $\hat{a}$ of game $\hat{g}$ and profile $a = \rho(\hat{a})$ of game g. Then $\hat{a}$ is a no-worst-response profile of $\hat{g}$ if and only if a is a pure Nash equilibrium of g.

Proof of Lemma 4. Suppose a is a Nash equilibrium of g. We claim that $\hat{a}$ is a pure Nash equilibrium of $\hat{g}$, and hence in particular is a no-worst-response action profile. Let $(i, \ell) \in \hat{N}$ be any player in $\hat{g}$, and consider any unilateral deviation from $\hat{a}$ to $\hat{a}' := (\hat{a}'_{i\ell}, \hat{a}_{-i\ell})$ by that player. By construction of ρ, we see that $\rho(\hat{a}')$ and $\rho(\hat{a}) = a$ can differ only in the ith component. Assuming they differ, there must exist some $a'_i \in A_i$ such that $\rho(\hat{a}') = (a'_i, a_{-i})$. Since a is a pure Nash equilibrium in g, no such deviation is profitable, so $u_i(\rho(\hat{a}')) \leq u_i(\rho(\hat{a})) = u_i(a)$. But by construction of the utilities in $\hat{g}$, this implies that

$$\hat{u}_{i\ell}(\hat{a}') = u_i(\rho(\hat{a}')) \leq u_i(\rho(\hat{a})) = \hat{u}_{i\ell}(\hat{a}).$$

So $\hat{a}$ is a pure Nash equilibrium, and thus a no-worst-response action profile, of $\hat{g}$.

Conversely, suppose that $\hat{a} \in \hat{A}$ is a no-worst-response action profile of $\hat{g}$ but, for the sake of contradiction, $a = \rho(\hat{a}) \in A$ is not a pure Nash equilibrium of g. Then there exists some player $i \in N$ and some action $a'_i \in A_i$ such that $u_i(a'_i, a_{-i}) > u_i(a) = u_i(a_i, a_{-i})$. By

property ($\star$), there is a player (i, ℓ) in group $[i]$ of the game $\hat{g}$ for whom playing any $\hat{a}'_{i\ell} \neq \hat{a}_{i\ell}$ yields

$$\hat{u}_{i\ell}(\hat{a}') = u_i(\rho(\hat{a}')) = u_i(a'_i, a_{-i}) > u_i(a) = u_i(\rho(\hat{a})) = \hat{u}_{i\ell}(\hat{a}).$$

So for the player (i, ℓ) in $\hat{g}$, *every* alternative action strictly increases her payoff. In other words, (i, ℓ) is worst-responding at $\hat{a}$ in $\hat{g}$. □

We address a subtlety in the reduction. For any constants m and $\hat{m}$, the function f in Lemma 2 can be implemented as a constant-size Boolean circuit. In our reduction, the number of actions $m = 2$ and $\hat{m}$ of players in games g and $\hat{g}$ are fixed, and for each possible $\hat{m}$ we can assume that the circuit gadget has constant size and is available as a precomputed circuit fragment of constant size. When our reduction thus constructs Boolean circuits for the payoff function $u_{i\ell}$ of each player (i, ℓ) in $\hat{g}$, it thus takes constant computation time to insert each circuit gadget (a total number of n times; namely, one per gadget) into the overall circuit.

Proof of Theorem 2. We first see that the reduction is polynomial time. The game $\hat{g}$ has nq players and $\hat{m}$ actions. As $\hat{m}$ is a fixed constant and $q = 6\hat{m}2^{\hat{m}}$ is a constant, the number of players of $\hat{g}$ is linear in the number n of players in the original game g. The discussion above on the constant-size circuit gadget also makes it straightforward to see that the Boolean circuits for payoff functions $u_{i\ell}$ of game $\hat{g}$ can be constructed in polynomial time. Lemmas 3 and 4 and Corollary 3 show that we have constructed a PLS-reduction, so we see that the search problem of finding a no-worst-response action profile in a potential game with at most $\hat{m}$ actions is PLS-hard.

Finally, since Lemma 1 established PLS membership, the problem is PLS-complete. □

Our reduction for Theorem 2 requires a bounded number of actions for each player in the target game $\hat{g}$; otherwise, the circuit gadgets cannot be assumed to be precomputed and have constant size.

5 Non-universal guarantees

Thus far we considered minimal rationality guarantees for all players simultaneously and established that finding no-worst-response action profiles is as "hard" as finding pure Nash equilibria. We thus turn to the question whether action profiles where only some fraction (< 1) of players are guaranteed to play "reasonably good" actions are computationally tractable.

We identify a threshold, where on one side, the problem becomes tractable and, on the other side, the problem remains intractable. For ease of exposition, we consider games in which each player has the same number of actions.

Theorem 3. *Fix $\alpha, \beta \in (0,1)$ and integer $m \geq 2$. The search problem m-$\langle \alpha, \beta \rangle^{\mathcal{Q}}$ of finding, in a given game $g \in \mathbf{G}$ with n players and m actions per player, an action profile in which a fraction at least α of players play a top β action is in* RP *if $\lceil \alpha n \rceil < \beta n$, and is* NP-hard *if $\alpha m > \lceil \beta m \rceil$.*

Note that for large numbers of players and actions, the two bounds converge to a unique threshold. By Theorem 3, on the one hand, finding an action profile where a fraction at least α of players best-respond is tractable if $\alpha < 1/m$; on the other hand, finding an action profile where a fraction α of players do not worst-respond is tractable if $\alpha < (m-1)/m$ (see Figure 3). This shows that our negative results regarding action profiles where no player worst-responds are in some sense knife-edge. If the number of actions per player is large, an action profile where almost nobody worst-responds is tractable. By contrast, an action profile where a constant fraction of players best-responds is intractable.

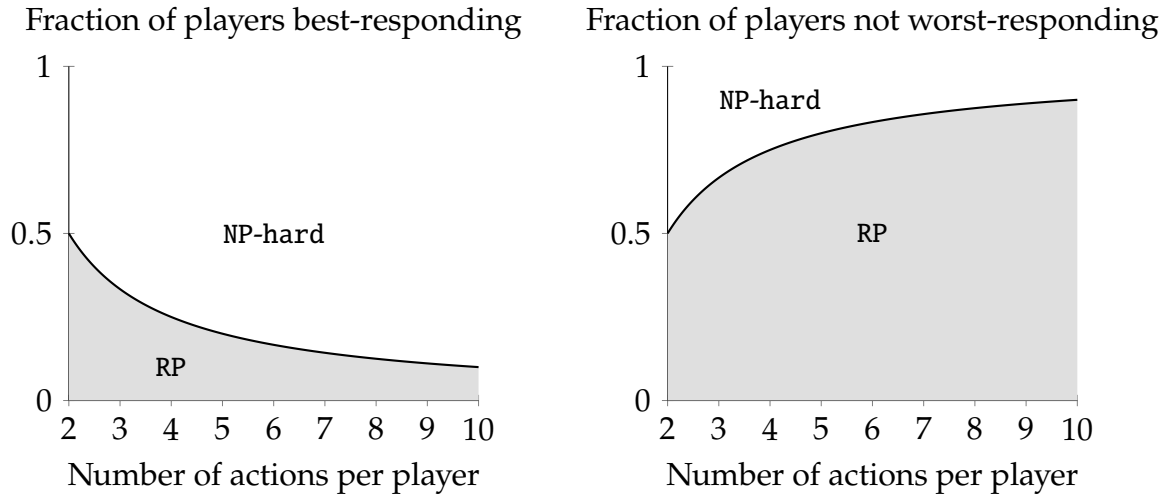


Figure 3: Schematic comparison of tractability frontier for best-responding and not worst-responding.

The proof of the first part of Theorem 3 follows from the more general Theorem 4 below. The second part of Theorem 3 is restated and proved in Theorem 5.

The intuition for the first (RP) part of Theorem 3 is the following. Consider the boundary where $\alpha \approx \beta$. If each player selects an action at random, with probability α they will play one of their approximately top α actions. Then, by linearity of expectation, when a random action profile is drawn, the expected number of players who are

using one of their approximately top $\alpha \cdot m$ actions is $\alpha \cdot n$. We formalize this intuition to arbitrary numbers of actions per player and for α, β that do not need to be fixed a priori.

Algorithm 1 A false-biased Monte Carlo algorithm.

Input: $g = (N, \{A_i\}_{i \in N}, \{u_i\}_{i \in N})$ with $n \coloneqq |N|$ and $m_i \coloneqq |A_i|$ for each $i \in N$.
Parameters $\alpha, \beta \in (0, 1)$ such that $\sum_{i \in N} \beta_i \geq \lceil \alpha n \rceil$ where $\beta_i \coloneqq \lceil \beta m_i \rceil / m_i$.

```
Let k = 1 and a be an arbitrary action profile.
while k ≤ n do
    Let c = 0.
    Let a be an action profile drawn uniformly at random from A.
    for each player i ∈ N do
        if a_i is a top β action in response to a_{-i} then
            Let c = c + 1
    if c ≥ αn then
        return action profile a.
    Let k = k + 1.
return action profile a from the loop's last iteration.
```

Proposition 4. For any game $g = (N, \{A_i\}_{i \in N}, \{u\}_{i \in N})$ with $n \coloneqq |N|$ players and $m_i \coloneqq |A_i|$ actions per player i, and for any $\alpha, \beta \in (0, 1)$ such that $\sum_{i \in N} \beta_i \geq \lceil \alpha n \rceil$ where $\beta_i \coloneqq \lceil \beta m_i \rceil / m_i$, Algorithm 1 returns an action profile in which at least $\lceil \alpha n \rceil$ players play a top β action with probability $\geq 1/2$.

Proof of Theorem 4. For any action profile $a \in A$ in g, let $\omega(a)$ denote the number of players whose action at a is a top β action. Since Algorithm 1 samples n action profiles, the probability that it returns a profile satisfying $\omega(a) \geq \lceil \alpha n \rceil$ is given by

$$1 - \mathbb{P}[\, \omega(a) < \lceil \alpha n \rceil \,]^n,$$

where $\mathbb{P}[\, \omega(a) < \lceil \alpha n \rceil \,]$ denotes the probability that $a \sim \text{Unif}[A]$ satisfies $\omega(a) < \lceil \alpha n \rceil$. But

$$\mathbb{P}[\, \omega(a) < \lceil \alpha n \rceil \,] = \mathbb{P}[\, \omega(a) \leq \lceil \alpha n \rceil - 1 \,] = \mathbb{P}[\, n - \omega(a) \geq n + 1 - \lceil \alpha n \rceil \,] \leq \frac{n - \mathbb{E}[\omega(a)]}{n + 1 - \lceil \alpha n \rceil},$$

where the last step is an application of Markov's inequality.

Now, since the probability that player i plays one of their top β_i actions is β_i, we have that $\mathbb{E}[\omega(a)] = \sum_{i \in N} \beta_i$. Moreover, since we assumed that $\sum_{i \in N} \beta_i \geq \lceil \alpha n \rceil$, we obtain

$$\frac{n - \mathbb{E}[\omega(a)]}{n + 1 - \lceil \alpha n \rceil} \leq \frac{n - \sum_{i \in N} \beta_i}{n + 1 - \sum_{i \in N} \beta_i} \leq \frac{n}{n + 1}.$$

It follows that the probability that Algorithm 1 returns an action profile a satisfying $\omega(a) \geq \lceil \alpha n \rceil$ is bounded below by

$$1 - \left(\frac{n}{n+1}\right)^n \geq \frac{1}{2},$$

for any $n \geq 1$, which completes the proof. □

We next turn to the hardness result in Theorem 3.

Proposition 5. Fix $\alpha, \beta \in (0,1)$ and integer $m \geq 2$ such that $\alpha > \lceil \beta m \rceil / m$. The decision problem $m\text{-}\langle \alpha, \beta \rangle^?$ of deciding whether a given game $g \in \mathbf{G}$ with m actions per player has an action profile in which a fraction at least α of players play a top β action is `NP-complete`.

To motivate the proof, consider, for any $n = m \geq 2$, a harmonic game with n players in which the players' actions are $\{0, \ldots, m-1\}$, and suppose that their utilities are given by $u_i(a) = \left(i + \sum_{j \in N} a_j\right) \bmod m$. At every action profile of the game, exactly κ players are using one of their κ best actions. It turns out that this property is preserved with the addition of the global term of (1) in Reduction 1.

For the proof of Theorem 5 and the discussion in the following, we define $k \in \mathbb{N}$ as the maximal value for which $\alpha > k/m$, and note that consequently $\beta \leq k/m$.

Membership in `NP` is immediate, as we can guess an action profile and check whether at least a fraction α of players are playing one of their top k actions in polynomial time. For `NP`-hardness, we argue that Reduction 1 is a polynomial-time reduction from `CircuitSAT`$^?$ to $m\text{-}\langle \alpha, \beta \rangle^?$. In the reduction, we choose m so that there exists k with $\alpha > k/m \geq \beta$. The main challenge to proving that Reduction 1 is a polynomial-time reduction to $m\text{-}\langle \alpha, \beta \rangle^?$, is to provide a more granular version of Observation 1.

Observation 2. If $f(a) = 1$, then every player is playing their best action. If $f(a) = 0$, then at most a fraction k/m of players are playing one of their k best actions.

Proof of Observation 2. If $f(a) = 1$, then Observation 1 tells us that every player is playing their best action. So in particular, at least αn players are using one of their k best actions.

So suppose from now on that $f(a) = 0$. If we can show that at most nk players are using one of their k best actions, we can apply the same reasoning as in the proof of Theorem 1 to show that Reduction 1 is a valid polynomial-time reduction from `CircuitSAT`$^?$ to $m\text{-}\langle \alpha, \beta \rangle^?$: the only action profiles a in which at least α fraction of players are playing one of their k best actions are those for which $f(a) = 1$, and these exist if and only if the circuit C is satisfiable.

Now suppose no player can unilaterally deviate to force $f(a'_{i\ell}, a_{-i\ell}) = 1$. Thus, by construction of the local term in (1), each player's mapping from actions to payoffs is a permutation on $\{0, \dots, m-1\}$. Secondly, within any group of players $(i,1), \dots, (i,m)$, for each possible payoff $p \in \{1, \dots, m\}$ there exists a player (i, ℓ) who achieves payoff p. It follows that exactly k players in each group play one of their k best actions. As there are n groups, nk players play one of their k best actions in total.

Finally, suppose there exists a player $(j, \hat{\ell})$ who can deviate to force $f(a'_{j\hat{\ell}}, a_{-j\hat{\ell}}) = 1$. We notice that there is always at most one group j which contains players who can achieve this. Also, by the same argument as above, exactly k players in each of the other groups $i \neq j$ play one of their k best actions. So we focus on counting the number of players the group $(j,1), \dots, (j,m)$ (which we call group j) who are playing one of their k best actions. If $m = 2$, if both players $(j,1), (j,2)$ may deviate to force $f(a'_{j\hat{\ell}}, a_{-j\hat{\ell}}) = 1$, then they are both worst responding and we are done. If instead $m \geq 3$, there is necessarily only one such player $(j, \hat{\ell})$. We thus focus on the case when there exists a unique such player $(j, \hat{\ell})$, and distinguish between two subcases. Recall that the local term of player $u_{j\ell}$'s payoffs is

$$\left(\ell + \sum_{\ell=1}^{m} a_{j\ell} \right) \bmod m,$$

and the value of this term can range from 0 to $m-1$ (cf. Equation (1) in Reduction 1).

Case 1: Suppose $\hat{\ell} + \sum_{\ell=1}^{m} a_{j\ell} < m - k$. Then exactly k of the remaining players in group j are playing one of their k best actions, by the same argument as above. Player $(j, \hat{\ell})$ can improve over her current action $a_{j\hat{\ell}}$ by playing the action that flips $f(a)$ to 1, or playing one of the k actions for which her local term is greater or equal to $m - k$. The flipping action and the k actions may intersect but, regardless, the action $a_{j\hat{\ell}}$ is not among her k best actions. It follows that exactly k players play one of their k best actions in action profile a.

Case 2: Suppose $\hat{\ell} + \sum_{\ell=1}^{m} a_{j\ell} \geq m - k$. By the same argument as above, there exist exactly $k-1$ other players in group i who are playing one of their k best actions. So, together with $(j, \hat{\ell})$, at most k players are playing one of their k best actions in profile a. □

With this observation, the proof of Theorem 5 is now straightforward.

Proof of Theorem 5. Let C be a Boolean circuit with n inputs. If C is satisfiable, then there exists an action profile for the game with $f(a) = 1$, and by Observation 2 all nm players play their best response (and so one of their k best responses). If C is unsatisfiable, then $f(a) = 0$ for all action profiles, and so Observation 2 tells us that at

most nk players play one of their k best actions. It follows that $C \in \texttt{CircuitSAT}^?$ if and only if $g \in m\text{-}\langle\alpha,\beta\rangle^?$. □

We can also extend the hardness result of Theorem 5 to the counting version of the problem.

Corollary 4. *Fix $\alpha, \beta \in (0,1)$ and integer $m \geq 2$ such that $\alpha > \lceil \beta m \rceil / m$. The counting problem $m\text{-}\langle\alpha,\beta\rangle^\#$ of determining, in a given game $g \in \mathbf{G}$ with m actions per player, the number of action profiles for which at least a fraction α of players play a top β action is* #P-hard.

Proof of Corollary 4. We use the same definition of k as for the proof of Theorem 5. We have observed above that an action profile in the reduced game g has at least αn players playing one of their k best actions if and only if $f(a) = 1$. It is straightforward that there is a one-to-one correspondence between satisfying truth assignments of C and action profiles a with $f(a) = 1$. Hence, $m\text{-}\langle\alpha,\beta\rangle^\#$ is #P-hard. □

6 Discussion

We have shown that even the weakest meaningful individual rationality requirement—that no player plays a worst response—is computationally intractable if required for all players. In general games, the decision, search, and counting problems for no-worst-response action profiles are computationally equivalent to their pure Nash equilibrium counterparts (NP-complete, NP-hard, and #P-complete, respectively). This hardness persists in potential games, where finding a no-worst-response action profile is PLS-complete despite the simple structure of potential games and the guaranteed existence of pure Nash equilibria and thus no-worst-response outcomes.

Our results suggest that, while not playing one's worst action may be straightforward in a single agent setting, requiring such behaviour for all players in a game turns out to be prohibitively difficult. This carries implications for what outcomes one can expect to arise in games. Solution concepts in games are usually founded on the premise that they would be played by rational players, or that an analyst recommends an action profile to which the players (after consideration) adhere to, or as an implicit or explicit agreement found by players after communication, or as arising from a decentralised learning procedure. Our results imply that, regardless of which of these foundations one considers, either some player or the analyst has to overcome a computationally intractable problem if the solution concept implements only no-worst-response outcomes. This computational intractability violates a basic plausibility requirement of decision-making (Camara, 2022).

From this perspective, our results are sobering. However, it is important to note that we considered worst-case complexity; that is, the maximal complexity of a given problem across all instances in a given game class. A natural route forward is to consider less penalizing notions such as average-case complexity, randomised solutions, as well as average-case dynamic foundations:

Average-case complexity. Our results establish worst-case hardness in general games and potential games. However, many games encountered in practice may not exhibit worst-case complexity. What can be said about the average-case complexity of finding action profiles with universal rationality guarantees? Are there natural distributions over games for which the problem becomes tractable with high probability?

Randomised solutions. We have focused on solutions in pure strategies. Does randomization, that is, considering mixed actions, reduce complexity? In this regard, the relationship between ϵ-equilibria and no-worst-response action profiles deserves further investigation.

Average-case dynamic foundations. A natural question is whether simple learning dynamics converge to no-worst-response action profiles more quickly than to Nash equilibria. Although our PLS-completeness result for potential games implies that worst-case convergence times for local improvement dynamics are exponential, average-case convergence behaviour may differ.

References

Elliot Anshelevich, Anirban Dasgupta, Jon Kleinberg, Éva Tardos, Tom Wexler, and Tim Roughgarden. The price of stability for network design with fair cost allocation. *SIAM Journal on Computing*, 38(4):1602–1623, 2008.

Florian M Artinger, Gerd Gigerenzer, and Perke Jacobs. Satisficing: Integrating two traditions. *Journal of Economic Literature*, 60(2):598–635, 2022.

Yakov Babichenko, Christos H. Papadimitriou, and Aviad Rubinstein. Can almost everybody be almost happy? pcp for ppad and the inapproximability of Nash. In Madhu Sudan, editor, *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science*, pages 1–9. ACM, 2016.

Modibo K. Camara. Computationally tractable choice. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, EC '22, page 28, 2022.

Ozan Candogan, Ishai Menache, Asuman Ozdaglar, and Pablo A Parrilo. Flows and decompositions of games: Harmonic and potential games. *Mathematics of Operations Research*, 36(3): 474–503, 2011.

Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Settling the complexity of computing two-player nash equilibria. *Journal of the ACM (JACM)*, 56(3):1–57, 2009.

Constantinos Daskalakis, Aranyak Mehta, and Christos Papadimitriou. A note on approximate Nash equilibria. In *International Workshop on Internet and Network Economics*, pages 297–306, 2006.

Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The complexity of computing a Nash equilibrium. *Communications of the ACM*, 52(2):89–97, 2009.

Geoffroy de Clippel and Kareen Rozen. Bounded rationality in choice theory: A survey. *Journal of Economic Literature*, 62(3):995–1039, 2024.

Alex Fabrikant, Christos H. Papadimitriou, and Kunal Talwar. The complexity of pure nash equilibria. In László Babai, editor, *Proceedings of the 36th Annual ACM Symposium on Theory of Computing*, pages 604–612. ACM, 2004.

Joan Feigenbaum, Daphne Koller, and Peter Shor. A game-theoretic classification of interactive complexity classes. In *Proceedings of Structure in Complexity Theory. Tenth Annual IEEE Conference*, pages 227–237, 1995.

Itzhak Gilboa and Eitan Zemel. Nash and correlated equilibria: Some complexity considerations. *Games and Economic Behavior*, 1(1):80–93, 1989.

Georg Gottlob, Gianluigi Greco, and Francesco Scarcello. Pure nash equilibria: Hard and easy games. In *Proceedings of the 9th Conference on Theoretical Aspects of Rationality and Knowledge*, pages 215–230, 2003.

Ehud Kalai and William Stanford. Finite rationality and interpersonal complexity in repeated games. *Econometrica*, 56(2):397–410, 1988.

Richard M. Karp. *Reducibility among Combinatorial Problems*, pages 85–103. Springer US, Boston, MA, 1972.

Michael Mitzenmacher and Eli Upfal. *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press, 2005.

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

Abraham Neyman. Bounded complexity justifies cooperation in the finitely repeated prisoners' dilemma. *Economics Letters*, 19(3):227–229, 1985.

Christos H. Papadimitriou. Algorithms, games, and the internet. In *Proceedings of the Thirty-Third Annual ACM Symposium on Theory of Computing (STOC)*, pages 749–753. ACM, 2001. doi: 10.1145/380752.380883.

Panos M Pardalos and Somesh Jha. Complexity of uniqueness and local search in quadratic 0–1 programming. *Operations Research Letters*, 11(2):119–123, 1992.

Bary SR Pradelski and Bassel Tarbush. Satisficing equilibrium. In *Proceedings of the 26th ACM Conference on Economics and Computation*, EC '25, page 442, 2025.

Roy Radner. Collusive behavior in noncooperative epsilon-equilibria of oligopolies with long but finite lives. *Journal of Economic Theory*, 22(2):136–154, 1980.

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

Ariel Rubinstein. Finite automata play the repeated prisoner's dilemma. *Journal of Economic Theory*, 39(1):83–96, 1986.

Aviad Rubinstein. Inapproximability of nash equilibrium. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, pages 409–418, 2015.

Yuval Salant. Procedural analysis of choice rules with applications to bounded rationality. *American Economic Review*, 101(2):724–748, 2011.

Alejandro A Schäffer and Mihalis Yannakakis. Simple local search problems that are hard to solve. *SIAM journal on Computing*, 20(1):56–87, 1991.

Grant R Schoenebeck and Salil Vadhan. The computational complexity of nash equilibria in concisely represented games. *ACM Transactions on Computation Theory (TOCT)*, 4(2):1–50, 2012.

Andrea Wilson. Bounded memory and biases in information processing. *Econometrica*, 82(6): 2257–2294, 2014.