

Computational Aspects of Shapley’s Saddles

Felix Brandt

Markus Brill

Felix Fischer

Paul Harrenstein

Institut für Informatik
Universität München
80538 München, Germany
{brandtf,brill,fischerf,harrenst}@tcs.ifi.lmu.de

ABSTRACT

Game-theoretic solution concepts, such as Nash equilibrium, are playing an ever increasing role in the study of systems of autonomous computational agents. A common criticism of Nash equilibrium is that its existence relies on the possibility of *randomizing* over actions, which in many cases is deemed unsuitable, impractical, or even infeasible. In work dating back to the early 1950s Lloyd Shapley proposed ordinal set-valued solution concepts for zero-sum games that he refers to as *strict and weak saddles*. These concepts are intuitively appealing, they always exist, and are unique in important subclasses of games. We initiate the study of computational aspects of Shapley’s saddles and provide polynomial-time algorithms for computing strict saddles in normal-form games and weak saddles in a subclass of symmetric zero-sum games. On the other hand, we show that certain problems associated with weak saddles in bimatrix games are NP-hard. Finally, we extend our results to mixed refinements of Shapley’s saddles introduced by Duggan and Le Breton.

Categories and Subject Descriptors

F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems; I.2.11 [Distributed Artificial Intelligence]: Multiagent Systems; J.4 [Computer Applications]: Social and Behavioral Sciences—*Economics*

General Terms

Theory, Algorithms, Economics

Keywords

Game Theory, Solution Concepts, Shapley’s Saddles

1. INTRODUCTION

Game-theoretic solution concepts, such as Nash equilibrium, are playing an ever increasing role in the study of systems of autonomous computational agents. A common criticism of Nash equilibrium is that its existence relies on the possibility of *randomizing* over actions, which has been attacked on various grounds [cf. 16, pp. 74-76]. Aumann chal-

Cite as: Computational Aspects of Shapley’s Saddles, Felix Brandt, Markus Brill, Felix Fischer, and Paul Harrenstein, *Proc. of 8th Int. Conf. on Autonomous Agents and Multiagent Systems (AA-MAS 2009)*, Decker, Sichman, Sierra and Castelfranchi (eds.), May, 10-15, 2009, Budapest, Hungary, pp. XXX-XXX.

Copyright © 2009, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

lenges the suitability of randomized strategies in one-shot games: “When randomized strategies are used in a strategic game, payoff must be replaced by expected payoff. Since the game is played only once, the law of large numbers does not apply, so it is not clear why a player would be interested specifically in the mathematical expectation of his payoff” [1, p. 63]. On top of that, players might simply be incapable of executing reliable randomizations. This is particularly true for games with more than two players, in which equilibrium probabilities may be *irrational* numbers [19].

Shapley has proposed ordinal set-valued solution concepts for zero-sum games that he refers to as *saddles* [21, 22, 23, 24]. What makes these concepts intuitively appealing is that they are based on the elementary notions of dominance and stability. Call a *generalized saddle point (GSP)* a tuple of subsets of actions for each player such that every action not contained in the GSP is dominated by some action in the GSP, given that the remaining players choose actions from the GSP. Then, 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 and weak saddles. Shapley [24] showed that every two-player zero-sum game admits a *unique* strict saddle. Duggan and Le Breton [12] proved that the same is true for the weak saddle in a certain subclass of symmetric two-player zero-sum games.

Despite the fact that Shapley’s saddles were devised as early as 1953 [21, 22] and are thus almost as old as Nash equilibrium [19], surprisingly little is known about their computational properties. In this paper, we provide polynomial-time algorithms for computing strict saddles in normal-form games (with any number of players) and weak saddles in a subclass of symmetric two-player zero-sum games introduced by Duggan and Le Breton [12]. On the other hand, we show that certain problems associated with weak saddles in bimatrix games, such as deciding whether there exists a weak saddle with at most k actions for some player, are NP-hard. Finally, we extend our results to mixed refinements of Shapley’s saddles that were introduced by Duggan and Le Breton [11].

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 [7, 10] and thus unlikely to admit a polynomial-time algorithm. Shapley’s saddles

are based on the notion of dominance, which has also been studied from a computational perspective, in particular in the form of *iterated dominance* [e.g., 15, 8, 9, 6]. Our algorithm for computing strict saddles is interesting insofar as most solution concepts are not known to be efficiently computable in general games, one of the few exceptions being iterated strict dominance. Strict saddles may be considered a “refinement” of iterated strict dominance as all strict saddles of a normal-form game are contained in the subgame that one obtains by iterated elimination of strictly dominated actions.

Another concept related to Shapley’s saddles are *CURB sets* [2], for which Benisch et al. [3] have proposed polynomial-time algorithms for bimatrix games. Both CURB sets and Shapley’s saddles are *set-valued* concepts. However, CURB sets are not *ordinal* as they are based on randomized strategies. Every strict saddle represents the support of a CURB set, and thus contains the support of a minimal CURB set. In confrontation games, as defined in Section 5.1, the support of a minimal CURB set and the strict saddle trivially coincide. Moreover, in this particular class of games, the strict mixed saddle is identical to the support of the minimal CURB set when only allowing pure strategies. There appears to be no such relationship between *weak* saddles and CURB sets.

3. PRELIMINARIES

An accepted way to model situations of strategic interactions is by means of a *normal-form game* [e.g., 20].

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 $p_i : (\times_{i \in N} A_i) \rightarrow \mathbb{R}$ is a function mapping each action profile (i.e., combination of actions) to a real-valued payoff for player i .

A two-player game $(\{1, 2\}, (A_1, A_2), (p_1, p_2))$ is alternatively called a *bimatrix game*, because it can be represented by two matrices M_1 and M_2 with rows and columns indexed by A_1 and A_2 , respectively, and $M_i(a_1, a_2) = p_i(a_1, a_2)$ for all $a_1 \in A_1, a_2 \in A_2$. A bimatrix game is called a *zero-sum* or *matrix game*, and represented by a single matrix M that just contains the payoffs for the first player, if $p_2(a, b) = -p_1(a, b)$ for all $(a, b) \in A_1 \times A_2$. Γ_M denotes the matrix game with matrix M . Finally, a bimatrix game is called *symmetric* if $A_1 = A_2$ and $p_1(a, b) = p_2(b, a)$ for all $a, b \in A_1$. Observe that Γ_M is symmetric if and only if M is *skew symmetric*, i.e., $M^T = -M$. We assume throughout the paper that games are given explicitly, i.e., as a table containing the payoffs for every possible action profile.

A *solution concept* identifies combinations of (sets of) strategies that are significant in some specified sense. Here, a strategy s_i for a player $i \in N$ is a probability distribution over his set of actions, i.e., $s_i \in \Delta(A_i)$. Actions can be identified with strategies that put probability 1 on that action, often called pure strategies. There are plenty of solution concepts for normal-form games, chief among them *Nash equilibrium* [18]. A Nash equilibrium is a combination of strategies, one for each player, such that no player can achieve a higher payoff by unilaterally changing his strategy. Formally, a vector $s = (s_1, s_2, \dots, s_n)$ is called a *strategy profile*

if $s_i \in \Delta(A_i)$ for all $i \in N$. For a strategy profile s , denote by s_{-i} the vector that contains the strategies of all players except player i , and by (s'_i, s_{-i}) the strategy profile where player i plays strategy s'_i and all other players play the same strategy as in s . Payoff functions can naturally be extended to strategy profiles s in terms of the *expected* payoff under the probability distribution generated by s_1, s_2, \dots, s_n .

DEFINITION 2. A Nash equilibrium is a strategy profile $s = (s_1, s_2, \dots, s_n)$ such that for all players $i \in N$ and all strategies $s'_i \in \Delta(A_i)$, $p_i(s_i, s_{-i}) \geq p_i(s'_i, s_{-i})$.

A well-known drawback of Nash equilibria is that their existence is not guaranteed if strategies are required to be pure. To illustrate this, define a *saddle point* of a matrix game Γ_M as a pair (i, j) of actions $i \in A_1, j \in A_2$ such that entry $M(i, j)$ is maximal in column j and minimal in row i . If such a saddle point exists, it is also a Nash equilibrium in pure strategies and constitutes a reasonable prediction of the outcome of the game. The problem is, of course, that there are matrix games without a saddle point, for example the game of *Matching Pennies* given by the matrix

$$\begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}.$$

The only Nash equilibrium of this game has both players pick one of their actions uniformly at random.

As pointed out in the introduction, requiring randomization in order to reach a stable outcome has been criticized for various reasons. A possible solution is to consider *set-valued* solution concepts that identify, for each player i , a subset $S_i \subseteq A_i$, such that the tuple (S_1, S_2, \dots, S_n) satisfies some notion of stability. Shapley’s saddles generalize saddle points by requiring that for every action a_i of a player $i \in N$ that is *not* included in S_i , there should be some reason for its exclusion, namely an action in S_i that is strictly better than a_i . To formalize this idea, we need some notation. Let $A = (A_1, A_2, \dots, A_n)$. For $S = (S_1, S_2, \dots, S_n)$, we write $S \subseteq A$ and say that S is a subset of A if $\emptyset \neq S_i \subseteq A_i$ for all $i \in N$. Further let $S_{-i} = (S_1, S_2, \dots, S_{i-1}, S_{i+1}, \dots, S_n)$. For a player $i \in N$ and two actions $a_i, b_i \in A_i$ we say that

- a_i *strictly dominates* b_i with respect to S_{-i} , denoted $a_i \gg_{S_{-i}} b_i$, if $p(a_i, s_{-i}) > p(b_i, s_{-i})$ for all $s_{-i} \in S_{-i}$, and that
- a_i *weakly dominates* b_i with respect to S_{-i} , denoted $a_i \succ_{S_{-i}} b_i$, if $p(a_i, s_{-i}) \geq p(b_i, s_{-i})$ for all $s_{-i} \in S_{-i}$, with at least one strict inequality.

Based on these notions of dominance, strict and weak saddles can be defined as follows.

DEFINITION 3 (STRICT AND WEAK SADDLE). Let $\Gamma = (N, (A_i)_{i \in N}, (p_i)_{i \in N})$ and $S = (S_1, S_2, \dots, S_n) \subseteq A$. Then, S is a generalized saddle point (GSP) of Γ if for each player $i \in N$ and

$$\forall a_i \in A_i \setminus S_i, \exists s_i \in S_i \text{ such that } s_i \gg_{S_{-i}} a_i. \quad (1)$$

A strict saddle is a GSP that contains no other GSP.

Similarly, S is a weak generalized saddle point (WGSP) of Γ if for each player $i \in N$ and

$$\forall a_i \in A_i \setminus S_i, \exists s_i \in S_i \text{ such that } s_i \succ_{S_{-i}} a_i. \quad (2)$$

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

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 dominates a_i , provided that the other players play only actions from their distinguished sets. Properties (1) and (2) are sometimes referred to as *external stability*. Using this terminology, a (W)GSP is a tuple S that is externally stable for every player. Since strict dominance implies weak dominance, every strict saddle is a WGSP and thus contains a weak saddle. Consider for example the matrix game Γ_M given by

$$M = \begin{pmatrix} 3 & 3 & 4 \\ 2 & 3 & 3 \\ 1 & 2 & 3 \\ 2 & 0 & 5 \end{pmatrix}.$$

Throughout the paper, the rows and columns of a matrix are indexed by r_1, r_2, \dots and c_1, c_2, \dots , respectively. The pair $S = (\{r_1, r_2\}, \{c_1, c_2\})$ is a strict saddle and a WGSP of Γ_M . Since r_1 weakly dominates r_2 with respect to $\{c_1, c_2\}$ and both c_1 and c_2 dominate c_3 with respect to $\{r_1\}$, the pair $S' = (\{r_1\}, \{c_1, c_2\})$ is also a WGSP. Indeed, S' is a weak saddle because it contains no smaller WGSP. Some reflection reveals that S and S' are in fact the *unique* strict and weak saddle of this game, respectively.

It is easy to see that every normal-form game has a strict and a weak saddle. By definition, the set A is a GSP. Furthermore, every GSP that is not a saddle must contain a GSP that is strictly smaller. Finiteness of A implies that there exists a minimal GSP, i.e., a strict saddle. An analogous argument applies to the weak saddle. Strict saddles are *unique* in matrix games but not in general games, whereas weak saddles are not even unique in matrix games. For examples see the coordination game in Section 4 and the matrix game Γ_D in Section 5.1. We finally note that both strict and weak saddle are *ordinal* solution concepts, i.e., they are invariant under order-preserving transformations of the payoff functions. This is in contrast to mixed-strategy Nash equilibrium, for which invariance holds only under positive *affine* transformations.

4. STRICT SADDLE

Shapley [24] has shown that every matrix game possesses a unique strict saddle, because the set of GSPs in such games is closed under intersection, and describes an algorithm, attributed to Harlan Mills, to compute this saddle. The idea behind this algorithm is that given a subset of the saddle, the saddle itself can be computed by iteratively adding actions that are maximal, i.e., not dominated with respect to the current subset of actions of the other player. Shapley further points out that a subset of the strict saddle can easily be found by taking all rows and columns that contain a *minimax* or a *maximin* point, i.e., an entry that is minimal among all column maxima or maximal among all row minima. This establishes that the strict saddle of a matrix game can be computed in polynomial time.

Observe, however, that being able to find a subset of the saddle is not crucial. Starting the above procedure from singleton sets of actions, and invoking it for every combination of such sets, yields a number of candidates for the strict saddle. The strict saddle can then be identified as the inclusion-minimal set among these candidates. The correctness of this procedure follows from the fact that every candidate set is a

Algorithm 1 Minimal GSP

```

procedure minGSP( $\Gamma, (S_1^0, S_2^0, \dots, S_n^0)$ )
  for all  $i \in N$  do
     $S_i \leftarrow S_i^0$ 
  end for
  repeat
    for all  $i \in N$  do
       $A'_i \leftarrow \{a_i \in A_i \setminus S_i : \nexists s_i \in A_i \text{ with } s_i \gg_{S_{-i}} a_i\}$ 
       $S_i \leftarrow S_i \cup A'_i$ 
    end for
  until  $\bigcup_{i=0}^n A'_i = \emptyset$ 
  return  $(S_1, S_2, \dots, S_n)$ 

```

GSP and that the unique strict saddle is contained in every GSP. Furthermore, the iterative procedure itself is invoked only a polynomial number of times.

In contrast to matrix games, strict saddles are no longer unique in general n -player games. For example, take the two-player coordination game given by matrices

$$M_1 = M_2 = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

This game has two strict saddles: One where both players play their first action, and one where both players play their second action.¹

From a computational point of view, however, the existence of multiple strict saddles does not have any serious consequences. Indeed, we proceed to show how Mills' algorithm can be generalized to efficiently compute *all* strict saddles of an arbitrary n -player game. To this end, recall that Mills' iterative procedure required as an input some non-empty subset of the strict saddle. Algorithm 1 is a straightforward generalization of this procedure to the n -player case. Given a tuple $S^0 = (S_1^0, S_2^0, \dots, S_n^0) \subseteq A$ as input, it computes the minimal GSP containing S^0 .

LEMMA 1. *Algorithm 1 computes the inclusion-minimal GSP containing a given input set S^0 .*

PROOF. Let S^{\min} be the minimal GSP containing S^0 . We show that during the execution of Algorithm 1, the set S is always a subset of S^{\min} . At the end of the algorithm, $\bigcup_{i=0}^n A'_i = \emptyset$ implies that S is a GSP, and the statement of the lemma follows.

We prove $S \subseteq S^{\min}$ by induction on $|S| = \sum_{i=1}^n |S_i|$. At the beginning of the algorithm, $S = S^0 \subseteq S^{\min}$ by definition of S^{\min} . Now assume that $S \subseteq S^{\min}$ at the beginning of a particular iteration. We have to show that for all $i \in N$, $A'_i \subseteq S_i^{\min}$. Let $a \in A'_i$, and assume for contradiction that $a \notin S_i^{\min}$. Since S^{\min} is a GSP, there exists $a^* \in S_i^{\min} \subseteq A_i$ with $a^* \gg_{S_{-i}^{\min}} a$. By the induction hypothesis, $S_{-i} \subseteq S_{-i}^{\min}$, which in turn implies $a^* \gg_{S_{-i}} a$. This contradicts the assumption that $a \in A'_i$. \square

Whenever S^0 is contained in a strict saddle, Algorithm 1 returns this strict saddle. This property can be used to construct an algorithm to compute all strict saddles of a game:

¹Also recall that strict saddles where every set S_i is a singleton are pure Nash equilibria. For the converse statement to be true we must require that the pure Nash equilibrium is *strict*, i.e., every player strictly loses when deviating from his equilibrium action.

Algorithm 2 Strict saddle

```

procedure StrictSaddle( $\Gamma$ )
  for all  $S^0 = (\{s_1\}, \{s_2\}, \dots, \{s_n\}) \subseteq A$  do
     $C \leftarrow C \cup \text{minGSP}(\Gamma, S^0)$ 
  end for
  return  $\{S \in C : S \text{ is inclusion-minimal in } C\}$ 

```

Call Algorithm 1 for every possible combination of singleton sets of actions of the different players. The result is a set of GSPs, and the strict saddles of the game are the inclusion-minimal elements of this set. Algorithm 2 implements this idea.

THEOREM 1. *All strict saddles of an n -player game can be computed in polynomial time.*

PROOF. We show that Algorithm 2 computes all strict saddles of game Γ and runs in time polynomial in the size of Γ . Correctness follows from Lemma 1. For the running time, observe that there are $|A| = \prod_{i=1}^n |A_i|$ calls to Algorithm 1, which clearly is polynomial in the size of Γ . Polynomial running time of Algorithm 2 now follows directly from the fact that at least one action is added in every iteration of Algorithm 1. \square

Observe that Theorem 1 directly implies that the number of strict saddles of a game is at most the number of action profiles.

5. WEAK SADDLE

The computation of weak saddles turns out to be significantly more complicated than that of strict saddles. Somewhat surprisingly, this even holds for matrix games. In particular, Mills' algorithm does not easily generalize to weak saddles. When it is invoked on the game Γ_M from Section 3, for example, and initialized with $S^0 = (\{r_1\}, \{c_2\})$, it might add action r_2 in the next step, which is not contained in a minimal WGSP containing S_0 . While the computational complexity of weak saddles of general matrix games remains an open problem, we propose a polynomial-time algorithm for finding the weak saddle in a subclass of *symmetric* matrix games that possess a unique weak saddle. Furthermore, we give evidence for the computational intractability of weak saddles in bimatrix games.

5.1 Confrontation Games

Duggan and Le Breton [12] have put forward a subclass of symmetric matrix games that is characterized by the fact that the two players get the same payoff if and only if they play the same action. Otherwise there will always be a winner and a loser, and the outcome would be reversed if players were to exchange actions. We therefore call these games *confrontation games*. Since this section is concerned exclusively with symmetric games, in which all players have the same set of actions, we slightly deviate from the notation used in the rest of the paper and denote this set by A for notational convenience.

DEFINITION 4. *Let $\Gamma = \Gamma_M$ be a symmetric matrix game, and denote by A the set of actions of Γ . Then, Γ is called confrontation game if for all $a, b \in A$, $M(a, b) = 0$ if and only if $a = b$.²*

²Duggan and Le Breton [12] refer to this property as the *off-diagonal property*.

Duggan and Le Breton [12] have shown that confrontation games have a unique weak saddle $S = (S_1, S_2)$, and that this weak saddle is symmetric, i.e., $S_1 = S_2$. In the following, we denote by $WS(\Gamma)$ the weak saddle of a confrontation game Γ . We proceed to show that $WS(\Gamma)$ can be computed in polynomial time. To this end, we leverage the concept of *quasi-strict* equilibrium proposed by Harsanyi [14], which refines the Nash equilibrium concept by requiring that actions played with positive probability must yield a *strictly* higher payoff than actions played with probability zero.

DEFINITION 5. *Let $\Gamma = (N, (A_i)_{i \in N}, (p_i)_{i \in N})$ be a normal-form game. A Nash equilibrium $s = (s_1, s_2, \dots, s_n)$ is called quasi-strict if for all players $i \in N$ and all $a, b \in A_i$ with $s_i(a) > 0$ and $s_i(b) = 0$, $p_i(a, s_{-i}) > p_i(b, s_{-i})$.*

Quasi-strict equilibrium is a very natural concept in that it requires *all* best responses to be played with positive probability. Brandt and Fischer [4] have shown that quasi-strict equilibria in matrix games have a unique support, and can be found efficiently by linear programming. The unique support in a *symmetric* matrix game Γ is the same for both players, and will henceforth be denoted by $QS(\Gamma)$.

The following lemma establishes that $QS(\Gamma)$, and thus the support of any Nash equilibrium, is contained in $WS(\Gamma)$ if Γ is a confrontation game. The proof is adapted from Dutta and Laslier [13], who show a slightly more general statement in the context of tournament solutions.

LEMMA 2. *Let Γ be a confrontation game. Then, $QS(\Gamma) \subseteq WS(\Gamma)$.*

PROOF. Let p be the payoff function of player 1 and A the set of actions available to the players. Denote by $N(\Gamma)$ the set of Nash equilibrium strategies of Γ . Since the set of equilibria of a matrix game is convex, it suffices to restrict attention to symmetric equilibria, i.e.,

$$N(\Gamma) = \{s \in \Delta(A) : (s, s) \text{ is a Nash equilibrium of } \Gamma\}.$$

For an action $a \in A$ and a strategy $s \in \Delta(A)$, denote by $p(a, s)$ the expected payoff from a if the opponent plays s . The proof then relies on the following three facts:

- (i) The support of a quasi-strict equilibrium contains exactly those actions that are played with positive probability in *some* Nash equilibrium, i.e.,

$$QS(\Gamma) = \{a \in A : s(a) > 0 \text{ for some } s \in N(\Gamma)\}$$

- (ii) $QS(\Gamma) = \{a \in A : p(a, s) = 0 \text{ for all } s \in N(\Gamma)\}$

- (iii) $N(\Gamma|_S) \subseteq N(\Gamma)$, where $S = WS(\Gamma)$.

(i) and (ii) were shown by Brandt and Fischer [4] and Dutta and Laslier [13], respectively. For (iii), let $s \in N(\Gamma|_S)$. In order to establish that s is a Nash equilibrium of Γ , it suffices to show that $p(a, s) \leq 0$ for all actions $a \in A$. This is obvious for actions in S , since s is a Nash equilibrium in $\Gamma|_S$. Thus consider an action $a \in A \setminus S$. Since $S = WS(\Gamma)$, there exists $a^* \in S$ with $a^* \geq_s a$. Since s places positive probability only on actions in S , it follows that $p(a, s) \leq p(a^*, s) \leq 0$, as desired.

We now show that $QS(\Gamma) \subseteq WS(\Gamma)$. Assume for contradiction that there exists an action a that is contained in $QS(\Gamma)$ but not in $S = WS(\Gamma)$. Since S is the weak saddle of Γ , there exists some $a^* \in S$ such that $a^* \geq_s a$. We distinguish two different cases:

Algorithm 3 Weak saddle of a confrontation game

procedure WeakSaddle(Γ) $S \leftarrow QS(A)$ **repeat** $A' \leftarrow \{a \in A \setminus S : \nexists s \in S \text{ with } s \gg_S a\}$ $S \leftarrow S \cup QS(A')$ **until** $A' = \emptyset$ **return** (S, S)

If $a^* \in QS(\Gamma|_S)$, consider a Nash equilibrium strategy $s \in N(\Gamma|_S)$ of $\Gamma|_S$ in which a^* is played with a positive probability. Such an equilibrium is guaranteed to exist by (i). Since Γ is a confrontation game, $p(a, a^*) \neq 0 = p(a^*, a^*)$, and thus $p(a, s) < p(a^*, s)$. By (iii) s is also a Nash equilibrium of Γ , and thus $p(a, s) < p(a^*, s) = 0$, which together with (ii) contradicts the assumption that $a \in QS(\Gamma)$.

If, on the other hand, $a^* \notin QS(\Gamma|_S)$, there has to be some $s \in N(\Gamma|_S)$ with $p(a, s) \leq p(a^*, s) < 0$, leading to the same contradiction as above. \square

We are now ready to describe Algorithm 3 for computing the weak saddle of a confrontation game. It is similar in spirit to Mills' algorithm in that it starts with a subset of the set to be computed, in this case with $QS(A)$, and iteratively adds actions that are not yet dominated. In contrast to the strict saddle, however, it is no longer obvious which actions to choose, because an action that is currently undominated might become dominated later on for a larger set of actions of the other player. As we will see, the latter can not happen for actions in the weak saddle of the *subgame* induced by the undominated actions. Since a non-empty subset of the weak saddle of any game can be found efficiently, this completes the algorithm.³

More formally, let Γ be a confrontation game with action set A . For a subset $C \subseteq A$, denote by $\Gamma|_C$ the induced subgame with action set C . Obviously, $\Gamma|_C$ is also a confrontation game. For notational convenience, we sometimes identify $\Gamma|_C$ and C , and write $QS(A) = QS(\Gamma)$ and $WS(A) = WS(\Gamma)$. The following is our key lemma.

LEMMA 3. *Let S be a subset of $WS(A)$, A' is the subset of actions that are not weakly dominated by S , i.e.,*

$$A' = \{a \in A \setminus S : \nexists s \in S \text{ with } s \gg_S a\}.$$

Then $WS(\Gamma|_{A'}) \subseteq WS(\Gamma)$.

PROOF. In order to prove the lemma, we first make the following observations. Let Γ be a confrontation game with actions A . Further let C_1 and C_2 be nonempty subsets of A , and $x, y \in A$. Then the following holds:

(i) if $x >_{C_1} y$ and $C_2 \subseteq C_1$ with $C_2 \cap \{x, y\} \neq \emptyset$, then $x >_{C_2} y$; and

(ii) if $x >_{C_1} y$, $y >_{C_2} z$, and $x \in C_1 \cap C_2$, then $x >_{C_1 \cap C_2} z$.

We can assume that A' is nonempty, since otherwise $WS(\Gamma|_{A'})$ is empty and there is nothing to prove.

Now, partition A' , the set of undominated elements, into two sets $C = A' \cap WS(A)$ and $C' = A' \setminus WS(A)$ of elements

³The same idea was used in an algorithm by Brandt and Fischer [5] to compute the minimal bidirectional covering set of an oriented graph.

contained in $WS(A)$ and elements not contained in $WS(A)$. We will show that C is a WGSP of the game $\Gamma|_{A'}$. This implies $WS(A') \subseteq C \subseteq WS(A)$, because $WS(A')$ is contained in every WGSP of $\Gamma|_{A'}$.

It suffices to show that C is a WGSP of $\Gamma|_{A'}$, i.e., that for all $y \in C' = A' \setminus C$, there exists $x \in C$ such that $x >_{A'} y$. Let $y \in C'$. Since $y \notin WS(A)$, there has to be some $x \in WS(A)$ that dominates y with respect to $WS(A)$, i.e., $x >_{WS(A)} y$. It is easy to see that $x \notin S$, since otherwise (i) would imply that $x >_S y$, contradicting the assumption that $y \in A'$. On the other hand, assume that $x \in WS(A) \setminus (S \cup C)$. Then there is some $s \in S$ such that $s >_S x$. However, according to (ii), $s >_S x$ and $x >_{WS(A)} y$ imply $s >_S y$, again contradicting the assumption that $y \in A'$. Thus $x \in C$, and using (i) again, $x >_{WS(A)} y$ and $y \in A'$ imply $x >_{A'} y$. Hence C is a WGSP of $\Gamma|_{A'}$. \square

THEOREM 2. *The weak saddle of a confrontation game can be computed in polynomial time.*

PROOF. We prove that Algorithm 3 computes the weak saddle and runs in time polynomial in the size of the game.

In each iteration, at least one action is added to the set S , so the algorithm is guaranteed to terminate after at most $|A|$ iterations. Each iteration consists of a single call to QS , which requires only polynomial time as shown by Brandt and Fischer [4].

As for correctness, we show by induction on the number of iterations that $S \subseteq WS(A)$ holds at any time. When the algorithm terminates, S is a WGSP, which, together with the induction hypothesis, implies that $S = WS(A)$. The base case follows directly from Lemma 2, i.e., from the fact that $QS(A) \subseteq WS(A)$. Now assume that $S \subseteq WS(A)$ at the beginning of a particular iteration. Then $S \cup QS(A') \subseteq S \cup WS(A') \subseteq WS(A)$, where the first inclusion is due to Lemma 2 and the second inclusion follows from Lemma 3 and the induction hypothesis. \square

In the remainder of this section, we present a family of symmetric matrix games that are not confrontation games and have an exponential number of weak saddles. Define two matrices D and $\mathbf{1}$ as

$$D = \begin{pmatrix} 0 & 1 & 0 & 0 \\ -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & -1 & 0 \end{pmatrix} \quad \text{and} \quad \mathbf{1} = \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{pmatrix}.$$

It is easily verified that each of the following pairs is a weak saddle of Γ_D : $(\{r_1, r_2\}, \{c_1, c_2\})$, $(\{r_3, r_4\}, \{c_3, c_4\})$, $(\{r_1, r_3\}, \{c_1, c_3\})$, $(\{r_2, r_3\}, \{c_1, c_4\})$, and $(\{r_1, r_4\}, \{c_2, c_3\})$.

For an odd integer $k \geq 1$, define M_k as the block matrix whose diagonal blocks are D and whose remaining blocks are arranged in a checker-board pattern consisting of $\mathbf{1}$ s and $-\mathbf{1}$ s, i.e.,

$$M_k = \begin{pmatrix} D & -\mathbf{1} & \mathbf{1} & -\mathbf{1} & \cdots & \mathbf{1} \\ \mathbf{1} & D & -\mathbf{1} & \mathbf{1} & & -\mathbf{1} \\ -\mathbf{1} & \mathbf{1} & D & -\mathbf{1} & & \mathbf{1} \\ \mathbf{1} & -\mathbf{1} & \mathbf{1} & D & & -\mathbf{1} \\ \vdots & & & & \ddots & \vdots \\ -\mathbf{1} & \mathbf{1} & -\mathbf{1} & \mathbf{1} & \cdots & D \end{pmatrix}.$$

For any ordered multiset of k weak saddles of Γ_D , consider the sets of rows and columns of M_k containing for each $i \leq k$,

the rows and columns of M_k obtained by identifying the i th weak saddle in the set in the i th diagonal block of M_k . We leave it to the reader to verify that the latter forms a weak saddle of Γ_{M_k} , such that total number of weak saddles of Γ_{M_k} is at least 5^k . An immediate consequence of this example is that computing *all* weak saddles of a game requires exponential time in the worst case, even for matrix games.

5.2 Bimatrix Games

In this section, we establish a relationship between weak saddles of bimatrix games and inclusion-maximal cliques of undirected graphs. Our construction is inspired by McLennan and Tourky [17] and will be used to derive results concerning the computational hardness of weak saddles.

Let $G = (V, E)$ be an undirected graph. A *clique* in a graph G is a subset $C \subseteq V$ of vertices such that $(i, j) \in E$ for all $i, j \in C$. Define the bimatrix game Γ_G where both players have V as their set of actions and payoffs are given by

$$p_1(i, j) = \begin{cases} 1 & \text{if } \{i, j\} \in E \\ 0 & \text{if } i = j \\ -1 & \text{otherwise,} \end{cases} \quad p_2(i, j) = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{otherwise.} \end{cases}$$

THEOREM 3. *A pair (S_1, S_2) is a weak saddle in Γ_G if and only if $S_1 = S_2$ and S_1 is an inclusion-maximal clique in G .*

The proof consists of three lemmas. Recall that a WGSP is a pair of subsets of V that is externally stable for both players. For $v \in V$ and $S \subseteq V$, define $p_i(v, S) = (p_i(v, s))_{s \in S}$ as the vector of payoffs for player i if he plays v and the other player plays some $s \in S$.

LEMMA 4. *(S_1, S_2) is externally stable for player 2 if and only if $\emptyset \neq S_1 \subseteq S_2$.*

The proof of Lemma 4 is straightforward and omitted due to space constraints.

LEMMA 5. *If S is a maximal clique in G , then (S, S) is a weak saddle in Γ_G .*

PROOF. We have to show that (S, S) is a WGSP, i.e., externally stable for both players, and that there is no WGSP strictly contained in (S, S) .

External stability for player 2 follows from Lemma 4. For external stability for player 1, consider any $v \in V \setminus S$. Since S is a maximal clique, there must exist some $s \in S$ with $(s, v) \notin E$ or, equivalently, $p_1(s, v) = -1$. Then, $s \succ_S v$ because $p_1(s, S) = (1, \dots, 1, 0, 1, \dots, 1)$ with entry 0 at position s .

Now assume for contradiction that there exists a WGSP (S'_1, S'_2) with $S'_1 \subseteq S$ and $S'_2 \subseteq S$, such that at least one inclusion is strict. By Lemma 4, $S'_1 \subseteq S'_2$, which means that S'_1 must be a strict subset of S_1 , because otherwise $(S'_1, S'_2) = (S, S)$. Consider some $s \in S \setminus S'_1$. Since (S'_1, S'_2) is a WGSP, there must exist some $s^* \in S'_1$ with $s^* \succ_{T'} s$. This is a contradiction, since $s^* \in S'_2$ and $p_1(s^*, s^*) = 0 < 1 = p_1(s, s^*)$, where the last equality is due to the fact that both s and s^* are in the clique S . \square

LEMMA 6. *If (S_1, S_2) is a weak saddle of Γ_G , then $S_1 = S_2$ and S_1 is a maximal clique in G .*

PROOF. Let (S_1, S_2) be a weak saddle in Γ_G . Let C be an inclusion-maximal clique in the induced subgraph $G|_{S_1}$ of G with vertex set S_1 . We claim that C is also inclusion-maximal in G .

Assume for contradiction that there exists some $v \in V \setminus C$ that is connected to every vertex in C , i.e., $p_1(v, C) = (1, \dots, 1)$. Since (S_1, S_2) is a weak saddle, there exists $s \in S_1$ with $s \succ_{S_2} v$. In particular, $p_1(s, C) = (1, \dots, 1)$, implying that $s \notin C$ and that s is connected to all vertices in C . This obviously contradicts the assumption that C is an inclusion-maximal clique in G .

Thus, C is a maximal clique in G and Lemma 5 implies that (C, C) is a weak saddle. Furthermore, by Lemma 4, $S_1 \subseteq S_2$. From the inclusion-minimality of saddles and from $C \subseteq S_1 \subseteq S_2$, we conclude that $(S_1, S_2) = (C, C)$. \square

This completes the proof of Theorem 3. The main result of this section now follows as a corollary.

COROLLARY 1. *Deciding whether there exists a weak saddle with any of the following properties is NP-hard, even in bimatrix games with only three different payoffs:*

- at most k actions for some player,
- at least k actions for some player, or
- an average payoff of at least p for a particular player.

PROOF SKETCH. It is not hard to see that the first problem is equivalent to the second one under polynomial-time Turing reductions, which in turn is equivalent to the problem of deciding the existence of a clique of size at least k in an undirected graph. NP-hardness of the former under polynomial-time *many-one* reductions can be shown *via* a reduction from the *exact cover* problem, which we omit due to space constraints. Hardness of the third problem follows by observing that the average payoff of player 1 in our construction is a strictly increasing function of the size of the weak saddle. \square

6. MIXED REFINEMENTS OF SADDLES

Duggan and Le Breton [11] introduce refinements of Shapley's saddles, motivated by the possibility that players may use *randomized* strategies. For an action to be excluded from a mixed saddle, it suffices to find a mixture of saddle actions that dominates it.

DEFINITION 6 (STRICT MIXED SADDLE). *A Mixed Generalized Saddle Point (MGSP) of the game $(N, (A_i)_{i \in N}, (p_i)_{i \in N})$ is a tuple $(S_1, S_2, \dots, S_n) \subseteq A$, such that for each player $i \in N$,*

$$\forall a_i \in A_i \setminus S_i, \exists s_i \in \Delta(S_i) \text{ such that } s_i \gg_{S_{-i}} a_i.$$

A strict mixed saddle is a MGSP that contains no other MGSP.

Weak mixed generalized saddle points and weak mixed saddles are defined analogously, replacing strict by weak domination. Unlike strict and weak saddles, mixed saddles are not ordinal solution concepts. They are, however, invariant under positive affine transformations of the payoff functions and we can therefore restrict our attention to games in which all payoffs are positive.

Algorithm 4 Minimal MGSP

```

procedure minMGSP( $\Gamma, (S_1^0, S_2^0, \dots, S_n^0)$ )
  for all  $i \in N$  do
     $S_i \leftarrow S_i^0$ 
  end for
  repeat
    for all  $i \in N$  do
       $A'_i \leftarrow \{a_i \in A_i \setminus S_i : \nexists s_i \in \Delta(A_i) \text{ with } s_i \gg_{S_{-i}} a_i\}$ 
       $S_i \leftarrow S_i \cup A'_i$ 
    end for
  until  $\bigcup_{i=0}^n A'_i = \emptyset$ 
  return  $(S_1, S_2, \dots, S_n)$ 

```

6.1 Strict Mixed Saddle

Since every strict saddle contains a strict *mixed* saddle, strict mixed saddles are not unique in non-zero-sum games. Nevertheless, we present an algorithm that computes all strict mixed saddles of an arbitrary n -player game. Similar to Algorithm 2 in Section 4, we use as a subroutine an algorithm that computes the minimal MGSP that contains a given subset.

LEMMA 7. *Algorithm 4 computes the inclusion-minimal MGSP containing a given input set S^0 .*

PROOF. The following geometric interpretation will be useful. For an action a_i of player $i \in N$, define $p_i(a_i, S_{-i}) = (p_i(a_i, s_{-i}))_{s_{-i} \in S_{-i}}$ as the vector of possible payoffs for player i if he plays a_i and the other players play some $s_{-i} \in S_{-i}$. For a set $B_i \subseteq A_i$ of actions of player i , denote by $p_i(B_i, S_{-i}) = \cup_{b_i \in B_i} p_i(b_i, S_{-i})$ the union of all such vectors, and write $m = |S_{-i}|$ for their dimension. For a set of vectors $V \subseteq \mathbb{R}_{\geq 0}^m$, define $L(V)$ to be the *lower contour set of conv(V)*, i.e.,

$$L(V) = \bigcup \{x \in \mathbb{R}_{\geq 0}^m : \exists v \in \text{conv}(V) \text{ with } v \geq x\},$$

where $v \geq x$ is to be read componentwise.

The underlying intuition is that each action whose vector of payoffs lies in the *interior* of $L(V)$ is strictly dominated by some strategy in $\Delta(V)$. More formally, a_i is strictly dominated by S_i with respect to S_{-i} if and only if $p_i(a_i, S_{-i}) \in L(p_i(S_i, S_{-i}))$.

Let S^{\min} be the minimal MGSP containing S^0 . It suffices to show that (i) during the execution of Algorithm 4, the set S is always a subset of S^{\min} , and that (ii) upon termination of the algorithm, S is a MGSP.

For (i), perform an induction on the size of S . Initially, $S = S^0 \subseteq S^{\min}$ by definition of S^{\min} . Now assume that $S \subseteq S^{\min}$ at the beginning of a particular iteration. We have to show that for all $i \in N$, $A'_i \subseteq S_i^{\min}$. Let $a \in A'_i$ and assume for contradiction that $a \notin S_i^{\min}$. Since S^{\min} is an MGSP, there exists some $a^* \in \Delta(S_i^{\min}) \subseteq \Delta(A_i)$ with $a^* \gg_{S_{-i}^{\min}} a$.

By the induction hypothesis, $S_{-i} \subseteq S_{-i}^{\min}$, which in turn implies $a^* \gg_{S_{-i}} a$. This contradicts the assumption that $a \in A'_i$.

For (ii), observe that upon termination of the algorithm, $\bigcup_{i=0}^n A'_i = \emptyset$, and thus $A'_i = \emptyset$ for all $i \in N$. We need to show that S is a MGSP, i.e., that for all $i \in N$ and for all $a_i \in A_i \setminus S_i$, there exists $s_i \in \Delta(S_i)$ with $s_i \gg_{S_{-i}} a_i$. Since $A'_i = \emptyset$, we know that there must be some $s_i \in \Delta(A_i)$ with $s_i \gg_{S_{-i}} a_i$. It thus suffices to show that $L(p_i(S_i, S_{-i})) = L(p_i(A_i, S_{-i}))$, such that an action is strictly dominated by

Algorithm 5 Strict Mixed Saddle

```

procedure StrictMixedSaddle( $\Gamma$ )
  for all  $S^0 = (\{s_1\}, \{s_2\}, \dots, \{s_n\}) \subseteq A$  do
     $C \leftarrow C \cup \text{minMGSP}(\Gamma, S^0)$ 
  end for
  return  $\{S \in C : S \text{ is inclusion-minimal in } C\}$ 

```

some strategy in $\Delta(S_i)$ if and only if it is strictly dominated by some strategy in $\Delta(A_i)$.

The inclusion from right to left is trivial since $S_i \subseteq A_i$. For the inclusion from left to right, recall *Minkowski's Theorem*, which states that a convex and compact set in \mathbb{R}^m is equal to the convex hull of the set of its extreme points. As both $L(p_i(A_i, S_{-i}))$ and $L(p_i(S_i, S_{-i}))$ are compact and convex, it remains to be shown that no point in $p_i(A_i \setminus S_i, S_{-i})$ is an extreme point of $L(p_i(A_i, S_{-i}))$. This follows from the fact that $A'_i = \emptyset$, which means that for all $a_i \in A_i \setminus S_i$, there exists $a_i^* \in \Delta(A_i)$ with $a_i^* \gg_{S_{-i}} a_i$. \square

Whenever S^0 is contained in a strict mixed saddle, Algorithm 4 returns a strict mixed saddle. If we call Algorithm 4 for every possible combination of singleton sets of actions of the different players, we get as a result a set of MGSPs. The strict mixed saddles of the game are the inclusion-minimal elements of this set. We thus obtain the main result of this section.

THEOREM 4. *All strict mixed saddles of an n -player game can be computed in polynomial time.*

PROOF. We show that Algorithm 5 computes all strict mixed saddles of an n -player game Γ and runs in time polynomial in the size of Γ . Correctness follows from Lemma 7. Concerning time complexity, observe that the number of calls to Algorithm 4 is $|A| = \prod_{i=1}^n |A_i|$, which obviously is polynomial in the size of the game. Furthermore, at least one action is added in every iteration of Algorithm 4, and each iteration takes only polynomial time because the set of undominated actions can be computed efficiently by linear programming (see, e.g., Proposition 1 by Conitzer and Sandholm [8]). \square

6.2 Weak Mixed Saddle

It turns out that some of the results we obtained for weak saddles can be extended to weak *mixed* saddles. For example, in confrontation games where payoffs are restricted to $\{-1, 0, 1\}$, the possibility of mixing does not affect the set of dominated actions. As a consequence, the weak mixed saddle and the weak saddle coincide in such games. In general confrontation games, it is still true that a subset of a weak mixed saddle can be found efficiently, namely the *sign essential set* introduced by Dutta and Laslier [13]. Whether this property can be used to efficiently construct a weak mixed saddle remains an open problem.

On the other hand, it can be shown that weak mixed saddles and weak saddles coincide in all games Γ_G used in Section 5.2. All hardness results for weak saddles in bimatrix games thus also apply to weak mixed saddles.

7. CONCLUSION

We have initiated the study of computational aspects of Shapley's saddles – ordinal set-valued solution concepts dating back to the early 1950s – by proposing polynomial-time

algorithms for computing pure and mixed strict saddles in general normal-form games and pure weak saddles in a subclass of symmetric two-player zero-sum games. The latter algorithm is highly non-trivial and surprisingly relies on linear programs that determine the support of Nash equilibria in certain subgames of the original game. We also showed that, in general bimatrix games, natural problems associated with weak (pure or mixed) saddles, such as deciding the existence of a weak saddle with at most k actions for some player, are NP-hard. Several open questions with respect to weak saddles remain. In particular, it is not known whether weak saddles can be computed efficiently in general two-player zero-sum games. Furthermore, the aforementioned NP-hardness results do not imply that finding an arbitrary weak saddle is NP-hard.

All of our results apply to games with a constant number of players and many actions. It is an interesting question whether strict saddles can still be computed efficiently in certain classes of games that allow for a compact representation when the number of players is unbounded. Similarly, one might ask for compact classes of games where computing weak saddles becomes tractable.

8. ACKNOWLEDGMENTS

This material is based upon work supported by the Deutsche Forschungsgemeinschaft under grants BR-2312/6-1 (within the European Science Foundation's EUROCORES program LogICCC) and BR 2312/3-2. We further thank Jan Hoffmann for helpful discussions on the hardness of weak saddles.

References

1. R. J. Aumann. Game theory. In J. Eatwell, M. Milgate, and P. Newman, editors, *The New Palgrave: A Dictionary of Economics*, volume 2, pages 460–482. MacMillan, 1987.
2. K. Basu and J. Weibull. Strategy subsets closed under rational behavior. *Economics Letters*, 36:141–146, 1991.
3. M. Benisch, G. B. Davis, and T. Sandholm. Algorithms for rationalizability and CURB sets. In *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI)*, pages 598–604. AAAI Press, 2006.
4. F. Brandt and F. Fischer. On the hardness and existence of quasi-strict equilibria. In B. Monien and U.-P. Schroeder, editors, *Proceedings of the 1st International Symposium on Algorithmic Game Theory (SAGT)*, volume 4997 of *Lecture Notes in Computer Science (LNCS)*, pages 291–302. Springer-Verlag, 2008.
5. F. Brandt and F. Fischer. Computing the minimal covering set. *Mathematical Social Sciences*, 56(2):254–268, 2008.
6. F. Brandt, F. Fischer, P. Harrenstein, and Y. Shoham. Ranking games. *Artificial Intelligence*, 173(2):221–239, 2009.
7. X. Chen and X. Deng. Settling the complexity of 2-player Nash-equilibrium. In *Proceedings of the 47th Symposium on Foundations of Computer Science (FOCS)*, pages 261–272. IEEE Press, 2006.
8. 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.
9. V. Conitzer and T. Sandholm. A generalized strategy eliminability criterion and computational methods for applying it. In *Proceedings of the 20th National Conference on Artificial Intelligence (AAAI)*, pages 483–488. AAAI Press, 2005.
10. C. Daskalakis, P. Goldberg, and C. Papadimitriou. The complexity of computing a Nash equilibrium. In *Proceedings of the 38th Annual ACM Symposium on the Theory of Computing (STOC)*, pages 71–78. ACM Press, 2006.
11. J. Duggan and M. Le Breton. Mixed refinements of Shapley's saddles and weak tournaments. *Social Choice and Welfare*, 18(1):65–78, 2001.
12. J. Duggan and M. Le Breton. Dutta's minimal covering set and Shapley's saddles. *Journal of Economic Theory*, 70:257–265, 1996.
13. B. Dutta and J.-F. Laslier. Comparison functions and choice correspondences. *Social Choice and Welfare*, 16(4):513–532, 1999.
14. J. C. Harsanyi. Oddness of the number of equilibrium points: A new proof. *International Journal of Game Theory*, 2:235–250, 1973.
15. D. E. Knuth, C. H. Papadimitriou, and J. N. Tsitsiklis. A note on strategy elimination in bimatrix games. *Operations Research Letters*, 7:103–107, 1988.
16. R. D. Luce and H. Raiffa. *Games and Decisions: Introduction and Critical Survey*. John Wiley & Sons Inc., 1957.
17. A. McLennan and R. Tourky. Simple complexity from imitation games. Unpublished manuscript, 2005.
18. J. Nash. Equilibrium points in n -person games. *Proceedings of the National Academy of Sciences (PNAS)*, 36:48–49, 1950.
19. J. F. Nash. Non-cooperative games. *Annals of Mathematics*, 54(2):286–295, 1951.
20. M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, 1994.
21. L. Shapley. Order matrices. I. Technical Report RM-1142, The RAND Corporation, 1953.
22. L. Shapley. Order matrices. II. Technical Report RM-1145, The RAND Corporation, 1953.
23. L. Shapley. A condition for the existence of saddle-points. Technical Report RM-1598, The RAND Corporation, 1955.
24. 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.