Game Transformations That Preserve Nash Equilibria or Best-Response Sets*

Emanuel Tewolde and **Vincent Conitzer**
Foundations of Cooperative AI Lab (FOCAL)
Computer Science Department, Carnegie Mellon University, USA
emanueltewolde@cmu.edu, conitzer@cs.cmu.edu

Abstract

In this paper, we investigate under which conditions normal-form games are (guaranteed to be) strategically equivalent. First, we show for N-player games ($N \geq 3$) that

(A) it is NP-hard to decide whether a given strategy is a best response to some strategy profile of the opponents, and that

(B) it is co-NP-hard to decide whether two games have the same best-response sets.

Combining that with known results from the literature, we move our attention to equivalence-preserving game transformations.

It is a widely used fact that a positive affine (linear) transformation of the utility payoffs neither changes the best-response sets nor the Nash equilibrium set. We investigate which other game transformations also possess either of the following two properties when being applied to an arbitrary N-player game ($N \geq 2$):

(i) The Nash equilibrium set stays the same;

(ii) The best-response sets stay the same.

For game transformations that operate player-wise and strategy-wise, we prove that (i) implies (ii) and that transformations with property (ii) must be positive affine. The resulting equivalence chain highlights the special status of positive affine transformations among all the transformation procedures that preserve key game-theoretic characteristics.

1 Introduction

Motivation

When faced with a strategic interaction with other agents, it can be computationally useful for AI systems – as we will discuss further down – to detect when the current situation can be treated in the same way as another strategic game that has already been dealt with in the past. This problem can also be critical for the robustness and generalizability of our AI systems. AI agents will often need to act in new environments, possibly with multiple equilibria, and yet be predictable to each other and to us to avoid bad outcomes. Recognizing a new environment as strategically equivalent to a previously encountered one can help tremendously with this, providing a precedent for action that ensures that everyone's expectations on behavior are well calibrated. Oesterheld and Conitzer [2022], for example, build on that in order to achieve Pareto-improved outcomes in games played by AI representatives (where games can be transformed by reprogramming the AIs' rewards).

Not least for these various reasons, even the simplest class of games – 2-player normal-form with 2 actions per player – has been studied extensively in order to obtain a complete taxonomy for them; see for example [Robinson and Goforth, 2005; Rapoport *et al.*, 1976; Borm, 1987]. With it, it is easy to recognize when a 2×2 game contains traits of competition, cooperation, coordination, etc. [Bruns and Kimmich, 2021]. Such interpretation tools are also being developed for more complex strategic situations [Marris *et al.*, 2023], but this still remains an important venue of further research. One challenge is that larger games become prohibitively complex to compare directly: Du [2008] show that deciding whether two 2-player normal-form games share Nash equilibria is a computationally hard task. We will show in Section 3 that this task is also computationally hard for the case of best response sets and at least three players.

One classic tool that emerged in the beginnings of game theory has been to transform a given game into other strategically equivalent games that are easier to analyze [von Neumann and Morgenstern, 1944]. Positive affine (linear) transformations (PATs) have been particularly useful in that regard [Aumann, 1961; Adler *et al.*, 2009]. To illustrate PATs, consider any 2-player normal-form game in which the players' utilities are measured in dollars. Then, the best-response strategies of player 1 do not change if her utility payoffs are multiplied by a factor of 5. Moreover, they also do not change if 10 dollars are added to all outcomes that involve player 2 playing his, say, third strategy. More generally, PATs have the power to rescale the utility payoffs of each player and to add constant terms to the utility payoffs of a player i for each strategy choice $\mathbf{k}_{-i}$ of her opponents.

Through leveraging PATs, previous work significantly ex-

*An earlier extended abstract of this paper can be found in the Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024).

tended the applicability of efficient Nash equilibrium solvers [von Neumann, 1928; Dantzig, 1951; Adler, 2013; Adsul *et al.*, 2021] to classes beyond those of zero-sum and rank-1 games[1] [Moulin and Vial, 1978; Kontogiannis and Spirakis, 2012; Heyman and Gupta, 2023].

PATs are also popular in mechanism design and e-commerce: *Affine maximizer auctions* are PAT transformations of the classic VCG mechanism, and as such, inherit strategy-proofness and individual rationality by the strategy-preserving nature of PATs. They play a key role in finding revenue-maximizing mechanisms (both with classical optimizers [Likhodedov and Sandholm, 2004] and deep learning [Curry *et al.*, 2023; Curry *et al.*, 2024], and in improving welfare in redistribution mechanisms [Guo and Conitzer, 2010] and advertisement auctions [Deng *et al.*, 2021].

The versatility of PATs is based on their well-known property that they do not change preferences, best responses, or Nash equilibria, when being applied to an arbitrary game. In a very precise sense, PATs are also *the only* game transformations that do not change preferences; cf. Section 7. The main result in this paper addresses the question of whether there are other (efficiently computable) game transformations that do not change best responses or Nash equilibria.

Overview

Sections 2 and 4 provide some background on game-theoretic concepts that are relevant to understanding and deriving our main results. In Section 3, we develop computational hardness results for deciding whether a strategy in a game ever constitutes a best response and for deciding whether two games have the same best-response sets. We believe these results are of independent interest. However, they are also important for Section 5, in which we discuss why we will henceforth restrict our attention to game transformations that transform utilities player-wise and strategy-wise (called *separability*). In Section 6, we proceed to characterize all separable game transformations that preserve the Nash equilibrium set – or, alternatively, the best response sets – when being applied to an arbitrary N-player game. Last but not least, Section 7 puts our results into context with further related work.

To illustrate the insights of Section 6 on an example, consider H_{Ex} that takes any 2-player 2×2 normal-form game with payoff matrices

$$A = \begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix} \quad , \quad B = \begin{pmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{pmatrix}$$

and transforms it into the game $H_{\text{Ex}}(A, B) := (A', B')$ that is defined as

$$A' = \begin{pmatrix} -2a_{11} + 10 & a_{12}^5 \\ e^{a_{21}} & 0 \end{pmatrix}, B' = \begin{pmatrix} |b_{11}| & \text{sign}(b_{12}) \\ \sqrt{|b_{21}|} & \arctan(b_{22}) \end{pmatrix}$$

As one can see with the sign function in B', it is noteworthy to highlight that our notion of a game transformation allows for non-continuous functions. With Theorem 2, we will show that there must exist 2×2 games $(\bar{A}, \bar{B})$ for which H_{Ex} does not preserve their Nash equilibrium set or - respectively - their best-response sets. More generally, we derive that *universally* preserving the Nash equilibrium set implies that the best-response sets always have to be preserved as well; and that the latter property is only satisfied by game transformations H with the very restricted structure of a PAT. In the example of H_{Ex}, each transformation map within it single-handedly already violates a PAT structure.

Full proofs for statements in this paper can be found in the full version of this paper.

2 Normal-Form Games

Notation-wise, we denote $[n] := \{1, \ldots, n\}$ for any $n \in \mathbb{N}$. A normal-form multiplayer game G specifies

(a) the number of players $N \in \mathbb{N}, N \geq 2$,

(b) a set of pure strategies $S^i = [m_i]$ for each player i where $m_i \in \mathbb{N}, m_i \geq 2$, and

(c) the utility payoffs for each player i given as a function $u_i : S^1 \times \ldots \times S^N \to \mathbb{R}$.

Denote the set of strategy profiles in G as $S := S^1 \times \ldots \times S^N$. Throughout this paper, all considered multiplayer games shall have the same number of players N and the same set of strategy profiles S. Hence, any game G will be determined by its utility functions $\{u_i\}_{i \in [N]}$. The players choose their strategies simultaneously and they cannot communicate with each other. A utility function u_i can be summarized by its pure strategy outcomes for player i, captured as an N-dimensional tensor or array $\{u_i(\mathbf{k})\}_{\mathbf{k} \in S}$.

As usual, we allow the players to randomize over their pure strategies, called mixed strategies. Then, player i's strategy space extends to the set of probability distributions $\Delta(S^i) := \{s^i = (s^i_k)_k \in \mathbb{R}^{m_i}_{\geq 0} : \sum_{k \in [m_i]} s^i_k = 1\}$ over S^i. A tuple $\mathbf{s} = (s^1, \ldots, s^N) \in \Delta(S^1) \times \ldots \times \Delta(S^N) =: \Delta(S)$ is called a mixed strategy profile[2] in G. The utility payoff of player i under profile $\mathbf{s}$ is defined as the player's utility payoff in expectation $u_i(\mathbf{s}) := \sum_{\mathbf{k} \in S} s^1_{k_1} \cdot \ldots \cdot s^N_{k_N} \cdot u_i(\mathbf{k})$. The goal of each player is to maximize her utility.

We will abbreviate with S^{-i} the set that consists of all possible pure strategy choices $\mathbf{k}_{-i} = (k_1, \ldots, k_{i-1}, k_{i+1}, \ldots, k_N)$ of the opponent players (resp. $\Delta(S^{-i})$ for the set of mixed strategy choices $\mathbf{s}^{-i} = (s^1, \ldots, s^{i-1}, s^{i+1}, \ldots, s^N)$). We will also use $u_i(k_i, \mathbf{k}_{-i})$ instead of $u_i(\mathbf{k})$ to stress how player i can only influence her own strategy when it comes to her payoff (resp. $u_i(s^i, \mathbf{s}^{-i})$ instead of $u_i(\mathbf{s})$).

Definition 2.1. The best-response set of player i to the opponents' strategy choices $\mathbf{s}^{-i}$ is defined as
$\text{BR}_{u_i}(\mathbf{s}^{-i}) := \text{argmax}_{t^i \in \Delta(S^i)} \{ u_i(t^i, \mathbf{s}^{-i}) \}$.

Best-response strategies capture the idea of optimal play against the other player's strategy choices. The most popular equilibrium concept in non-cooperative games is based on best responses.

[1] A 2-player game, represented by its payoff matrices $A, B \in \mathbb{R}^{m \times n}$, is said to have rank 1 if $\text{rank}(A + B) = 1$.

[2] Not to be confused with a correlated strategy: In our notation, $\Delta(S)$ itself is not a simplex of high dimension but only the product of N lower-dimensional simplices.

Definition 2.2. A strategy profile $\mathbf{s} \in \Delta(S)$ to a game $G = \{u_i\}_{i \in [N]}$ is called a Nash equilibrium if for every player $i \in [N]$ we have $s^i \in \mathrm{BR}_{u_i}(\mathbf{s}^{-i})$.

By a result of Nash [1950], any such multiplayer game G admits at least one Nash equilibrium.

3 Decision Problems about Best Responses

In this section we show that two decision problems about best responses are hard for N-player games, when $N \geq 3$. To our knowledge, these results are novel.

For computational problems involving N-player games G with strategy sets $(S^i)_{i \in [N]}$ and utility functions $(u_i)_{i \in [N]}$, we are interested in their computational complexities in terms of $|S|$ and the binary encoding of all utility payoffs $(u_i(\mathbf{s}))_{\mathbf{s} \in S, i \in [N]}$. For that, we require that utility payoffs take on rational values only.

First, we consider the problem of deciding whether a mixed strategy of a player is ever a best response to some mixed strategy profile of the opponent players. In its computationally easiest form, we may formulate it as the following.

Definition 3.1 (CHECKIFEVERBR)**.** Given a 3-player normal-form game, do there exist mixed strategies $\mathbf{r} \in \Delta(S^2)$ of PL2 and $\mathbf{s} \in \Delta(S^3)$ of PL3 such that pure strategy 1 of PL1 is a best response to $(\mathbf{r}, \mathbf{s})$?

This is different from determining the best responses of a player to a given strategy profile of the opponents, a task that can be solved in polynomial time. Our problem is related to *rationalizable* strategies [Bernheim, 1984; Pearce, 1984] - a concept that is based on the idea that a rational player can and should eliminate any strategy that is not a best response to some belief over what her opponents may play.

Proposition 3.2. CHECKIFEVERBR *is* $\mathbf{NP}$*-hard.*

The analogous formulation of CHECKIFEVERBR for the case of 2-player games can be efficiently decided by solving a system of linear (in-)equalities. We can recover polynomial-time solvability for many-player games if we allow the opponents to play in a coordinated fashion (cf. correlated strategies). On a related note, Pearce [1984][Lemma 3] shows that a strategy s^* is a best response to some correlated strategy of the opponents if and only if s^* is not strictly dominated by a mixed strategy.

We prove Proposition 3.2 by a reduction from the Balanced Complete Bipartite Subgraph problem. This decision problem asks whether a given weighted bipartite graph $G = (V \cup W, E)$ has subsets $V^* \subseteq V$ and $W^* \subseteq W$ of given size $K \in \mathbb{N}$ that are fully connected, that is, $(v, w) \in E$ for all $v \in V^*, w \in W^*$. This problem is known to be $\mathbf{NP}$-complete [Garey and Johnson, 1990][GT24].

Proof sketch of Proposition 3.2. Given an instance $G = (V \cup W, E)$ and K of the Balanced Complete Bipartite Subgraph problem, construct a three player game where PL2 has strategy set V and PL3 has strategy set W. PL1 will have the following strategies: Strategy "1" which will be the subject of interest in CHECKIFEVERBR, one strategy for each node in G, and one strategy for each edge $(v, w) \in V \times W$ that is not present in G. The utility payoffs of PL1 will be carefully constructed such that strategy 1 is a best response to mixed strategies $(\mathbf{r}, \mathbf{s})$ of PL2 and PL3 if and only if the support of $\mathbf{r}$ and $\mathbf{s}$ form subsets V^* and W^* that make a balanced complete bipartite subgraph of G. To that end, we make strategy v (resp. w) of PL1 very attractive for PL1 in the case that PL2 (resp. PL3) plays their corresponding strategy v (resp. w) with too much probability. Moreover, we make a strategy $(v, w) \notin E$ of PL1 very attractive for PL1 in the case that PL2 and PL3 both play their corresponding strategies v and w with any significant probability at all. Intuitively, these two conditions accomplish that in any potential certificate $(\mathbf{r}, \mathbf{s})$, PL2 and PL3 will mix over at least K strategies and, moreover, they will only put non-negligible weight on strategies v and w if $(v, w) \in E$. □

Based on the hardness of CHECKIFEVERBR, we can prove $\mathbf{co}$-$\mathbf{NP}$-hardness of deciding best-response equivalence.

Definition 3.3 (CHECKIFSAMEBRS)**.** Given two 3-player normal-form games with strategy set $S^1 \times S^2 \times S^3$, do they have the same best-response sets?

Theorem 1. CHECKIFSAMEBRS *is* $\mathbf{co}$-$\mathbf{NP}$*-hard.*

Proof sketch. Given a game instance G of CHECKIFEVERBR, construct another game G' by changing the utility that PL1 receives from playing strategy 1 to something worse than the lowest payoff present in G. If a best-response set changed from G to G', then it must also be the case that strategy 1 for PL1 was added or removed from that best-response set. The former cannot happen because strategy 1 is strictly dominated for PL1 in G' which prevents it from ever being a best response. Thus, G and G' will have the same best-response sets if and only if strategy 1 is never a best-response strategy in G. □

Together with prior work found in the literature, Theorem 1 will guide us in the next sections when it comes to the types of game transformations that we may consider for preserving key game-theoretic characteristics. We believe, however, that Proposition 3.2 and Theorem 1 are also of independent interest for algorithmic game theory and AI research.

4 Preliminaries on Game Transformations

4.1 Positive Affine Transformations

The following lemma (or restricted versions of it) is a well-known result for 2-player games.[3] Here, the notation $\mathbf{1}_n \in \mathbb{R}^n$ stands for the vector with all ones as its entries.

Lemma 4.1. *Let (A, B) be an $m_1 \times m_2$ bimatrix game and take arbitrary scalars $\alpha_1, \alpha_2 > 0$ and vectors $c^1 \in \mathbb{R}^{m_2}, c^2 \in \mathbb{R}^{m_1}$. Define*

$$A' = \alpha_1 A + \mathbf{1}_{m_1} (c^1)^T \quad \text{and} \quad B' = \alpha_2 B + c^2 \mathbf{1}_{m_2}^T .$$

Then (A', B') has the same best-response sets as (A, B). Thus, both games have the same Nash equilibrium set.

[3]See Heyman and Gupta [2023][Lemma 2.1], Maschler *et al.* [2013][Theorem 5.35], [Moulin and Vial, 1978][Theorem 1], Harsanyi and Selten [1988][Chapter 3] or Başar and Olsder [1998][Proposition 3.1].

The game transformations in Lemma 4.1 are called (2-player) positive affine transformations (PATs). An explicit example of a 2-player PAT is one that transforms a 2×2 game (A, B) into

$$A' = \begin{pmatrix} 2a_{11} + 10 & 2a_{12} - 5 \\ 2a_{21} + 10 & 2a_{22} - 5 \end{pmatrix},$$

$$B' = \begin{pmatrix} \frac{1}{2}b_{11} & \frac{1}{2}b_{12} \\ \frac{1}{2}b_{21} - \sqrt{3} & \frac{1}{2}b_{22} - \sqrt{3} \end{pmatrix}.$$

The intuition behind Lemma 4.1 is as follows: PL1 wants to maximize her utility given the strategy of PL2. A positive rescaling of u_1 will change the utility payoffs but not the utility-maximizing strategies. The same holds true if we add utility payoffs to u_1 that are only dependent on the strategy choice of her opponent PL2, because that would make a constant shift in terms of the decision variables of PL1.
Let us generalize PATs to multiplayer games.

Definition 4.2. A positive affine transformation (PAT) specifies for each player i a scaling parameter $\alpha^i \in \mathbb{R}, \alpha^i > 0$, and translation constants $C^i := (c^i_{\mathbf{k}_{-i}})_{\mathbf{k}_{-i} \in S^{-i}}$ for each choice of pure strategies from the opponents. The PAT $H_{\text{PAT}} = \{\alpha^i, C^i\}_{i \in [N]}$ then takes any game $G = \{u_i\}_{i \in [N]}$ as an input and returns the transformed game $H_{\text{PAT}}(G) = \{u'_i\}_{i \in [N]}$ with utility functions

$$u'_i : S \to \mathbb{R}, \ \mathbf{k} \mapsto \alpha_i \cdot u_i(\mathbf{k}) + c^i_{\mathbf{k}_{-i}}. \tag{1}$$

Multiplayer PATs also preserve the best-response sets and Nash equilibrium set, which we prove in the full version of this paper for completeness.

Lemma 4.3. *Take a PAT $H_{\text{PAT}} = \{\alpha^i, C^i\}_{i \in [N]}$ and any game $G = \{u_i\}_{i \in [N]}$. Then, the transformed game $H_{\text{PAT}}(G) = \{u'_i\}_{i \in [N]}$ has the same best-response sets as the original game G. Consequently, $H_{\text{PAT}}(G)$ also has the same Nash equilibrium set as G.*

PATs have found much success as a tool for simplifying a given game precisely because of this property. We want to investigate which other game transformations also preserve the best-response sets or the Nash equilibrium set. If we found more of these transformations, we could use them to, e.g., further increase the class of efficiently solvable games.

4.2 Separable Game Transformations

In this paper, we will focus on the following space of game transformations. We discuss in Section 5 why this forms a maximally large search space within which we may still reasonably hope to find game transformation that are equivalence-preserving and efficiently computable.

Definition 4.4. A separable game transformation $H = \{H^i\}_{i \in [N]}$ specifies for each player i a collection of functions $H^i := \{h^i_{\mathbf{k}} : \mathbb{R} \to \mathbb{R}\}_{\mathbf{k} \in S}$, indexed by the different pure strategy profiles $\mathbf{k}$.
The transformation H can then be applied to any N-player game $G = \{u_i\}_{i \in [N]}$ with strategy set S to construct the transformed game $H(G) = \{H^i(u_i)\}_{i \in [N]}$ where
$H^i(u_i) : S \to \mathbb{R}, \quad \mathbf{k} \mapsto h^i_{\mathbf{k}}(u_i(\mathbf{k}))$.

Observe that the utility payoff of player i in the transformed game $H(G)$ from the pure strategy outcome $\mathbf{k}$ is only a function of the utility payoff from that same player in that same pure strategy outcome of the original game G.

We extend the utility functions $H^i(u_i)$ to mixed strategy profiles $\mathbf{s} \in \Delta(S)$ as usual through
$H^i(u_i)(\mathbf{s}) := \sum_{\mathbf{k} \in S} s^1_{k_1} \cdot \ldots \cdot s^N_{k_N} \cdot h^i_{\mathbf{k}}(u_i(\mathbf{k}))$.
To simplify future notation, we will often use $h^i_{k_i, \mathbf{k}_{-i}}$ to refer to $h^i_{\mathbf{k}}$.

Remark 4.5. A multiplayer positive affine transformation $H_{\text{PAT}} = \{\alpha^i, C^i\}_{i \in [N]}$ makes a separable game transformation $H = \{H^i\}_{i \in [N]}$ by setting
$h^i_{\mathbf{k}} : \mathbb{R} \to \mathbb{R} \ , \ z \mapsto \alpha^i \cdot z + c^i_{\mathbf{k}_{-i}}$.

In the following Definitions 4.6 and 4.7, we define the universally preserving characteristics that we are interested in.

Definition 4.6. Let $H = \{H^i\}_{i \in [N]}$ be a separable game transformation. Then we say that H universally preserves Nash equilibrium sets if for all games $G = \{u_i\}_{i \in [N]}$ the transformed game $H(G) = \{H^i(u_i)\}_{i \in [N]}$ has the same Nash equilibrium set as G.

Definition 4.7. Let map H^i come from a separable game transformation H. Then we say that H^i universally preserves best responses if for all utility functions $u_i : S \to \mathbb{R}$ and for all opponents' strategy choices $\mathbf{s}^{-i} \in \Delta(S^{-i})$:

$$\begin{aligned} \text{BR}_{H^i(u_i)}(\mathbf{s}^{-i}) &= \operatorname*{argmax}_{t^i \in \Delta(S^i)} \{H^i(u_i)(t^i, \mathbf{s}^{-i})\} \\ &= \operatorname*{argmax}_{t^i \in \Delta(S^i)} \{u_i(t^i, \mathbf{s}^{-i})\} = \text{BR}_{u_i}(\mathbf{s}^{-i}). \end{aligned}$$

Lemma 4.3 states that the maps H^i of a PAT universally preserve best responses. Note, moreover, that by definition of a Nash equilibrium, a game transformation $H = \{H^i\}_{i \in [N]}$ will universally preserve Nash equilibrium sets if for every player i the map H^i universally preserves best responses. Therefore, being a PAT implies Definition 4.7 implies Definition 4.6. In Section 6 we will show the reverse implication chain for game transformations that are separable.

5 Discussion of Restrictions

The space of separable game transformations forms a vast landscape in which we may search for universally preserving transformations. This can be seen from the game transformation example H_{Ex} of Section 1. However, one might still ask why this paper does not expand its attention to non-separable game transformations. We will discuss that in this section.

For example, consider a game transformation that introduces or removes duplicate strategies or dummy players. Note that this would require the transformations to have the power to change the strategy sets and player set. Nonetheless, these specific examples are well-behaved in the sense that they alter the Nash equilibrium set (or best responses) in an easily describable manner. Abdou *et al.* [2022], for example, managed to characterize how selected examples of these transformations interact with various methods of decomposing a game. Transformations that change the strategy sets and player set also appear in the literature under the term *Nash homomorphism*, and they have been of use for complexity-theoretic studies, e.g., of win-lose games [Abbott *et al.*, 2005] or ranking games [Brandt *et al.*, 2006]. Suffice to say, once we allow for game transformations to arbitrarily change the game structure, i.e. the player set and strategy sets, it is not straightforward to define anymore under what conditions two games of different game structure should be considered "strategically equivalent". This makes such general game transformations prohibitively complex (or impossible) to analyze beyond a case by case basis. Therefore, and in accordance with most of the literature on strategic equivalence between games [Moulin and Vial, 1978; Morris and Ui, 2004; Du, 2008; Liu, 1996], we restrict our attention to games whose game structures are directly comparable.

Indeed, game transformations that preserve the player set and the strategy sets form an interesting search space because Definitions 4.6 and 4.7 can be directly extended to it and because within that search space, some of our following results will not hold true anymore. Compare the Prisoner's Dilemma with the Quality game, as presented by von Stengel [2022]:

$$\begin{pmatrix} 2,2 & 0,3 \\ 3,0 & 1,1 \end{pmatrix} \text{ and } \begin{pmatrix} 2,2 & 0,1 \\ 3,0 & 1,1 \end{pmatrix}.$$

Both games have the same unique Nash equilibrium, namely, where PL1 plays the bottom row and PL2 plays the right column. But the best response of PL2 to PL1 playing the top row is different in the two games. This example illustrates the fact that strictly dominated strategies will never be a best response, and so they will never appear in a Nash equilibrium (nor in a best-response set). Therefore, we can think of a game transformation procedure that iteratively detects strictly dominated strategies and sets their payoffs to a large negative number. This transformation universally preserves Nash equilibria, but it does not universally preserve best-response sets. Note that this game transformation is not separable because its maps $h^i_{\mathbf{k}}$ now need to take all utility payoffs of the game into consideration, and not only what utility player i receives from strategy profile $\mathbf{k}$.

In a similar fashion, one may think of best-response-preserving transformations that are not PATs. This was studied extensively by Liu [1996], who discusses the following example of 3×2 payoff matrices of PL1 in 2-player games:

$$A = \begin{pmatrix} 6 & 0 \\ 0 & 6 \\ 4 & 4 \end{pmatrix} \text{ and } A' = \begin{pmatrix} 6 & 0 \\ 2 & 5 \\ 4 & 4 \end{pmatrix}. \tag{2}$$

As visualized by Figure 1, the best responses of PL1 to any mixed strategy of PL2 are the same in A and A'. However, A' cannot be obtained from A through a PAT: If there were such a PAT, then the payoff from profile $(2, 1)$ requires a shift of $c^1_1 = 2$. Hence, the payoff from profile $(1, 1)$ requires a scaling of $\alpha^1 = \frac{2}{3}$. But these components of a positive affine transformation do not work out for the payoff from profile $(3, 1)$, leaving us with a contradiction.

Liu [1996] develops a polynomial-time method, called *bi-affine* transformation, that determines whether two 2-player normal-form games have the same best-response sets. The procedure detects which strategies and strategy pairs are essential, and derives that only the essential pairs need to be in a positive affine relationship. Hence, the method includes PATs, but it is also more powerful than that. Liu's dissertation (1995) extends those ideas to N-player games ($N \geq 3$). But in N-player games, the method downgrades to a sufficient condition: Two N-player games ($N \geq 3$) may have the same best-response sets while not being a *quasi-affine* transformation of each other. Furthermore, their method becomes computationally inefficient. In fact, we have shown in Section 3 more generally that determining whether two 3-players games have the same BR sets is $\mathbf{co\text{-}NP}$-hard.

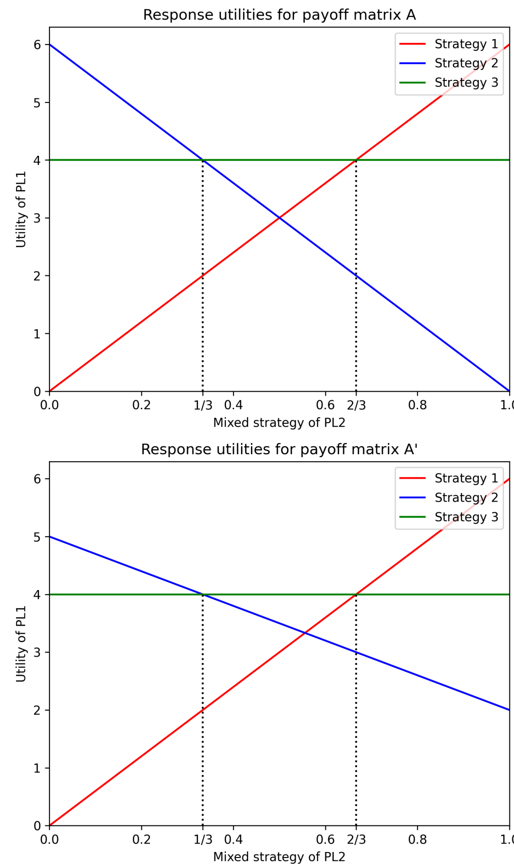

Figure 1: The utility payoffs of each pure strategy $1, 2, 3$ of PL1 in response to the mixed strategy of PL2 that plays 1 with probability x. Plots correspond to matrices A and A' from (2). The best-response set to a strategy $(x, 1 - x)$ of PL2 will be all convex combinations of pure strategies of PL1 that are maximal at x in the respective plot.

Liu concludes with an immediate open problem for future work: to characterize games with the same Nash equilibria. To that end, Du [2008] proves that it is $\mathbf{NP}$-complete to decide whether two 2-player games share a common Nash equilibrium, and that it is $\mathbf{co\text{-}NP}$-hard to decide whether two 2-player games have the same Nash equilibrium set.

In light of these negative results about characterizing best-response equivalence and Nash equilibrium equivalence in full generality - assuming the well-accepted complexity belief $\mathbf{co\text{-}NP} \neq \mathbf{P}$ - we restrict our focus to a subclass of equivalence-preserving transformations based on separability. We argue that among naturally defined subclasses, separable game transformations constitute a maximal subclass for which it is still open whether it contains tractable and equivalence-preserving transformations aside from PATs.

6 Transformations that Preserve Nash Equilibrium Sets or Best-Response Sets

To our knowledge, the results of this section are all novel unless explicitly stated otherwise. They can be summarized in the following statement.

Theorem 2. *Let $H = \{H^i\}_{i\in[N]}$ be a separable game transformation. Then:*

H universally preserves Nash equilibrium sets (i)

$\iff$ *for each player i, map H^i universally preserves best responses* (ii)

$\iff$ *H is a positive affine transformation.* (iii)

Lemma 4.3 gives (iii) $\implies$ (i), and so the novel part of Theorem 2 is the implication chain (i) $\implies$ (ii) $\implies$ (iii). The key property that enables us to develop this chain is that we require the separable game transformations $H = \{H^i\}_{i\in[N]}$ to be *universally* applicable, no matter the game $G = \{u_i\}_{i\in[N]}$ we have at hand.

We shall state two algorithmic consequences of Theorem 2.

Corollary 6.1. *Given two normal-form games, we can decide within polynomial time whether one is a transform of the other through an equivalence-preserving separable game transformation.*

This is because deciding whether a game is a PAT transform of another reduces to solving a linear (in-)equation system for the variables $\{\alpha^i, C^i\}_{i\in[N]}$. A case distinction is needed for solution points that take on values $\alpha^i = 0$.

Corollary 6.2. *Given a 2-player normal-form game G, we can find a transform G' of it (if it exists) via an equivalence-preserving separable game transformation, such that G' is a zero-sum or rank-1 game. With that, a Nash equilibrium for G can be computed subsequently. Both take polynomial time.*

This follows from the results in [Heyman and Gupta, 2023; Adsul *et al.*, 2021].

Before tackling Theorem 2, let us characterize a special property that a game transformation can satisfy in which the strategy choice of player i does not influence the map that is being used to transform her utilities.

Definition 6.3. Let map H^i come from a separable game transformation H. Then we say that H^i only depends on the strategy choices of the opponents if for all pure strategy choices $\mathbf{k}_{-i} \in S^{-i}$ of the opponents, we have the map identities $h^i_{1,\mathbf{k}_{-i}} = \ldots = h^i_{m_i,\mathbf{k}_{-i}} : \mathbb{R} \to \mathbb{R}$.

Next, we can show (i) $\implies$ (ii).

Proposition 6.4. *Let $H = \{H^i\}_{i\in[N]}$ be a separable game transformation that universally preserves Nash equilibrium sets and consider the map H^i of a player i. Then H^i only depends on the strategy choices of the opponents. Furthermore, H^i universally preserves best responses.*

Proof sketch.

1. Such a universally preserving transformation H should in particular not change the Nash equilibrium set for a trivial game in which all players receive the same constant utility $z \in \mathbb{R}$ from all strategy profiles. In such a game, the whole strategy set S will make the Nash equilibrium set. For that to also be the case in the transformed game, we show for every player i, that the transformations maps $h^i_{1,\mathbf{k}_{-i}}, \ldots, h^i_{m_i,\mathbf{k}_{-i}}$ must all evaluate the same on any input value z.

2. Let u_i be an arbitrary utility function of player i. Complete u_i to a full game G by setting the utilities of all other players to the constant payoff of 0. This makes any strategy s^j of another player $j \neq i$ always a best-response strategy in G. We can then show that this must also hold in the transformed game $H(G)$, using the first conclusion. Therefore, we get the following equivalence chain:

(a) a strategy s^i of player i is a best response to a profile s^{-i} of the opponent players and with respect to u_i if and only if

(b) (s^i, s^{-i}) is a Nash equilibrium of G if and only if

(c) (s^i, s^{-i}) is a Nash equilibrium of $H(G)$ if and only if

(d) s^i a best response to s^{-i} with respect to $H^i(u_i)$.

□

The first conclusion captures the intuition that if the maps $h^i_{\mathbf{k}}$ from H^i would depend on the strategy choice of player i, then in the transformed game $H(G)$, player i may need to adjust her strategy choice to those $h^i_{\mathbf{k}}$ that map payoffs to high values. This would affect the strategic decision making of player i and therefore the Nash equilibrium set. Similar reasoning provides us with a related (but independent) result:

Lemma 6.5. *If map H^i universally preserves best responses, then H^i only depends on the strategy choices of the opponents.*

Due to Proposition 6.4, we can transition to the analysis of transformation maps H^i that universally preserve best responses. Thus from now on, our results also become relevant to game theory research that focuses on best-response sets, such as best-response dynamics or fictitious play.

Proposition 6.4 moreover allows us to restrict our analysis to the map H^1 for PL1 w.l.o.g. because any results for H^1 will analogously also hold for maps $H^2, \ldots, H^N$. By Lemma 6.5, we can also drop the dependence of H^1 on k_1 and write $H^1 := \left\{h^1_{\mathbf{k}_{-1}} : \mathbb{R} \to \mathbb{R}\right\}_{\mathbf{k}_{-1}\in S^{-1}}$.

For each pure-strategy map $h^1_{\mathbf{k}_{-1}}$ we introduce its *distance distortion* function which takes two utility values and measures their distance after a $h^1_{\mathbf{k}_{-1}}$-transformation:

$$\Delta h^1_{\mathbf{k}_{-1}} : \mathbb{R} \times \mathbb{R} \to \mathbb{R}\,,\ (z, w) \mapsto h^1_{\mathbf{k}_{-1}}(z) - h^1_{\mathbf{k}_{-1}}(w) \quad (3)$$

The following lemma reveals an important preliminary observation on how the distance distortion functions $\Delta h^1_{\mathbf{k}_{-1}}$ relate to each other. It highlights how the distorted utility distances are connected upon a strategy change of a player $j \neq 1$ from, e.g., some pure strategy $k_j \neq 1$ to their first pure strategy $1 \in [m_j]$. It is again crucial that H^1 preserves best responses *universally* in order to deduce these global properties of and connections between the maps within H^1.

Lemma 6.6. *Suppose transformation map H^1 universally preserves best responses. Take a player $j \in [N] \setminus \{1\}$ and profile $\mathbf{k}_{-1} \in S^{-1}$ with $k_j \neq 1$. Define $\mathbf{k}'_{-1} \in S^{-1}$ to be the same as $\mathbf{k}_{-1}$ except for player j's choice which shall be set to $k'_j = 1$. Then, for all $z, z', w, w' \in \mathbb{R}$:*

$$z - w \geq z' - w' \iff \Delta h^1_{\mathbf{k}_{-1}}(z, w) \geq \Delta h^1_{\mathbf{k}'_{-1}}(z', w').$$

Proof sketch. Construct a utility function u^1 for each set of values for $j, \mathbf{k}_{-1}, z, z', w$, and w'. Namely, set $u_1(1, \mathbf{k}_{-1}) := z$ and $u_1(1, \mathbf{k}'_{-1}) := w'$, and for all strategies $l \in [m_1] \setminus \{1\}$ set $u_1(l, \mathbf{k}_{-1}) := w$ and $u_1(l, \mathbf{k}'_{-1}) := z'$. Observe that uniformly randomizing over $\mathbf{k}_{-1}$ and $\mathbf{k}'_{-1}$ not only makes a correlated strategy of the opponents, but also a valid mixed strategy profile. Hence, the left hand side of the equivalence can be reinterpreted as strategy $1 \in [m_1]$ performing better for player 1 than any other of her strategies $l \in [m_1] \setminus \{1\}$ if player j uniformly mixes over strategies k_j and k'_j and if all other players $r \notin \{1, j\}$ play their respective strategy $k_r \in [m_r]$. We then derive the equivalence by using that H^1 preserves strategy 1 being such a best response and by using Lemma 6.5. □

Next, observe that by definition, these distance distortion functions are skew-symmetric, that is, $\forall\, z, w \in \mathbb{R}$: $\Delta h^1_{\mathbf{k}_{-1}}(z, w) = -\Delta h^1_{\mathbf{k}_{-1}}(w, z)$.

With the following lemma, we further tighten the connection between the pure-strategy maps $h^1_{\mathbf{k}_{-1}}$ through their distance distortion functions. Last but not least, we shine some light on how those maps $h^1_{\mathbf{k}_{-1}}$ behave individually in the subsequent lemma.

Lemma 6.7. *If map H^i universally preserves best responses, then the pure-strategy maps in H^1 equally distort distances. That is, $\forall\, \mathbf{k}_{-1} \in S^{-1} : \Delta h^1_{\mathbf{k}_{-1}} = \Delta h^1_{\mathbf{1}_{-1}}$, where $\mathbf{1}_{-1} := (1, \ldots, 1) \in S^{-1}$.*

Proof sketch. Make iterative use of Lemma 6.6 for all other players $j \neq 1$, and make use of the skew-symmetry. □

Lemma 6.8. *If map H^i universally preserves best responses, then for all $\mathbf{k}_{-1} \in S^{-1}$:*

1. *map $h^1_{\mathbf{k}_{-1}}$ is strictly increasing, and*
2. *map $h^1_{\mathbf{k}_{-1}}$ distorts distances independently of their reference points:*

$$\forall z, z', \lambda \in \mathbb{R} \, : \, \Delta h^1_{\mathbf{k}_{-1}}(z + \lambda, z) = \Delta h^1_{\mathbf{k}_{-1}}(z' + \lambda, z')\,.$$

Proof sketch. For the first conclusion make use of Lemma 6.6 for values $z' = w'$, and of Lemma 6.7. For the second conclusion, utilize skew-symmetry together with the same two lemmata. □

With Lemmata 6.7 and 6.8, we can finally show that positive affine transformations are the only game transformations that universally preserve best responses. Intuitively speaking, the second conclusion of Lemma 6.8 states that taking a step of length λ in the domain space consistently maps to taking a step of some other length in the range space, independently of the base point z from which we take such a step.

This brings us to two known results from the analysis literature. Recall that a function $h : \mathbb{R} \to \mathbb{R}$ is called linear if there exists some $a \in \mathbb{R}$ such that $\forall z \in \mathbb{R} \, : \, h(z) = az$. A function $h : \mathbb{R} \to \mathbb{R}$ is said to be additive if it satisfies $\forall x, y \in \mathbb{R} \, : \, h(x + y) = h(x) + h(y)$.

Lemma 6.9 ([Darboux, 1875; Reem, 2017])**.** *If a map $h : \mathbb{R} \to \mathbb{R}$ is monotone and additive, then it is also linear.*

Corollary 6.10. *Let $h : \mathbb{R} \to \mathbb{R}$ be monotone and satisfy for all $z, z', \lambda \in \mathbb{R} : h(z + \lambda) - h(z) = h(z' + \lambda) - h(z')$. Then h is affine linear, i.e., there exist some $a, c \in \mathbb{R}$ such that for all $\forall z \in \mathbb{R} \, : \, h(z) = az + c$.*

This brings us to the completion of this section.

Proof sketch of Theorem 2.
Implication (iii) $\Longrightarrow$ (i) follows from Lemma 4.3, and implication (i) $\Longrightarrow$ (ii) follows from Proposition 6.4. For (ii) $\Longrightarrow$ (iii), recall that by symmetry, our results for H^1 hold analogously for all maps H^i. By Lemmata 6.5 and 6.8, the maps $h^i_{\mathbf{k}} = h^i_{\mathbf{k}_{-i}}$ satisfy the conditions of Corollary 6.10. Thus, there exist parameters $a^i_{\mathbf{k}_{-i}}, c^i_{\mathbf{k}_{-i}} \in \mathbb{R}$ for each $\mathbf{k}_{-i} \in S^{-i}$ such that $\forall z \in \mathbb{R} \, : \, h^i_{\mathbf{k}_{-i}}(z) = a^i_{\mathbf{k}_{-i}} \cdot z + c^i_{\mathbf{k}_{-i}}$.
Lemma 6.7 implies $a^i_{\mathbf{k}_{-i}} = a^i_{\mathbf{1}_{-i}}$ for all $\mathbf{k}_{-i} \in S^{-i}$. Therefore, we only have to keep track of one scaling parameter α^i for all the maps within H^i. With the first conclusion of Lemma 6.8, we obtain $\alpha^i > 0$. Putting everything together, we have shown that $H = (H^1, \ldots, H^N)$ makes a positive affine transformation.

□

Theorem 2 gives two novel equivalent characterizations of PATs that highlight their special status among game transformations: PATs are the only separable game transformations that always preserve the Nash equilibrium set or, respectively, the best-response sets.

One way to circumvent this result is to focus on game transformations that we only care to apply on particular subclasses of N-player games. Preferably, the game properties defining such a subclass would be generic enough to still contain "most" games. On the other hand, one may instead also consider non-separable game transformations as discussed in Section 5.

7 Further Related Literature

Much work has gone into identifying when two games can be considered strategically equivalent.

Strategic similarity, for example, is an important aspect of Potential Games (cf. Monderer and Shapley [1996]). Morris and Ui [2004] noted that a game G is a weighted potential game if and only if it is the PAT transformation of an identical interest game[4]. They also characterized when two given games are *best-response equivalent*, *better-response equivalent* or *von Neumann-Morgenstern equivalent*. The former and latter are directly tied to our concepts of preserving best-response sets and to PATs. Unfortunately, we were not able to

[4]Identical interest game: Given an action profile s, each player shall receive the same utility from s.

base the second part of Theorem 2 on the insights from Morris and Ui because their characterization for best-response equivalence only holds for games that satisfy specific properties.

Hammond [2005] described that the strategic decision-making in a game in mixed strategies does not depend on the player's numerical utility values, but solely on the preferences that the utility functions induce over the strategies. In the full version of this paper, we give some further background on utility theory in order to put Hammond's work into our context. Using the Expected Utility Theorem – cf., e.g., Mas-Colell *et al.* [1995] – Hammond deduced that utility functions that induce the same preferences can only differ up to a positive affine transformation. Note that the property of preserving the player's preferences is, in general, strictly harder to satisfy than preserving best responses (and, hence, Nash equilibria). Thus, our Theorem 2 generalizes their result to the broader question of strategic equivalence.

Moving to more broadly related work, Gabarró *et al.* [2011; 2013] gave several complexity-theoretic results for the problem of deciding whether two pure strategy games are *isomorphic* w.r.t. a notion of game transformation that can help us understand the symmetries within a game [Harsanyi and Selten, 1988, Chapter 3]. McKinsey [1951] and Chang and Tijs [2006] studied two notions of game equivalency specific to cooperative games.

Finally, there are other lines of related research that work more explicitly with different notions of transforming a game and preserving strategic features [Thompson, 1952; Kohlberg and Mertens, 1986; Elmes and Reny, 1994; Casajus, 2003; Pottier and Nessah, 2014; Wu *et al.*, 2022].

8 Conclusion

In this paper, we first gave hardness results about deciding whether a strategy constitutes a best response or whether two games have the same best-response sets. Next, we introduced separable game transformations for multiplayer games, and define the properties (i) universally preserving Nash equilibrium sets and (ii) universally preserving best responses. It is well-known that PATs universally preserve Nash equilibrium sets. We showed that separable game transformations which universally preserve Nash equilibrium sets also universally preserve best responses. In the subsequent results, we derived further that if a separable game transformation universally preserves best responses then it is a positive affine transformation.

When faced with a strategic interaction it can be highly beneficial to consider equivalent variations of it that are easier to analyze. In this paper, we shed light on why PATs have become the go-to transformation method for that purpose, reinforcing their standing as the standard off-the-shelf approach. The current literatures on game theory and on decision making in AI are lacking methods to detect or generate strategically equivalent games, and we hope that our results can serve as guidance to the development of any such detection or generation toolkit.

Acknowledgments

We want to thank Sebastian van Strien for the valuable discussions related to this project and to thank Bernhard von Stengel whose feedback inspired the development of Sections 4 and 6. We are also grateful to Vojtěch Kovařík, Rinor Cakaj, and the anonymous reviewers for their valuable improvement suggestions for this paper. Finally, we thank the Cooperative AI Foundation, Polaris Ventures (formerly the Center for Emerging Risk Research) and Jaan Tallinn's donor-advised fund at Founders Pledge for financial support.

References

[Abbott *et al.*, 2005] T. Abbott, D. Kane, and P. Valiant. On the complexity of two-player win-lose games. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS'05)*, pages 113–122, 2005.

[Abdou *et al.*, 2022] Joseph M. Abdou, Nikolaos Pnevmatikos, Marco Scarsini, and Xavier Venel. Decomposition of games: Some strategic considerations. *Math. Oper. Res.*, 47(1):176–208, 2022.

[Adler *et al.*, 2009] Ilan Adler, Constantinos Daskalakis, and Christos H. Papadimitriou. A note on strictly competitive games. In Stefano Leonardi, editor, *Internet and Network Economics, 5th International Workshop, WINE 2009, Rome, Italy, December 14-18, 2009. Proceedings*, pages 471–474. Springer, 2009.

[Adler, 2013] Ilan Adler. The equivalence of linear programs and zero-sum games. *Int. J. Game Theory*, 42(1):165–177, 2013.

[Adsul *et al.*, 2021] Bharat Adsul, Jugal Garg, Ruta Mehta, Milind A. Sohoni, and Bernhard von Stengel. Fast algorithms for rank-1 bimatrix games. *Oper. Res.*, 69(2):613–631, 2021.

[Aumann, 1961] Robert J. Aumann. Almost strictly competitive games. *Journal of The Society for Industrial and Applied Mathematics*, 9:544–550, 1961.

[Başar and Olsder, 1998] Tamer Başar and Geert Jan Olsder. *Dynamic Noncooperative Game Theory*. Society for Industrial and Applied Mathematics, 2nd. edition, 1998.

[Bernheim, 1984] B. Douglas Bernheim. Rationalizable strategic behavior. *Econometrica*, 52(4):1007–1028, 1984.

[Borm, 1987] P.E.M. Borm. A classification of 2x2 bimatrix games. *Cahiers du C.E.R.O.*, 29(1-2):69–84, 1987. Pagination: 16.

[Brandt *et al.*, 2006] Felix Brandt, Felix A. Fischer, and Yoav Shoham. On strictly competitive multi-player games. In *Proceedings, The Twenty-First National Conference on Artificial Intelligence and the Eighteenth Innovative Applications of Artificial Intelligence Conference, July 16-20, 2006, Boston, Massachusetts, USA*, pages 605–612. AAAI Press, 2006.

[Bruns and Kimmich, 2021] Bryan Bruns and Christian Kimmich. Archetypal games generate diverse models of power, conflict, and cooperation. *Ecology and Society*, 26(4):2, 2021.

[Casajus, 2003] André Casajus. Weak isomorphism of extensive games. *Math. Soc. Sci.*, 46(3):267–290, 2003.

[Chang and Tijs, 2006] Chih Chang and Stef Tijs. A note on isomorphism and strategic equivalence of cooperative games. *TOP*, 14:333–342, 2006.

[Curry *et al.*, 2023] Michael J. Curry, Tuomas Sandholm, and John P. Dickerson. Differentiable economics for randomized affine maximizer auctions. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI 2023, 19th-25th August 2023, Macao, SAR, China*, pages 2633–2641. ijcai.org, 2023.

[Curry *et al.*, 2024] Michael J. Curry, Vinzenz Thoma, Darshan Chakrabarti, Stephen McAleer, Christian Kroer, Tuomas Sandholm, Niao He, and Sven Seuken. Automated design of affine maximizer mechanisms in dynamic settings. In Michael J. Wooldridge, Jennifer G. Dy, and Sriraam Natarajan, editors, *Thirty-Eighth AAAI Conference on Artificial Intelligence, AAAI 2024, Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence, IAAI 2024, Fourteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2014, February 20-27, 2024, Vancouver, Canada*, pages 9626–9635. AAAI Press, 2024.

[Dantzig, 1951] George B Dantzig. A proof of the equivalence of the programming problem and the game problem. In Tjalling C. Koopmans, editor, *Activity analysis of production and allocation*, pages 330–335. Cowles Commission Monograph No.13, 1951.

[Darboux, 1875] Gaston Darboux. Sur la composition des forces en statique. *Bulletin des Sciences Mathématiques et Astronomiques*, 9:281–288, 1875.

[Deng *et al.*, 2021] Yuan Deng, Jieming Mao, Vahab S. Mirrokni, and Song Zuo. Towards efficient auctions in an auto-bidding world. In Jure Leskovec, Marko Grobelnik, Marc Najork, Jie Tang, and Leila Zia, editors, *WWW '21: The Web Conference 2021, Virtual Event / Ljubljana, Slovenia, April 19-23, 2021*, pages 3965–3973. ACM / IW3C2, 2021.

[Du, 2008] Ye Du. On the complexity of deciding bimatrix games similarity. *Theor. Comput. Sci.*, 407(1-3):583–586, 2008.

[Elmes and Reny, 1994] Susan Elmes and Philip J. Reny. On the strategic equivalence of extensive form games. *Journal of Economic Theory*, 62(1):1–23, 1994.

[Gabarró *et al.*, 2011] Joaquim Gabarró, Alina García, and Maria J. Serna. The complexity of game isomorphism. *Theor. Comput. Sci.*, 412(48):6675–6695, 2011.

[Gabarró *et al.*, 2013] Joaquim Gabarró, Alina García, and Maria J. Serna. On the hardness of game equivalence under local isomorphism. *RAIRO Theor. Informatics Appl.*, 47(2):147–169, 2013.

[Garey and Johnson, 1990] Michael R. Garey and David S. Johnson. *Computers and Intractability; A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., USA, 1990.

[Guo and Conitzer, 2010] Mingyu Guo and Vincent Conitzer. Optimal-in-expectation redistribution mechanisms. *Artif. Intell.*, 174(5-6):363–381, 2010.

[Hammond, 2005] Peter J. Hammond. Utility invariance in non-cooperative games. In Ulrich Schmidt and Stefan Traub, editors, *Advances in Public Economics: Utility, Choice and Welfare*, pages 31–50. Springer, Boston, MA, 2005.

[Harsanyi and Selten, 1988] J.C. Harsanyi and R. Selten. *A General Theory of Equilibrium Selection in Games*. MIT Press Classics. MIT Press, 1988.

[Heyman and Gupta, 2023] Joseph L. Heyman and Abhishek Gupta. Rank reduction in bimatrix games. *IGTR*, 25(1), 2023.

[Kohlberg and Mertens, 1986] Elon Kohlberg and Jean-Francois Mertens. On the strategic stability of equilibria. *Econometrica*, 54(5):1003–1037, 1986.

[Kontogiannis and Spirakis, 2012] Spyros C. Kontogiannis and Paul G. Spirakis. On mutual concavity and strategically-zero-sum bimatrix games. *Theor. Comput. Sci.*, 432:64–76, 2012.

[Likhodedov and Sandholm, 2004] Anton Likhodedov and Tuomas Sandholm. Methods for boosting revenue in combinatorial auctions. In Deborah L. McGuinness and George Ferguson, editors, *Proceedings of the Nineteenth National Conference on Artificial Intelligence, Sixteenth Conference on Innovative Applications of Artificial Intelligence, July 25-29, 2004, San Jose, California, USA*, pages 232–237. AAAI Press / The MIT Press, 2004.

[Liu, 1995] Luchuan A. Liu. *On Correlated Equilibrium and Invariance of Best Reply Correspondence*. PhD thesis, State University of New York at Stony Brook, August 1995. UMI number: 9605308.

[Liu, 1996] Luchuan A. Liu. The invariance of best reply correspondences in two-player games. 93, City University of Hong Kong, Faculty of Business, Department of Economics, 1996.

[Marris *et al.*, 2023] Luke Marris, Ian Gemp, and Georgios Piliouras. Equilibrium-invariant embedding, metric space, and fundamental set of 2×2 normal-form games, 2023.

[Mas-Colell *et al.*, 1995] Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory*. Oxford University Press, 1995.

[Maschler *et al.*, 2013] Michael Maschler, Eilon Solan, and Shmuel Zamir. *Game Theory*. Cambridge University Press, 2013.

[McKinsey, 1951] J. C. C. McKinsey. Isomorphism of games, and strategic equivalence. In Harold William Kuhn and Albert William Tucker, editors, *Contributions to the Theory of Games (AM-24), Volume I*, pages 117–130. Princeton University Press, Princeton, 1951.

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

[Morris and Ui, 2004] Stephen Morris and Takashi Ui. Best response equivalence. *Games Econ. Behav.*, 49(2):260–287, 2004.

[Moulin and Vial, 1978] H. Moulin and J.P. Vial. Strategically zero-sum games: The class of games whose completely mixed equilibria cannot be improved upon. *International Journal of Game Theory*, 7:201–221, 1978.

[Nash, 1950] John F. Nash. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences*, 36(1):48–49, 1950.

[Oesterheld and Conitzer, 2022] Caspar Oesterheld and Vincent Conitzer. Safe pareto improvements for delegated game playing. *Auton. Agents Multi Agent Syst.*, 36(2):46, 2022.

[Pearce, 1984] David G. Pearce. Rationalizable strategic behavior and the problem of perfection. *Econometrica*, 52(4):1029–1050, 1984.

[Pottier and Nessah, 2014] Antonin Pottier and Rabia Nessah. Berge-vaisman and nash equilibria: Transformation of games. *IGTR*, 16(4), 2014.

[Rapoport *et al.*, 1976] Anatol Rapoport, Melvin Guyer, and David G. Gordon. *The 2 X 2 game*. University of Michigan Press Ann Arbor, 1976.

[Reem, 2017] Daniel Reem. Remarks on the cauchy functional equation and variations of it. *Aequationes mathematicae*, 91:237–264, 2017.

[Robinson and Goforth, 2005] David Robinson and David Goforth. *The Topology of the 2×2 Games*. Routledge, 2005.

[Thompson, 1952] F. B. Thompson. *Equivalence of Games in Extensive Form*. RAND Corporation, Santa Monica, CA, 1952.

[von Neumann and Morgenstern, 1944] John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1944.

[von Neumann, 1928] John von Neumann. Zur theorie der gesellschaftsspiele. *Mathematische Annalen*, 100:295–320, 1928.

[von Stengel, 2022] Bernhard von Stengel. *Game Theory Basics*. Cambridge University Press, 2022.

[Wu *et al.*, 2022] Yuhu Wu, Shuting Le, Kuize Zhang, and Xi-Ming Sun. Agent transformation of bayesian games. *IEEE Transactions on Automatic Control*, 67(11):5793–5808, 2022.