

Computing Dominance-Based Solution Concepts

Felix Brandt
TU München
85748 Munich, Germany
brandtf@in.tum.de

Markus Brill
University of Oxford
Oxford OX1 3QD, UK
mbrill@cs.ox.ac.uk

Two common criticisms of Nash equilibrium are its dependence on very demanding epistemic assumptions and its computational intractability. We study the computational properties of less demanding set-valued solution concepts that are based on varying notions of dominance. These concepts are intuitively appealing, always exist, and admit unique minimal solutions in important subclasses of games. Examples include Shapley’s saddles, Harsanyi and Selten’s primitive formations, Basu and Weibull’s CURB sets, and Dutta and Laslier’s minimal covering set. Based on a unifying framework proposed by Duggan and Le Breton, we formulate two generic algorithms for computing these concepts and investigate for which classes of games and which properties of the underlying dominance notion the algorithms are sound and efficient. We identify two sets of conditions that are sufficient for polynomial-time computability and show that the conditions are satisfied, for instance, by saddles and primitive formations in normal-form games, minimal CURB sets in two-player games, and the minimal covering set in symmetric matrix games. Our positive algorithmic results explain regularities observed in the literature, but also apply to several solution concepts whose computational complexity was unknown.

1 Introduction

Saddle points, i.e., combinations of actions such that no player can gain by deviating, are among the earliest solution concepts considered in game theory (see, e.g., von Neumann and Morgenstern, 1944). In two-player zero-sum games (henceforth *matrix games*), every saddle point happens to coincide with the 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, Borel (1921) introduced *mixed*—i.e., randomized—strategies and von Neumann (1928) 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 normal-form games by Nash (1951), at the expense of its normative foundation. Since then, Nash equilibrium has commonly been criticized for resting on very demanding epistemic assumptions.¹

Shapley (1953a,b) showed that the 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.² Shapley defines a *generalized saddle point (GSP)* to be a pair of subsets of actions for each player that satisfies a simple external stability condition: Every action not contained in a player’s subset is dominated by some action in the set, given that the other player chooses actions from his set. A GSP is minimal if it does not contain another GSP. Minimal GSPs, which Shapley calls *saddles*, also satisfy internal stability in the sense that no two actions within a set dominate each other, given that the other player chooses actions from his set. While Shapley was the first to conceive GSPs, he was not the only one. Apparently unaware of Shapley’s work, Samuelson (1992) 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 (1976) in the context of spatial voting games and the *minimal covering set* as defined by Dutta (1988) in the context of majority tournaments are GSPs (Duggan and Le Breton, 1996). In a regrettably unpublished paper, Duggan and Le Breton (2014)³ extend Shapley’s approach to normal-form games and define a *D-solution* as a tuple of sets that is internally and externally stable with respect to a so-called *dominance structure D*. Depending on *D*, a number of different solution concepts can be defined. The framework is rich enough to not only cover Shapley’s saddles, but also other common set-valued solution concepts such as *rationalizability* (Bernheim, 1984; Pearce, 1984) and *CURB sets* (Basu and Weibull, 1991); see Section 3.2 for details.

We are mainly interested in *D-cores*, which are (*inclusion-*)*minimal D-solutions*. For the case of *strict dominance (S)*, Shapley (1964) showed that every matrix game admits a *unique S-core*. Duggan and Le Breton (2014) extend this uniqueness result to other dominance structures and to a larger class of games by showing, among other things, that *equilibrium safe* games (a class of *n*-player games that includes matrix games) have a unique core with respect to strict dominance, *mixed strict dominance (S*)*, and *Börgers dominance (B)* (Börgers, 1993). Furthermore, Duggan and Le Breton (1996) proved uniqueness of the core with respect to *weak dominance (W)* and *very weak dominance (V)* in a subclass of

¹See, e.g., Luce and Raiffa (1957, pp. 74–76), Fishburn (1978), Bernheim (1984), Pearce (1984), Myerson (1991, pp. 88–91), Börgers (1993), Aumann and Brandenburger (1995), Perea (2007), Jungbauer and Ritzberger (2011), Barelli (2009) and Bach and Tsakas (2014).

²The main results of the 1953 reports later reappeared in revised form (Shapley, 1964).

³An earlier version of the paper by Duggan and Le Breton has been circulating since 1996 under the title “Dominance-based Solutions for Strategic Form Games.”

| | S | B | S^* | C_M | C_D | V | V^* |
|------------------------|--------|--------|--------|--------|--------|-----|-------|
| normal-form games | poly | poly | poly | | | | |
| matrix games | unique | unique | unique | | | | |
| symmetric matrix games | | | | unique | unique | exp | |
| confrontation games | | | | unique | | | exp |
| tournament games | | unique | | unique | | | |

Table 1: Summary of results. For a given dominance structure D and a class of games (ordered by set inclusion), the table shows bounds on the *asymptotic number* of D -cores (unique, polynomial, or exponential). If a cell is highlighted in dark gray, the greedy algorithm finds all D -cores in the given class in polynomial time. If it is highlighted in light gray, the analogous statement holds for the sophisticated algorithm. If a cell spans several columns, the corresponding D -cores coincide within the respective class of games. Covering (C_M) and deep covering (C_D) are only defined for symmetric matrix games.

symmetric matrix games that we refer to as confrontation games. While it is easy to see that a matrix game can have multiple V -cores, Brandt et al. (2016) showed that V -cores in matrix games are—just like pure and mixed saddle points—interchangeable and equivalent.

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 (Chen et al., 2009; Daskalakis et al., 2009) and thus unlikely to admit an efficient algorithm. This result has sparked the search for alternative, computationally tractable, solution concepts. Despite the fact that Shapley’s saddles were devised as early as 1953 (Shapley, 1953a,b) and are thus almost as old as Nash equilibrium (Nash, 1951), surprisingly little is known about their computational properties. Common notions of dominance have widely been studied from a computational perspective in the form of *iterated dominance* (Knuth et al., 1988; Gilboa et al., 1993; Conitzer and Sandholm, 2005; Brandt et al., 2011a). D -cores are “refinements” of iterated D dominance and *cannot* be found by iterated elimination of dominated actions.

In this paper, we propose two generic algorithms (a greedy and a sophisticated one) for computing D -cores and study their soundness and efficiency for various dominance structures D and subclasses of games. In addition to the dominance structures mentioned above, we study their mixed counterparts (denoted by D^* for a given dominance structure D) (Duggan and Le Breton, 2014), *covering* (C_M) (McKelvey, 1986; Dutta and Laslier,

1999), and *deep covering* (C_D) (Duggan, 2013). We then define abstract properties that, when satisfied by a dominance structure within a particular class of games, allow for our algorithms to be sound and efficient. Our results yield

- greedy algorithms for computing all S -cores (aka saddles), S^* -cores (aka primitive formations or mixed saddles and equivalent to CURB sets in two-player games), and B -cores of a given normal-form game, and
- sophisticated algorithms for computing the unique C_M -core and the unique C_D -core of a given symmetric matrix game. Within the subclass of confrontation games, these algorithms coincide and also yield the W -core and the V -core (aka weak saddle).

Our algorithms subsume existing algorithms for computing saddles in matrix games (Shapley, 1964), the (unique) minimal covering set in symmetric win-lose-tie games (Brandt and Fischer, 2008), and minimal CURB sets in two-player games (Benisch et al., 2010). Interestingly, the sophisticated algorithms rely on the repeated computation of Nash equilibria via linear programming, even though most of the corresponding solution concepts are purely ordinal. For the remaining combinations of dominance structures and classes of games, we show that these classes admit an exponential number of D -cores. This renders the computation of *all* D -cores infeasible.⁴ Our results are summarized in Table 1.

2 Preliminaries

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

A two-player game $\Gamma = (\{1, 2\}, (A_1, A_2), (u_1, u_2))$ is a *matrix game* (or *zero-sum game*) if $u_1(a_1, a_2) + u_2(a_1, a_2) = 0$ for all $(a_1, a_2) \in A_1 \times A_2$. It is *symmetric* if $A_1 = A_2$ and $u_1(a_1, a_2) = u_2(a_2, a_1)$ for all $(a_1, a_2) \in A_1 \times A_1$. Whenever we are concerned with symmetric matrix games, we slightly deviate from the notation used in the rest of the paper for notational convenience: The set $A_1 = A_2$ of actions is denoted by A , the utility function of player 1 is denoted by u , and the game itself is denoted by (A, u) . A symmetric matrix game can be conveniently represented by a skew-symmetric matrix containing the utilities of player 1. For a subset $B \subseteq A$ of actions, $\Gamma|_B$ denotes the *subgame* of $\Gamma = (A, u)$ restricted to B , i.e., $\Gamma|_B$ is the symmetric matrix game $(B, u|_{B \times B})$.

Confrontation games are symmetric matrix games characterized by the fact that the two players receive the same utility if and only if they play the same action (Duggan and Le Breton, 1996). Formally, a symmetric matrix game $\Gamma = (A, u)$ is a confrontation game if

⁴In the case of very weak dominance, it has also been shown that finding *some* V -core is computationally intractable in two-player games (Brandt et al., 2011b). Whether some V -core (or some V^* -core) can be computed efficiently in matrix games remains an open problem.

for all $a, b \in A$, $u(a, b) = 0$ if and only if $a = b$.⁵ A *win-lose-tie game* is a matrix game in which all payoffs are either -1 , 0 , or 1 . A confrontation game that is also a win-lose-tie game is called a *tournament game*.⁶ Tournament games generalize the well-known game of rock-paper-scissors and are surprisingly rich (see, e.g., Fisher and Ryan, 1992; Laffond et al., 1993; Fisher and Ryan, 1995; Fisher and Reeves, 1995). Many interesting game-theoretic phenomena already appear in tournament games. In fact, finding Nash equilibria in tournament games may be as hard as finding Nash equilibria in general zero-sum games.⁷

Let $\Delta(M)$ denote the set of all probability distributions over a finite set M . A (mixed) *strategy* of a player $i \in N$ is an element of $\Delta(A_i)$. Utility functions are extended to profiles of strategies in the usual way. A *Nash equilibrium* is a strategy profile such that no player can benefit by unilaterally deviating from his strategy. The *essential set* $ES(\Gamma)$ is the set of all actions that are played with positive probability in some Nash equilibrium of Γ (Dutta and Laslier, 1999). As every normal-form game contains a Nash equilibrium (Nash, 1951), the essential set is never empty.

3 Dominance-Based Solution Concepts

In this section, we formally define the dominance structures and solution concepts considered in this paper. Furthermore, we introduce a number of properties that will be critical for our algorithmic results.

3.1 Dominance Structures

The following notation will be used throughout the paper. Let A_N denote the n -tuple (A_1, \dots, A_n) containing all action sets. An n -tuple $X = (X_1, \dots, X_n)$ is said to be *nonempty*, denoted $X \neq \emptyset$, if $X_i \neq \emptyset$ for all $i \in N$. For a nonempty n -tuple $X = (X_1, \dots, X_n)$, we write $X \subseteq A_N$ if $X_i \subseteq A_i$ for all $i \in N$. To simplify the exposition, we will frequently abuse terminology and refer to an n -tuple $X \subseteq A_N$ as a “set.” For every player i , we furthermore let X_{-i} denote the set $\prod_{j \in N \setminus \{i\}} X_j$ of all opponent action profiles in which each opponent $j \in N \setminus \{i\}$ plays an action from $X_j \subseteq A_j$.

Consider a player $i \in N$. Whether an action (or a set of actions) in A_i dominates another action $a_i \in A_i$ naturally depends on which actions the other players have at their disposal. This is reflected in the following definition, in which a dominance structure D is defined as a mapping from sets X_{-i} of opponent action profiles to *dominance relations* $D(X_{-i})$ for player i . A dominance relation for a player relates *subsets of actions* of this

⁵Duggan and Le Breton (1996) refer to this property as the *off-diagonal property*.

⁶The term *tournament game* refers to the fact that such a game $\Gamma = (A, u)$ can be represented by a tournament graph with vertex set A and edge set $\{(a, b) : u(a, b) = 1\}$. In a similar fashion, a confrontation game can be represented by a *weighted* tournament graph.

⁷Brandt and Fischer (2008) pointed out that computing Nash equilibria in symmetric win-lose-tie games—which are slightly more general than tournament games—is P-complete (under log-space reductions) and therefore at least as hard as any problem in P.

player to individual actions of the same player, i.e., $D(X_{-i}) \subseteq 2^{A_i} \times A_i$. Intuitively, $(X_i, a_i) \in D(X_{-i})$ signifies that subset $X_i \subseteq A_i$ is “preferable” to action a_i given X_{-i} . We require that if X_i is preferable to a_i , then so is every superset of X_i .⁸

Definition 1. Let $\Gamma = (N, (A_i)_{i \in N}, (u_i)_{i \in N})$ be a game in normal form and $X \subseteq A_N$. For each player $i \in N$, a dominance structure D maps X_{-i} to a subset of $2^{A_i} \times A_i$ such that $(X_i, a_i) \in D(X_{-i})$ implies $(Y_i, a_i) \in D(X_{-i})$ for all Y_i with $X_i \subseteq Y_i \subseteq A_i$.

For $X_i \subseteq A_i$ and $a_i \in A_i$, we write $X_i D(X_{-i}) a_i$ if $(X_i, a_i) \in D(X_{-i})$. In this case, we say that X_i *D-dominates* a_i with respect to X_{-i} . If X_i consists of a single action x_i , we write $x_i D(X_{-i}) a_i$ instead of $\{x_i\} D(X_{-i}) a_i$ to avoid cluttered notation.

We go on to define the main dominance structures considered by Duggan and Le Breton (2001, 2014).⁹

Definition 2. Let $\Gamma = (N, (A_i)_{i \in N}, (u_i)_{i \in N})$ be a game in normal form and $i \in N$. Furthermore, let $X \subseteq A_N$ and $a_i \in A_i$.

- strict dominance (S): $X_i S(X_{-i}) a_i$ if there exists $x_i \in X_i$ with $u_i(x_i, x_{-i}) > u_i(a_i, x_{-i})$ for all $x_{-i} \in X_{-i}$.
- weak dominance (W): $X_i W(X_{-i}) a_i$ if there exists $x_i \in X_i$ with $u_i(x_i, x_{-i}) \geq u_i(a_i, x_{-i})$ for all $x_{-i} \in X_{-i}$ and the inequality is strict for at least one $x_{-i} \in X_{-i}$.
- very weak dominance (V): $X_i V(X_{-i}) a_i$ if there exists $x_i \in X_i$ with $u_i(x_i, x_{-i}) \geq u_i(a_i, x_{-i})$ for all $x_{-i} \in X_{-i}$.
- Börgers dominance (B): $X_i B(X_{-i}) a_i$ if $X_i W(Y_{-i}) a_i$ for all $\emptyset \neq Y_{-i} \subseteq X_{-i}$.¹⁰
- mixed strict dominance (S^*): $X_i S^*(X_{-i}) a_i$ if there exists $s_i \in \Delta(X_i)$ with $u_i(s_i, x_{-i}) > u_i(a_i, x_{-i})$ for all $x_{-i} \in X_{-i}$.
- mixed weak dominance (W^*): $X_i W^*(X_{-i}) a_i$ if there exists $s_i \in \Delta(X_i)$ with $u_i(s_i, x_{-i}) \geq u_i(a_i, x_{-i})$ for all $x_{-i} \in X_{-i}$ and the inequality is strict for at least one $x_{-i} \in X_{-i}$.
- mixed very weak dominance (V^*): $X_i V^*(X_{-i}) a_i$ if there exists $s_i \in \Delta(X_i)$ with $u_i(s_i, x_{-i}) \geq u_i(a_i, x_{-i})$ for all $x_{-i} \in X_{-i}$.

⁸We remark that our definition of a dominance structure differs from the one by Duggan and Le Breton (2014), who define dominance structures as mappings from X to binary relations on A_i , and *mixed* dominance structures as mappings from X to relations from $\Delta(A)$ to A . Importantly, for all dominance structures D considered in this paper, all definitions give rise to the same notion of a D -solution.

⁹Duggan and Le Breton (2014) use a different terminology: strict, weak, and very weak dominance are referred to as *Shapley*, *weak Shapley*, and *Nash* dominance, respectively.

¹⁰More formally, $X_i W(Y_{-i}) a_i$ needs to hold for all $Y \subseteq A_N$ with $\emptyset \neq Y_j \subseteq X_j$ for all $j \in N \setminus \{i\}$.

Börger's dominance has a mixed counterpart as well, requiring that $X_i W^*(Y_{-i}) a_i$ for all $Y_{-i} \subseteq X_{-i}$. However, mixed Börger's dominance coincides with mixed strict dominance (Duggan and Le Breton, 2014).

The following dominance structures are only defined for symmetric matrix games; we sometimes refer to them as *symmetric dominance structures*.

Definition 3. Let (A, u) be a symmetric matrix game, $X, Y \subseteq A$, and $a \in A$.

- covering (C_M): $X C_M(Y)$ if there exists $x \in X \cap Y$ with
 - $u(x, a) > 0$ and
 - $u(x, y) \geq u(a, y)$ for all $y \in Y$.
- deep covering (C_D): $X C_D(Y)$ if there exists $x \in X \cap Y$ with
 - $u(x, a) > 0$,
 - $u(x, y) \geq u(a, y)$ for all $y \in Y$, and
 - $u(x, y) > u(a, y)$ for all $y \in Y$ with $u(a, y) = 0$.

Covering was introduced by McKelvey (1986) and later generalized by Dutta and Laslier (1999), and deep covering is a generalization of a notion by Duggan (2013). Covering and deep covering are equivalent in confrontation games.

For two dominance structures D and D' , we write $D \subseteq D'$ and say that D is *coarser* than D' and that D' is *finer* than D , if $D(X_{-i}) \subseteq D'(X_{-i})$ for all $X \subseteq A_N$. The following relations follow immediately from the respective definitions:

$$S \subseteq B \subseteq W \subseteq V, \quad C_D \subseteq C_M, \quad \text{and} \quad D \subseteq D^* \text{ for all } D \in \{S, W, V\}.$$

3.2 D -Solutions and D -Cores

Generalizing a classic cooperative solution concept by von Neumann and Morgenstern (1944), a set of actions X can be said to be *stable* if it consists precisely of those alternatives not dominated by X (see also (Wilson, 1970)). This fixed-point characterization can be split into two conditions of internal and external stability: First, there should be no reason to restrict the selection by excluding some action from it; second, there should be an argument against each proposal to include an outside action into the selection.

Definition 4. Let D be a dominance structure and $\Gamma = (N, (A_i)_{i \in N}, (u_i)_{i \in N})$ a game in normal form. A tuple $X \subseteq A_N$ is a D -solution in Γ if for every $i \in N$,

$$X_i \setminus \{x_i\} D(X_{-i}) x_i \text{ for no } x_i \in X_i, \text{ and} \tag{1}$$

$$X_i D(X_{-i}) a_i \text{ for all } a_i \in A_i \setminus X_i. \tag{2}$$

We refer to (1) and (2) as *internal* and *external D -stability*, respectively. We are mainly interested in *inclusion-minimal D -solutions*. Following Duggan and Le Breton (2014), these will be called *D -cores*.

| | a_1 | a_2 | a_3 | a_4 | a_5 | a_6 | a_7 | a_8 |
|-------|-------|-------|-------|-------|-------|-------|-------|-------|
| a_1 | 0 | 2 | -2 | 1 | 2 | 3 | | |
| a_2 | -2 | 0 | 2 | 1 | 2 | | | |
| a_3 | 2 | -2 | 0 | 1 | 1 | | | |
| a_4 | -1 | -1 | -1 | 0 | 1 | | | |
| a_5 | -2 | -2 | -1 | -1 | 0 | | | |
| a_6 | -3 | | | | | 0 | 3 | -3 |
| a_7 | | | | | | -3 | 0 | 3 |
| a_8 | | | | | | 3 | -3 | 0 |

| | b_1 | b_2 | b_3 | b_4 | b_5 | b_6 | b_7 | b_8 |
|-------|-------|-------|-------|-------|-------|-------|-------|-------|
| b_1 | 0 | 0 | 0 | 0 | 1 | 1 | 1 | 1 |
| b_2 | 0 | 0 | 1 | -1 | 0 | 1 | 1 | -1 |
| b_3 | 0 | -1 | 0 | 1 | 1 | 1 | -1 | 1 |
| b_4 | 0 | 1 | -1 | 0 | 0 | -1 | 1 | 1 |
| b_5 | -1 | 0 | -1 | 0 | 0 | 1 | 1 | 1 |
| b_6 | -1 | -1 | -1 | 1 | -1 | 0 | -1 | 1 |
| b_7 | -1 | -1 | 1 | -1 | -1 | 1 | 0 | -1 |
| b_8 | -1 | 1 | -1 | -1 | -1 | -1 | 1 | 0 |

| D | D -core |
|-----------------|--|
| S | $(\{a_1, a_2, a_3, a_4, a_5\}, \{a_1, a_2, a_3, a_4, a_5\})$ |
| W, V, B | $(\{a_1, a_2, a_3, a_4\}, \{a_1, a_2, a_3, a_4\})$ |
| S^*, W^*, V^* | $(\{a_1, a_2, a_3\}, \{a_1, a_2, a_3\})$ |

| D | D -core |
|-------|--|
| C_D | $(\{b_1, b_2, b_3, b_4, b_5\}, \{b_1, b_2, b_3, b_4, b_5\})$ |
| C_M | $(\{b_1, b_2, b_3, b_4\}, \{b_1, b_2, b_3, b_4\})$ |
| V | $(\{b_1\}, \{b_1\})$ |

Figure 1: Example games with unique D -cores for several dominance structures D . The game on the left is a confrontation game and the game on the right is a symmetric win-lose-tie game.

Definition 5. A D -core is a D -solution X such that there does not exist a D -solution Y with $Y \subseteq X$ and $Y \neq X$.

Figure 1 contains examples of D -cores for all dominance structures considered in this paper.

Various set-valued solution concepts that have been proposed in the literature can be characterized as D -cores for some dominance structure D . An action profile $(a_1, \dots, a_n) \in \prod_{i \in N} A_i$ is a *Nash equilibrium in pure strategies* if and only if $(\{a_1\}, \dots, \{a_n\})$ is a V -core. Shapley's (1964) *saddles* and *weak saddles* for matrix games correspond to S - and V -cores, respectively, Dutta and Laslier's (1999) *minimal covering set* for symmetric matrix games corresponds to the C_M -core, and Duggan's (2013) *deep covering set* for symmetric win-lose-tie games corresponds to the C_D -core. Furthermore, mixed refinements of Shapley's saddles, as proposed by Duggan and Le Breton (2001) for symmetric win-lose-tie games, correspond to S^* - and W^* -cores.

Two further solution concepts that fit into this framework are Harsanyi and Selten's (1988) *formations* and Basu and Weibull's (1991) *CURB sets*. The respective dominance structures are defined in terms of best response sets. An action a_i is *rationally dominated* with respect to a set X_{-i} of opponent action profiles if it is *not* a best response to any mixed opponent strategy with support in X_{-i} . A subtle difference occurs if there are more than two players (and therefore more than one opponent). While in *correlated rational dominance* (R_c), opponents are allowed to play joint, i.e., correlated, mixtures (and thus

to act like a single opponent), *uncorrelated rational dominance* (R_u) restricts opponents to independent mixtures.

A tuple of sets is a CURB set if and only if it is externally stable with respect to R_u , and *minimal CURB sets* coincide with R_u -cores. Similarly, a tuple of sets is a formation if and only if it is externally stable with respect to R_c , and *primitive formations* are R_c -cores. Since it is well known that an action is not a best response to some correlated opponent strategy if and only if it is strictly dominated by a mixed strategy (see, e.g., Pearce, 1984, Lemma 3), the dominance structures R_c and S^* coincide. As a consequence, all our results concerning S^* -cores directly apply to primitive formations as well. The same is true for minimal CURB sets in two-player games, due to the equivalence of R_c and R_u for $n = 2$.¹¹

3.3 Properties of Dominance Structures

We now define a number of properties in order to formalize for which dominance structures, D -cores can be computed efficiently. An action $a_i \in A_i$ is said to be D -maximal with respect to X_{-i} if it is not D -dominated by A_i .

Definition 6. *Let D be a dominance structure and $X \subseteq A_N$. The set of D -maximal elements of A_i with respect to X_{-i} is defined as*

$$\max(D(X_{-i})) = A_i \setminus \{a_i \in A_i : A_i D(X_{-i}) a_i\}.$$

Definition 7. *Let $X \subseteq A_N$ and $a_i \in A_i$. A dominance structure D satisfies*

- *monotonicity (MON) if $X_i D(X_{-i}) a_i$ implies $X_i D(Y_{-i}) a_i$ for all $\emptyset \neq Y_{-i} \subseteq X_{-i}$,*
- *computational tractability (COM) if $X_i D(X_{-i}) a_i$ can be checked in polynomial time,*
- *maximal domination (MAX) if $A_i D(X_{-i}) a_i$ implies $\max(D(X_{-i})) D(X_{-i}) a_i$, and*
- *singularity (SING) if $X_i D(X_{-i}) a_i$ implies the existence of an action $x_i \in X_i$ with $x_i D(X_{-i}) a_i$.*

It is easily seen that S , B , and V are monotonic, and that W is not. S and W satisfy maximal domination because the relations $S(X_{-i})$ and $W(X_{-i})$ —restricted to pairs of singletons—are transitive and irreflexive. On the other hand, V violates MAX because $\max(V(X_{-i}))$ may be empty. It directly follows from the definitions that S , W , V , C_M , and C_D are singular.

Computational tractability of dominance structures is a mild requirement. Indeed, if a dominance structure does not satisfy COM, there is no hope for computing D -solutions efficiently. As shown in Sections 5 and 6, all the dominance structures defined in Section 3.1 satisfy COM.

The following properties of dominance structures are defined in the context of symmetric matrix games.

¹¹In games with more than two players, CURB sets are computationally intractable (Hansen et al., 2010). In fact, even checking uncorrelated rational dominance is coNP-hard.

Definition 8. Let (A, u) be a symmetric matrix game. Let furthermore $X, X' \subseteq A$ and $a, b, c \in A$. A dominance structure D satisfies

- weak monotonicity (*weak MON*) if $a D(X) b$ implies $a D(X') b$ for all $X' \subseteq X$ with $a \in X'$,
- transitivity (*TRA*) if $a D(X) b$, $b D(X') c$, and $a \in X \cap X'$ imply $a D(X \cap X') c$,¹²
- computational tractability of finding subsets (*SUB-COM*) if a nonempty subset of a D -core can be computed in polynomial time, and
- uniqueness in symmetric matrix games (*UNI*) if every symmetric matrix game has a unique D -core.

In contrast to COM, property SUB-COM is rather demanding. However, there is a useful sufficient condition involving the essential set. Since the essential set of a game can be computed efficiently via linear programming, a dominance structure D satisfies SUB-COM if every D -core of a game Γ contains the essential set $ES(\Gamma)$.

Monotonicity turns out to be sufficient for the *existence of solutions*: If a dominance structure D satisfies MON, a D -solution can be constructed by iteratively eliminating actions that are D -dominated (Duggan and Le Breton, 2014). Once the elimination process terminates, MON ensures that the resulting set is externally D -stable. Note, however, that these solutions need not be minimal (see, e.g., Figure 1).¹³ In symmetric matrix games, the same is true for dominance structures satisfying weak monotonicity. As weak dominance and mixed weak dominance are not monotonic, the above argument does not apply to those dominance structures. In fact, there are games without any W - or W^* -solution (see Figure 2 for an example). For this reason, we do not consider W - and W^* -solutions in this paper.¹⁴

Monotonicity of a dominance structure is also an important ingredient in several results on the *uniqueness* of (minimal or maximal) D -solutions (Duggan and Le Breton, 2014). However, those uniqueness results do not apply to the dominance structures C_M and C_D (which only satisfy weak monotonicity). Results on the uniqueness of D -cores with respect to these dominance structures are given in Section 6.

¹² $X \cap X'$ is to be read componentwise. Hence, $X \cap X' \neq \emptyset$ if and only if $X_i \cap X'_i \neq \emptyset$ for all $i \in N$.

¹³Under fairly general conditions, D -solutions obtained by iterated elimination of D -dominated actions are *maximal* (Duggan and Le Breton, 2014). The maximal S^* -solution of a two-player game, for instance, consists of all *rationalizable* actions (Pearce, 1984; Bernheim, 1984).

¹⁴The fact that W -solutions may fail to exist was first observed by Samuelson (1992). There are at least three approaches to restore the existence of W -solutions. First, one can ignore internal stability and define W -solutions as externally W -stable sets (Duggan and Le Breton, 2001). The properties of W -cores defined in this way are similar to those of V -cores: The number of W -cores may be exponential, even in symmetric matrix games, and a number of natural problems concerning W -cores are computationally intractable (Brandt et al., 2011b). Second, one can look for restricted classes of games in which W -solutions are guaranteed to exist. One such class is the class of confrontation games, where the W -core is unique and coincides with the V -core. Third, one can consider the so-called *monotonic kernel* of W , which turns out to be identical to B (Duggan and Le Breton, 2014).

$$\begin{array}{ccc}
& a_1 & a_2 & a_3 \\
a_1 & \begin{array}{|c|} \hline 0 \\ \hline \end{array} & \begin{array}{|c|} \hline 1 \\ \hline \end{array} & \begin{array}{|c|} \hline 0 \\ \hline \end{array} \\
a_2 & \begin{array}{|c|} \hline -1 \\ \hline \end{array} & \begin{array}{|c|} \hline 0 \\ \hline \end{array} & \begin{array}{|c|} \hline 1 \\ \hline \end{array} \\
a_3 & \begin{array}{|c|} \hline 0 \\ \hline \end{array} & \begin{array}{|c|} \hline -1 \\ \hline \end{array} & \begin{array}{|c|} \hline 0 \\ \hline \end{array}
\end{array}$$

Figure 2: Symmetric matrix game without W - and W^* -solutions.

Another beneficial property of (weakly) monotonic dominance structures is that *minimal* externally stable sets also happen to be internally stable. This is again due to the fact that the iterative elimination of dominated actions preserves external stability.

Proposition 1. (i) Let $\Gamma = (N, (A_i)_{i \in N}, (u_i)_{i \in N})$ be a game in normal form and D a dominance structure satisfying *MON*. Then, a set $X \subseteq A_N$ is a D -core of Γ if and only if it is a minimal externally D -stable set.

(ii) Let $\Gamma = (A, u)$ be a symmetric matrix game and D a dominance structure satisfying weak *MON*. Then, a set $X \subseteq A$ is a D -core of Γ if and only if it is a minimal externally D -stable set.

The first part of Proposition 1 follows from Proposition 2 of Duggan and Le Breton (2014).¹⁵ The proof of the second part is completely analogous.

Our proofs will frequently exploit the equivalence of D -cores and minimal externally D -stable sets. In particular, we will make use of the following easy corollary of Proposition 1.

Corollary 1. (i) Let $\Gamma = (N, (A_i)_{i \in N}, (u_i)_{i \in N})$ be a game in normal form and D a dominance structure satisfying *MON*. If $X \subseteq A_N$ is externally D -stable, then there exists a D -core Y with $Y \subseteq X$.

(ii) Let $\Gamma = (A, u)$ be a symmetric matrix game and D a dominance structure satisfying weak *MON*. If $X \subseteq A$ is externally D -stable, then there exists a D -core Y with $Y \subseteq X$.

4 General Results

We will now study how to compute minimal solutions for the dominance structures introduced in the previous section. To this end, we introduce two generic algorithms: a *greedy* and a *sophisticated* one. In principle, these algorithms can be applied to any game and any of the dominance structures introduced in Section 3.1. The goal of this section is to identify, for each algorithm, which properties of a dominance structure guarantee that the

¹⁵Note that Proposition 2 of Duggan and Le Breton (2014) also assumes Duggan and Le Breton's version of *transitivity*, but this is satisfied by all the dominance structures of Definition 2 when translated into their binary framework.

algorithm is sound and efficient. In addition, we will construct families of games that admit an exponential number of minimal solutions.

4.1 Generic Greedy Algorithm

Shapley (1964) has shown that every matrix game possesses a unique S -core and described an algorithm, attributed to Harlan Mills, to compute this set. The idea behind this algorithm is that given a nonempty subset of the S -core, the S -core 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. We generalize Mills' algorithm in two directions. First, we identify general conditions on a dominance structure D that ensure that this greedy approach works. Second, we consider arbitrary n -player normal-form games, thereby losing uniqueness of D -cores, and devise an algorithm that computes *all* D -cores of such games in polynomial time.

We start by looking at some structural properties of externally stable sets. For monotonic dominance structures satisfying maximal domination, externally stable sets are closed under intersection.¹⁶

Proposition 2. *Let D be a dominance structure satisfying MON and MAX. If X and Y are externally D -stable and $X \cap Y \neq \emptyset$, then $X \cap Y$ is externally D -stable.*

Proof. Suppose that X and Y are externally D -stable and $X \cap Y \neq \emptyset$. In order to show that $X \cap Y$ is externally D -stable, fix $i \in N$ and consider $a_i \in A_i \setminus (X_i \cap Y_i)$. Without loss of generality, assume that $a_i \notin X_i$. As X is externally D -stable, $X_i \not D(X_{-i}) a_i$, and thus $A_i \not D(X_{-i}) a_i$. Now MON implies $A_i \not D(X_{-i} \cap Y_{-i}) a_i$. Since $a_i \in A_i \setminus (X_i \cap Y_i)$ was chosen arbitrarily, no action in $A_i \setminus (X_i \cap Y_i)$ is D -maximal with respect to $X_{-i} \cap Y_{-i}$. Therefore, $\max(D(X_{-i} \cap Y_{-i})) \subseteq X_i \cap Y_i$. Moreover, MAX implies $\max(D(X_{-i} \cap Y_{-i})) \not D(X_{-i} \cap Y_{-i}) a_i$, which finally yields $(X_i \cap Y_i) \not D(X_{-i} \cap Y_{-i}) a_i$. \square

One particularly useful consequence of Proposition 2 is the uniqueness of minimal externally D -stable sets containing given sets of actions.

Corollary 2. *Let D be a dominance structure satisfying MON and MAX. For any $X^0 \subseteq A_N$, the minimal externally D -stable set containing X^0 is unique: If Y and Z are externally D -stable with $X^0 \subseteq Y$ and $X^0 \subseteq Z$, then $Y \subseteq Z$ or $Z \subseteq Y$.*

Proof. Let $X^0 \subseteq A_N$. Assume for contradiction that both Y and Z are minimal among all externally D -stable sets containing X^0 , and that neither $Y \subseteq Z$ nor $Z \subseteq Y$. As both Y and Z contain X^0 , $Y \cap Z$ is nonempty and Proposition 2 implies that $Y \cap Z$ is externally D -stable. This contradicts minimality of both Y and Z . \square

If D moreover satisfies computational tractability, the minimal externally D -stable set containing X^0 can be computed efficiently by greedily adding D -maximal actions.

¹⁶ $X \cap Y$ is to be read componentwise. Hence, $X \cap Y \neq \emptyset$ if and only if $X_i \cap Y_i \neq \emptyset$ for all $i \in N$.

Algorithm 1 Minimal externally D -stable set containing X^0

```

procedure min_ext( $\Gamma, (X_1^0, \dots, X_n^0)$ )
  for all  $i \in N$  do
     $X_i \leftarrow X_i^0$ 
  end for
  repeat
    for all  $i \in N$  do
       $Y_i \leftarrow \max(D(X_{-i})) \setminus X_i$ 
       $X_i \leftarrow X_i \cup Y_i$ 
    end for
  until  $\bigcup_{i=0}^n Y_i = \emptyset$ 
  return  $(X_1, \dots, X_n)$ 

```

Proposition 3. *Let $X^0 \subseteq A_N$. If D satisfies MON, MAX, and COM, the minimal externally D -stable set containing X^0 can be computed in polynomial time.*

Proof. We show that Algorithm 1 computes the minimal externally D -stable set containing X^0 and runs in polynomial time. Algorithm 1 starts with X^0 and iteratively adds all actions that are maximal with respect to the current set X_{-i} of opponent action profiles. As D satisfies COM, these actions can be computed efficiently. Moreover, the number of loops is bounded by $\sum_{i=1}^n |A_i|$.

Let X^{min} be the minimal externally D -stable set containing X^0 . We show that during the execution of Algorithm 1, the set X is always a subset of X^{min} . At the end of the algorithm, $\bigcup_{i=0}^n Y_i = \emptyset$ implies that $\max(D(X_{-i})) \subseteq X_i$ for all $i \in N$. As D satisfies MAX, this shows that X is externally D -stable.

We prove $X \subseteq X^{min}$ by induction on $|X| = \sum_{i=1}^n |X_i|$. At the beginning of the algorithm, $X = X^0 \subseteq X^{min}$ by definition of X^{min} . Now assume that $X \subseteq X^{min}$ at the beginning of a particular iteration. We have to show that for all $i \in N$, $Y_i \subseteq X_i^{min}$. Let $a_i \in Y_i = \max(D(X_{-i})) \setminus X_i$, and assume for contradiction that $a_i \notin X_i^{min}$. Since X^{min} is externally D -stable, $X_i^{min} \not D(X_{-i}^{min}) a_i$. By the induction hypothesis, $X_{-i} \subseteq X_{-i}^{min}$, which together with MON implies $X_i^{min} \not D(X_{-i}) a_i$. It follows that $A_i \not D(X_{-i}) a_i$, contradicting the assumption that $a_i \in \max(D(X_{-i}))$. \square

Whenever X^0 is contained in a D -core, Algorithm 1 returns this D -core. This property can be used to construct an algorithm to compute all D -cores of a game: Call Algorithm 1 for every possible combination of singleton sets of actions of the different players. The result is a set of externally D -stable sets, and the D -cores of the game are the inclusion-minimal elements of this set. This idea is made precise in Algorithm 2.

Theorem 1. *If D satisfies MON, MAX, and COM, all D -cores of a normal-form game can be computed in polynomial time.*

Algorithm 2 All D -cores

procedure $D_cores(\Gamma)$ $C \leftarrow \emptyset$ **for all** $(a_1, \dots, a_n) \in \prod_{i \in N} A_i$ **do** add $\min_ext(\Gamma, (\{a_1\}, \dots, \{a_n\}))$ to set C **end for****return** $\{X \in C : \text{there does not exist } X' \in C \text{ with } X \neq X' \text{ and } X' \subseteq X\}$

| | a_1 | a_2 | a_3 | a_4 | a_5 | a_6 |
|-------|-------|-------|-------|-------|-------|-------|
| a_1 | 0 | 1 | -1 | 1 | 1 | -1 |
| a_2 | -1 | 0 | 1 | 1 | -1 | 1 |
| a_3 | 1 | -1 | 0 | -1 | 1 | 1 |
| a_4 | -1 | -1 | 1 | 0 | 1 | -1 |
| a_5 | -1 | 1 | -1 | -1 | 0 | 1 |
| a_6 | 1 | -1 | -1 | 1 | -1 | 0 |

Figure 3: Tournament game with unique V -core $(\{a_1, a_2, a_3\}, \{a_1, a_2, a_3\})$. A naive adaptation of Algorithm 2—starting with a pair $(\{a_i\}, \{a_j\})$ and iteratively adding all actions that are not V -dominated—results in proper supersets of the V -core.

Proof. Let $\Gamma = (N, (A_i)_{i \in N}, (u_i)_{i \in N})$ be a game in normal form. We show that Algorithm 2 computes all D -cores of Γ and runs in polynomial time. Polynomial running time follows immediately because Algorithm 1 is invoked $\prod_{i \in N} |A_i|$ times, and inclusion-minimality can be checked easily.

As for correctness, we first note that every D -core X is an element of the set C . To see this, note that Proposition 2 implies that X is the minimal externally D -stable set containing $(\{x_1\}, \dots, \{x_n\})$, for every $(x_1, \dots, x_n) \in \prod_{i=1}^n X_i$. To show that all inclusion-minimal elements of C are D -cores, observe that all elements of C are externally D -stable. Thus, the first part of Corollary 1 implies that every element of C contains a D -core. Since all D -cores are elements of C , the inclusion-minimal elements of C are exactly the D -cores of Γ . \square

4.2 Generic Sophisticated Algorithm for Symmetric Matrix Games

Algorithm 2 is not sound for all dominance structures considered in this paper. For instance, very weak dominance violates maximal domination and therefore does not satisfy the conditions of Theorem 1. The example given in Figure 3 shows that, even in tournament games, where V -cores are generally unique, Algorithm 2 fails to find the V -core. The failure of the greedy algorithm can be traced back to the following problem: Since

the set of D -maximal actions with respect to a given set X_{-i} may be empty, it is no longer obvious which actions should be added to a strict subset of a D -core. In particular, adding all actions that are not D -dominated with respect to the current subset does not work (see Figure 3). We will identify conditions on dominance structures D that allow for the following sophisticated method: Instead of adding all D -undominated actions, merely add actions contained in a D -core of the *subgame* induced by the D -undominated actions. This immediately gives rise to a recursive algorithm, whose running time may, however, be exponential. Whenever a nonempty subset of a D -core can be found efficiently, which fortunately is the case for many symmetric dominance structures we consider, an efficient algorithm can be constructed.

In this section, we will only be concerned with symmetric matrix games and dominance structures D satisfying uniqueness in symmetric matrix games (UNI). If a symmetric matrix game Γ has a unique D -core (X_1, X_2) , it is easily verified that $X_1 = X_2$.¹⁷ In this case, the set $X_1 = X_2$ will be denoted by $S_D(\Gamma)$. The following lemma is the key ingredient for the sophisticated algorithm. Given a game Γ and a subset X of the D -core of Γ , it identifies sufficient conditions for the D -core of $\Gamma|_{A'}$ to be contained in the D -core of Γ , where $\Gamma|_{A'}$ is the subgame of Γ that is induced by the D -undominated actions with respect to X .

Lemma 1. *Let Γ be a symmetric matrix game and D a dominance structure satisfying weak MON, TRA, SING, and UNI. Let furthermore $X \subseteq S_D(\Gamma)$ and let A' denote the set of actions that are neither contained in X nor D -dominated with respect to X , i.e., $A' = \{a \in A \setminus X : \neg(X D(X) a)\}$. Then, $S_D(\Gamma|_{A'}) \subseteq S_D(\Gamma)$.*

Proof. Let $X \subseteq S_D(\Gamma)$ and $A' = \{a \in A \setminus X : \neg(X D(X) a)\}$. We can assume that A' is nonempty, as otherwise $S_D(\Gamma|_{A'})$ is empty and there is nothing to prove.

Now partition the set A' into two sets $C = A' \cap S_D(\Gamma)$ and $C' = A' \setminus S_D(\Gamma)$ of actions contained in $S_D(\Gamma)$ and actions not contained in $S_D(\Gamma)$. We will show that (C, C') is externally D -stable in $\Gamma|_{A'}$. Then, the second part of Corollary 1 and UNI imply that $S_D(\Gamma|_{A'}) \subseteq C$ and, therefore, $S_D(\Gamma|_{A'}) \subseteq S_D(\Gamma)$.

In order to show that (C, C') is externally D -stable in $\Gamma|_{A'}$, consider some $z \in C'$. Since $z \notin S_D(\Gamma)$, SING implies that there exists $y \in S_D(\Gamma)$ with $y D(S_D(\Gamma)) z$. We show that $y \in C$. It is easy to see that $y \notin X$, since otherwise weak MON would imply that $y D(X) z$, contradicting the assumption that $z \in A'$. On the other hand, assume that $y \in S_D(\Gamma) \setminus (X \cup C)$. Then there is some $x \in X$ such that $x D(X) y$. However, according to TRA, $x D(X) y$ and $y D(S_D(\Gamma)) z$ imply $x D(X) z$, again contradicting the assumption that $z \in A'$. Thus $y \in C$, and using weak MON again, $y D(S_D(\Gamma)) z$ and $z \in A'$ imply $y D(A') z$. Hence (C, C') is externally D -stable in $\Gamma|_{A'}$. \square

Two further properties are required to turn the insight of Lemma 1 into an efficient algorithm: First, we need a polynomial-time algorithm to compute a nonempty subset of the unique D -core; second, the dominance structure D itself must be computationally tractable.

¹⁷Suppose $X_1 \neq X_2$. Then, (X_2, X_1) is another D -core in Γ .

Algorithm 3 D -core of a symmetric matrix game

```
procedure D_core_symm( $\Gamma$ )
   $X \leftarrow$  subset of  $S_D(\Gamma)$ 
  repeat
     $A' \leftarrow \{a \in A \setminus X : \neg(X \ D(X) \ a)\}$ 
     $X' \leftarrow$  subset of  $S_D(\Gamma|_{A'})$ 
     $X \leftarrow X \cup X'$ 
  until  $\max(D(X)) \setminus X = \emptyset$ 
  return  $(X, X)$ 
```

Theorem 2. *If D satisfies weak MON, TRA, SING, UNI, SUB-COM, and COM, the D -core of a symmetric matrix game can be computed in polynomial time.*

Proof. Let $\Gamma = (A, u)$ be a symmetric matrix game. We show that Algorithm 3 computes $S_D(\Gamma)$ and runs in polynomial time. In each iteration, at least one action is added to the set X , so the algorithm is guaranteed to terminate after at most $|A|$ iterations. Each iteration consists of (1) computing the set A' of D -undominated actions and (2) finding a subset X' of $S_D(\Gamma|_{A'})$. Since D satisfies COM and SUB-COM, both tasks can be performed in polynomial time.

As for correctness, we show by induction on the number of iterations that $X \subseteq S_D(\Gamma)$ holds at any time. When the algorithm terminates, X is externally D -stable, which together with the induction hypothesis implies that $X = S_D(\Gamma)$. The base case is trivial. Now assume that $X \subseteq S_D(\Gamma)$ at the beginning of a particular iteration. Then $X \cup X' \subseteq X \cup S_D(\Gamma|_{A'}) \subseteq S_D(\Gamma)$, where the first inclusion is due to $X' \subseteq S_D(\Gamma|_{A'})$ and the second inclusion follows from Lemma 1 and the induction hypothesis. \square

4.3 Games with an Exponential Number of D -Cores

Our algorithms do not apply to all dominance structures considered in this paper. In fact, some dominance structures give rise to an *exponential* number of D -cores, even in symmetric matrix games. We need the following lemma, which is easily established.

Lemma 2. *Let $\Gamma = (A, u)$ be a symmetric matrix game. Define Γ' as the matrix game that is identical to Γ except that player 1 has an additional action \hat{a} that always yields a utility of 1. That is, $\Gamma' = (\{1, 2\}, (A \cup \{\hat{a}\}, A), (u_1, u_2))$ with $u_1(a, b) = u(a, b)$ for all $a, b \in A$, $u_1(\hat{a}, a) = 1$ for all $a \in A$, and $u_2 = -u_1$. Then, there exists no subset $X \subseteq A$ such that $X \ V^*(X) \ \hat{a}$.*

Proof. Assume for contradiction that $X \ V^*(X) \ \hat{a}$ for some $X \subseteq A$ and let $s \in \Delta(X)$ be a strategy that strictly dominates \hat{a} with respect to X . Consider the matrix game $\Gamma|_X$. In this game, playing strategy s guarantees a payoff of at least 1. However, the game $\Gamma|_X$ is symmetric and thus has a value of zero, meaning that no player can guarantee a strictly positive payoff. \square

A D -core (X_1, X_2) is *symmetric* if $X_1 = X_2$. It is straightforward to verify that every symmetric matrix game has a symmetric D -core for all dominance structures D considered in this paper. For $D \in \{V, V^*\}$, any symmetric matrix game with multiple symmetric D -cores can be used to show that the number of D -cores may be exponential in general. Let $c(\Gamma)$ denote the number of symmetric D -cores of game Γ .

Theorem 3. *Let $D \in \{V, V^*\}$. For every symmetric matrix game $\Gamma = (A, u)$, there exists a family $(\Gamma_k)_{k \in \mathbb{N}}$ of symmetric matrix games with the following properties:*

(i) *For all $k \in \mathbb{N}$, the game $\Gamma_k = (A^k, u^k)$ satisfies $|A^k| = 3^{k-1}|A|$ and*

$$c(\Gamma_k) = c(\Gamma)^{3^{k-1}} = c(\Gamma)^{|A^k|/|A|}.$$

In particular, $c(\Gamma_k)$ is exponential in $|A^k|$ whenever $c(\Gamma) \geq 2$.

(ii) *If Γ is a confrontation game, then so is Γ_k , for all $k \in \mathbb{N}$.*

Proof. Let $D \in \{V, V^*\}$ and consider a symmetric matrix game $\Gamma = (A, u)$. We now construct the family $(\Gamma_k)_{k \in \mathbb{N}}$ with $\Gamma_k = (A^k, u^k)$ for all $k \in \mathbb{N}$. Let $\Gamma_1 = \Gamma$. For $k \geq 1$, $\Gamma_{k+1} = (A^{k+1}, u^{k+1})$ is defined inductively as follows.

$$A^{k+1} = A^{k,0} \cup A^{k,1} \cup A^{k,2},$$

where for each $\ell \in \{0, 1, 2\}$, $A^{k,\ell}$ is a copy of A^k . For $a \in A^{k,\ell}$ and $b \in A^{k,\ell'}$, the utility function u^{k+1} is defined by

$$u^{k+1}(a, b) = \begin{cases} u^k(a, b) & \text{if } \ell = \ell', \\ -1 & \text{if } \ell' = \ell + 1, \\ 1 & \text{if } \ell' = \ell + 2, \end{cases}$$

where $\ell + r$ should be understood to mean $\ell + r \pmod{3}$. If M_k is the matrix representing Γ_k , $\mathbf{1}$ is the $|A^k| \times |A^k|$ matrix containing only ones, and $-\mathbf{1}$ is $(-1) \cdot \mathbf{1}$, then the game Γ_{k+1} is represented by the block matrix

$$M_{k+1} = \begin{pmatrix} M_k & -\mathbf{1} & \mathbf{1} \\ \mathbf{1} & M_k & -\mathbf{1} \\ -\mathbf{1} & \mathbf{1} & M_k \end{pmatrix}.$$

Property (ii) follows immediately from the construction. For (i), we will show that, for each $k \geq 1$, the symmetric D -cores of Γ_{k+1} can be characterized in terms of the symmetric D -cores of Γ_k . The following notation will be useful. For $X \subseteq A^{k+1}$ and $\ell \in \{0, 1, 2\}$, let $X_\ell = X \cap A^{k,\ell}$ denote the part of X that lies in block ℓ . We claim that for each $k \geq 1$,

$$\begin{aligned} (X, X) \text{ is a symmetric } D\text{-core in } \Gamma_{k+1} & \quad \text{if and only if} \\ (X_\ell, X_\ell) \text{ is a symmetric } D\text{-core in } \Gamma_k & \text{ for all } \ell \in \{0, 1, 2\}. \end{aligned} \tag{3}$$

Before proving this equivalence, we make the following observation.

$$\text{If } (X, X) \text{ is a } D\text{-core in } \Gamma_{k+1}, \text{ then } X_\ell \neq \emptyset \text{ for all } \ell \in \{0, 1, 2\}. \quad (4)$$

To see this, let $x \in X$ be an arbitrary action in X and choose ℓ such that $x \in X_\ell$. Consider the game where the actions of player 2 are restricted to X_ℓ . As $u^{k+1}(a, b) = 1$ for all $a \in X_{\ell+1}$ and $b \in X_\ell$, Lemma 2 implies that no action in $X_{\ell+1}$ is D -dominated by X_ℓ . Therefore, at least one of the actions in $X_{\ell+1}$ has to be contained in X , i.e., $X_{\ell+1} \neq \emptyset$. Repeating the argument, $X_{\ell+1} \neq \emptyset$ implies $X_{\ell+2} \neq \emptyset$, which proves (4).

We are now ready to prove the equivalence (3). For the direction from left to right, assume that (X, X) is a D -core in Γ_{k+1} and let $\ell \in \{0, 1, 2\}$. We need to show that (X_ℓ, X_ℓ) is a D -core in Γ_k . By (4), we know that $X_\ell \neq \emptyset$. To show that (X_ℓ, X_ℓ) is externally D -stable, consider some $a \in A^{k, \ell} \setminus X_\ell$. As (X, X) is externally D -stable in Γ_{k+1} , $X D(X) a$. However, the definition of u^{k+1} ensures that none of the actions in $X_{\ell+1} \cup X_{\ell+2}$ is involved in the domination of a , and that actually $X_\ell D(X) a$. Monotonicity of D finally yields $X_\ell D(X_\ell) a$. For minimality of (X_ℓ, X_ℓ) , note that the existence of an externally D -stable set $(X', Y') \neq (X_\ell, X_\ell)$ in Γ_k with $X', Y' \subseteq X_\ell$ would contradict the minimality of (X, X) in Γ_{k+1} .

For the direction from right to left, (X, X) is externally D -stable in Γ_{k+1} because each (X_ℓ, X_ℓ) is externally D -stable in Γ_k . Furthermore (X, X) is minimal, as a proper subset of (X, X) that is externally D -stable in Γ_{k+1} would yield an externally D -stable subset of (X_ℓ, X_ℓ) for some $\ell \in \{0, 1, 2\}$, contradicting the minimality of (X_ℓ, X_ℓ) in Γ_k . We have thus proven (3).

Finally, let $c_k = c(\Gamma_k)$. It follows from (3) that $c_{k+1} = c_k^3$ for all $k \geq 1$. As $c_1 = c(\Gamma)$, this yields $c_k = c(\Gamma)^{3^{k-1}}$. As $|A^k| = 3^{k-1}|A|$, we have $c_k = c(\Gamma)^{|A^k|/|A|}$. In particular, c_k is exponential in $|A^k|$ whenever $c(\Gamma) \geq 2$. \square

The construction used in the proof of Theorem 3 also works for weak dominance and mixed weak dominance.

5 Greedy Algorithms

In this section, we investigate the consequences of Theorem 1 on S -, B -, and S^* -cores.

Corollary 3. *All S -cores of a normal-form game can be computed in polynomial time.*

Proof. According to Theorem 1, it suffices to show that S satisfies MON, MAX, and COM. It can easily be verified that S satisfies MON and MAX. Furthermore, S satisfies COM because $x_i S(X_{-i}) a_i$ can be checked efficiently by simply comparing $u_i(x_i, x_{-i})$ and $u_i(a_i, x_{-i})$ for each $x_{-i} \in X_{-i}$. \square

The same is true for Börgers dominance.

Corollary 4. *All B -cores of a normal-form game can be computed in polynomial time.*

Algorithm 4 Checking Börgers dominance

```
procedure Boergers_dom( $\Gamma, (X_i)_{i \in N}, a_i$ )  
   $Y_{-i} \leftarrow X_{-i}$   
  repeat  
    if  $X_i W(Y_{-i}) a_i$  then  
      choose  $x_i \in X_i$  such that  $x_i W(Y_{-i}) a_i$   
       $C(x_i) \leftarrow \{y_{-i} \in Y_{-i} : u_i(x_i, y_{-i}) > u_i(a_i, y_{-i})\}$   
       $Y_{-i} \leftarrow Y_{-i} \setminus C(x_i)$   
    else  
      return “no”  
    end if  
  until  $Y_{-i} = \emptyset$   
  return “yes”
```

Proof. According to Theorem 1, it suffices to show that B satisfies MON, MAX, and COM. As was the case for S , it can easily be checked that B satisfies MON and MAX.

It remains to be shown that B satisfies COM. Consider the following procedure, formalized in Algorithm 4, which checks whether $X_i B(X_{-i}) a_i$ holds. First check whether X_i weakly dominates a_i . If no, then X_i does not B -dominate a_i either. If yes, we can find an action $x_i \in X_i$ with $u_i(x_i, x_{-i}) \geq u_i(a_i, x_{-i})$ for all $x_{-i} \in X_{-i}$. Define $C(x_i)$ as the set of all tuples $x_{-i} \in X_{-i}$ for which the latter inequality is strict. $C(x_i)$ is nonempty by definition of W . It follows that $x_i W(Y_{-i}) a_i$ for all Y_{-i} with $Y_{-i} \cap C(x_i) \neq \emptyset$. We can therefore restrict attention to subsets of $Y_{-i} \setminus C(x_i)$ and “recursively” check whether $X_i B(Y_{-i} \setminus C(x_i)) a_i$.

It is easily verified that this procedure runs in polynomial time and correctly checks Börgers dominance. For the former, observe that the number of loops is bounded by $|X_{-i}|$. For the latter, there are two cases to consider. If Algorithm 4 returns “no”, then there is $Y_{-i} \subseteq X_{-i}$ such that X_i does not weakly dominate a_i with respect to Y_{-i} . Therefore, $X_i B(X_{-i}) a_i$ does not hold. If, on the other hand, Algorithm 4 returns “yes”, then every $x_{-i} \in X_{-i}$ is contained in $C(x_i)$ for some $x_i \in X_i$. In order to show that $X_i B(Y_{-i}) a_i$, consider an arbitrary subset $Y_{-i} \subseteq X_{-i}$. We need to show that $x_i W(Y_{-i}) a_i$ for some $x_i \in X_i$. Consider the first time Algorithm 4 chooses an action x_i with $C(x_i) \cap Y_{-i} \neq \emptyset$. Denote this action by x_i^* and let Y_{-i}^* be the corresponding set of opponent action profiles such that $x_i^* W(Y_{-i}^*) a_i$. By choice of x_i^* , we have $Y_{-i} \subseteq Y_{-i}^*$. It follows that $u_i(x_i^*, y_{-i}) \geq u_i(a_i, y_{-i})$ for all $y_{-i} \in Y_{-i}$. And since $C(x_i) \cap Y_{-i} \neq \emptyset$, the inequality is strict for at least one $y_{-i} \in Y_{-i}$. \square

The requirements for the greedy algorithm are also met by mixed strict dominance.

Corollary 5. *All S^* -cores of a normal-form game can be computed in polynomial time.*

Proof. According to Theorem 1, it suffices to show that S^* satisfies MON, MAX, and COM. It is easily verified that S^* satisfies MON. Furthermore, S^* satisfies COM because

$X_i S^*(X_{-i}) a_i$ can be checked efficiently with the help of a linear program (see, e.g., Conitzer and Sandholm, 2005, Proposition 1).

We now show that S^* satisfies MAX.¹⁸ We use the following notation. For a strategy $s \in \Delta(A_i)$ and an action $a \in A_i$, we write $s >_{X_{-i}} a$ if s strictly dominates a with respect to X_{-i} , i.e., if $u(s, x_{-i}) > u(a, x_{-i})$ for all $x_{-i} \in X_{-i}$. We simply write $s > a$ if X_{-i} is clear from the context. We call an action *dominated* if there exists a strategy $s \in \Delta(A_i)$ with $s > a$.

Consider a strategy $s \in \Delta(A_i)$ and an action $a \in A_i$. For a strategy $t \in \Delta(A_i)$ with $t(a) < 1$, we denote by $s^{a \rightarrow t}$ the strategy in which action a gets probability zero and the probability mass from action a is redistributed among other alternatives according to t . Formally, $s^{a \rightarrow t}(a) = 0$ and $s^{a \rightarrow t}(b) = s(b) + s(a) \frac{t(b)}{\sum_{c \neq a} t(c)}$ for all $b \in A_i \setminus \{a\}$.

Consider a game Γ and a set X_{-i} of opponent action profiles. Let $a \in A_i$ be an action such that $A_i S^*(X_{-i}) a$. We call an action $b \in A_i$ *necessary for the domination of a* if every strategy $s \in \Delta(A_i)$ with $s >_{X_{-i}} a$ satisfies $s(b) > 0$. In other words, b is necessary for the domination of a if $A_i S^*(X_{-i}) a$ but not $A_i \setminus \{b\} S^*(X_{-i}) a$. We will use the following observations:

- (i) For every dominated action $a \in A_i$, a is not necessary for the domination of a .
- (ii) If b is necessary for the domination of a , then $b \in \max(S^*(X_{-i}))$.

To see (i), consider a dominated action $a \in A_i$ and $s \in \Delta(A_i)$ with $s > a$. If $s(a) = 0$, there is nothing to show. Otherwise, define $s' = s^{a \rightarrow s}$. It is easily verified that $s' > a$. Since $s'(a) = 0$, we have that a is not necessary for the domination of a .

To see (ii), consider a dominated action $a \in A_i$ and an action $b \in A_i$ that is necessary for the domination of a . Assume for contradiction that $b \notin \max(S^*(X_{-i}))$. Then there exists a strategy $s_b \in \Delta(A_i)$ with $s_b > b$. By (i), we can assume that $s_b(b) = 0$. Let furthermore $s_a \in \Delta(A_i)$ be a strategy with $s_a > a$. Since b is necessary for the domination of a , it must hold that $s_a(b) > 0$. Define $s' = s_a^{b \rightarrow s_b}$. It is easily verified that $s' > a$. But $s'(b) = 0$, contradicting the assumption that b is necessary for the domination of a .

We are now ready to prove that S^* satisfies MAX. Consider a game Γ , a player $i \in N$, and a set X_{-i} of opponent action profiles. Let $M = \max(S^*(X_{-i}))$ and $L = A_i \setminus M$ denote the maximal and dominated actions of player i , respectively. We need to show that $M S^*(X_{-i}) a$ for all $a \in L$. Our proof is by induction on the number $k = |L|$ of dominated strategies.

If $k = 0$, there is nothing to show, and for $k = 1$ the statement follows from observation (i). Now assume $k \geq 2$ and consider two dominated actions $a, b \in L$. By observation (ii), we know that b is not necessary for the domination of a . Thus, a is still dominated in the subgame Γ' of Γ in which player i is restricted to action set $A_i \setminus \{b\}$. By induction, a_i is dominated by the set M' of maximal actions of player i in Γ' . Clearly, M' is identical to the set M of maximal actions of player i in Γ . Thus $M S^*(X_{-i}) a$. Since the argument works for every dominated action $a \in L$, we have shown that S^* satisfies MAX. \square

¹⁸We are grateful to an anonymous reviewer for pointing out this argument.

6 Sophisticated Algorithms for Symmetric Matrix Games

In this section, we investigate the consequences of Theorems 2 and 3 on C_{D^-} , C_{M^-} , V^- , and V^* -cores.

Let us first consider C_M and C_D , which are only defined in symmetric matrix games. For this class of games, it turns out that both dominance structures yield unique minimal solutions. In fact, we can show the following more general result.

Proposition 4. *Let D be a dominance structure satisfying weak MON and TRA. If D is coarser than C_M (i.e., $D \subseteq C_M$), then every symmetric matrix game has a unique D -core.*

For the case $D = C_M$, it was already known that every symmetric matrix game has a unique *symmetric* C_M -core (Dutta and Laslier, 1999, Theorem 4.2). Furthermore, Lemmata 1 and 2 of Duggan and Le Breton (1996) imply that every confrontation game has a unique C_M -core. (See Footnote 21 for details.)

It is unknown whether $D \subseteq C_M$ is *necessary* for the uniqueness of D -cores in symmetric matrix games. Various dominance structures that are finer than C_M have been shown to admit disjoint minimal solutions, sometimes involving rather elaborate combinatorial arguments. Examples include *unidirectional covering* (Brandt and Fischer, 2008) and *extending* (Brandt, 2011; Brandt et al., 2013).

In order to prove Proposition 4, we need the following two lemmata.¹⁹

Lemma 3. *Let $D \subseteq C_M$ be a dominance structure satisfying weak MON and TRA. Let furthermore $\Gamma = (A, u)$ be a symmetric matrix game and $X, Y \subseteq A$. If (X, X) and (Y, Y) are externally D -stable sets, then $(X \cap Y, X \cap Y)$ is externally D -stable.*

Proof. We first show that $X \cap Y \neq \emptyset$. Assume for contradiction that $X \cap Y = \emptyset$ and let $a_0 \in Y$. As (X, X) is externally D -stable, there exists $a_1 \in X$ with $a_1 D(X) a_0$. As (Y, Y) is externally D -stable, there exists $a_2 \in Y$ with $a_2 D(Y) a_1$. Repeatedly applying these arguments yields an infinite sequence (a_0, a_1, a_2, \dots) such that

- for all *even* $i \geq 0$, $a_i \in Y$ and $a_{i+1} D(X) a_i$, and
- for all *odd* $i \geq 1$, $a_i \in X$ and $a_{i+1} D(Y) a_i$.

By the assumption that $D \subseteq C_M$ and the definition of C_M , it follows that

- $u(a_{i+1}, a_i) > 0$ and $u(a_{i+1}, x) \geq u(a_i, x)$ for all *even* $i \geq 0$ and for all $x \in X$, and
- $u(a_{i+1}, a_i) > 0$ and $u(a_{i+1}, y) \geq u(a_i, y)$ for all *odd* $i \geq 1$ and for all $y \in Y$.

¹⁹These lemmata are adapted from Duggan and Le Breton (1996), who proved analogous statements for the special case of W -cores in confrontation games.

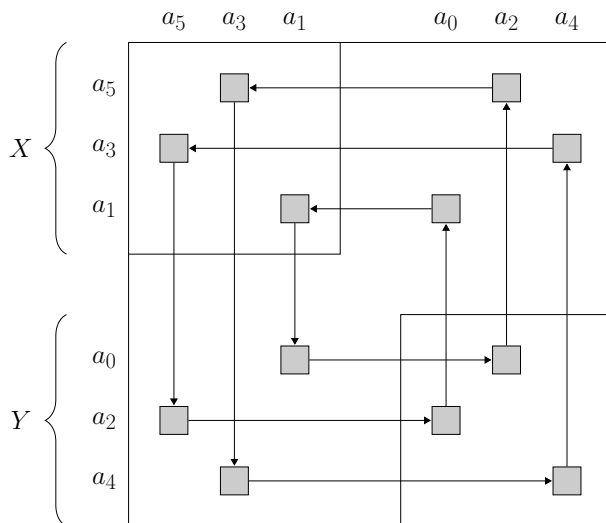


Figure 4: Construction in the proof of Lemma 3. An arrow from a profile (a_i, a_j) to another profile $(a_{i'}, a_{j'})$ indicates that $u(a_i, a_j) \geq u(a_{i'}, a_{j'})$.

Since A is finite, the sequence (a_0, a_1, a_2, \dots) must contain repetitions. Without loss of generality, assume that $a_k = a_0$ for some even $k > 0$. We can therefore construct the following chain of inequalities (consult Figure 4 for an example with $k = 6$):

$$\begin{aligned}
 u(a_0, a_1) &\leq u(a_1, a_1) \leq u(a_1, a_0) \leq u(a_2, a_0) \leq u(a_2, a_{k-1}) \\
 &\leq u(a_3, a_{k-1}) \leq u(a_3, a_{k-2}) \leq u(a_4, a_{k-2}) \leq \dots \\
 &\leq u(a_0, a_2) \leq u(a_0, a_1)
 \end{aligned}$$

It follows that all utilities in this chain of inequalities are equal. Moreover, since Γ is a symmetric matrix game, we have that $u(a_1, a_1) = 0$ and hence that all utilities in this chain are zero. In particular, $u(a_1, a_0) = 0$, which contradicts the assumption that $a_1 D(Y) a_0$ (because we have observed that the latter implies $u(a_1, a_0) > 0$). This proves that $X \cap Y \neq \emptyset$.²⁰

In order to show that $(X \cap Y, X \cap Y)$ is externally D -stable, take an arbitrary $a_0 \notin X \cap Y$. Without loss of generality, assume that $a_0 \notin X$. As (X, X) is externally D -stable and $D \subseteq C_M$, there exists $a_1 \in X$ with $a_1 C_M(X) a_0$. If $a_1 \notin Y$, there exists $a_2 \in Y$ with $a_2 C_M(X) a_1$. This construction finally yields some $a_k \in X \cap Y$, for otherwise we have a contradiction as in the first part of the proof. Repeated application of weak MON now yields $a_k D(X \cap Y) a_i$ for all $i < k$. In particular, $a_k D(X \cap Y) a_0$, as desired. \square

The proof of the following lemma is similar to that of Lemma 3 and we omit it.

²⁰ $X \cap Y \neq \emptyset$ also follows from Claim 2 in the the proof of Theorem 4.2 of Dutta and Laslier (1999).

Lemma 4. *Under the assumptions of Lemma 3, if (X, Y) is externally D -stable, then $(X \cap Y, X \cap Y)$ is externally D -stable.*

It is now easy to prove Proposition 4.

Proof of Proposition 4. Lemma 4 implies that every D -core (X, Y) satisfies $X = Y$, as otherwise $(X \cap Y, X \cap Y)$ would be a smaller externally D -stable set. Similarly, Lemma 3 implies that there cannot exist two D -cores (X, X) and (Y, Y) with $X \neq Y$. \square

Since C_D is coarser than C_M and both C_D and C_M are weakly monotonic and transitive, it immediately follows that C_D and C_M satisfy uniqueness in symmetric matrix games, which is one of the main ingredients of the sophisticated algorithm.

Corollary 6. *Every symmetric matrix game has a unique C_D -core and a unique C_M -core.*

We now show that the other requirements for the sophisticated algorithm are satisfied as well.

Corollary 7. *The C_D -core and the C_M -core of a symmetric matrix game can be computed in polynomial time.*

Proof. According to Theorem 2, it is sufficient to show that both C_M and C_D satisfy weak MON, TRA, SING, UNI, SUB-COM, and COM.

It is easily verified that both C_M and C_D satisfy weak MON, TRA, SING, and COM. UNI was shown in Corollary 6. Finally, Dutta and Laslier (1999) have shown that the essential set $ES(\Gamma)$ is a (nonempty) subset of $S_{C_M}(\Gamma)$. Since the essential set can be computed in polynomial time using linear programming (see Brandt and Fischer, 2008), this proves that C_M satisfies SUB-COM. The same is true for C_D because $S_{C_M}(\Gamma) \subseteq S_{C_D}(\Gamma)$. \square

Let us now turn to very weak dominance. In contrast to S -cores, V -cores are not unique in matrix games. It is in fact easily seen that even a symmetric matrix game can have multiple very weak saddles: If all action profiles yield the same utility, then every single profile constitutes a V -core. Therefore, Theorem 3 implies that there are symmetric matrix games with an exponential number of V -cores. The following corollary is an immediate consequence.

Corollary 8. *Computing all V -cores of a game requires exponential time in the worst case, even for symmetric matrix games.*

We emphasize that this result does not preclude the existence of an efficient algorithm that finds a succinct representation of V -cores for a given matrix game. In fact, the recent finding that V -cores in matrix games are interchangeable and equivalent raises hope that this might indeed be possible (Brandt et al., 2016). It is even open whether a single V -core of a matrix game can be computed efficiently. For non-zero-sum games, Brandt et al. (2011b) have shown that a number of natural problems like *finding* V -cores, checking

| | a_1 | a_2 | a_3 | a_4 | a_5 | a_6 |
|-------|-------|-------|-------|-------|-------|-------|
| a_1 | 0 | 3 | -3 | -1 | -1 | 2 |
| a_2 | -3 | 0 | 3 | -1 | 2 | -1 |
| a_3 | 3 | -3 | 0 | 2 | -1 | -1 |
| a_4 | 1 | 1 | -2 | 0 | 3 | -3 |
| a_5 | 1 | -2 | 1 | -3 | 0 | 3 |
| a_6 | -2 | 1 | 1 | 3 | -3 | 0 |

Figure 5: Confrontation game with two symmetric V^* -cores: $(\{a_1, a_2, a_3\}, \{a_1, a_2, a_3\})$ and $(\{a_4, a_5, a_6\}, \{a_4, a_5, a_6\})$. In the former V^* -core, a_4 is dominated by $\frac{2}{3}a_1 + \frac{1}{3}a_3$, a_5 is dominated by $\frac{1}{3}a_2 + \frac{2}{3}a_3$, and a_6 is dominated by $\frac{1}{3}a_1 + \frac{2}{3}a_2$. In the latter V^* -core, a_1 is dominated by $\frac{2}{3}a_5 + \frac{1}{3}a_6$, a_2 is dominated by $\frac{2}{3}a_4 + \frac{1}{3}a_5$, and a_3 is dominated by $\frac{1}{3}a_4 + \frac{2}{3}a_6$.

whether a given action is contained in *some* V -core, or deciding whether there is a unique V -core are computationally intractable.

In confrontation games, the picture is different: Duggan and Le Breton (1996) have shown that these games have a unique V -core, which moreover coincides with the (unique) C_M -core.²¹ The following positive result now follows from Corollary 7.

Corollary 9. *The unique V -core of a confrontation game can be computed in polynomial time.*

V^* -cores, on the other hand, are not even unique in confrontation games, as witnessed by the game in Figure 5. Applying Theorem 3 again, we get the following.

Corollary 10. *Computing all V^* -cores of a game requires exponential time in the worst case, even for confrontation games.*

In order to guarantee the uniqueness of V^* -cores, we thus have to restrict the class of games even further. Duggan and Le Breton (2001) have shown that tournament games have a unique V^* -core, which moreover coincides with the (unique) C_M -core.

Corollary 11. *The unique V^* -core of a tournament game can be computed in polynomial time.*

²¹To be precise, Duggan and Le Breton (1996) have shown that the W -core of a confrontation game is unique and coincides with the C_M -core. However, it can be shown that the proofs carry over for very weak dominance. As a consequence, V -cores, W -cores, and C_M -cores all coincide in confrontation games.

Acknowledgments

This material is based on work supported by the Deutsche Forschungsgemeinschaft under grants BR 2312/7-1, BR 2312/7-2, and BR 2312/10-1, by a Feodor Lynen research fellowship of the Alexander von Humboldt Foundation, and by ERC-StG 639945 (ACCORD). We thank Felix Fischer and Paul Harrenstein for valuable discussions and their contribution to an earlier related paper that was presented at the 8th International Joint Conference on Autonomous Agents and Multi-Agent Systems (Budapest, May 2009). Moreover, this work was presented at the 13th ACM Conference on Electronic Commerce (Valencia, June 2012), at the 4th World Congress of the Game Theory Society (Istanbul, July 2012), and at the CS-ECON seminar series at Duke University (Durham, October 2014). We thank the anonymous reviewers for helpful comments.

References

- R. J. Aumann and A. Brandenburger. 1995. Epistemic Conditions for Nash Equilibrium. *Econometrica* 63, 5 (1995), 1161–1180.
- C. W. Bach and E. Tsakas. 2014. Pairwise epistemic conditions for Nash equilibrium. *Games and Economic Behavior* 85 (2014), 48–59.
- P. Borelli. 2009. Consistency of beliefs and epistemic conditions for Nash and correlated equilibria. *Games and Economic Behavior* 67, 2 (2009), 363–375.
- K. Basu and J. Weibull. 1991. Strategy Subsets Closed under Rational Behavior. *Economics Letters* 36 (1991), 141–146.
- M. Benisch, G. B. David, and T. Sandholm. 2010. Algorithms for Closed Under Rational Behavior (CURB) Sets. *Journal of Artificial Intelligence Research* 38 (2010), 513–534.
- B. D. Bernheim. 1984. Rationalizable Strategic Behavior. *Econometrica* 52, 4 (1984), 1007–1028.
- E. Borel. 1921. La théorie du jeu et les équations intégrales à noyau symétrique. *Comptes Rendus de l'Académie des Sciences* 173 (1921), 1304–1308.
- T. Börgers. 1993. Pure Strategy Dominance. *Econometrica* 61, 2 (1993), 423–430.
- F. Brandt. 2011. Minimal Stable Sets in Tournaments. *Journal of Economic Theory* 146, 4 (2011), 1481–1499.
- F. Brandt, M. Brill, F. Fischer, and P. Harrenstein. 2011a. On The Complexity of Iterated Weak Dominance in Constant-Sum Games. *Theory of Computing Systems* 49, 1 (2011), 162–181.

- F. Brandt, M. Brill, F. Fischer, and J. Hoffmann. 2011b. The Computational Complexity of Weak Saddles. *Theory of Computing Systems* 49, 1 (2011), 139–161.
- F. Brandt, M. Brill, and W. Suksompong. 2016. An Ordinal Minimax Theorem. *Games and Economic Behavior* 95 (2016), 107–112.
- F. Brandt, M. Chudnovsky, I. Kim, G. Liu, S. Norin, A. Scott, P. Seymour, and S. Thomassé. 2013. A counterexample to a conjecture of Schwartz. *Social Choice and Welfare* 40, 3 (2013), 739–743.
- F. Brandt and F. Fischer. 2008. Computing the Minimal Covering Set. *Mathematical Social Sciences* 56, 2 (2008), 254–268.
- X. Chen, X. Deng, and S.-H. Teng. 2009. Settling the Complexity of Computing Two-Player Nash Equilibria. *J. ACM* 56, 3 (2009).
- V. Conitzer and T. Sandholm. 2005. Complexity of (Iterated) Dominance. In *Proc. of 6th ACM-EC Conference*. ACM Press, 88–97.
- C. Daskalakis, P. Goldberg, and C. Papadimitriou. 2009. The complexity of computing a Nash equilibrium. *SIAM J. Comput.* 39, 1 (2009), 195–259.
- J. Duggan. 2013. Uncovered Sets. *Social Choice and Welfare* 41, 3 (2013), 489–535.
- J. Duggan and M. Le Breton. 1996. Dutta’s Minimal Covering Set and Shapley’s Saddles. *Journal of Economic Theory* 70, 1 (1996), 257–265.
- J. Duggan and M. Le Breton. 2001. Mixed Refinements of Shapley’s Saddles and Weak Tournaments. *Social Choice and Welfare* 18, 1 (2001), 65–78.
- J. Duggan and M. Le Breton. 2014. Choice-theoretic Solutions for Strategic Form Games. (2014). Mimeo. Available at <http://www.johnduggan.net>.
- B. Dutta. 1988. Covering Sets and A New Condorcet Choice Correspondence. *Journal of Economic Theory* 44, 1 (1988), 63–80.
- B. Dutta and J.-F. Laslier. 1999. Comparison Functions and Choice Correspondences. *Social Choice and Welfare* 16, 4 (1999), 513–532.
- P. C. Fishburn. 1978. Non-Cooperative Stochastic Dominance Games. *International Journal of Game Theory* 7, 1 (1978), 51–61.
- D. C. Fisher and R. B. Reeves. 1995. Optimal Strategies for Random Tournament Games. *Linear Algebra Appl.* 217 (1995), 83–85.
- D. C. Fisher and J. Ryan. 1992. Optimal Strategies for a Generalized “Scissors, Paper, and Stone” Game. *Amer. Math. Monthly* 99, 10 (1992), 935–942.

- D. C. Fisher and J. Ryan. 1995. Tournament Games and Positive Tournaments. *Journal of Graph Theory* 19, 2 (1995), 217–236.
- I. Gilboa, E. Kalai, and E. Zemel. 1993. The Complexity of Eliminating Dominated Strategies. *Mathematics of Operations Research* 18, 3 (1993), 553–565.
- K. A. Hansen, P. B. Miltersen, and T. B. Sørensen. 2010. The Computational Complexity of Trembling Hand Perfection and Other Equilibrium Refinements. In *Proc. of 3rd International Symposium on Algorithmic Game Theory (SAGT) (LNCS)*, Vol. 6386. Springer, 198–209.
- J. C. Harsanyi and R. Selten. 1988. *A General Theory of Equilibrium Selection in Games*. MIT Press.
- T. Jungbauer and K. Ritzberger. 2011. Strategic games beyond expected utility. *Economic Theory* 48, 2–3 (2011), 377–398.
- D. E. Knuth, C. H. Papadimitriou, and J. N. Tsitsiklis. 1988. A Note on Strategy Elimination in Bimatrix Games. *Operations Research Letters* 7, 3 (1988), 103–107.
- G. Laffond, J.-F. Laslier, and M. Le Breton. 1993. The Bipartisan Set of a Tournament Game. *Games and Economic Behavior* 5, 1 (1993), 182–201.
- R. D. Luce and H. Raiffa. 1957. *Games and Decisions: Introduction and Critical Survey*. John Wiley & Sons Inc.
- R. D. McKelvey. 1986. Covering, Dominance, and Institution-Free Properties of Social Choice. *American Journal of Political Science* 30, 2 (1986), 283–314.
- R. D. McKelvey and P. C. Ordeshook. 1976. Symmetric Spatial Games Without Majority Rule Equilibria. *The American Political Science Review* 70, 4 (1976), 1172–1184.
- R. B. Myerson. 1991. *Game Theory: Analysis of Conflict*. Harvard University Press.
- J. F. Nash. 1951. Non-cooperative games. *Annals of Mathematics* 54, 2 (1951), 286–295.
- D.G. Pearce. 1984. Rationalizable Strategic Behavior and the Problem of Perfection. *Econometrica* 52, 4 (1984), 1029–1050.
- A. Perea. 2007. A One-Person Doxastic Characterization of Nash strategies. *Synthese* 158 (2007), 251–271.
- L. Samuelson. 1992. Dominated Strategies and Common Knowledge. *Games and Economic Behavior* 4, 2 (1992), 284–313.
- L. Shapley. 1953a. *Order Matrices. I*. Technical Report RM-1142. The RAND Corporation.

- L. Shapley. 1953b. *Order Matrices. II*. Technical Report RM-1145. The RAND Corporation.
- L. Shapley. 1964. Some Topics in Two-Person Games. In *Advances in Game Theory*, M. Dresher, L. S. Shapley, and A. W. Tucker (Eds.). Annals of Mathematics Studies, Vol. 52. Princeton University Press, 1–29.
- J. von Neumann. 1928. Zur Theorie der Gesellschaftspiele. *Math. Ann.* 100, 1 (1928), 295–320.
- J. von Neumann and O. Morgenstern. 1944. *Theory of Games and Economic Behavior*. Princeton University Press.
- R. B. Wilson. 1970. The Finer Structure of Revealed Preference. *Journal of Economic Theory* 2, 4 (1970), 348–353.