

The Computational Complexity of Weak Saddles

Felix Brandt, Markus Brill, Felix Fischer, and Jan Hoffmann

Institut für Informatik, Ludwig-Maximilians-Universität München
80538 München, Germany
{brandtf,brill,fischerf,hoffmann}@tcs.ifi.lmu.de

Abstract. We continue the recently initiated study of the computational aspects of weak saddles, an ordinal set-valued solution concept proposed by Shapley. Brandt et al. gave a polynomial-time algorithm for computing weak saddles in a subclass of matrix games, and showed that certain problems associated with weak saddles of bimatrix games are NP-complete. The important question of whether weak saddles can be *found* efficiently was left open. We answer this question in the negative by showing that finding weak saddles of bimatrix games is NP-hard, under polynomial-time Turing reductions. We moreover prove that recognizing weak saddles is coNP-complete, and that deciding whether a given action is contained in some weak saddle is hard for parallel access to NP and thus not even in NP unless the polynomial hierarchy collapses. Our hardness results are finally shown to carry over to a natural weakening of weak saddles.

1 Introduction

Saddle points, i.e., combinations of actions such that no player can gain by deviating, are one of the earliest solutions suggested in game theory (see, e.g., [23]). In two-player zero-sum games (henceforth *matrix games*), every saddle point happens to coincide with an optimal outcome both players can guarantee in the worst case and thus enjoys a very strong normative foundation. Unfortunately, however, not every matrix game possesses a saddle point. In order to remedy this situation, von Neumann [22] considered *mixed*, i.e., randomized, strategies and proved that every matrix game contains a mixed saddle point (or equilibrium) that moreover maintains the appealing normative properties of saddle points. The existence result was later generalized to arbitrary general-sum games by Nash [15], at the expense of its normative foundation. Since then, Nash equilibrium has commonly been criticized for its need for randomization, which may be deemed unsuitable, impractical, or even infeasible (see, e.g., [13, 14, 5]).

In two papers from 1953, Lloyd Shapley showed that existence of saddle points (and even uniqueness in the case of matrix games) can also be guaranteed by moving to *minimal sets* of actions rather than randomizations over them [19, 20].¹ Shapley defines a *generalized saddle point (GSP)* to be a tuple of subsets of actions of each player, such that every action not contained in the GSP is

¹ The main results of the 1953 reports later reappeared in revised form [21].

dominated by some action in the GSP, given that the remaining players choose actions from the GSP. A *saddle* is an inclusion-minimal GSP, i.e., a GSP that contains no other GSP. Depending on the underlying notion of dominance, one can define strict, weak, and very weak saddles. Shapley [21] showed that every matrix game admits a *unique* strict saddle. Duggan and Le Breton [9] proved that the same is true for the weak saddle in a certain subclass of symmetric matrix games that we refer to as *confrontation games*. While Shapley was the first to conceive weak GSPs, he was not the only one. Apparently unaware of Shapley’s work, Samuelson [18] uses the very related concept of a *consistent pair* to point out epistemic inconsistencies in the concept of iterated weak dominance. Also, *weakly admissible sets* as defined by McKelvey and Ordeshook [14] in the context of spatial voting games are identical to weak GSPs. Other common *set-valued* concepts in game theory include *rationalizability* [3, 17] and *CURB sets* [1].

In this paper we continue the recently initiated study of the computational aspects of Shapley’s saddles. Brandt et al. [5] gave polynomial-time algorithms for computing strict saddles in general games and weak saddles in confrontation games. Although it was shown that certain problems associated with weak saddles in bimatrix games are NP-complete, the question of whether weak saddles can be found efficiently was left open. We answer this question in the negative by showing that finding weak saddles is NP-hard. Moreover, we prove that recognizing weak saddles is coNP-complete, and that deciding whether an action is contained in a weak saddle of a bimatrix game is complete for parallel access to NP and thus not even in NP unless the polynomial hierarchy collapses. We finally demonstrate that our hardness results carry over to very weak saddles.

2 Related Work

In recent years, the computational complexity of game-theoretic solution concepts has come under increasing scrutiny. One of the most prominent results in this stream of research is that the problem of finding Nash equilibria in bimatrix games is PPAD-complete [6, 8], and thus unlikely to admit a polynomial-time algorithm. PPAD is a subclass of FNP, and it is obvious that Nash equilibria can be recognized in polynomial time. Interestingly, our results imply that this is not the case for weak saddles unless $P=NP$.

Weak saddles rely on the elementary concept of weak dominance, whose computational aspects have been studied extensively in the form of *iterated* weak dominance [11, 7]. In contrast to iterated dominance, saddles are based on a notion of stability reminiscent of Nash equilibrium and its various refinements. Weak saddles are also related to minimal covering sets, a concept that has been proposed independently in social choice theory [10, 9] and whose computational complexity has recently been analyzed [4, 2].

Brandt et al. [5] constructed a class of games that established a strong relationship between weak saddles and inclusion-maximal cliques in undirected graphs. Based on this construction and a reduction from the NP-complete problem CLIQUE, they showed that deciding whether there exists a weak saddle

with a *certain number of actions* is NP-hard. This construction, however, did not permit any statements about the more important problems of *finding* a weak saddle, *recognizing* a weak saddle, or *deciding whether a certain action is contained* in some weak saddle.

3 Preliminaries

An accepted way to model situations of strategic interactions is by means of a *normal-form game* (see, e.g., [13]).

Definition 1 (Normal-Form Game). A (*finite*) game in normal-form is a tuple $\Gamma = (N, (A_i)_{i \in N}, (p_i)_{i \in N})$ where $N = \{1, 2, \dots, n\}$ is a set of players and for each player $i \in N$, A_i is a nonempty finite set of actions available to player i , and $p_i : (\prod_{i \in N} A_i) \rightarrow \mathbb{R}$ is a function mapping each action profile (i.e., combination of actions) to a real-valued payoff for player i .

A *subgame* of a (normal-form) game $\Gamma = (N, (A_i)_{i \in N}, (p_i)_{i \in N})$ is a game $\Gamma' = (N, (A'_i)_{i \in N}, (p'_i)_{i \in N})$ where, for each $i \in N$, A'_i is a nonempty subset of A_i and $p'_i(a') = p_i(a')$ for all $a' \in A'_1 \times \dots \times A'_n$. Γ' is then called a *subgame* of Γ .

In order to formally define Shapley's weak saddles, we need some additional notation. Let $A_N = (A_1, \dots, A_n)$. For a tuple $S = (S_1, \dots, S_n)$, write $S \subseteq A_N$ and say that S is a subset of A_N if $\emptyset \neq S_i \subseteq A_i$ for all $i \in N$. Further let $S_{-i} = (S_1, \dots, S_{i-1}, S_{i+1}, \dots, S_n)$. For a player $i \in N$ and two actions $a_i, b_i \in A_i$ say that a_i *weakly dominates* b_i with respect to S_{-i} , denoted $a_i >_{S_{-i}} b_i$, if $p_i(a_i, s_{-i}) \geq p_i(b_i, s_{-i})$ for all $s_{-i} \in S_{-i}$, with at least one strict inequality.

Definition 2 (Weak Saddle). Let $\Gamma = (N, (A_i)_{i \in N}, (p_i)_{i \in N})$ be a game and $S = (S_1, \dots, S_n) \subseteq A_N$. Then, S is a weak generalized saddle point (WGSP) of Γ if for each player $i \in N$ the following holds:

$$\text{For every } a_i \in A_i \setminus S_i \text{ there exists } s_i \in S_i \text{ such that } s_i >_{S_{-i}} a_i. \quad (1)$$

A weak saddle is a WGSP that contains no other WGSP.

An example game with two weak saddles is given in Figure 1. The interpretation of this definition is the following: Every player i has a distinguished set S_i of actions such that for every action a_i that is not in the set S_i , there is some action in S_i that weakly dominates a_i , provided that the other players play only actions from their distinguished sets. Condition (1) will be called *external stability* in the following. A WGSP thus is a tuple S that is externally stable for each player. Observe that the tuple A_N of all actions is always a WGSP, thereby guaranteeing existence of a weak saddle in every game. As the game in Figure 1 illustrates, weak saddles do not have to be unique. It is also not very hard to see that weak saddles are invariant under order-preserving transformations of the payoff functions and that every weak saddle contains a (mixed) Nash equilibrium.

	b_1	b_2	b_3
a_1	0	1	0
a_2	1	1	1

Fig. 1. Example game with two weak saddles: $(\{a_1\}, \{b_1, b_2\})$ and $(\{a_1, a_2\}, \{b_2\})$.

In the remainder of the paper we will concentrate on two-player games.² For such games, we can simplify notation and write $\Gamma = (A, B, p)$, where A is the set of actions of player 1, B is the set of actions of player 2, and $p : A \times B \rightarrow \mathbb{R} \times \mathbb{R}$ is the payoff function on the understanding that $p(a, b) = (p_1(a, b), p_2(a, b))$ for all $(a, b) \in A \times B$. A two-player game is often called a *bimatrix* game, as it can conveniently be represented as a $|A| \times |B|$ bimatrix M , i.e., a matrix with rows indexed by A , columns indexed by B and $M(a, b) = p(a, b)$ for every action profile $(a, b) \in A \times B$. We will commonly refer to actions of players 1 and 2 by the rows and columns of this matrix, respectively. When representing bimatrix game graphically, we follow the convention to write player 1’s payoffs in the lower left corner and player 2’s payoff in the upper right corner of the corresponding matrix cell (see Figure 1 for an example).

For an action a and a weak saddle $S = (S_1, S_2)$, we will sometimes slightly abuse notation and write $a \in S$ if $a \in (S_1 \cup S_2)$. In such cases, whether a is a row or a column should be either clear from the context or irrelevant for the argumentation. This partial identification of S and $S_1 \cup S_2$ is also reflected in referring to S as a “set” rather than a “pair” or “tuple.” When reasoning about the structure of the saddles of game, the following definition will be useful

Definition 3. Let $\Gamma = (A, B, p)$ be a game and $x, y \in A \cup B$ two actions. We say that x compels y , denoted $x \rightsquigarrow y$, if every weak saddle containing x also contains y .

Observe that \rightsquigarrow as a relation on $(A \cup B) \times (A \cup B)$ is transitive. We now identify two sufficient conditions for $x \rightsquigarrow y$ to hold.

Fact 1. Let $\Gamma = (A, B, p)$ be a two-player-game, $b \in B$ an action of player 2, and $a \in A$ an action of player 1. Then $b \rightsquigarrow a$ if one of the following two conditions holds:³

- (i) a is the unique action that maximizes $p_1(\cdot, b)$, i.e., $\{a\} = \arg \max_{a' \in A} p_1(a', b)$.
- (ii) a maximizes $p_1(\cdot, b)$ and all actions maximizing $p_1(\cdot, b)$ yield identical payoffs for all opponent actions, i.e., $a \in \arg \max_{a' \in A} p_1(a', b)$ and $p_1(a_1, b') = p_1(a_2, b')$ for all $a_1, a_2 \in \arg \max_{a' \in A} p_1(a', b)$ and all $b' \in B$.

² Naturally, all hardness results carry over to the general n -player case by adding an arbitrary number of “dummy” players that always receive the same payoff.

³ The statement remains true if the roles of the two players are reversed.

Part (i) of the statement above can be generalized in the following way. An action a is in the weak saddle if it is the unique best response to a subset of saddle actions: if $\{b_1, \dots, b_t\} \subset S$ and $a \succ_{\{b_1, \dots, b_t\}} a'$ for all $a' \in A \setminus \{a\}$, then $a \in S$. In this case, we write $\{b_1, \dots, b_t\} \rightsquigarrow a$. Moreover, for two sets of actions X and Y , we write $X \rightsquigarrow Y$ if $X \rightsquigarrow y$ for all $y \in Y$. For example, in the game in Figure 1, $b_1 \rightsquigarrow a_1 \rightsquigarrow b_2$, $\{b_2, b_3\} \rightsquigarrow a_2$ and $\{a_1, a_2\} \rightsquigarrow \{b_1, b_3\}$.

We assume throughout the paper that games are given explicitly, i.e., as tables containing the payoffs for every possible action profile. We will be interested in the following computational problems for a given game Γ :

- FINDWEAKSADDLE: Find a weak saddle of Γ .
- ISWEAKSADDLE: Is a given collection (S_1, \dots, S_n) of subsets of actions for each player a weak saddle of Γ ?
- UNIQUEWEAKSADDLE: Does Γ contain exactly one weak saddle?
- INWEAKSADDLE: Is a given action a contained in a weak saddle of Γ ?
- INALLWEAKSADDLES: Is a given action a contained in *every* weak saddle of Γ ?
- NONTRIVIALWEAKSADDLE: Does Γ contain a weak saddle that does *not* consist of all actions?

We assume the reader to be familiar with the basic notions of complexity theory, such as polynomial-time many-one reductions and Turing reductions, and the related notions of hardness and completeness, and with standard complexity classes such as P, NP, and coNP (see, e.g., [16]). We will further use the complexity classes Σ_2^p and Θ_2^p . $\Sigma_2^p = \text{NP}^{\text{NP}}$ is the second level of the polynomial hierarchy and consists of all problem that can be solved on a non-deterministic Turing machine with access to an NP oracle. $\Theta_2^p = \text{P}_{\parallel}^{\text{NP}}$ consists of all problems that can be solved on a deterministic Turing machine with parallel (non-adaptive) access to an NP oracle.

4 Hardness Results for Weak Saddles

We will now derive various hardness results for weak saddles. We begin by presenting a general construction that transforms a Boolean formula φ into a bimatrix game Γ_φ , such that the existence of certain weak saddles in Γ_φ depends the satisfiability of φ . This construction will be instrumental for each of the hardness proofs given in the sequel.

4.1 A General Construction

Let $\varphi = C_1 \wedge \dots \wedge C_m$ be a Boolean formula in conjunctive normal form (CNF) over a finite set $V = \{v_1, \dots, v_n\}$ of variables. Denote by $L = \bigcup_{v \in V} \{\{v, \bar{v}\} : v \in V\}$ the set of all *literals*, where a literal is either a variable or its negation. Each *clause* C_j is a set of literals. An *assignment* $\alpha : L \rightarrow \{0, 1\}$ is a function mapping each literal to either 1 (“true”) or 0 (“false”). Assignment α is *valid* if $\alpha(v) \neq \alpha(\bar{v})$ for all $v \in V$. For a valid assignment α , denote by $L^\alpha = \{\ell \in L : \alpha(\ell) = 1\}$

	b^*	v_1	\bar{v}_1	v_2	\bar{v}_2	\dots	v_n	\bar{v}_n
a^*	1	0	0	0	0	\dots	0	0
d^*	0	1	1	1	1	\dots	1	1
C_1	1	0	0	0	0	\dots	0	0
C_2	0	0	1	1	0	\dots	1	1
\vdots								
C_m	1	0	0	0	0	\dots	0	0
	0	1	1	1	1	\dots	1	1

Fig. 2. Subgame of Γ_φ for a formula $\varphi = C_1 \wedge \dots \wedge C_m$ with $v_1, \bar{v}_2 \in C_1$ and $\bar{v}_1, v_n \in C_2$.

the set of literals that are set to true under α . We say that α *satisfies* a clause C_j if $C_j \cap L^\alpha \neq \emptyset$. Finally, formula φ is *satisfiable* if there is an assignment that satisfies each of its clauses. We assume without loss of generality that φ does not contain *trivial* clauses, i.e., clauses that contain a literal ℓ as well as its negation $\bar{\ell}$. The game $\Gamma_\varphi = (A, B, p)$ is defined in three steps.

Step 1. Player 1 has actions $\{a^*, d^*\} \cup C$, where $C = \{C_1, \dots, C_m\}$ is the set of clauses of φ . Player 2 has actions $B = \{b^*\} \cup L$, where L is the set of literals.⁴ Payoffs are given by

- $p(a^*, b^*) = (1, 1)$,
- $p(d^*, \ell) = (1, 1)$ for all $\ell \in L$,
- $p(C_j, b^*) = (0, 1)$ for all $j \in [m]$,
- $p(C_j, \ell) = (1, 0)$ if and only if $\ell \notin C_j$,
- $p(a, b) = (0, 0)$ otherwise.

An example of such a game is shown in Figure 2. Observe that (a^*, b^*) is a weak saddle and thus no strict superset can be a weak saddle. Furthermore, row d^* dominates row C_j with respect to a set of columns $\{\ell_1, \dots, \ell_t\} \subseteq L$ if and only if $\ell_i \in C_j$ for some $i \in [t]$.⁵ In particular, $d^* >_{L^\alpha} C_j$ if and only if α satisfies C_j . Another noteworthy property of this game is the fact that no weak saddle contains any of the rows C_j , because $C_j \rightsquigarrow b^* \rightsquigarrow a^*$ for each $j \in [m]$.

The basic idea behind this construction is the following. We want to have an “assignment saddle” $S^\alpha = (S_1, S_2)$ with $d^* \in S_1$ and $S_2 = L^\alpha$ if and only if α is satisfiable. For the direction from left to right, we have to ensure that S^α cannot

⁴ There shall be no confusion by identifying literals with corresponding actions of player 2, which will henceforth be called “literal actions” (or “literal columns”).

⁵ For $n \in \mathbb{N}$, we write $[n] = \{1, 2, \dots, n\}$.

be a weak saddle if α does not satisfy φ or if α is not a valid assignment. This is achieved by means of additional actions, for which the payoffs are defined in such a way that every “wrong” (i.e., unsatisfying or invalid) assignment yields a set containing both a^* and b^* . Obviously, such a set can never be a weak saddle, because it contains the weak saddle (a^*, b^*) as a proper subset. In fact, (a^*, b^*) will be the unique weak saddle in cases where there is no satisfying assignment.

Step 2. We augment the action sets of both players. Player 1 has one additional row ℓ' for each literal $\ell \in L$.⁶ Player 2 has one additional column y_i for each variable $v_i \in V$. Payoffs for profiles involving new actions are defined as follows:

- $p(a^*, y_i) = (1, 0)$ for all $i \in [n]$,
- $p(\ell', \ell) = (2, 1)$ if $\ell' = \ell$,
- $p(\ell', y_i) = (0, 1)$ if $\ell' \in \{v_i, \bar{v}_i\}$,
- $p(a, b) = (0, 0)$ otherwise.

Observe that, by Fact 1 and the following discussion, $\ell \rightsquigarrow \ell'$, $\{\ell', \bar{\ell}'\} \rightsquigarrow y_i$ and $y_i \rightsquigarrow a^* \rightsquigarrow b^*$. This means that no assignment saddle can contain both ℓ as well its negation $\bar{\ell}$.

There only remains one subtlety to be dealt with. In the game defined so far, there are weak saddles containing row d^* , whose existence is independent of the satisfiability of φ , namely $(\{d^*, \ell'\}, \{\ell\})$ for each $\ell \in L$. We destroy these saddles using additional rows.

Step 3. We introduce new rows $r_1, \bar{r}_1, \dots, r_n, \bar{r}_n$, one for each literal, with the property that $r_i \rightsquigarrow b^*$, and that r_i or \bar{r}_i can only be weakly dominated (by v_i and \bar{v}_i , respectively) if at least one literal column other than v_i or \bar{v}_i is in the saddle. For this, we define

- $p(r_i, b^*) = p(\bar{r}_i, b^*) = (0, 1)$ for all $i \in [n]$,
- $p(r_i, v_i) = p(\bar{r}_i, \bar{v}_i) = (2, 0)$,
- $p(r_i, \ell) = p(\bar{r}_i, \ell) = (-1, 0)$ if $\ell \in \{v_{i+1}, \bar{v}_{i+1}\}$,
- $p(a, b) = (0, 0)$ otherwise.

The game Γ_φ now has action sets $A = \{a^*, d^*\} \cup C \cup L \cup \{r_1, \dots, \bar{r}_n\}$ for player 1 and $B = \{b^*\} \cup L \cup \{y_1, \dots, y_n\}$ for player 2. The size of Γ_φ thus is clearly polynomial in the size of φ . A complete example of such a game is given in Figure 4 in the appendix.

For an assignment α , define the assignment saddle S^α as $S^\alpha = (\{d^*\} \cup L^\alpha, L^\alpha)$. It should be clear from the argumentation above that S^α is a weak saddle of Γ_φ if and only if α satisfies φ . To show that membership of a given action in a weak saddle is NP-hard, it suffices to show that there are no other weak saddles containing row d^* . We do so in the following section.

⁶ Action ℓ' of player 1 and action ℓ of player 2 refer to the same literal, but we name them differently to avoid confusion.

4.2 Membership is NP-hard

We show NP-hardness of the membership problem via a reduction from SAT. Given a CNF formula φ , we show that the game Γ_φ defined in Section 4.1 has a weak saddle containing action d^* if and only if φ is satisfiable. A detailed proof of the following theorem is given in the appendix.

Theorem 1. *INWEAKSADDLE is NP-hard.*

4.3 Membership is coNP-hard

We have just seen that it is NP-hard to decide whether there exists a weak saddle containing a given action. In order to show that this problem is also coNP-hard, we first show the following: given a game and an action c , it is possible to augment the game with additional actions such that every weak saddle of the augmented game that contains c contains all actions of this game.

Lemma 1. *Let $\Gamma = (A, B, p)$ be a two-player game, $c \in A \cup B$ and action of Γ . Then there exists a supergame $\Gamma^c = (A', B', p')$ of Γ with the following properties:*

- (i) *If S is a weak saddle of Γ^c containing c , then $S = (A', B')$.*
- (ii) *If S is a weak saddle of Γ that does not contain c , then S is a weak saddle of Γ^c .*
- (iii) *The size of Γ^c is polynomial in the size of Γ .*

The game Γ^c is sketched in Figure 3. Briefly, we add new actions such that $c \rightsquigarrow (A' \setminus A) \rightsquigarrow (B' \setminus \{c\}) \rightsquigarrow A$. A detailed proof of Lemma 1 can be found in the appendix.

Theorem 2. *INWEAKSADDLE is coNP-hard.*

Proof. We give a reduction from UNSAT. For a given CNF formula φ , consider the game $\Gamma_\varphi^{b^*}$ obtained by augmenting the game Γ_φ defined in Section 4.1 in such a way that every weak saddle containing action b^* in fact contains all actions. We show that $\Gamma_\varphi^{b^*}$ has a weak saddle containing b^* if and only if φ is unsatisfiable.

For the direction from left to right, assume that there exists a weak saddle $S = (S_1, S_2)$ with $b^* \in S_2$. By Lemma 1, S is trivial, i.e., equals the set of all actions. Furthermore, S must be the unique weak saddle of $\Gamma_\varphi^{b^*}$, because any other weak saddle would violate minimality of S . In particular, S^α cannot be a saddle for any assignment α , which by the discussion in Section 4.1 means that φ is unsatisfiable.

For the direction from right to left, assume that φ is unsatisfiable. It is not very hard to see that every weak saddle $S = (S_1, S_2)$ contains at least one column not corresponding to a literal, i.e., $S_2 \not\subseteq L$ (otherwise, S would be an assignment saddle). However, since $a^* \rightsquigarrow b^*$ and $b \rightsquigarrow a^*$ for every non-literal column $b \in B \setminus L$, we have that $b^* \in S_2$ for every weak saddle S . \square

The proof of Theorem 2 implies several other hardness results.




Fig. 3. Construction used in the proof of Lemma 1. Payoffs are $(0, 0)$ unless specified otherwise, k is chosen to maximize $p_1(\cdot, c)$. Every weak saddle containing column c then equals the set of all actions.

Corollary 1. *The following holds:*

- ISWEAKSADDLE is coNP-complete.
- INALLWEAKSADDLES is coNP-complete.
- UNIQUEWEAKSADDLE is coNP-hard.

Proof. Recall the definition of the game $\Gamma_\varphi^{b^*}$ used in the proof of Theorem 2. It is easily verified that the following statements are equivalent: formula φ is unsatisfiable, $\Gamma_\varphi^{b^*}$ has a trivial weak saddle, the unique weak saddle of $\Gamma_\varphi^{b^*}$ is the trivial one, and b^* is contained in all weak saddles of $\Gamma_\varphi^{b^*}$.

Membership of INALLWEAKSADDLES in coNP holds because any externally stable set that does not contain the action in question serves as a witness that this actions is *not* contained in every weak saddle. For membership of ISWEAKSADDLE, consider a tuple S of actions that is *not* a weak saddle. Then either S itself is not externally stable, or a proper subset of S is. For both cases there exists a witness of polynomial size. \square

4.4 Finding a Saddle is NP-hard

A particularly interesting consequence of Theorem 2 concerns the existence of a nontrivial weak saddle. As we will see, hardness of deciding the latter can be used to obtain a result about the complexity of the search problem.

Corollary 2. NONTRIVIALWEAKSADDLE is NP-complete.

Proof. For membership in NP, observe that proving the existence of a nontrivial weak saddle is tantamount to finding a proper subset of the set of all actions that

is externally stable. By definition, every such subset is guaranteed to contain a weak saddle. Obviously, external stability can be checked in polynomial time.

Hardness is again straightforward from the proof of Theorem 2, since the game $\Gamma_\varphi^{b^*}$ has a nontrivial weak saddle if and only if formula φ is satisfiable. \square

Corollary 3. FINDWEAKSADDLE is NP-hard under polynomial-time Turing reductions.

Proof. Suppose there exists an algorithm that computes some weak saddle of a game in time polynomial in the size of the game. Such an algorithm could obviously be used to solve the NP-hard problem NONTRIVIALWEAKSADDLE in polynomial time. Just run the algorithm once. If it returns a nontrivial saddle, the answer is “yes.” Otherwise the set of all actions must be the unique weak saddle of the game, and the answer is “no.” \square

4.5 Membership is Θ_2^p -hard

Now that we have established that INWEAKSADDLE is both NP-hard and coNP-hard, we will raise the lower bound to Θ_2^p . Wagner provided a sufficient condition for Θ_2^p -hardness that turned out to be very useful (see, e.g., [12]).

Lemma 2 (Wagner [24]). Let S be an NP-complete problem, and let T be any set. Further let f be a polynomial-time computable function such that the following holds for all $k \geq 1$ and all strings x_1, x_2, \dots, x_{2k} satisfying $x_{j-1} \in S$ whenever $x_j \in S$ for every j with $1 < j \leq 2k$:

$$\|\{i : x_i \in S\}\| \text{ is odd} \iff f(x_1, x_2, \dots, x_{2k}) \in T. \quad (2)$$

Then T is Θ_2^p -hard.

The following statement is shown by applying Wagner’s Lemma to the NP-complete problem $S = \text{SAT}$ and to $T = \text{INWEAKSADDLE}$. The proof is given in the appendix.

Theorem 3. INWEAKSADDLE is Θ_2^p -hard.

We conclude this section by showing that Σ_2^p is an upper bound for the membership problem.

Proposition 1. INWEAKSADDLE is in Σ_2^p .

Proof. Let $\Gamma = (N, (A_i)_{i \in N}, (p_i)_{i \in N})$ be a game, $d^* \in \bigcup_i A_i$ a designated action. First observe that we can verify in polynomial time whether a subset of A_N is externally stable. We can guess a weak saddle S containing $d^* \in S$ in nondeterministic polynomial time and verify its minimality by checking that none of its subsets are externally stable. This places INWEAKSADDLE in NP^{coNP} and thus in Σ_2^p . \square

5 Very Weak Saddles

A natural weakening of weak dominance is *very weak dominance*, which does not require a strict inequality in addition to the weak inequalities [13]. Thus, in particular, two actions that always yield the same payoff may very weakly dominate each other. Formally, for a player $i \in N$ and two actions $a_i, b_i \in A_i$ we say that a_i *very weakly dominates* b_i with respect to S_{-i} , denoted $a_i \geq_{S_{-i}} b_i$, if $p_i(a_i, s_{-i}) \geq p_i(b_i, s_{-i})$ for all $s_{-i} \in S_{-i}$. Based on this modified notion of dominance, one can define the very weak analog of the weak saddle.

Definition 4 (Very Weak Saddle). *Let $\Gamma = (N, (A_i)_{i \in N}, (p_i)_{i \in N})$ be a game, $S = (S_1, \dots, S_n) \subseteq A_N$. Then, S is a very weak generalized saddle point (VWGSP) of Γ if for each player $i \in N$ the following condition holds:*

$$\text{For every } a_i \in A_i \setminus S_i \text{ there exists } s_i \in S_i \text{ such that } s_i \geq_{S_{-i}} a_i.$$

A very weak saddle is a VWGSP that contains no other VWGSP.

Computational problems for very weak saddles are defined analogously to their counterparts for weak saddles. It turns out that most of our results for the latter can be transferred to the former.

Theorem 4. *The following holds:*

- INVVERYWEAKSADDLE is NP-hard.
- INVVERYWEAKSADDLE is coNP-hard.
- ISVERYWEAKSADDLE is coNP-complete.
- INALLVERYWEAKSADDLES is coNP-complete.
- UNIQUEVERYWEAKSADDLE is coNP-hard.
- NONTRIVIALVERYWEAKSADDLE is NP-complete
- FINDVERYWEAKSADDLE is NP-hard.

It should be noted that the hardness results for very weak saddles do not follow in an obvious way from the corresponding results for weak saddles, or vice versa. While the proofs are based on the same general idea, and again on one core construction, there are some significant technical differences. The proofs of all results are given in the appendix.

6 Conclusion

In the early 1950s, Shapley proposed an ordinal set-valued solution concept known as the weak saddle. We have shown that weak saddles are intractable in bimatrix games. As it turned out, not only *finding* but also *recognizing* weak saddles is computationally hard. This distinguishes weak saddles from Nash equilibrium, iterated dominance, and any other game-theoretic solution concept we are aware of. Three of the most challenging remaining problems are to study the complexity of weak saddles in matrix games, to close the gap between Θ_2^P and Σ_2^P for INVWEAKSADDLE, and to completely characterize the complexity of FINDWEAKSADDLE.

References

1. K. Basu and J. Weibull. Strategy subsets closed under rational behavior. *Economics Letters*, 36:141–146, 1991.
2. D. Baumeister, F. Brandt, F. Fischer, and J. Rothe. Deciding membership in minimal upward covering sets is hard for parallel access to NP. Technical report, <http://arxiv.org/abs/0901.3692>, 2009.
3. B. Bernheim. Rationalizable strategic behavior. *Econometrica*, 52(4):1007–1028, 1984.
4. F. Brandt and F. Fischer. Computing the minimal covering set. *Mathematical Social Sciences*, 56(2):254–268, 2008.
5. F. Brandt, M. Brill, F. Fischer, and P. Harrenstein. Computational aspects of Shapley’s saddles. In *Proc. of 8th AAMAS Conference*, 2009. To Appear.
6. X. Chen and X. Deng. Settling the complexity of 2-player Nash-equilibrium. In *Proc. of 47th FOCS Symposium*, pages 261–272. IEEE Press, 2006.
7. V. Conitzer and T. Sandholm. Complexity of (iterated) dominance. In *Proc. of 6th ACM-EC Conference*, pages 88–97. ACM Press, 2005.
8. C. Daskalakis, P. Goldberg, and C. Papadimitriou. The complexity of computing a Nash equilibrium. In *Proc. of 38th STOC*, pages 71–78. ACM Press, 2006.
9. J. Duggan and M. Le Breton. Dutta’s minimal covering set and Shapley’s saddles. *Journal of Economic Theory*, 70:257–265, 1996.
10. B. Dutta. Covering sets and a new Condorcet choice correspondence. *Journal of Economic Theory*, 44:63–80, 1988.
11. I. Gilboa, E. Kalai, and E. Zemel. The complexity of eliminating dominated strategies. *Mathematics of Operations Research*, 18(3):553–565, 1993.
12. E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Exact analysis of Dodgson elections: Lewis Carroll’s 1876 voting system is complete for parallel access to NP. *Journal of the ACM*, 44(6):806–825, 1997.
13. R. D. Luce and H. Raiffa. *Games and Decisions: Introduction and Critical Survey*. John Wiley & Sons Inc., 1957.
14. R. D. McKelvey and P. C. Ordeshook. Symmetric spatial games without majority rule equilibria. *The American Political Science Review*, 70(4):1172–1184, 1976.
15. J. F. Nash. Non-cooperative games. *Annals of Mathematics*, 54(2):286–295, 1951.
16. C. H. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.
17. D. Pearce. Rationalizable strategic behavior and the problem of perfection. *Econometrica*, 52(4):1029–1050, 1984.
18. L. Samuelson. Dominated strategies and common knowledge. *Games and Economic Behavior*, 4:284–313, 1992.
19. L. Shapley. Order matrices. I. Technical Report RM-1142, The RAND Corporation, 1953.
20. L. Shapley. Order matrices. II. Technical Report RM-1145, The RAND Corporation, 1953.
21. L. Shapley. Some topics in two-person games. In M. Dresher, L. S. Shapley, and A. W. Tucker, editors, *Advances in Game Theory*, volume 52 of *Annals of Mathematics Studies*, pages 1–29. Princeton University Press, 1964.
22. J. von Neumann. Zur Theorie der Gesellschaftspiele. *Mathematische Annalen*, 100:295–320, 1928.
23. J. von Neumann and O. Morgenstern. *The Theory of Games and Economic Behavior*. Princeton University Press, 1944.
24. K. Wagner. More complicated questions about maxima and minima, and some closures of NP. *Theoretical Computer Science*, 51:53–80, 1987.

A Proofs for Weak Saddles

Theorem 1. INWEAKSADDLE is NP-hard.

Proof. We give a reduction from SAT. For a CNF formula φ , we show that the game Γ_φ , defined in Section 4.1, has a weak saddle that contains action d^* if and only if φ is satisfiable. The direction from left to right is straightforward. If α is a satisfying assignment, then S^α is a weak saddle that contains d^* .

For the other direction, we will show that all weak saddles containing d^* are (essentially) assignment saddles. Let $S = (S_1, S_2)$ be a weak saddle of Γ_φ such that $d^* \in S_1$. We can assume that $S_2 \subseteq L$. If this was not the case, i.e., if there was a column $c \in \{b^*, y_1, \dots, y_n\}$ with $c \in S_2$, then $c \rightsquigarrow a^* \rightsquigarrow b^*$, and (a^*, b^*) would be a smaller saddle contained in S , a contradiction. We will now show that

- (i) $|S_2| \geq 2$,
- (ii) $|\{v_i, \bar{v}_i\} \cap S_2| \leq 1$ for all $i \in [n]$, and
- (iii) $C \cap S_1 = \emptyset$.

For (i), suppose that $S_2 = \{\ell\}$, where $\ell = v_i$ or $\ell = \bar{v}_i$. Then, both ℓ and r_i have to be in S_1 , as they are maximal with respect to $\{\ell\}$. Together with $r_i \rightsquigarrow b^*$, this however contradicts the fact that $b^* \notin S_2$.

For (ii), suppose that there exists $i \in [n]$ with $\{v_i, \bar{v}_i\} \subseteq S_2$. Then at least one of the rows v_i or r_i and at least one of the rows \bar{v}_i or \bar{r}_i is in the set S_1 . Since $r_i \rightsquigarrow b^*$ as well as $\bar{r}_i \rightsquigarrow b^*$, and since $b^* \notin S_2$, we know that $\{v_i, \bar{v}_i\} \subseteq S_1$. On the other hand, $\{v_i, \bar{v}_i\} \rightsquigarrow y_i$, again contradicting $S_2 \subseteq L$.

For (iii), merely observe that $C_i \rightsquigarrow b^*$ for all $i \in [m]$.

From (iii) and the fact that $S = (S_1, S_2)$ is a weak saddle, we know that for each $j \in [m]$, there exists a saddle row $s_j \in S_1$ with $s_j >_{S_2} C_j$. From (i) and (ii) we further know that there are (at least) two distinct saddle columns $\ell_1, \ell_2 \in L$ with $p_1(C_j, \ell_1) = p_1(C_j, \ell_2) = 1$. By definition of p_1 , d^* is the only row that can weakly dominate C_j , and therefore $s_j = d^*$ for all $j \in [m]$.

The fact that $d^* >_{S_2} C_j$ implies that there exists $\ell \in S_2$ with $p_1(d^*, \ell) > p_1(C_j, \ell)$. This can only be the case if $p_1(d^*, \ell) = 1$ and $p_1(C_j, \ell) = 0$, where the latter equality means that $\ell \in C_j$. Define an assignment α for φ such that $\alpha(\ell) = 1$ if and only if $\ell \in S_2$. Note that by (ii), α is well-defined. We now have that for each clause C_j , there is a literal $\ell \in C_j$ with $\alpha(\ell) = 1$, i.e., α satisfies φ . \square

	b^*	v_1	\bar{v}_1	v_2	\bar{v}_2	\dots	v_n	\bar{v}_n	y_1	y_2	\dots	y_n
a^*	1 1					\dots			0 1	0 1	\dots	0 1
d^*		1 1	1 1	1 1	1 1	\dots	1 1	1 1			\dots	
C_1	1 0		0 1	0 1		\dots	0 1	0 1			\dots	
C_2	1 0	0 1		0 1	0 1	\dots		0 1			\dots	
\vdots						\vdots						
C_m	1 0	0 1	0 1	0 1	0 1	\dots	0 1	0 1			\dots	
v'_1		1 2				\dots			1 0		\dots	
\bar{v}'_1			1 2			\dots			1 0		\dots	
v'_2				1 2		\dots				1 0	\dots	
\bar{v}'_2					1 2	\dots				1 0	\dots	
\vdots						\vdots						
v'_n						\dots	1 2				\dots	1 0
\bar{v}'_n						\dots		1 2			\dots	1 0
r_1	1 0	0 2		0 -1	0 -1	\dots					\dots	
\bar{r}_1	1 0		0 2	0 -1	0 -1	\dots					\dots	
r_2	1 0			0 2		\dots					\dots	
\bar{r}_2	1 0				0 2	\dots					\dots	
\vdots						\vdots						
r_n	1 0	0 -1	0 -1			\dots	0 2				\dots	
\bar{r}_n	1 0		0 -1	0 -1		\dots		0 2			\dots	

Fig. 4. Game Γ_φ used in the proof of Theorem 1. Payoffs equal $(0,0)$ unless specified otherwise. S^α is a weak saddle of Γ_φ if and only if φ is satisfiable, while (a^*, b^*) always is a weak saddle. For improved readability, thick lines are used to separate different types of actions.

Lemma 1. *For every two-player game $\Gamma = (A, B, p)$ and every action $c \in A \cup B$, there exists a supergame $\Gamma^c = (A', B', p')$ with the following properties:*

- (i) *If S is a weak saddle of Γ^c containing c , then $S = (A', B')$.*
- (ii) *If S is a weak saddle of Γ not containing c , then S is a weak saddle of Γ^c .*
- (iii) *The size of Γ^c is polynomial in the size of Γ .*

Proof. Let $n = |A|$ and $m = |B|$. Without loss of generality, we may assume that all payoffs in Γ are positive and that c is a column, i.e., $p_\ell(a, b) > 0$ for all $(a, b) \in A \times B$, $\ell \in [2]$ and $c \in B$. Define λ to be greater than the maximum payoff to player 1 in column c , e.g., $\lambda = \max_{a \in A} p_1(a, c) + 1$.

Let $\Gamma^c = (A', B', p')$ be the supergame of Γ with $n+m+1$ additional rows and n additional columns, given by $A' = A \cup \{a'_1, \dots, a'_{n+m-1}\}$, $B' = B \cup \{b'_1, \dots, b'_n\}$, and $p'|_{A \times B} = p$. Payoffs for action profiles not in $A \times B$ are shown in Figure 3.

For (i), let $S = (S_1, S_2)$ be a weak saddle of Γ^c with $c \in S_2$. By definition of λ , all new rows a'_1, \dots, a'_{n+m-1} are maximal with respect to column c . Since all these rows are identical for player 1, it follows that *all* of them have to be included in the saddle, i.e., $(A' \setminus A) \subseteq S_1$. The definition of u' now ensures that all columns other than c are in the saddle, again by maximality. Invoking the argument a third time, we finally find that all rows in A are contained in S_1 , because the i th row of A is maximal with respect to column b'_i .

For (ii) observe that our assumption concerning the payoffs in Γ implies that all additional actions are dominated by each of the original actions, as long as c is not contained in the weak saddle.

Finally, (iii) follows directly from the definition of Γ^c . □

Theorem 3. *INWEAKSADDLE is Θ_2^p -hard.*

Proof. We apply Wagner’s Lemma with the NP-complete problem $S = \text{SAT}$ and with $T = \text{INWEAKSADDLE}$. Fix an arbitrary $k \geq 1$ and let $\varphi_1, \dots, \varphi_{2k}$ be $2k$ boolean formulas such that satisfiability of φ_j implies satisfiability of φ_{j-1} , for each j , $1 < j \leq 2k$.

We will now define a polynomial-time computable function f which maps the given $2k$ boolean formulas to an instance of **FINDWEAKSADDLE** such that the requirements of (2) are satisfied. For *odd* $i \in [2k]$, let $\Gamma_i = (A_i, B_i, p_i)$ be a game as defined in the proof of Theorem 1. Recall that this game has a weak saddle containing a certain action d_i if and only if φ_i is satisfiable. Analogously, for *even* $i \in [2k]$, let $\Gamma_i = (A_i, B_i, p_i)$ be a game defined in the proof of Theorem 2, which has a weak saddle containing a certain action d_i if and only if φ_i is *unsatisfiable*. Without loss of generality, we may assume that all payoffs in Γ_i are positive and bounded from above by some $K \in \mathbb{N}$, and that the decision action d_i of game Γ_i is a row, i.e., $0 < p_\ell(a, b) < K$ for all $(a, b) \in A_i \times B_i$, $\ell \in [2]$ and $d_i \in A_i$ for all $i \in [2k]$.⁷

⁷ Adding a positive number to every payoff does not change the dominance relation between the actions. As the minimum payoff in Γ_i is -1 , adding a number greater than 1 suffices. If d_i is a column, as in the proof of Theorem 2, we can simply transpose the game by exchanging the two players.

Now define a game Γ as the union of the games Γ_i , $i \in [2k]$, with one additional row r_i and two additional columns c_i^1 and c_i^2 for each formula φ_i , as well as a decision row d^* , i.e., $\Gamma = (A, B, p)$ where $A = \bigcup_{i=1}^{2k} A_i \cup \{r_1, \dots, r_{2k}\} \cup \{d^*\}$ and $B = \bigcup_{i=1}^{2k} B_i \cup \bigcup_{i=1}^{2k} \{c_i^1, c_i^2\}$. Payoffs $p(a, b)$ for $a \in \bigcup A_i$ and $b \in \bigcup B_i$ are defined as in the games Γ_i . If $a \in A_i$, $b \in B_j$ for $i \neq j$, let $p(a, b) = (0, 0)$. Furthermore, let $p(r_i, b) = (0, 1)$ for all $i \in [2k]$ and $b \in \bigcup B_i$. The definition of p on profiles containing a new column c_i^ℓ , $i \in [2k]$, $\ell \in [2]$ is quite complicated, and we recommend consulting Figure 5 for an overview. Player 2 has only two distinct payoffs for these columns:

$$p_2(a, c_i^\ell) = \begin{cases} K & \text{if } a = d_i \\ 0 & \text{otherwise.} \end{cases}$$

Recall that all payoffs in the games Γ_i are bounded by K , such that the payoff for player 2 in the case $a = d_i$ is maximal in Γ .

The payoffs for player 1 are defined in order to connect the games Γ_{2i} and Γ_{2i+1} , for each i . We need some notation. For $i \in [2k]$, let i° be $i + 1$ if i is even and $i - 1$ if i is odd. Thus, each pair $\{i, i^\circ\}$ is of the form $\{2j, 2j + 1\}$ for some j . For $a \in \bigcup A_i$, define

$$p_1(a, c_i^\ell) = \begin{cases} 1 & \text{if } \ell = 1 \text{ and } a \in A_i \\ 2 & \text{if } \ell = 1 \text{ and } a \in A_{i^\circ} \\ 0 & \text{otherwise,} \end{cases} \quad \text{and}$$

$$(p_1(z_j, c_i^1), p_1(z_j, c_i^2)) = \begin{cases} (1, 1) & \text{if } j = i \\ (0, 0) & \text{if } j = i^\circ \\ (0, 1) & \text{otherwise.} \end{cases}$$

Finally, let $p_1(d^*, c_i^1) = 0$ and $p_1(d^*, c_i^2) = 1$.

An example of the game Γ for the case $k = 2$ is depicted in Figure 5, where we assume without loss of generality that d_i is the first row of Γ_i .

The following facts are readily appreciated.

Fact 2. *If S is a weak saddle of Γ_i not containing d_i , then S is also a weak saddle of Γ .*

For a weak saddle S of Γ and $i \in [2k]$, define $S_i = S \cap (A_i, B_i)$ as the intersection of S with Γ_i .

Fact 3. *If S is a weak saddle of Γ , then S_i is either a weak saddle of Γ_i or empty.*

For Fact 2 it suffices to check external stability. For Fact 3, observe that our assumption that $p_\ell(a, b) > 0$ implies that weak domination with respect to a subset of $A_i \cup B_i$ can only occur among actions belonging to $A_i \cup B_i$. Therefore, if some action profile in $A_i \times B_i$ is contained in a weak saddle, all actions of Γ_i not contained in the saddle must be dominated by some saddle action of the same subgame Γ_i .

		c_1^1	c_1^2	c_2^1	c_2^2	c_3^1	c_3^2	c_4^1	c_4^2
d_1	Γ_1	4	4	0					
		1	0	2					
		0		0					
		1		2					
d_2	Γ_2	0		4	4				
		2		1	0				
		0		0					
		2		1					
d_3	Γ_3					4	4	0	
						1	0	2	
						0		0	
						1		2	
d_4	Γ_4					0		4	4
						2		1	0
						0		0	
						2		1	
d^*	0		0		0		0		0
			1		1		1		1
			0	0				0	
z_1			1	1			1		1
z_2				0	0		0		0
z_3			1	1		1		1	
z_4		0		0	0	0			
		1		1		1	1		
			0		0		0	0	
		1		1			1	1	

Fig. 5. Game Γ used in the proof of Theorem 3. Payoffs are $(0, 0)$ unless specified otherwise. Γ has a weak saddle containing row d^* if and only if both Γ_1 and Γ_2 or both Γ_3 and Γ_4 have a weak saddle containing their respective decision rows d_i .

In order to be able to apply Lemma 2, we now prove (2), which here amounts to showing the following equivalence:

$$\|\{i : \varphi_i \in \text{SAT}\}\| \text{ is odd} \iff \Gamma \text{ has a weak saddle } S \text{ with } d^* \in S \quad (3)$$

For the direction from left to right, let S_i be a weak saddle of Γ_i containing d_i , and let S_{i° be a weak saddle of Γ_{i° containing d_{i° . The existence of these weak saddles is guaranteed by construction of Γ and the fact that φ_i is satisfiable and φ_{i° is unsatisfiable. Now let $S = S_i \cup S_{i^\circ} \cup (\{d^*, z_1, \dots, z_{2k}\}, \{c_i^1, c_i^2, c_{i^\circ}^1, c_{i^\circ}^2\})$. We claim that S is a weak saddle of Γ . The proof consists of two parts.

First, we have to show that S is externally stable, i.e., all actions not in the saddle have to be weakly dominated by saddle actions. To see this, let $a \in A_j$ be a row that is not in S . If $j \notin \{i, i^\circ\}$, then a is weakly dominated by every saddle row because it yields payoff 0 to player 1 against any saddle columns. If, on the other hand, $j \in \{i, i^\circ\}$, then a is weakly dominated by the same row that weakly dominates it in the subgame Γ_j . The argument for non-saddle columns

$b \in \bigcup_i B_i$ is analogous. Moreover, every column c_j^ℓ with $j \notin \{i, i^\circ\}$ is weakly dominated by each of the saddle columns $c_i^1, c_i^2, c_{i^\circ}^1, c_{i^\circ}^2$.

Second, we have to show that S is inclusion-minimal, i.e., that no proper subset of S is a weak saddle of Γ . Assume for contradiction that such a subset $S' \subset S$ exists. By Fact 3, we know that $S'_i = S_i$, as otherwise inclusion-minimality of S_i would be violated. In particular, $d_i \in S'_i$, which implies that $\{c_i^1, c_i^2\} \subseteq S'$. The same reasoning for i° shows that $S'_{i^\circ} = S_{i^\circ}$ and $\{c_{i^\circ}^1, c_{i^\circ}^2\} \subseteq S'$. Then, both z_i and z_{i° have to be in S' , because they are both uniquely maximal with respect to $\{c_i^1, c_i^2, c_{i^\circ}^1, c_{i^\circ}^2\}$. Furthermore, all rows z_j with $j \notin \{i, i^\circ\}$, as well as d^* , are in S' , because they are all identical and maximal with respect to S' . Here, maximality is due to the fact that they are the only rows that yield a positive payoff to player 1 against both saddle columns c_i^2 and $c_{i^\circ}^2$. Thus $S' = S$, meaning that S is indeed inclusion-minimal.

For the direction from right to left, let S be a weak saddle of Γ with $d^* \in S$. From the definition of $p_2(d^*, \cdot)$, we infer that $S \cap \bigcup_i B_i \neq \emptyset$, which in turn implies that there is at least one column $c_i^\ell \in S$. Otherwise, row d^* would always yield 0 against all saddle actions and thus would be weakly dominated by all saddle rows in $\bigcup_i A_i$. Now observe that for any $i \in [2k]$, c_i^1 and c_i^2 are identical for player 2, which implies that every weak saddle of Γ contains either none or both of them. We thus have that $\{c_i^1, c_i^2\} \subseteq S$. It then has to be the case that $z_i \in S$, because this row is maximal with respect to $\{c_i^1, c_i^2\}$. However, z_i must not weakly dominate d^* with respect to S , because d^* is itself a saddle action. This means there has to be a saddle column $c \in S$ with $p_1(z_i, c) < p_1(d^*, c)$. The only column satisfying this property is $c_{i^\circ}^2$, which means that both $c_{i^\circ}^2$ and, by the same argument as above, $c_{i^\circ}^1$ are contained in S . Now that both c_i^1 and $c_{i^\circ}^1$ are in S , at least one row from each of the games Γ_i and Γ_{i° has to be a saddle action, i.e., $S \cap A_i \neq \emptyset$ and $S \cap A_{i^\circ} \neq \emptyset$. By Fact 3, we conclude that S_i and S_{i° are weak saddles of Γ_i and Γ_{i° , respectively.

It remains to be shown that $d_i \in S_i$ and $d_{i^\circ} \in S_{i^\circ}$. If $d_i \notin S_i$, then by Fact 2 $S_i \subset S$ would be a weak saddle of Γ , contradicting inclusion-minimality of S . The argument for S_{i° is analogous. It finally follows from the construction that φ_i is satisfiable and φ_{i° is unsatisfiable,⁸ which completes the proof of (3). By Lemma 2, FINDWEAKSADDLE is Θ_2^P -hard. \square

B Proofs for Very Weak Saddles

As in the case of weak saddles, we begin by defining, for each Boolean formula φ , a two-player game Γ_φ that admits a certain type of very weak saddles if and only if φ is satisfiable. Let $\varphi = C_1 \wedge \dots \wedge C_m$ be a 3-CNF formula⁹ over

⁸ Here we have assumed without loss of generality that $i < i^\circ$, i.e., i is even and $i^\circ = i + 1$ is odd.

⁹ A formula in 3-CNF is a CNF formula where every clause consists of exactly three literals. Recall that SAT is NP-complete even for this restricted class of formulas. While the construction works for arbitrary CNF formulas, we employ 3-CNFs for ease of notation.

	$\ell_{1,1}$	$\ell_{1,2}$	$\ell_{1,3}$	$\ell_{2,1}$	$\ell_{2,2}$	$\ell_{2,3}$	\dots	$\ell_{3,1}$	$\ell_{3,2}$	$\ell_{3,3}$
C_1	1	1	1				\dots	0	0	0
C_2	0	0	0	1	1	1	\dots			
C_3				0	0	0	\dots			
\vdots				1	1	1	\dots			
C_m							\dots	1	1	1
$[\ell_{1,1}, \ell_{2,3}]$	0					0	\dots	0	0	0
$[\ell_{1,2}, \ell_{2,1}]$	1					1	\dots			
$[\ell_{1,1}, \ell_{m,1}]$		0		0			\dots			
\vdots	1						\dots	0		
\vdots		1		1			\dots			
\vdots							\dots	1		
\vdots							\dots			

Fig. 6. Game Γ_φ for a formula φ with $C_1 = v_1 \vee \bar{v}_2 \vee v_3$, $C_2 = v_2 \vee v_4 \vee \bar{v}_1$ and $C_m = \bar{v}_1 \vee \bar{v}_2 \vee v_4$.

variables v_1, \dots, v_n , where $C_i = \{\ell_{i,1}, \ell_{i,2}, \ell_{i,3}\}$. Call a pair $\{\ell_{i,j}, \ell_{i',j'}\}$ of variable occurrences a *conflicting pair* if $i \neq i'$ and $\ell_{i,j} = \bar{\ell}_{i',j'}$. Conflicting pairs are denoted $[\ell_{i,j}, \ell_{i',j'}]$.

Define the bimatrix game $\Gamma_\varphi = (A, B, p)$ as follows. The set A of actions of player 1 comprises the set $C = \{C_1, \dots, C_m\}$ of clauses of φ as well as one additional action for each conflicting pair $[\ell_{i,j}, \ell_{i',j'}]$ of literals. The set B of actions of player 2 equals the set $O = \bigcup_{j=1}^m \{\ell_{j,1}, \ell_{j,2}, \ell_{j,3}\}$ of all literal occurrences. Payoffs are given by

$$p(C_i, \ell_{j,k}) = \begin{cases} (0, 1) & \text{if } j = i, \\ (1, 0) & \text{if } j = i + 1 \pmod{m}, \\ (0, 0) & \text{otherwise, and} \end{cases}$$

$$p([\ell_{i,j}, \ell_{i',j'}], \ell_{p,q}) = \begin{cases} (1, 0) & \text{if } i = p \text{ and } j = q, \\ (1, 0) & \text{if } i' = p \text{ and } j' = q, \\ (0, 0) & \text{otherwise.} \end{cases}$$

An example of a game Γ_φ is shown in Figure 6. In the following, we will exploit three key properties of Γ_φ .

Let $S = (S_1, S_2)$ be a very weak saddle of Γ_φ . Then the following properties hold:

- (i) If $C_i \in S_1$ for some $i \in [m]$, then $\ell_{i,j} \in S_2$ for some $j \in [3]$.
- (ii) If $\ell_{i,j} \in S_2$ for some $i \in [m]$ and $j \in [3]$, then $C_{i+1} \in S_1$ or $[\ell_{i,j}, \ell_{i',j'}] \in S_1$ for some $i' \in [m]$ and $j' \in [3]$.
- (iii) For two conflicting literals $\ell_{i,j} = \bar{\ell}_{i',j'}$ we have $\{\ell_{i,j}, \ell_{i',j'}\} \rightsquigarrow [\ell_{i,j}, \ell_{i',j'}]$.

The idea underlying this construction is formalized in the following lemma.

Lemma 3. *The game Γ_φ has a very weak saddle $S = (S_1, S_2)$ with $S_1 = C$ if and only if φ is satisfiable.*

Proof. For the direction from left to right, consider a saddle $S = (S_1, S_2)$ as in the statement of the lemma. By (iii), S_2 does not include any conflicting literals and thus defines a valid assignment α for φ . Moreover, (i) ensures that $|\{\ell_{i,1}, \ell_{i,2}, \ell_{i,3}\} \cap S_2| \geq 1$ for each $i \in [m]$, meaning that α satisfies φ .

For the direction from right to left, let α be a satisfying assignment of φ and $f : [m] \rightarrow [3]$ be a function such that $\alpha(\ell_{i,f(i)}) = 1$ for all $i \in [m]$. It is then easily verified that $S = (C, \{\ell_{i,f(i)} : i \in [m]\})$ is a very weak saddle of Γ_φ . \square

In the following we define two bimatrix games Γ'_φ and Γ_φ^+ that extend Γ_φ with new actions such that properties (i), (ii), and (iii) still hold for the extended games. In particular, Lemma 3 still holds for Γ'_φ and Γ_φ^+ . The game Γ'_φ is then used to prove the NP-hardness of INVERYWEAKSADDLE, while Γ_φ^+ is used in the proofs of all other hardness results. Both extensions are independent of the initial formula φ .

The game Γ'_φ is defined by adding a column d to Γ_φ . Payoffs for new profiles are defined as

$$\begin{aligned} p(C_i, d) &= (0, 0) \text{ for all } i \in [m], \\ p([\ell_{i,j}, \ell_{i',j'}], d) &= (1, 1) \text{ for each conflicting pair.} \end{aligned}$$

Lemma 4. *Γ'_φ has a very weak saddle $S = (S_1, S_2)$ with $C_1 \in S_1$ if and only if φ is satisfiable.*

Proof. By Lemma 3, there is a very weak saddle S with $S_1 = C$ in the game Γ_φ if and only if φ is satisfiable. Since $p(C_i, d) = (0, 0)$ for all $i \in [m]$, this property still holds for Γ'_φ .

It remains to show that if (S_1, S_2) is a very weak saddle with $C_1 \in S_1$ in Γ'_φ , then $S_1 = C$. But this is true since property (ii) holds for Γ'_φ , $[\ell_{i,j}, \ell_{i',j'}] \rightsquigarrow d$ for every $[\ell_{i,j}, \ell_{i',j'}]$, and $(\{[\ell_{i,j}, \ell_{i',j'}]\}, \{d\})$ is a very weak saddle. Thus, $(\{[\ell_{i,j}, \ell_{i',j'}]\}, \{d\})$ is the only very weak saddle containing $[\ell_{i,j}, \ell_{i',j'}]$. \square

To show the remaining hardness results, we define the bimatrix Γ_φ^+ that is another extension of the basic game Γ_φ . In addition to the properties (i), (ii), and (iii) we will have the following new property in Γ_φ^+ :

(iv) For every action of the form $[\ell_{i,j}, \ell_{i',j'}]$ it is true that $[\ell_{i,j}, \ell_{i',j'}] \rightsquigarrow a$ for every action a of Γ_φ^+ .

Recall that $\Gamma_\varphi = (A, B, p)$ and label actions $A = \{a_1, \dots, a_r\}$ such that $C_i = a_i$ for all $i \in [m]$ and let $B = \{b_1, \dots, b_{3m}\}$. To obtain Γ_φ^+ , we add s new columns d_1, \dots, d_s and s new lines f_1, \dots, f_s to Γ_φ where $s = \max(|A|, |B|) + 1$. Payoffs for new profiles are defined as

$$\begin{aligned} p(f_i, d_j) &= \begin{cases} (2, 0) & \text{if } j = i \\ (0, 2) & \text{if } j = i + 1 \pmod s \\ (0, 0) & \text{otherwise} \end{cases} \\ p(f_i, b_j) &= \begin{cases} (0, 1) & \text{if } i = j \text{ or } i = j + 1 \\ (0, 0) & \text{otherwise} \end{cases} \\ p([\ell_{i,j}, \ell_{i',j'}], d_1) &= (0, 1) \text{ for all conflicting pairs} \\ p(C_i, d_1) &= \begin{cases} (1, 0) & \text{if } 1 \leq i \leq 2 \\ (0, 0) & \text{otherwise} \end{cases} \\ p(a_i, d_j) &= \begin{cases} (1, 0) & \text{if } j > 1 \text{ and } j \in \{i, i + 1\} \\ (0, 0) & \text{if } j > 1 \text{ and } j \notin \{i, i + 1\} \end{cases} \end{aligned}$$

The game Γ_φ^+ is shown in Figure 7. Note that (i), (ii), and (iii) hold for the Γ_φ^+ . This is the case since we have $a >_B f_i$ for all $a \in A$, $i \in [s]$ as well as $b >_A d_i$ for $b \in B$ and all $i \in [s]$. Therefore we can show the following lemma analogously to Lemma 3.

Lemma 5. *The game Γ_φ^+ has a very weak saddle $S = (S_1, S_2)$ with $S_1 = C$ if and only if φ is satisfiable.*

To prove (iv), note that $[\ell_{i,j}, \ell_{i',j'}] \rightsquigarrow d_1$ for every conflicting pair $[\ell_{i,j}, \ell_{i',j'}]$. Furthermore we have $d_i \rightsquigarrow f_i$ for every $i \in [s]$ and $f_j \rightsquigarrow d_{j+1}$ for every $j \in [s-1]$. So it follows from the transitivity of \rightsquigarrow that $[\ell_{i,j}, \ell_{i',j'}] \rightsquigarrow d_k$ and $[\ell_{i,j}, \ell_{i',j'}] \rightsquigarrow f_k$ for every $[\ell_{i,j}, \ell_{i',j'}] \in A$ and all $k \in [s]$. Finally it follows directly from the construction that $\{d_i, d_{i+1}\} \rightsquigarrow a_i$ and $\{f_i, f_{i+1}\} \rightsquigarrow b_i$ for all $1 \leq i < s$. Since $s > \max(|A|, |B|)$ this shows (iv).

Lemma 6. *The game Γ_φ^+ has a nontrivial very weak saddle if and only if the formula φ is satisfiable.*

Proof. If φ is satisfiable, there is a nontrivial very weak saddle by Lemma 5. Now assume that φ is unsatisfiable. From (iv) we know that there is no nontrivial saddle (S_1, S_2) with $[\ell_{i,j}, \ell_{i',j'}] \in S_1$ for any conflicting pair $[\ell_{i,j}, \ell_{i',j'}]$. By Lemma 5, there is no saddle (S_1, S_2) with $S_1 = C$. And it follows from (ii) and (iv) that there cannot be a very weak saddle (S_1, S_2) with $S_1 \subset C$. It remains to show that a nontrivial very weak saddle cannot contain any of the new actions f_i or d_j . As mentioned above, $d_i \rightsquigarrow f_i$ and $f_j \rightsquigarrow d_{j+1}$. But by the construction we also have $f_s \rightsquigarrow d_1$. Hence, we can conclude— analogously to the proof of (iv)—that $d_i \rightsquigarrow a$ and $f_i \rightsquigarrow a$ for every action a and, therefore, that d_i and f_i are not part of a nontrivial saddle for every $i \in [s]$. \square

	$\ell_{1,1}$...	$\ell_{m,3}$	d_1	d_2	d_3	d_4	...	d_s	d_s
C_1	1	...	0	0	0			...		
	0	...	1	1	1			...		
C_2	0	...			0	0		...		
	1	...			1	1		...		
C_3		...				0	0	...		
						1	1	...		
\vdots										
C_m		...	1					...		
			0					...		
$\{\ell_{i,j}, \ell_{i',j'}\}$	0	...		1				...		
	1	...		0				...		
\vdots										
$\{\ell_{p,q}, \ell_{p',q'}\}$...		1				...	0	0
				0				...	1	1
f_1	1	...		0	2			...		
	0	...		2	0			...		
f_2	1	...			0	2		...		
	0	...			2	0		...		
f_3		...				0	2	...		
						2	0	...		
f_4		...					0	...		
							2	...		
\vdots										
f_s		...		2				...		0
				0				...		2

Fig. 7. The game Γ_φ^+ . Payoffs are $(0, 0)$ unless specified otherwise.

Theorem 4. *The following holds:*

- (i) *INVERYWEAKSADDLE is NP-hard.*
- (ii) *INVERYWEAKSADDLE is coNP-hard.*
- (iii) *ISVERYWEAKSADDLE is coNP-complete.*
- (iv) *INALLVERYWEAKSADDLES is coNP-complete.*
- (v) *UNIQUEVERYWEAKSADDLE is coNP-hard.*
- (vi) *NONTRIVIALVERYWEAKSADDLE is NP-complete*
- (vii) *FINDVERYWEAKSADDLE is NP-hard.*

Proof. Let φ be a Boolean formula and let Γ_φ^+ be the game defined above.

- (i) By a reduction from 3-SAT. See Lemma 4.
- (ii) It follows directly from Lemma 5 and Lemma 6 that an action $[\ell_{i,j}, \ell_{i',j'}]$ is in a very weak saddle (namely the trivial saddle) if and only if the formula is unsatisfiable. (One can assume w.l.o.g. that φ has a pair of conflicting literals.)
- (iii) To see that ISVERYWEAKSADDLE is in coNP, note that a minor modification of the coNP algorithm for ISWEAKSADDLE is a coNP algorithm for ISVERYWEAKSADDLE. The problem is hard for coNP since the set of all actions of Γ_φ^+ is a very weak saddle if and only if φ is unsatisfiable (Lemma 6).
- (iv) The above discussion also shows that INALLVERYWEAKSADDLES is coNP-complete.
- (v) One can assume w.l.o.g. that φ has more than one satisfying assignment. (Just add a new clause $\hat{v}_1 \vee \hat{v}_2 \vee \hat{v}_3$ with fresh variables.) Then it follows from the proof of Lemma 5 that there are multiple very weak saddles $S = (S_1, S_2)$ with $S_1 = C$, each one corresponding to a satisfying assignment. On the other hand it follows from Lemma 6 that Γ_φ^+ has only the trivial saddle if φ is unsatisfiable. That shows that UNIQUEVERYWEAKSADDLE is coNP-hard.
- (vi) The proof of NP-membership of NONTRIVIALVERYWEAKSADDLE is similar to the proof of NP-membership of NONTRIVIALWEAKSADDLE. The NP-hardness of the problem follows directly from Lemma 6.
- (vii) Analogous to the proof of the NP-hardness of FINDWEAKSADDLE.

□