

On the Complexity of Iterated Weak Dominance in Constant-Sum Games

Felix Brandt · Markus Brill · Felix Fischer · Paul Harrenstein

Abstract In game theory, a player's action is said to be weakly dominated if there exists another action that, with respect to what the other players do, is never worse and sometimes strictly better. We investigate the computational complexity of the process of iteratively eliminating weakly dominated actions (IWD) in two-player constant-sum games, i.e., games in which the interests of both players are diametrically opposed. It turns out that deciding whether an action is eliminable via IWD is feasible in polynomial time whereas deciding whether a given subgame is reachable via IWD is NP-complete. The latter result is quite surprising, as we are not aware of other natural computational problems that are intractable in constant-sum normal-form games. Furthermore, we slightly improve on a result of Conitzer and Sandholm by showing that typical problems associated with IWD in win-lose games with at most one winner are NP-complete.

Keywords Game Theory · Constant-Sum Games · Solution Concepts · Iterated Weak Dominance · Computational Complexity

1 Introduction

A simple and indisputable conviction in game theory is that a player need not bother to consider an action that yields less payoff than some other action *no matter* what all the other players do (see, e.g., [11]). In the language of game theory, such an action is called *strictly dominated*. Similarly, one says that an action is *weakly dominated* if there exists another action that, with respect to what the other players do, is never worse and sometimes strictly better. An action that is not weakly dominated is also said to be *admissible*. When a (strictly or weakly) dominated action is eliminated from a player's consideration, it may be possible that a previously undominated action of

F. Brandt · M. Brill · P. Harrenstein
Institut für Informatik, Technische Universität München, 85748 Garching, Germany
E-mail: {brandtf,brill,harrenst}@in.tum.de

F. Fischer
Harvard School of Engineering and Applied Sciences, Cambridge, MA 02138, USA
E-mail: fischerf@seas.harvard.edu

another player becomes dominated. Thus, based on the mutual rational belief that (some) dominated actions will not be played, one can define an iterative process of eliminating actions. It is well-known that this process invariably leads to the same subgame no matter in which order *strictly* dominated actions are eliminated, whereas the same is not the case for weak dominance (see, e.g., [1, 19]). The dependence on the order of elimination gives rise to some combinatorial difficulties, as witnessed by the NP-completeness of various computational problems related to iterated weak dominance [8, 6]. By contrast, the corresponding problems for iterated strict dominance are computationally tractable. This disparity has also become apparent in the complexity analysis of other solution concepts based on dominance [4].

We investigate the computational complexity of *iterated weak dominance (IWD)*—or *iterated admissibility*—in *two-player constant-sum* games, i.e., games in which the interests of both players are diametrically opposed. Our analysis is restricted to dominance by pure strategies, but most of our results readily apply to mixed strategies as well (see Section 6). It turns out that deciding whether an action is eliminable via IWD is feasible in polynomial time, whereas deciding whether a given subgame is reachable via IWD is NP-complete. The latter result is quite surprising, as we are not aware of other natural computational problems that are intractable in normal-form constant-sum games.¹ Furthermore, we slightly improve on a result of Conitzer and Sandholm [6] by showing that typical problems associated with IWD in win-lose games *with at most one winner* are NP-complete.

Iterated weak and strict dominance are well-established solution concepts, which have a long history and appear in virtually every textbook on game theory. The work of Bernheim [2] and Pearce [16] has instigated a renewed discussion concerning the formal and intuitive connections of iterated dominance with rationalizability and the epistemic foundations of solution concepts [17, 3], the stability of equilibria [9], and backward induction solutions [7, 18]. It cannot be said that iterated weak dominance has left the arena entirely unscathed. Unlike iterated strict dominance, proper epistemic foundations for iterated weak dominance are pretty hard to come by. In particular, Samuelson [17] showed that common knowledge of admissibility does not imply iterated weak dominance. Nevertheless, IWD has its place as a tool in the analysis of games (see, e.g., [13, 14], for discussions).

Our aim in this paper is not to pass judgement on iterated weak dominance as a solution concept. Rather, our focus is on the computational aspects of IWD in two-player constant-sum games and win-lose games with at most one winner. As mentioned above, the fact that some of these problems turn out to be NP-hard is interesting and surprising in its own right.

After having introduced our formal framework in Section 2, we propose the auxiliary concept of a regionalized game in Section 3 and show that this concept may be used as a convenient tool in the proofs of our hardness results. In Section 4 we deal with the computational complexity of reachability and eliminability problems in two-player constant-sum games. Finally, in Section 5, we address the same problems for win-lose games that allow at most one winner.

¹ Hard problems are known in the context of *extensive-form* constant-sum games. For instance, Koller and Megiddo [10] show that finding maximin behavior strategies in extensive-form constant-sum games without perfect recall is NP-hard.

2 Preliminaries

A *two-player game* $\Gamma = (A_1, A_2, u)$ is given by a finite set A_1 of actions of player 1, a finite set A_2 of actions of player 2, and a utility function $u : A_1 \times A_2 \rightarrow \mathbb{R} \times \mathbb{R}$. We assume each action to be indexed by its respective player, such that $A_1 \cap A_2 = \emptyset$, but sometimes omit the indices to avoid cluttered notation. We have A denote $A_1 \cup A_2$ and write $u_1(a, b) = x$ and $u_2(a, b) = y$ if $u(a, b) = (x, y)$. Both players are assumed to choose one of their actions simultaneously. If player 1 chooses a and player 2 chooses b , their payoffs will be $u_1(a, b)$ and $u_2(a, b)$, respectively.

A two-player game is called *constant-sum game* if $u_1(a, b) + u_2(a, b) = u_1(c, d) + u_2(c, d)$ for all $a, c \in A_1$ and $b, d \in A_2$. Such a game can conveniently be represented by writing down the payoffs of player 1 in a matrix with rows indexed by the actions of player 1 and columns indexed by the actions of player 2.

Let $\Gamma = (A_1, A_2, u)$ be a two-player game and $a, b \in A_1$ two actions of player 1. Then, a is said to *weakly dominate* b at $c \in A_2$ in Γ if $u_1(a, c) > u_1(b, c)$ and for all $d \in A_2$, $u_1(a, d) \geq u_1(b, d)$. More generally, a is said to weakly dominate b if a weakly dominates b at c for some $c \in A_2$. The weak dominance relation is both asymmetric—i.e., if action a weakly dominates action b , action b does not weakly dominate action a —and transitive—i.e., if action a weakly dominates action b and action b weakly dominates action c , then action a also weakly dominates action c . For a game $\Gamma' = (A'_1, A'_2, u)$ with $A'_1 \subseteq A_1$ and $A'_2 \subseteq A_2$, we say further that an action $c \in A_2$ *backs* the elimination of $b \in A_1$ by $a \in A_1$ in Γ' if $a, b, c \in A'_1 \cup A'_2$ and $u_1(a, c) > u_1(b, c)$, and *blocks* the elimination of b by a in Γ' if $a, b, c \in A'_1 \cup A'_2$ and $u_1(a, c) < u_1(b, c)$. Dominance, backing, and blocking for actions of player 2 are defined analogously. Note that an action is dominated by another action of the same player if some action of the other player backs the elimination and none of them block it. As the remainder of this paper only concerns (iterated) weak dominance, we will drop the qualification ‘weak’ and by ‘dominance’ understand weak dominance.

An *elimination sequence* of a game $\Gamma = (A_1, A_2, u)$ is a finite sequence $\Sigma = (\Sigma_1, \dots, \Sigma_n)$ of pairwise disjoint subsets of actions in A . For a game $\Gamma = (A_1, A_2, u)$ and an elimination sequence $\Sigma = (\Sigma_1, \dots, \Sigma_n)$ of Γ we have $\Gamma(\Sigma)$ denote the subgame where the actions in $\Sigma_1 \cup \dots \cup \Sigma_n$ have been removed, i.e., $\Gamma(\Sigma) = (A'_1, A'_2, u')$ where $A'_1 = A_1 \setminus (\Sigma_1 \cup \dots \cup \Sigma_n)$, $A'_2 = A_2 \setminus (\Sigma_1 \cup \dots \cup \Sigma_n)$, and u' is the restriction of u to $A'_1 \times A'_2$. The *validity* of elimination sequences is then defined inductively: the empty sequence ϵ is valid for every game, and an elimination sequence $(\Sigma_1, \dots, \Sigma_n, \Sigma_{n+1})$ is valid for Γ if $(\Sigma_1, \dots, \Sigma_n)$ is valid for Γ and every action $a \in \Sigma_{n+1}$ is dominated in $\Gamma(\Sigma_1, \dots, \Sigma_n)$. If every Σ_i is a singleton, we say the elimination sequence $(\Sigma_1, \dots, \Sigma_n)$ is *simple*. Simple elimination sequences we usually write as sequences $\sigma = (\sigma_1, \dots, \sigma_n)$ of actions in A .

An action a is called *eliminable by* b at c in a game Γ if there exists a valid elimination sequence Σ such that a is dominated by b at c in $\Gamma(\Sigma)$. Action a is *eliminable* in Γ if there are actions b and c such that a is eliminable by b at c . A subgame Γ' of Γ is *reachable* from Γ if there exists a valid elimination sequence Σ such that $\Gamma(\Sigma) = \Gamma'$. Furthermore Γ is called *solvable* if some subgame $\Gamma' = (A'_1, A'_2, u')$ with $|A'_1| = |A'_2| = 1$ is reachable from Γ . Finally, we say that Γ is *irreducible* if none of its actions are dominated.

We assume the reader to be familiar with the theory of computational complexity, in particular with the complexity classes P and NP and the NP-complete problem *3SAT* (see, e.g., [15]).

3 Regions and Regionalized Games

An essential building block of our hardness proofs are *regionalized games*, i.e., games in which the action set A_i of each player i is divided up into *regions*. Intuitively, the regions prevent eliminations of actions by actions from other regions. We assume for each player i that the regions constitute a partition of A_i , i.e., a set of non-empty and pairwise disjoint subsets of A_i the union of which exhausts A_i . More formally, a *regionalized two-player game* is a tuple (Γ, X_1, X_2) consisting of a two-player game $\Gamma = (A_1, A_2, u)$, a partition X_1 of A_1 , and a partition X_2 of A_2 . The elements of X_1 and X_2 are called *regions*.

For regionalized games, the concept of a valid elimination sequence is modified so as to allow only eliminations of actions that are dominated by other actions in the same region. A *valid elimination sequence* for a regionalized game (Γ, X_1, X_2) is a sequence $\Sigma = (\Sigma_1, \dots, \Sigma_n)$ for Γ such that for each i with $1 \leq i \leq n$ and each $a \in \Sigma_i$, there is some action b and some $x \in X_1 \cup X_2$ such that $a, b \in x$ and b dominates a in $\Gamma(\Sigma_1, \dots, \Sigma_{i-1})$. With a slight abuse of notation we will use $(\Gamma(\Sigma_1, \dots, \Sigma_n), X_1, X_2)$ to refer to the regionalized game $(\Gamma(\Sigma_1, \dots, \Sigma_n), X'_1, X'_2)$ where $X'_1 = \{x \setminus (\Sigma_1 \cup \dots \cup \Sigma_n) : x \in X_1\} \setminus \{\emptyset\}$ and $X'_2 = \{x \setminus (\Sigma_1 \cup \dots \cup \Sigma_n) : x \in X_2\} \setminus \{\emptyset\}$.

We now prove a useful lemma: any regionalized two-player game can be transformed in polynomial time into a non-regionalized two-player game with the same valid elimination sequences. It follows that for two of the three computational problems we consider—reachability of (irreducible) subgames and eliminability—we can thus restrict ourselves to regionalized games, which are often more practical for and afford more insight into the constructions used in our hardness proofs than games without regions.

Lemma 1 *For each regionalized game (Γ, X_1, X_2) with $\Gamma = (A_1, A_2, u)$, there is a game $\Gamma' = (A'_1, A'_2, u')$ computable in polynomial time such that the valid elimination sequences for Γ' and (Γ, X_1, X_2) coincide. Moreover, $u'(a, b) \in \{(0, 1), (1, 0)\}$ for all $(a, b) \in (A'_1 \times A'_2) \setminus (A_1 \times A_2)$.*

Proof The game Γ' is constructed from Γ by adding actions that impose the same restrictions on the elimination of actions as the regions did in (Γ, X_1, X_2) . More auxiliary actions are then added to ensure that all elimination sequences that are valid for (Γ, X_1, X_2) are still valid for Γ' while no new valid elimination sequences are created.

Formally, let $\Gamma' = (A'_1, A'_2, u')$ with $A'_1 = A_1 \cup X_2 \cup Y_1$ and $A'_2 = A_2 \cup X_1 \cup Y_2$, where $Y_1 = \{y_1^1, y_1^2, y_1^3, y_1^4\}$ and $Y_2 = \{y_2^1, y_2^2, y_2^3, y_2^4\}$ are sets of actions disjoint from $A_1 \cup X_2$ and $A_2 \cup X_1$. Observe that the regions in (Γ, X_1, X_2) correspond to actions of the respective other player in Γ' . Further define the utility function u' such that $u'(a, b) = u(a, b)$ for all $(a, b) \in A_1 \times A_2$. For every $(a, x) \in A_1 \times X_1$ and $(x, b) \in X_2 \times A_2$, let

$$u'(a, x) = \begin{cases} (1, 0) & \text{if } a \in x, \\ (0, 1) & \text{otherwise,} \end{cases} \quad \text{and} \quad u'(x, b) = \begin{cases} (0, 1) & \text{if } b \in x, \\ (1, 0) & \text{otherwise.} \end{cases}$$

Without loss of generality we may assume that $|X_1| = |X_2| = k$ for some index $k \geq 0$, as we can always introduce copies of actions to the game. Thus, let $X_1 = \{x_1^1, \dots, x_1^k\}$ and $X_2 = \{x_1^1, \dots, x_1^k\}$ and define for all indices i and j with $1 \leq i, j \leq k$,

$$u'(x_1^i, x_2^j) = \begin{cases} (1, 0) & \text{if } i = j, \\ (0, 1) & \text{otherwise.} \end{cases}$$

					X_1						
		a_2^1	\dots	a_2^m	x_2^1	\dots	x_2^k	y_2^1	y_2^2	y_2^3	y_2^4
a_1^1	\vdots	Γ			\cdot	\dots	\cdot	$(0,1)$	$(1,0)$	$(0,1)$	$(1,0)$
					\vdots	\ddots	\vdots	\vdots	\vdots	\vdots	\vdots
					\cdot	\dots	\cdot	$(0,1)$	$(1,0)$	$(0,1)$	$(1,0)$
a_1^ℓ	\vdots				\cdot	\dots	\cdot	$(0,1)$	$(1,0)$	$(0,1)$	$(1,0)$
					\vdots	\ddots	\vdots	\vdots	\vdots	\vdots	\vdots
					\cdot	\dots	\cdot	$(0,1)$	$(1,0)$	$(0,1)$	$(1,0)$
X_2	$\left\{ \begin{array}{l} x_1^1 \\ \vdots \\ x_1^k \end{array} \right.$	\cdot	\dots	\cdot	$(1,0)$	\dots	$(0,1)$	$(1,0)$	$(0,1)$	$(1,0)$	$(0,1)$
		\vdots	\ddots	\vdots	\vdots	\ddots	\vdots	\vdots	\vdots	\vdots	\vdots
		\cdot	\dots	\cdot	$(0,1)$	\dots	$(1,0)$	$(1,0)$	$(0,1)$	$(1,0)$	$(0,1)$
y_1^1	$(1,0)$	\dots	$(1,0)$	$(0,1)$	\dots	$(0,1)$	$(1,0)$	$(1,0)$	$(0,1)$	$(0,1)$	
y_1^2	$(0,1)$	\dots	$(0,1)$	$(1,0)$	\dots	$(1,0)$	$(1,0)$	$(1,0)$	$(0,1)$	$(0,1)$	
y_1^3	$(1,0)$	\dots	$(1,0)$	$(0,1)$	\dots	$(0,1)$	$(0,1)$	$(0,1)$	$(1,0)$	$(1,0)$	
y_1^4	$(0,1)$	\dots	$(0,1)$	$(1,0)$	\dots	$(1,0)$	$(0,1)$	$(0,1)$	$(1,0)$	$(1,0)$	

Fig. 1 Game Γ' used in the proof of Lemma 1

The payoffs for the remaining action profiles are depicted in Figure 1. Obviously, Γ' can be obtained from (Γ, X_1, X_2) in polynomial time.

Before we show that the valid elimination sequences for (Γ, X_1, X_2) and Γ' coincide, we note that the utility function u' is chosen so as to ensure that none of the first player's actions in $X_2 \cup Y_1$ nor any of the second player's actions in $X_1 \cup Y_2$ appear in any valid elimination sequence for Γ' . To see this, observe that for each action $a \in X_2 \cup Y_1$ and each action $b \in A_1'$ there is some action $X_1 \cup Y_2$ that blocks the elimination of a by b in Γ' . For instance, x_2^1 blocks the elimination of y_2^1 by y_1^1 . Moreover, for $a \in A_1$ and $b \in X_2 \cup Y_1$, there is some action in $X_1 \cup Y_2$ blocking the elimination of a by b in Γ' . It follows that for every valid elimination sequence Σ for Γ' , if $a \in A_1'$ is dominated by $b_1 \in A_1'$ in $\Gamma'(\Sigma)$, then $a, b \in A_1$. By symmetrical arguments, an analogous statement holds for actions $a, b \in A_2'$.

Now consider an arbitrary valid elimination sequence $\Sigma = (\Sigma_1, \dots, \Sigma_n)$ for Γ' . Then, $\Sigma_1 \cup \dots \cup \Sigma_n \subseteq A_1 \cup A_2$. Also consider an arbitrary index i with $1 \leq i \leq n$ and an arbitrary action $a \in \Sigma_i$. Without loss of generality we may assume that $a \in A_1$ and that there are actions $b \in A_1$ and $c \in A_2'$ such that a is dominated by b at c in $\Gamma'(\Sigma_1, \dots, \Sigma_{i-1})$. Let $x \in X_1$ be the region of (Γ, X_1, X_2) with $b \in x$. It follows that $a \in x$ as well, otherwise the elimination of a by b would be blocked by x in $\Gamma'(\Sigma_1, \dots, \Sigma_{i-1})$. With this being the case, observe that $u_1'(a, z) = u_1'(b, z)$ for all $z \in X_1 \cup Y_2$, i.e., no $z \in X_1 \cup Y_2$ backs the elimination of a by b . Hence, $c \in A_2 \setminus (\Sigma_1 \cup \dots \cup \Sigma_{i-1})$. It follows that a is dominated by b at c in $\Gamma(\Sigma_1, \dots, \Sigma_{i-1})$ and, because a and b are in the same region $x \in X_1$, a is dominated by b at c in $(\Gamma(\Sigma_1, \dots, \Sigma_{i-1}), X_1, X_2)$ as well. We may conclude that Σ is also a valid elimination sequence for (Γ, X_2, X_2) .

Finally, consider a valid elimination sequence $\Sigma = (\Sigma_1, \dots, \Sigma_n)$ for (Γ, X_1, X_2) , an index i with $1 \leq i \leq n$, and an action $a \in \Sigma_i$. Without loss of generality we may assume that $a \in A_1$. Suppose that a is dominated by b at c in $(\Gamma(\Sigma_1, \dots, \Sigma_{i-1}), X_1, X_2)$ for

	a_2	b_2
a_1	(1, 0)	(1, 0)
b_1	(0, 1)	(1, 0)
c_1	(1, 0)	(0, 1)

Fig. 2 IWD is order dependent.

some actions $b \in A_1$ and $c \in A_2$. Obviously, a is dominated by b at c in $\Gamma(\Sigma_1, \dots, \Sigma_{i-1})$ as well. Moreover, a and b belong to the same region $x \in X_1$. Accordingly, no action $z \in X_1 \cup Y_2$ blocks the elimination of a by b in $\Gamma'(\Sigma_1, \dots, \Sigma_{i-1})$. It follows that a is dominated by b at c in $\Gamma'(\Sigma_1, \dots, \Sigma_{i-1})$ and that Σ is a valid elimination sequence for Γ' . \square

4 Two-Player Constant-sum Games

We will now show that subgame reachability is NP-complete even in games that only allow the outcomes (0, 1) and (1, 0). This may be attributed to the order dependence of IWD. For example, (b_1, a_2) is a valid elimination sequence for the game in Figure 2. However, if one eliminates row c_1 first, column a_2 is no longer eliminable.

In Section 4.2 we will find that for two-player constant-sum games a weak form of order independence can be salvaged, which allows us to formulate an efficient algorithm for the eliminability problem. Our first observation is that in the case of two-player constant-sum games we can restrict our attention to *simple* elimination sequences.

Lemma 2 *Let $\Gamma = (A_1, A_2, u)$ be a two-player constant-sum game and $\Sigma = (\Sigma_1, \dots, \Sigma_m)$ a valid elimination sequence. Then, there exists a simple elimination sequence $\sigma = (\sigma_1, \dots, \sigma_n)$ with $\{\sigma_1, \dots, \sigma_n\} = \Sigma_1 \cup \dots \cup \Sigma_m$ that is valid for Γ .*

Proof Let X be a non-empty subset of A . It suffices to show that validity of the one-element sequence (X) for Γ implies the existence of some $x \in X$ such that the sequence $(X \setminus \{x\}, \{x\})$ is valid for Γ as well.

Assume for contradiction that (X) is valid but, for any $x \in X$, $(X \setminus \{x\}, \{x\})$ is not valid. Consider an arbitrary $x \in X$ and assume without loss of generality that $x \in A_1$. Then, x is dominated by some $x' \in A_1$ at some $y \in A_2$, i.e., $u_1(x', y) > u_1(x, y)$. Recall that the dominance relation is asymmetric and transitive and that X is finite. Hence, without loss of generality, we may assume that $x' \notin X$.² By contrast, $y \in X$. To see this, observe that $(X \setminus \{x\})$ is valid for Γ . Moreover, as no action blocks the elimination of x by x' in Γ , neither is this the case for $\Gamma(X \setminus \{x\})$. If $y \notin X$, then $y \notin X \setminus \{x\}$, and x' would dominate x at y for $\Gamma(X \setminus \{x\})$. Consequently, $(X \setminus \{x\}, \{x\})$ would be valid for Γ , a contradiction.

Now, since $y \in X$, there must be some $y' \in A_2$ dominating y in Γ . By asymmetry and transitivity of the dominance relation, we may assume that $y' \notin X$. Moreover, there are no actions blocking the elimination of y by y' in Γ . Having assumed, however, that

² As X is finite, by asymmetry and transitivity of the dominance relation there is a maximal sequence x_1, \dots, x_k of pairwise distinct actions in X , such that $x = x_1$ and x_{i+1} dominates x_i for each i with $1 \leq i < k$. By assumption there also has to be some action x' in A that dominates x_k . By maximality of x_1, \dots, x_k , we have $x' \notin X$. Finally, by transitivity of the dominance relation, it follows that x' also dominates x .

	y	y'
x'	$u_1(x', y)$	$u_1(x', y')$
	\vee	\wedge
x	$u_1(x, y)$	$u_1(x, y')$

Fig. 3 Diagram illustrating the proof of Lemma 2

	y	y'
x	$(0, 0)$	$(0, 1)$
x'	$(1, 0)$	$(0, 0)$

Fig. 4 Game with weakly dominated actions x and y and a valid elimination sequence $(\{x, y\})$. The simple elimination sequences (x, y) and (y, x) are not valid.

$(X \setminus \{y\}, \{y\})$ is not valid for Γ , it follows that x' does not back the elimination of y by y' in Γ , i.e., $u_2(x', y') \leq u_2(x', y)$. As Γ is a constant-sum game, $u_1(x', y') \geq u_1(x', y)$. Similarly, there is no action blocking the elimination of x by x' in Γ , whereas $(X \setminus \{x\}, \{x\})$ is not valid for Γ . Hence, y' does not back the elimination of x by x' in Γ , i.e., $u_1(x, y') \geq u_1(x, y)$. This situation is illustrated in Figure 3. It now follows that $u_1(x, y') > u_1(x, y)$ and, since Γ is a constant-sum game, $u_2(x, y') < u_2(x, y)$, contradicting the assumption that y' dominates y in Γ . \square

As a corollary of Lemma 2 we find that a subgame of a two-player constant-sum game is reachable if and only if it is reachable by a simple elimination sequence. Analogous statements also hold for eliminability and solvability. Lemma 2 however does not hold for general strategic games. In fact, it already fails for games with outcomes in $\{(0, 0), (0, 1), (1, 0)\}$, as Figure 4 illustrates.

4.1 Reachability

We are now ready to show that subgame reachability in constant-sum games is computationally intractable.

Theorem 1 *Given constant-sum games Γ and Γ' , deciding whether Γ' is reachable from Γ is NP-complete, even if Γ only has outcomes $(0, 1)$ and $(1, 0)$ and Γ' is irreducible.*

Proof For membership in NP consider arbitrary constant-sum games Γ and Γ' . Given an elimination sequence σ , it can clearly be decided in polynomial time whether Σ is a valid elimination sequence for (Γ, X_1, X_2) such that $\Gamma(\Sigma) = \Gamma'$.

The proof of hardness proceeds by a reduction from 3SAT. By virtue of Lemma 1 it suffices to give a reduction for regionalized games. Consider an arbitrary 3CNF $\varphi = C_1 \wedge \dots \wedge C_k$, where each $C_i = (\lambda_i^1 \vee \lambda_i^2 \vee \lambda_i^3)$ is a clause and each λ_i^j is a literal, for $1 \leq i \leq k$ and $1 \leq j \leq 3$. Without loss of generality, we may assume that all clauses in φ are distinct. Define a regionalized game $(\Gamma_\varphi, X_1, X_2)$, with $\Gamma_\varphi = (A_1, A_2, u)$ as

follows.

$$\begin{aligned}
A_1 &= \{p, \neg p, p\downarrow : p \text{ a variable in } \varphi\} \\
&\quad \cup \{C_i, (\lambda_i^1, i), (\lambda_i^2, i), (\lambda_i^3, i) : C_i \text{ a clause in } \varphi\} \\
&\quad \cup \{e\} \\
A_2 &= \{p, \neg p : p \text{ a variable in } \varphi\} \cup \{a, b\} \\
X_1 &= \{\{p, \neg p, p\downarrow\} : p \text{ a variable in } \varphi\} \\
&\quad \cup \{\{C_i, (\lambda_i^1, i), (\lambda_i^2, i), (\lambda_i^3, i)\} : C_i \text{ a clause in } \varphi\} \\
&\quad \cup \{\{e\}\} \\
X_2 &= \{\{p, \neg p : p \text{ a variable in } \varphi\} \cup \{a, b\}\} = \{A_2\}
\end{aligned}$$

For each variable p in φ , the payoffs in rows p , $\neg p$ and $p\downarrow$ are defined as in the following table, where q is a typical variable in φ distinct from p .

	p	$\neg p$	q	$\neg q$	a	b
p	(1, 0)	(0, 1)	(0, 1)	(0, 1)	(1, 0)	(0, 1)
$\neg p$	(0, 1)	(1, 0)	(0, 1)	(0, 1)	(1, 0)	(0, 1)
$p\downarrow$	(1, 0)	(1, 0)	(0, 1)	(0, 1)	(0, 1)	(0, 1)

Due to the regionalization, row $p\downarrow$ can be eliminated only by row p or row $\neg p$. Column a is the only action backing such an elimination. Intuitively, removing column p means setting variable p to false, removing column $\neg p$ setting variable p to true, thus choosing an assignment. Row $p\downarrow$ can thus be eliminated only after one of these columns has been removed, i.e., after an assignment for p has been chosen.

For each i with $1 \leq i \leq k$, the payoffs in rows C_i , (λ_i^1, i) , (λ_i^2, i) , (λ_i^3, i) depend on the literals occurring in C_i . In the following table, $\bar{\lambda}_i^j = \neg p$ if $\lambda_i^j = p$, and $\bar{\lambda}_i^j = p$ if $\lambda_i^j = \neg p$. We further assume $i \neq m$.

	λ_i^1	$\bar{\lambda}_i^1$	λ_i^2	$\bar{\lambda}_i^2$	λ_i^3	$\bar{\lambda}_i^3$	λ_m^j	$\bar{\lambda}_m^j$	a	b
(λ_i^1, i)	(1, 0)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 1)
(λ_i^2, i)	(0, 1)	(0, 1)	(1, 0)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 1)
(λ_i^3, i)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(1, 0)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 1)
C_i	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(1, 0)	(0, 1)

Thus, the only columns backing the elimination of C_i are λ_i^1 , λ_i^2 , and λ_i^3 . Also note that column a blocks the elimination of C_i . On the other hand, as we saw above, column a is essential to the elimination of the rows $p\downarrow$. Intuitively, this means that an assignment needs to be chosen before any of the rows C_i is eliminated.

Finally, we let $u(e, y) = (1, 0)$ if $y \neq b$, and $u(e, b) = (0, 1)$:

	λ_1^1	$\bar{\lambda}_1^1$	\dots	λ_k^3	$\bar{\lambda}_k^3$	a	b
e	(1, 0)	(1, 0)	\dots	(1, 0)	(1, 0)	(1, 0)	(0, 1)

Observe that row e is the only action in its region and as such cannot be eliminated, and that it backs the elimination of every column by b .

Now define $(\Gamma'_\varphi, X'_1, X'_2)$ with $\Gamma'_\varphi = (A'_1, A'_2, u')$ such that

$$\begin{aligned} A'_1 &= \{p, \neg p : p \text{ a variable in } \varphi\} \\ &\quad \cup \{(\lambda_i^1, i), (\lambda_i^2, i), (\lambda_i^3, i) : C_i \text{ a clause in } \varphi\} \\ &\quad \cup \{e\}, \\ A'_2 &= \{b\}, \end{aligned}$$

and the utility function u' and the partitions X'_1 and X'_2 are restricted appropriately to A'_1 and A'_2 , i.e., $u' = u|_{A'_1 \times A'_2}$, $X'_1 = \{x \cap A'_1 : x \in X_1\} \setminus \{\emptyset\}$ and $X'_2 = \{x \cap A'_2 : x \in X_2\} \setminus \{\emptyset\}$. It is readily appreciated that no actions can be eliminated in $(\Gamma'_\varphi, X'_1, X'_2)$, i.e., that $(\Gamma'_\varphi, X'_1, X'_2)$ is irreducible.

We now prove that φ is satisfiable if and only if $(\Gamma'_\varphi, X'_1, X'_2)$ is reachable from $(\Gamma_\varphi, X_1, X_2)$.

For the direction from left to right, assume that φ is satisfiable and consider a satisfying assignment v . Start by eliminating, using column b , each column corresponding to a literal that is set to false by v . Subsequently, for each variable p , eliminate row $p \downarrow$ by row p or row $\neg p$. This is possible since either column p or column $\neg p$ have been eliminated in the first step. Next eliminate column a by column b . Since v is a satisfying assignment, there remains for each clause $C_i = (\lambda_i^1 \vee \lambda_i^2 \vee \lambda_i^3)$ a column λ_i^j , which now backs the elimination of row C_i by row (λ_i^j, i) . Finally eliminating by column b all other remaining columns, we reach subgame $(\Gamma'_\varphi, X'_1, X'_2)$.

For the direction from right to left, assume that $(\Gamma'_\varphi, X'_1, X'_2)$ is reachable from (Γ, X_1, X_2) . Observe that this specifically requires the elimination of row $p \downarrow$ for each variable p occurring in φ , and recall that for this to be possible at least one of the columns p and $\neg p$ needs to be eliminated while column a is still present to back the elimination. Furthermore, row C_i must be eliminated for each $1 \leq i \leq k$, which can only take place by some row (λ_i^j, i) and backed by column λ_i^j , and only when column a is no longer present to block the elimination. We can thus define an assignment v^* that satisfies exactly those literals λ_i^j corresponding to columns present when column a is eliminated. It is readily appreciated that v^* is well-defined and satisfies φ . \square

Solvability is a special case of subgame reachability, and is tractable for two-player single-winner games, i.e., for constant-sum games which only allow outcomes $(0, 1)$ and $(1, 0)$ [5]. Whether solvability is tractable in general constant-sum games remains an open question.

4.2 Eliminability

As we have seen, the iterated elimination of weakly dominated actions may depend on the order in which actions are eliminated. If an elimination sequence $\sigma = (\sigma_1, \dots, \sigma_n)$ is valid for a game Γ , it does not automatically follow that σ remains valid if some dominated action d different from σ_1 is removed first. Consider for example the game Γ depicted in Figure 5 and the elimination sequence (x, v, y, u, a) , which is valid for this game. Action d , which is itself dominated by action c , is the only action backing the

	c	d	u	v	w
a	(0, 2)	(2, 0)	(1, 1)	(2, 0)	(0, 2)
b	(1, 1)	(2, 0)	(0, 2)	(2, 0)	(0, 2)
x	(0, 2)	(0, 2)	(0, 2)	(2, 0)	(1, 1)
y	(0, 2)	(1, 1)	(0, 2)	(2, 0)	(1, 1)
z	(0, 2)	(2, 0)	(2, 0)	(1, 1)	(1, 1)

Fig. 5 Constant-sum game Γ illustrating that an elimination sequence need not remain valid if an action is eliminated. The elimination sequence (x, v, y, u, a) is valid for Γ , but the elimination sequence (d, x, v, y, u, a) is not.

elimination of x in Γ . Thus the elimination sequence (x, v, y, u, a) is no longer valid when d is eliminated first.

It turns out, however, that by delaying the elimination of x until y has been eliminated one can obtain an elimination sequence, viz. (v, y, x, u, a) , that is valid for $\Gamma(d)$. We will see presently that this is just an example of a more general property of elimination sequences in two-player constant-sum games: given a valid elimination sequence σ and a dominated action d , one can carry out the elimination of d early and still find a valid elimination sequence that eliminates all the actions in σ , provided that one is prepared to postpone the elimination of some of these actions. This insight will be instrumental to the proof of Theorem 2, which states that the eliminability problem for two-player constant-sum games can be solved efficiently.

We need some auxiliary terminology and notation. Fix a game $\Gamma = (A_1, A_2, u)$, and let $\sigma = (\sigma_1, \dots, \sigma_n)$ be a sequence of actions. For σ to be a valid elimination sequence, there has to exist, for each action σ_i , an action δ_i of the same player and an action γ_i of the other player, both of which have not yet been eliminated, such that δ_i dominates σ_i at γ_i . Let $\delta = (\delta_1, \dots, \delta_n)$ and $\gamma = (\gamma_1, \dots, \gamma_n)$ be sequences of actions of Γ . We say that σ is valid for Γ with respect to δ and γ if, for each i with $1 \leq i \leq n$, action δ_i dominates σ_i at γ_i in $\Gamma(\sigma_1, \dots, \sigma_{i-1})$. We call an action σ_i an *obstacle* in σ with respect to δ and γ in Γ if δ_i does not dominate σ_i at γ_i in $\Gamma(\sigma_1, \dots, \sigma_{i-1})$. Obviously, there are no obstacles in σ with respect to δ and γ if and only if σ is valid with respect to δ and γ . An elimination sequence $\sigma = (\sigma_1, \dots, \sigma_n)$ will be called *weakly valid* with respect to an action sequence $\delta = (\delta_1, \dots, \delta_n)$ if, for all i with $1 \leq i \leq n$, it is the case that $\delta_i \in A \setminus \{\sigma_1, \dots, \sigma_{i-1}\}$ and no action in $A \setminus \{\sigma_1, \dots, \sigma_{i-1}\}$ blocks the elimination of σ_i by δ_i in $\Gamma(\sigma_1, \dots, \sigma_{i-1})$.

We will show that for any constant-sum game Γ every elimination sequence $(\sigma_1, \dots, \sigma_n)$ that is weakly valid with respect to $(\delta_1, \dots, \delta_n)$ can be transformed into a valid elimination sequence, provided that the last action is not an obstacle, i.e., that there is an action actually backing the elimination of σ_n by δ_n in $\Gamma(\sigma_1, \dots, \sigma_{n-1})$. As a first step, the following lemma specifies a sufficient condition for the removal of an action from a weakly valid elimination sequence such that the sequence remains weakly valid and no new obstacles are created. Intuitively, this condition requires that if not eliminated, the action in question does not block any eliminations appearing later in the sequence.

Lemma 3 *Let $\Gamma = (A_1, A_2, u)$ be a two-player game, $\sigma = (\sigma_1, \dots, \sigma_n)$ and $\delta = (\delta_1, \dots, \delta_n)$ two action sequences such that σ is weakly valid for Γ with respect to δ . Let i be an index with $1 \leq i \leq n$ such that σ_i does not block the elimination of σ_j in $\Gamma(\sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_{j-1})$ for any j with $i < j \leq n$. Then,*

$(\sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_n)$ is weakly valid with respect to $(\delta_1, \dots, \delta_{i-1}, \delta_{i+1}, \dots, \delta_n)$. Moreover, the following holds for every index k with $1 \leq k \leq n$ and $k \neq i$, and for every action sequence $\gamma = (\gamma_1, \dots, \gamma_n)$: if σ_k is not an obstacle in σ with respect to δ and γ , then σ_k is not an obstacle in $(\sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_n)$ with respect to $(\delta_1, \dots, \delta_{i-1}, \delta_{i+1}, \dots, \delta_n)$ and $(\gamma_1, \dots, \gamma_{i-1}, \gamma_{i+1}, \dots, \gamma_n)$ either.

Proof Consider an arbitrary index m with $1 \leq m \leq n$ and $m \neq i$. First consider the case when $m < i$. Then, since σ is weakly valid with respect to δ , it follows immediately that no action in $A \setminus \{\sigma_1, \dots, \sigma_{m-1}\}$ blocks the elimination of σ_m by δ_m . Now assume that $m > i$. Then, $\delta_m \in A \setminus \{\sigma_1, \dots, \sigma_{m-1}\}$, and thus $\delta_m \in A \setminus \{\sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_{m-1}\}$. Moreover, since $m > i$, $\delta_m \neq \sigma_i$. It follows that no action in $A \setminus \{\sigma_1, \dots, \sigma_{m-1}\}$ blocks the elimination of σ_m by δ_m in $\Gamma(\sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_{m-1})$. By assumption, σ_i does not block this elimination either, and we may conclude that $(\sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_n)$ is weakly valid with respect to $(\delta_1, \dots, \delta_{i-1}, \delta_{i+1}, \dots, \delta_n)$.

For the second part of the claim, assume that σ_k is not an obstacle in σ with respect to δ and γ , i.e., δ_k dominates σ_k at γ_k in $\Gamma(\sigma_1, \dots, \sigma_{k-1})$. Observe that $\gamma_k \in A \setminus \{\sigma_1, \dots, \sigma_{k-1}\}$. The case when $k < i$ is trivial, so assume that $k > i$. We have already seen that $(\sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_n)$ is weakly valid with respect to $(\delta_1, \dots, \delta_{i-1}, \delta_{i+1}, \dots, \delta_n)$. Moreover, $A \setminus \{\sigma_1, \dots, \sigma_{k-1}\} \subseteq A \setminus \{\sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_{k-1}\}$. With action γ_k still available, δ_k dominates σ_k at γ_k in $\Gamma(\sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_{k-1})$. Thus, σ_k is not an obstacle. \square

A corollary of Lemma 3 is that a valid elimination sequence remains valid after the removal of an action that blocks no other elimination if it remains in the game. Moreover, if an obstacle of an elimination sequence is moved to a position where it blocks no additional eliminations but where it can itself be eliminated, the number of obstacles in the sequence strictly decreases. As we will see next, this can be used to transform a weakly valid elimination sequence into a valid one, given that the last element of the former is not an obstacle.

Lemma 4 *Let $\Gamma = (A_1, A_2, u)$ be a constant-sum game. Let a, b , and c be distinct actions in $A_1 \cup A_2$, and $\sigma = (\sigma_1, \dots, \sigma_n)$ and $\delta = (\delta_1, \dots, \delta_n)$ action sequences with $\sigma_n = a$ and $\delta_n = b$. If σ is weakly valid with respect to δ in Γ and b dominates a at c in $\Gamma(\sigma_1, \dots, \sigma_{n-1})$, then a is eliminable by b at c in Γ .*

Proof Assume that σ is weakly valid with respect to δ and that b dominates a at c in $\Gamma(\sigma_1, \dots, \sigma_{n-1})$. Let $\gamma = (\gamma_1, \dots, \gamma_n)$ be an arbitrary action sequence with $\gamma_n = c$, and assume for contradiction that a is not eliminable by b at c in Γ . Note that we may assume without loss of generality that σ, δ , and γ minimize the number of obstacles among all triples of action sequences with the above properties. We will derive a contradiction by showing that there exists a triple with strictly fewer obstacles.

Clearly, σ cannot be valid for Γ with respect to δ and γ , so there exists a smallest index i with $1 \leq i \leq n$ such that σ_i is an obstacle in σ with respect to δ and γ . By assumption, σ_n is not an obstacle with respect to δ and γ , and thus $i \neq n$. We distinguish two cases.

First assume that there is no index j with $i < j \leq n$ such that σ_i blocks the elimination of σ_j by δ_j in $\Gamma(\sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_{j-1})$. Then, by Lemma 3, $(\sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_n)$ is weakly valid with respect to $(\delta_1, \dots, \delta_{i-1}, \delta_{i+1}, \dots, \delta_n)$ and contains fewer obstacles with respect to $(\delta_1, \dots, \delta_{i-1}, \delta_{i+1}, \dots, \delta_n)$ and

$(\gamma_1, \dots, \gamma_{i-1}, \gamma_{i+1}, \dots, \gamma_n)$ than σ does with respect to δ and γ . Moreover, since $i \neq n$, a is still dominated by b at c in $\Gamma(\sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_{n-1})$, a contradiction.

For the remainder of the proof we will thus assume that σ_i blocks the elimination of σ_j by δ_j in $\Gamma(\sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_{j-1})$ for some index j with $i < j \leq n$. Without loss of generality we may also assume that j is the smallest such index, and that $\sigma_i \in A_1$. Accordingly, $\delta_j, \sigma_j \in A_2$ and $u_2(\sigma_i, \delta_j) < u_2(\sigma_i, \sigma_j)$. It also holds that $\sigma_i, \sigma_j \in A \setminus \{\sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_{j-1}\}$, otherwise σ_i could not block the elimination of σ_j by δ_j in $\Gamma(\sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_{j-1})$. As σ is weakly valid with respect to δ , it follows that γ_i does not back the elimination of σ_i by δ_i in $\Gamma(\sigma_1, \dots, \sigma_{i-1})$. We will see, however, that there exists an index k with $i \leq k < j$ such that δ_k dominates σ_i at σ_j in $\Gamma(\sigma_1, \dots, \sigma_k)$, and that by delaying the elimination of σ_i until σ_k has been removed, σ_i ceases to be an obstacle while no additional ones are being created.

Define B as the smallest subset of $A \setminus \{\sigma_1, \dots, \sigma_{i-1}\}$ such that (i) $\sigma_i \in B$, and (ii) $\delta_k \in B$ whenever $\sigma_k \in B$ and δ_k blocks the elimination of σ_j by δ_j in $\Gamma(\sigma_1, \dots, \sigma_{k-1})$. Obviously, B is non-empty and finite. We may also assume that $B = \{\sigma_{i_1}, \dots, \sigma_{i_m}\}$, where $\sigma_{i_1} = \sigma_i$ and $\sigma_{i_k} = \delta_{i_{k-1}}$, for all k with $1 \leq k \leq m$. Further observe that by weak validity of σ with respect to δ , all actions in B must be eliminated before σ_j is, i.e., $i_m < j$.

Now consider the sequences

$$\begin{aligned} (\sigma'_1, \dots, \sigma'_n) &= (\sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_{i_m}, \sigma_i, \sigma_{i_m+1}, \dots, \sigma_n), \\ (\delta'_1, \dots, \delta'_n) &= (\delta_1, \dots, \delta_{i-1}, \delta_{i+1}, \dots, \delta_{i_m}, \delta_{i_m}, \delta_{i_m+1}, \dots, \delta_n), \text{ and} \\ (\gamma'_1, \dots, \gamma'_n) &= (\gamma_1, \dots, \gamma_{i-1}, \gamma_{i+1}, \dots, \gamma_{i_m}, \sigma_j, \gamma_{i_m+1}, \dots, \gamma_n). \end{aligned}$$

We will show that $(\sigma'_1, \dots, \sigma'_n)$ is weakly valid with respect to $(\delta'_1, \dots, \delta'_n)$ and, moreover, contains fewer obstacles in Γ with respect to $(\delta'_1, \dots, \delta'_n)$ and $(\gamma'_1, \dots, \gamma'_n)$ than σ does with respect to δ and γ . This yields a contradiction, because b also dominates a at c in $\Gamma(\sigma'_1, \dots, \sigma'_{n-1})$. To appreciate the latter, simply observe that the games $\Gamma(\sigma'_1, \dots, \sigma'_{n-1})$ and $\Gamma(\sigma_1, \dots, \sigma_{n-1})$ are identical and, since $i_m < j \leq n$, $\sigma'_n = \sigma_n = a$, $\delta'_n = \delta_n = b$, and $\gamma'_n = \gamma_n = c$.

By Lemma 3 and the assumptions about σ_i , $(\sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_{i_m})$ is a weakly valid elimination sequence with respect to $(\delta_1, \dots, \delta_{i-1}, \delta_{i+1}, \dots, \delta_{i_m})$ in Γ . Moreover, for every index k with $i_m < k \leq n$, $\Gamma(\sigma_1, \dots, \sigma_{k-1})$ and $\Gamma(\sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_{i_m}, \sigma_i, \sigma_{i_m+1}, \dots, \sigma_{k-1})$ are the same game, and in this game no elimination of σ_k by δ_k is blocked. To show that $(\sigma'_1, \dots, \sigma'_n)$ is weakly valid with respect to $(\delta'_1, \dots, \delta'_n)$ and contains fewer obstacles in Γ with respect to $(\delta'_1, \dots, \delta'_n)$ and $(\gamma'_1, \dots, \gamma'_n)$ than σ with respect to δ and γ , it thus suffices to show that σ_i is dominated by δ_{i_m} at σ_j in $\Gamma(\sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_{i_m})$.

Since Γ is a constant-sum game, $u_2(\sigma_i, \delta_j) < u_2(\sigma_i, \sigma_j)$ implies that $u_1(\sigma_i, \delta_j) > u_1(\sigma_i, \sigma_j)$. Furthermore, by definition of B , $u_1(\delta_{i_m}, \delta_j) \geq u_1(\sigma_{i_m}, \delta_j)$ and $u_1(\sigma_{i_{k+1}}, \delta_j) \geq u_1(\sigma_{i_k}, \delta_j)$ for every k with $1 \leq k < m$. Since i_m is the largest index for which $\sigma_{i_m} \in B$, it follows that δ_{i_m} does not block the elimination of σ_j by δ_j . Now recall that σ is weakly valid with respect to δ . Thus, $\delta_{i_m} \in A \setminus \{\sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_{i_m}\}$ and $\delta_j \in A \setminus \{\sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_j\}$. This implies $u_2(\delta_{i_m}, \delta_j) \geq u_2(\delta_{i_m}, \sigma_j)$ and, since Γ is a constant-sum game, $u_1(\delta_{i_m}, \delta_j) \leq u_1(\delta_{i_m}, \sigma_j)$. The resulting situation is depicted in Figure 6, from which it can easily be read off that

$$u_1(\sigma_i, \sigma_j) < u_1(\sigma_i, \delta_j) \leq u_1(\delta_i, \delta_j) \leq \dots \leq u_1(\delta_{i_m}, \delta_j) \leq u_1(\delta_{i_m}, \sigma_j).$$




Fig. 6 Diagram illustrating the proof of Lemma 5

In particular, $u_1(\delta_{i_m}, \sigma_j) > u_1(\sigma_i, \sigma_j)$, i.e., σ_j backs the elimination of σ_i by δ_{i_m} in $\Gamma(\sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_{i_m})$. Since σ is weakly valid with respect to δ , none of the actions in $A \setminus \{\sigma_1, \dots, \sigma_{i_m}\}$ block the elimination σ_i by δ_i in $\Gamma(\sigma_1, \dots, \sigma_{i-1})$. By transitivity of the dominance relation, the same is true for the elimination of σ_i by δ_{i_k} in $\Gamma(\sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_{i_k})$ for any k with $1 \leq k \leq i_m$. It follows that σ_i is dominated by δ_{i_m} at σ_j in $\Gamma(\sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_{i_m})$, which completes the proof. \square

We have seen in the beginning of this section that the elimination of an action can turn a valid elimination sequence into one that is only weakly valid. Using Lemma 4, we will now show that the *existence* of an elimination sequence ending with a particular action a is not affected by such an earlier elimination, given that the eliminated action is not directly involved in the elimination of a .

Lemma 5 *Let $\Gamma = (A_1, A_2, u)$ be a constant-sum game. Let a , b , and c be distinct actions in $A_1 \cup A_2$, and σ a valid elimination sequence for Γ not containing a , b , or c . Then, if a is eliminable by b at c in Γ , a is still eliminable by b at c in $\Gamma(\sigma)$.*

Proof Assume that a is eliminable by b at c in Γ . Then there exist action sequences $\sigma' = (\sigma'_1, \dots, \sigma'_n)$, $\delta' = (\delta'_1, \dots, \delta'_n)$, and $\gamma' = (\gamma'_1, \dots, \gamma'_n)$ with $\sigma'_n = a$, $\delta'_n = b$, and $\gamma'_n = c$ such that σ' is valid with respect to δ' and γ' . Let $\sigma = (\sigma_1, \dots, \sigma_m)$, and let $\delta = (\delta_1, \dots, \delta_m)$ and $\gamma = (\gamma_1, \dots, \gamma_m)$ be action sequences such that σ is valid with respect to δ and γ . By transitivity of the dominance relation, we may further assume for each i with $1 \leq i \leq m$ that $\sigma_i \notin \{\delta'_1, \dots, \delta'_n\}$.³ Since σ' is valid with respect to δ' and γ' and σ is valid with respect to δ and γ , it follows that $(\sigma_1, \dots, \sigma_m, \sigma'_1, \dots, \sigma'_n)$ is weakly

³ To appreciate this, suppose that $\sigma_i = \delta'_j$ for some j with $1 \leq j \leq n$, i.e., σ_i dominates σ'_j in $\Gamma(\sigma'_1, \dots, \sigma'_{j-1})$. Define for each k with $1 \leq k \leq n$ action δ''_k as follows:

$$\delta''_k = \begin{cases} \delta_i & \text{if } k = 1, \\ \delta'_{k-1} & \text{if } \delta''_{k-1} = \sigma_{k-1}, \\ \delta''_{k-1} & \text{otherwise.} \end{cases}$$

Now observe that generally $\delta''_k \in A \setminus \{\sigma'_1, \dots, \sigma'_{k-1}\}$. Moreover, by transitivity of the dominance relation, σ'_j is also dominated by δ''_j in $\Gamma(\sigma'_1, \dots, \sigma'_{j-1})$ and can act as a proxy for δ_i .

valid with respect to $(\delta_1, \dots, \delta_m, \delta'_1, \dots, \delta'_n)$. Moreover, by the assumption that $a, b, c \notin \{\sigma_1, \dots, \sigma_m\}$, action a is still dominated by b at c in $\Gamma(\sigma_1, \dots, \sigma_m, \sigma'_1, \dots, \sigma'_{n-1})$. Lemma 4 now gives us the desired result. \square

Intuitively, Lemma 5 says the following: to eliminate a particular action a by b backed by c , one can eliminate dominated actions more or less in an arbitrary way; one just has to be careful not to eliminate actions b and c . On the basis of this observation, we obtain the main result of this section.

Theorem 2 *The problem of deciding whether a given action of a constant-sum game is eliminable can be solved in polynomial time.*

Proof Let a be the action to be eliminated, and assume without loss of generality that $a \in A_1$. Consider the algorithm that performs the following steps:

- (i) Compose a list $(b^1, c^1), \dots, (b^k, c^k)$ of all pairs $(b^i, c^i) \in A_1 \times A_2$ such that c^i backs the elimination of a by b^i .
- (ii) For each i with $1 \leq i \leq k$, arbitrarily eliminate actions *distinct from b^i and c^i* until no more eliminations are possible. Let $\sigma^i = (\sigma_1^i, \dots, \sigma_{m_i}^i)$ denote the resulting valid elimination sequence.
- (iii) If for some i with $1 \leq i \leq k$, action a is eliminated in σ^i , i.e., $a \in \{\sigma_1^i, \dots, \sigma_{m_i}^i\}$, output “yes,” otherwise “no.”

Obviously, this algorithm runs in polynomial time. If action a is not eliminable, the algorithm cannot find a valid elimination sequence and will always output “no.” If, on the other hand, a is eliminable by b at c for some actions b and c , the algorithm will check this at some point. If it does so, it will make sure not to eliminate actions b and c . Thus, by Lemma 5, a will remain eliminable by b at c as more and more actions are eliminated. Since the overall number of actions is finite, a will at some point become dominated by b at c and can subsequently be eliminated. \square

5 Win-Lose Games

Conitzer and Sandholm [6] have shown that both subgame reachability and eliminability are NP-complete in *win-lose* games, i.e., games which only allow outcomes $(0, 0)$, $(0, 1)$, $(1, 0)$, and $(1, 1)$. As both win-lose and constant-sum games generalize single-winner games, it is interesting to compare these results with those for constant-sum games in the previous section. It turns out that the results of Conitzer and Sandholm even hold for win-lose games with at most one winner, i.e., for games with outcomes $(0, 0)$, $(0, 1)$, and $(1, 0)$. For subgame reachability, this follows from Theorem 1, which shows NP-completeness even for games with outcomes in $\{(0, 1), (1, 0)\}$. For eliminability, we modify the construction used in the proof of Theorem 1 to provide a reduction from *3SAT*.

Theorem 3 *Deciding whether a given action of a two-player game with outcomes in $\{(0, 0), (0, 1), (1, 0)\}$ is eliminable is NP-complete.*

Proof Membership in NP is obvious.

Hardness is shown using a reduction from *3SAT*. By Lemma 1, it suffices to give a reduction for regionalized games. Consider a *3CNF* φ , and recall the regionalized game

	p	$\neg p$	q	$\neg q$	r	$\neg r$	a	b	c	d^*
p	(1, 0)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(1, 0)	(0, 1)	(0, 0)	(0, 0)
$\neg p$	(0, 1)	(1, 0)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(1, 0)	(0, 1)	(0, 0)	(0, 0)
$p\downarrow$	(1, 0)	(1, 0)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 0)	(0, 1)
q	(0, 1)	(0, 1)	(1, 0)	(0, 1)	(0, 1)	(0, 1)	(1, 0)	(0, 1)	(0, 0)	(0, 0)
$\neg q$	(0, 1)	(0, 1)	(0, 1)	(1, 0)	(0, 1)	(0, 1)	(1, 0)	(0, 1)	(0, 0)	(0, 0)
$q\downarrow$	(0, 1)	(0, 1)	(1, 0)	(1, 0)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 0)	(0, 1)
r	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(1, 0)	(0, 1)	(1, 0)	(0, 1)	(0, 0)	(0, 0)
$\neg r$	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(1, 0)	(1, 0)	(0, 1)	(0, 0)	(0, 0)
$r\downarrow$	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(1, 0)	(1, 0)	(0, 1)	(0, 1)	(0, 0)	(0, 1)
$p \vee q \vee \neg r$	(1, 0)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 0)	(0, 0)
$\neg p$	(0, 1)	(1, 0)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 0)	(0, 0)
q	(0, 1)	(0, 1)	(1, 0)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 0)	(0, 0)
$\neg r$	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(1, 0)	(0, 1)	(0, 1)	(0, 0)	(0, 0)
$\neg p \vee q \vee \neg r$	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(1, 0)	(0, 1)	(0, 0)	(0, 1)
$\neg p$	(0, 1)	(1, 0)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 0)	(0, 0)
q	(0, 1)	(0, 1)	(1, 0)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 0)	(0, 0)
r	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(1, 0)	(0, 1)	(0, 1)	(0, 1)	(0, 0)	(0, 0)
$\neg p \vee q \vee r$	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(1, 0)	(0, 1)	(0, 0)	(0, 1)
$\neg p$	(0, 1)	(1, 0)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 0)	(0, 0)
$\neg q$	(0, 1)	(0, 1)	(0, 1)	(1, 0)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 0)	(0, 0)
$\neg r$	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(1, 0)	(0, 1)	(0, 1)	(0, 0)	(0, 0)
$\neg p \vee \neg q \vee \neg r$	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(1, 0)	(0, 1)	(0, 0)	(0, 1)
e	(1, 0)	(1, 0)	(1, 0)	(1, 0)	(1, 0)	(1, 0)	(1, 0)	(0, 1)	(0, 1)	(0, 0)

Fig. 7 Construction used in the proof of Theorem 3, example for the formula $(p \vee q \vee \neg r) \wedge (\neg p \vee q \vee r) \wedge (\neg p \vee \neg q \vee \neg r)$

$(\Gamma_\varphi, X_1, X_2)$ with $\Gamma_\varphi = (A_1, A_2, u)$ defined in the proof of Theorem 1. This game only involved the outcomes (0, 1) and (1, 0). Define a regionalized game $(\Gamma'_\varphi, X'_1, X'_2)$ such that $A'_1 = A_1$, $A'_2 = A_2 \cup \{c, d^*\}$, $X'_1 = X_1$, and $X'_2 = X_2 \cup \{c, d^*\}$. The utility function u' extends u , i.e., $u'(a, b) = u(a, b)$ for all $a \in A_1$ and $b \in A_2$. Payoffs for columns c and d^* are as follows:⁴

	c	d^*		c	d^*		c	d^*	
p	(0, 0)	(0, 0)	(λ_i^1, i)	(0, 0)	(0, 0)		e	(0, 1)	(0, 0)
$\neg p$	(0, 0)	(0, 0)	(λ_i^2, i)	(0, 0)	(0, 0)				
$p\downarrow$	(0, 0)	(0, 1)	(λ_i^3, i)	(0, 0)	(0, 0)				
			C_i	(0, 0)	(0, 1)				

An example of the resulting game is given in Figure 7.

Observe that for all actions $x \in A'_1$, $u_1(x, c) = u_1(x, d^*) = 0$. The additional actions c and d^* thus do not back or block any eliminations. Furthermore, c and d^*

⁴ By setting $u'_1(x, c) = u'_1(x, d^*) = 1$ instead, one obtains a construction proving the intractability of the eliminability problem for games with outcomes in (0, 1), (1, 0) and (1, 1).

	a_2^1	\dots	a_2^m	c	d^*	f	g^*	x_2^1	\dots	y_2^4	z_2^1	z_2^2	z_2^3
a_1^1	\cdot	\dots	\cdot	\cdot	\cdot	\cdot	\cdot	\cdot	\dots	\cdot	(0, 1)	(0, 1)	(0, 1)
\vdots	\vdots	\ddots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\ddots	\vdots	\vdots	\vdots	\vdots
a_1^n	\cdot	\dots	\cdot	\cdot	\cdot	\cdot	\cdot	\cdot	\dots	\cdot	(0, 1)	(0, 1)	(0, 1)
e	\cdot	\dots	\cdot	(0, 1)	(0, 0)	(0, 1)	(0, 0)	\cdot	\dots	\cdot	(0, 1)	(0, 1)	(0, 1)
x_1^1	\cdot	\dots	\cdot	\cdot	\cdot	\cdot	\cdot	\cdot	\dots	\cdot	(0, 1)	(0, 1)	(0, 1)
\vdots	\vdots	\ddots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\ddots	\vdots	\vdots	\vdots	\vdots
y_1^4	\cdot	\dots	\cdot	\cdot	\cdot	\cdot	\cdot	\cdot	\dots	\cdot	(0, 1)	(0, 1)	(0, 1)
z_1^1	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 0)	(0, 1)	(1, 0)	(1, 0)
z_1^2	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(1, 0)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(1, 0)
z_1^3	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(1, 0)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 0)	(1, 0)	(0, 1)
z_1^4	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(0, 1)	(1, 0)	(0, 1)	(0, 1)	(0, 1)	(0, 0)	(0, 1)	(1, 0)

Fig. 8 Construction used in the proof of Theorem 4

constitute a separate region and can therefore neither eliminate nor be eliminated by any of the actions in A_2 . Finally, column d^* is dominated by c at e if and only if action $p\downarrow$ for each variable p and action C_i for each clause C_i have been eliminated. By virtue of an argument analogous to the one used in the proof of Theorem 1, we find that action d^* is eliminable if and only if φ is satisfiable. This completes the proof. \square

Conitzer and Sandholm [6] use a reduction from eliminability to solvability to show intractability of the latter in win-lose games. Their construction, however, hinges on the presence of the outcome (1, 1). For the more restricted class of games without (1, 1) as an outcome we instead reduce directly from $3SAT$ and exploit the internal structure of the construction used in the proof of Theorem 3.

Theorem 4 *Deciding whether a two-player game with outcomes in $\{(0, 0), (0, 1), (1, 0)\}$ is solvable is NP-complete.*

Proof Membership in NP is straightforward.

Hardness is shown using a reduction from $3SAT$. Consider a $3CNF$ formula φ , and let $(\Gamma'_\varphi, X'_1, X'_2)$ with $\Gamma'_\varphi = (A'_1, A'_2, u')$ be the regionalized game with outcomes in $\{(0, 0), (0, 1), (1, 0)\}$ defined in Theorem 3, with additional copies f and g^* of the actions c and d^* such that $\{f, g^*\}$ constitutes a separate region. Thus, A'_1 and X'_1 are as before, while for the column player we have

$$A'_2 = \{p, \neg p : p \text{ a variable in } \varphi\} \cup \{a, b, c, d^*, f, g^*\},$$

$$X'_2 = \{\{p, \neg p : p \text{ a variable in } \varphi\} \cup \{a, b\}, \{c, d^*\}, \{f, g^*\}\},$$

and $u'(x, f) = u'(x, c)$ and $u'(x, g^*) = u'(x, d^*)$ for each $x \in A'_1$. By the same reasoning as in the proof of Theorem 3, both d^* and g^* are eliminable if and only if φ is satisfiable.

Now consider the game $\Gamma''_\varphi = (A''_1, A''_2, u'')$ without regions corresponding to $(\Gamma'_\varphi, X'_1, X'_2)$ as defined in the proof of Lemma 1, and define $\Gamma'''_\varphi = (A'''_1, A'''_2, u''')$

	$\{(0, 1), (1, 0)\}$	Constant-Sum	$\{(0, 0), (0, 1), (1, 0)\}$	Win-Lose	General
Subgame reachability	NP-c ^a	NP-c ^a	NP-c ^a	NP-c ^f	NP-c ^g
Eliminability	in P ^b	in P ^b	NP-c ^c	NP-c ^f	NP-c ^g
Solvability	in P ^d		NP-c ^e	NP-c ^f	NP-c ^g

^a Theorem 1

^b Theorem 2

^c Theorem 3

^d Brandt et al. [5]

^e Theorem 4

^f Conitzer and Sandholm [6]

^g Gilboa et al. [8]

Table 1 Computational complexity of IWD in two-player games

with $A_1''' = A_1'' \cup \{z_1^1, z_1^2, z_1^3, z_1^4\}$ and $A_2''' = A_2'' \cup \{z_2^1, z_2^2, z_2^3\}$. Let $u'''(x, y) = u''(x, y)$ for all $(x, y) \in A_1''' \times A_2'''$, the payoffs for the remaining action profiles in $A_1''' \times A_2'''$ are shown in Figure 8.

We make the following observations about the game Γ_φ''' .

- (i) As long as columns d^* and g^* are not eliminated, the actions in $\{z_1^1, z_1^2, z_1^3, z_1^4\} \cup \{z_2^1, z_2^2, z_2^3\}$ do not dominate and are not dominated by any action in the game.
- (ii) Actions z_1^2 and z_1^3 back the elimination of d^* by c , and z_1^4 backs the elimination of g^* by f . However, since action e also backs the same eliminations, and since action e itself is not eliminable in Γ_φ'' , this does not make any additional eliminations possible as long as d^* and g^* have not been eliminated.

We now claim that Γ_φ''' can be solved, with (z_1^1, z_2^1) as the remaining action profile, if and only if φ is satisfiable.

For the direction from left to right, assume that φ is unsatisfiable. Then, using the same arguments as in the proof of Theorem 3 actions d^* and g^* cannot be eliminated. Hence, by (i), the game Γ_φ''' is not solvable.

For the direction from right to left, assume that φ is satisfiable. Again by the same arguments as in the proof of Theorem 3, columns d^* and g^* can be eliminated. Then, rows z_1^2 , z_1^3 and z_1^4 can be eliminated by row z_1^1 , followed by the elimination of columns a_2^1 through y_2^4 and z_2^3 by column z_2^1 . Finally, row z_1^1 can eliminate all other remaining rows, and the elimination of column z_2^2 solves the game. \square

6 Conclusion

We have investigated the computational complexity of iterated weak dominance in two-player constant-sum games. In particular, we have shown that eliminability of an action can be decided in polynomial time, whereas deciding reachability of a given subgame is NP-complete. We have further shown the NP-completeness of typical problems associated with iterated dominance in win-lose games *with at most one winner*. Table 1 provides an overview of our results, and related results obtained earlier.

In win-lose games an action is dominated by a mixed strategy if and only if it is dominated by a pure strategy [6]. All of our results results apart from Theorem 2 thus immediately extend to dominance by mixed strategies.

Marx and Swinkels [12] identify a condition under which all subgames that are reachable via iterated weak dominance and cannot be reduced further are equivalent in terms of the payoff profiles that can be obtained, i.e., differ only by the addition or removal of identical actions and the renaming of actions. Since the condition is satisfied by constant-sum games, we can decide in polynomial time which payoff profiles of a constant sum game can still be obtained after the iterated removal of weakly dominated actions, by simply eliminating dominated actions arbitrarily. This, however, does not imply any of our results, because it does not discriminate between actions that yield identical payoffs for *some* reachable subgame. In fact, Theorem 1 tells us that reachability of a given subgame is NP-hard to decide even in constant-sum games. The conceptual difference between our work and that of Marx and Swinkels is thus tightly linked to the question whether one is interested in *action profiles* or *payoff profiles* as “solutions” of a game or, more generally, whether one champions a prescriptive or a descriptive interpretation of game theory. It may be argued that the computational gap between both concepts is of particular interest in this context.

Acknowledgements This material is based on work supported by the Deutsche Forschungsgemeinschaft under grants BR 2312/3-2, BR 2312/3-3, BR 2312/6-1, BR 2312/7-1, and FI 1664/1-1. We thank the anonymous reviewers for helpful comments.

References

1. K. R. Apt. Uniform proofs of order independence for various strategy elimination procedures. *Contributions to Theoretical Economics*, 4(1), 2004.
2. B. Bernheim. Rationalizable strategic behavior. *Econometrica*, 52(4):1007–1028, 1984.
3. A. Brandenburger, A. Friedenberg, and H. J. Keisler. Admissibility in games. *Econometrica*, 76(2):307–352, 2008.
4. F. Brandt, M. Brill, F. Fischer, P. Harrenstein, and J. Hoffmann. Computing Shapley’s saddles. *ACM SIGecom Exchanges*, 8(2), 2009.
5. F. Brandt, F. Fischer, P. Harrenstein, and Y. Shoham. Ranking games. *Artificial Intelligence*, 173(2):221–239, 2009.
6. V. Conitzer and T. Sandholm. Complexity of (iterated) dominance. In *Proceedings of the 6th ACM Conference on Electronic Commerce (ACM-EC)*, pages 88–97. ACM Press, 2005.
7. C. Ewerhart. Iterated weak dominance in strictly competitive games of perfect information. *Journal of Economic Theory*, 107:474–482, 2002.
8. I. Gilboa, E. Kalai, and E. Zemel. The complexity of eliminating dominated strategies. *Mathematics of Operations Research*, 18(3):553–565, 1993.
9. E. Kohlberg and J.-F. Mertens. On the strategic stability of equilibria. *Econometrica*, 54:1003–1037, 1986.
10. D. Koller and N. Megiddo. The complexity of two-person zero-sum games in extensive form. *Games and Economic Behavior*, 4(4):528–552, 1992.
11. R. D. Luce and H. Raiffa. *Games and Decisions: Introduction and Critical Survey*. John Wiley & Sons Inc., 1957.
12. L. M. Marx and J. M. Swinkels. Order independence for iterated weak dominance. *Games and Economic Behavior*, 18:219–245, 1997.

-
13. R. B. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, 1991.
 14. M. Osborne. *An Introduction to Game Theory*. Oxford University Press, 2004.
 15. C. H. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.
 16. D. Pearce. Rationalizable strategic behavior and the problem of perfection. *Econometrica*, 52(4):1029–1050, 1984.
 17. L. Samuelson. Dominated strategies and common knowledge. *Games and Economic Behavior*, 4:284–313, 1992.
 18. M. Shimoji. On the equivalence of weak dominance and sequential best response. *Games and Economic Behavior*, 48:385–402, 2004.
 19. Y. Shoham and K. Leyton-Brown. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press, 2009.