Proceedings of the Thirty-Fifth Australasian Computer Science Conference (ACSC 2012), Melbourne, Australia

On the parameterized complexity of dominant strategies

Vladimir Estivill-Castro

Mahdi Parsa

School of ICT, Griffith University, Queensland, Australia 4111 Email: {v.estivill-castro, m.parsa}@griffith.edu.au

Abstract

In game theory, a strategy for a player is dominant if, regardless of what any other player does, the strategy earns a better payoff than any other. If the payoff is strictly better, the strategy is named strictly dominant, but if it is simply not worse, then it is called weakly dominant.

We investigate the parameterized complexity of two problems relevant to the notion of domination among strategies. First, we study the parameterized complexity of the Minimum Mixed Dominating Strategy Set problem, the problem of deciding whether there exists a mixed strategy of size at most k that dominates a given strategy of a player. We show that the problem can be solved in polynomial time on win-lose games. Also, we show that it is a fixed-parameter tractable problem on rsparse games, games where the payoff matrices of players have at most r nonzero entries in each row and each column. Second, we study the parameterized complexity of the ITERATED WEAK DOMINANCE problem. This problem asks whether there exists a path of at most k-steps of iterated weak dominance that eliminates a given pure strategy. We show that this problem is W[2]-hard, therefore, it is unlikely to be a fixed-parameter tractable problem.

Keywords: Algorithm, Complexity, Computational game theory, dominant strategies, parameterized complexity

1 Introduction

Game theory is a mathematical framework for the study of conflict and cooperation between intelligent agents. This theory offers models to study decision-making situations and proposes several long-standing solution concepts.

A game consists of a set of players, a set of strategies for each player, and a specification of payoffs for each combination of strategies. Each single strategy in the set of strategies of a player is called a pure strategy. However, if a player randomly chooses a pure strategy, we say that the player is using a mixed strategy. In each game, players want to optimize their payoff which depends both on their own choices and also the choices of others.

Here, we use the Prisoners' Dilemma, a classical example in game theory, to introduce the concept of dominant strategy. In this game, two prisoners, the row player and the column player, are collectively charged with a crime and held in separate cells with no way of communicating. Each prisoner has two choices, cooperate (C) which means not defect his partner or defect (D), which

Copyright ©2012, Australian Computer Society, Inc. This paper appeared at the 35th Australasian Computer Science Conference (ACSC 2012), Melbourne, Australia, January-February 2012. Conferences in Research and Practice in Information Technology (CRPIT), Vol. 122, Mark Reynolds and Bruce Thomas, Ed. Reproduction for academic, not-for-profit purposes permitted provided this text is included.

Table 1: Payoff matrix of the players in Prisoners' Dilemma.

column player

row player
$$\begin{array}{c|cccc} & C & D \\ C & -1, -1 & -10, 0 \\ D & 0, -10 & -9, -9 \end{array}$$

means betray his partner. The punishment for the crime is ten years of prison. Betrayal gives a reduction of one year for the confessor. If a prisoner is not betrayed by its partner, he is convicted to one year for a minor offense. This situation can be summarized in Table 1. The numbers in the table represent the payoff for the players and there are two payoffs at each position: by convention the first number is the payoff for the row player and the second number is the payoff for the column player. In this game, the strategy defecting (D) gives a better payoff for both players no matter how that player's opponents may play. Strategies like the strategy (D) are called *dominant* strategies.

The notion of dominant strategies is a more elementary than the well-known solution concept Nash equilibrium. A Nash equilibrium is a set of strategies, one for each player, such that all players have no incentive to unilaterally change their decision. The elimination of dominated strategies can be used as a preprocessing technique for computing Nash equilibria. For example, in the Prisoners' Dilemma, after eliminating the dominated strategies C for each player, the remaining respective strategies D specify a Nash equilibrium.

Gilboa, Kalai and Zemel (Gilboa et al. 1993) used classical complexity theory and showed that many decision problems regarding to computation of dominant strategies are **NP**-complete for two-player games. Later, other researches (Brandt et al. 2009, Conitzer & Sandholm 2005) extended their hardness results to other classes of games such as win-lose games.

In this paper, we study two problems relevant to the notion of domination from the perspective of parameterized complexity. Hence, we are interested in algorithms that compute exact optimal solutions, while attempting to confine the inevitable exponential-running time of such algorithms to an input-length independent parameter.

First, we study the parameterized complexity of MIN-IMUM MIXED DOMINATING STRATEGY SET problem.

MINIMUM MIXED DOMINATING STRATEGY SET

Instance: Given the row player's payoffs of a two-player game \mathcal{G} and a distinguished pure strategy i of the row player.

Parameter: An integer k.

Question: Is there a mixed strategy x for the row player that places positive probability on at most k pure strategies, and dominates the pure strategy i?

A strategy may fail to be dominated by a pure strategy, but may be dominated by a mixed strategy. Here, we focus on specializations of two-player games (by revisiting the original NP-completeness proof (Conitzer & Sandholm 2005) we can discover that this is a parameterized reduction). Thus, the problem is W[2]-hard, and it is unlikely to be in FPT for general two-player games. First, we focus on win-lose games, games where the payoff values are limited to 0 and 1. We show that this problem can be solved in polynomial time on win-lose games (Lemma 3.4). Second, we investigate this problem on r-sparse games. Here the payoff matrices have at most r nonzero entries in each row and each column. We show that MINIMUM MIXED DOMINATING STRATEGY SET is fixed-parameter tractable for r-sparse games (Theorem 3.5).

Next, we study the parameterized complexity of the ITERATED WEAK DOMINANCE problem.

ITERATED WEAK DOMINANCE

Instance: A two-player game and a distinguished pure strategy i.

Parameter: An integer k.

Question: Is there a path of at most k steps of iterated weak dominance that eliminates the pure strategy i?

It is well-known that iterated strict dominance is path-independent, that is, the elimination process will always terminate at the same point, and the elimination procedure can be executed in polynomial time (Gilboa et al. 1993). In contrast, iterated weak dominance is path-dependent and it is known that whether a given strategy is eliminated in some path is NP-complete (Conitzer & Sandholm 2005). We show that this problem is W[2]-hard, therefore it is unlikely to be fixed-parameter tractable (Theorem 4.1).

The rest of the paper is organized as follows. In Section 2 we give formal definitions for games, and parameterized complexity theory. In Section 3, we show the fixed-parameter tractability results. In Section 4 we show our parameterized hardness results. In Section 5 we discuss further the implications of our results and some open problems.

2 Preliminaries

In this section, we review relevant concepts game theory, and computational complexity theory including parameterized complexity.

2.1 Parameterized complexity theory

Parameterized complexity aims at providing an alternative to exponential algorithms for **NP**-complete problems by identifying a formulation where the parameter would take small values in practice and shifting the exponential explosion to this parameter while the rest of the computation is polynomial in the size of the input.

A parameterized problem is a language $L\subseteq \Sigma^* \times \mathbb{N}$ where the second part of the problem is called the parameter. A parameterized problem L is fixed-parameter tractable if there is an algorithm that decides in $f(k)|x|^{O(1)}$ time whether $(x,k)\in L$, where f is an arbitrary computable function depending only on the parameter k. Such an algorithm is called FPT-time algorithm, and **FPT** denotes the complexity class that contains all fixed-parameter tractable problems.

In order to characterize those problems that do not seem to admit an FPT algorithm, Downey and Fellows (Downey & Fellows 1998) defined a *parameterized* reduction and a hierarchy of classes $W[1] \subseteq W[2] \subseteq \ldots$

including likely fixed parameter intractable problems. A (many-to-one) parameterized reduction from a parameterized problem L to a parameterized problem L' is an FPT-time mapping Φ that transforms an instance (x,k) of L into an instance (x',k') such that, $(x,k) \in L$ if and only if $(x',k') \in L'$, where k' bounded by some function depends only on the parameter k.

The class W[t] is defined to be the class of all problems that are reducible to a parameterized version of the satisfiability problem for Boolean circuits of weft t (see Downey and Fellows (Downey & Fellows 1998) for the exact definition).

The above classes may be equal (if **NP=P**, for example); however, there is evidence to suspect (Downey & Fellows 1998) that W[2]-completeness is a strong indication of intractability in the FPT sense. The best known algorithm for any W[2]-complete problem is still just the brute force algorithm of trying all k subsets which has a running time $O(n^{k+1})$.

2.2 Games and dominant strategies

A two-player normal form game \mathcal{G} consists of two matrices $A=(a_{ij})_{m\times n}$ and $B=(b_{ij})_{m\times n}$, where a_{ij} denotes the payoff for the first player and b_{ij} denotes the payoff for the second player when the first player plays his i-th strategy and the second player plays his j-th strategy. We identify the first player as the row player and the second player as the column player. Each single strategy in the set of strategies of a player is called a pure strategy. However, if a player randomly chooses a pure strategy, we say that the player is using a mixed strategy.

Definition 2.1 An ordered n-tuple $\mathbf{x} = (x_1, ..., x_n)$ with $\sum_{i=1}^{n} x_i = 1$ and $\mathbf{x} \ge 0$ is a mixed strategy.

Thus, a mixed strategy is a probability distribution over the pure strategy space. The *support* (denoted $supp(\mathbf{x})$) of a mixed strategy \mathbf{x} is the set of pure strategies which are played with positive probability, that is $\{i: 1 \leq i \leq n, x_i > 0\}$.

In a two-player game $\mathcal{G}=(A,B)$, a strategy i of the row player is said to weakly dominate a strategy i' of the row player if for every strategy j of the column player we have $a_{ij} \geq a_{i'j}$ and there exists a strategy j_0 of the column player that $a_{ij_0} > a_{i'j_0}$. The strategy i is said to strictly dominate strategy i' if for every strategy j of the column player $a_{ij} > a_{i'j}$. A similar definition is used to define the domination relation of the column player, but now using the payoff matrix B.

If a strategy is dominated, the game (and thus the problem) can be simplified by removing it. Eliminating a dominated strategy may enable elimination of another pure strategy that was not dominated at the outset, but is now dominated. The elimination of dominated strategies can be repeated until no pure strategies can be eliminated in this manner. In a finite game this will occur after a finite number of eliminations and will always leave at least one pure strategy remaining for each player. This process is called *iterated dominant strategies* (Gilboa et al. 1993).

Note that a strategy may fail to be strongly eliminated by a pure strategy, but may be dominated by a mixed strategy.

Definition 2.2 Consider strategy i of the row player in two-player game (A, B). We say that the strategy i is dominated by a mixed strategy $\mathbf{x} = (x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n)$ of the row player, if the following holds for every strategy j of the column player

$$\sum_{i'\neq i} x_{i'} a_{i'j} \ge a_{ij}.$$

3 FPT results on mixed strategy domination

Recall that sometimes a strategy is not dominated by any pure strategy, but it is dominated by some mixed strategies. Example 3.1 illustrates the differences between these two types of strategies.

Example 3.1 Consider the payoff matrix of the row player that is given as follows:

$$\left(\begin{array}{ccc} 4 & 0 & 2 \\ 0 & 4 & 0 \\ 1 & 1 & 0 \end{array}\right).$$

In this situation, no pure strategy can eliminate any other. However playing the first and the second strategy with probability 1/2, dominates the third strategy. Because, the expected payoff of those two strategies is equal to $1/2 \cdot (4,0,2) + 1/2 \cdot (0,4,0) = (2,0,1) + (0,2,0) = (2,2,1)$.

Moreover, we can test in polynomial time, whether a given strategy of a player is dominated by a mixed strategy of the same player. The following proposition shows the tractability of this issue.

Proposition 3.2 Consider a two-player game $\mathcal{G} = (A_{m \times n}, B_{m \times n})$, a subset S' of the row player's pure strategies, and a distinguished strategy i for the row player. We can determine in polynomial time (in the size of the game) whether there exists a mixed strategy x, that places positive probability only on strategies in S' and dominates the pure strategy i. Similarly, for the column player, a subset S' of the column player's pure strategies, and a distinguished strategy j for the column player. We can determine in polynomial time (in the size of the game) whether there exists a mixed strategy y, that places positive probability only on strategies in S' and dominates the pure strategy j. This applies both for strict and weak dominance (Conitzer & Sandholm 2005).

Nevertheless, finding such a mixed strategy that dominates a pure strategy with the smallest support size (MINIMUM MIXED DOMINATING STRATEGY SET) is computationally hard (NP-complete) (Conitzer & Sandholm 2005). Moreover, it is not hard to obtain a proof of W[2]-hardness for this problem. The original proof (Conitzer & Sandholm 2005) introduced a reduction from SET COVER, a W[2]-complete problem, to MINIMUM MIXED DOMINATING STRATEGY SET. We just need to verify that it is a parameterized reduction. Furthermore, this W-hardness result shows that it is unlikely to find an FPT algorithm for this problem by considering only the size of the support as the parameter.

Moreover, the review of the proof reveals that the constructed instances of the MINIMUM MIXED DOMINATING STRATEGY SET problem in the reduction have limited payoffs, which are $\{0,1,k+1\}$. Therefore, a natural question to ask next is whether it is possible to find an FPT algorithm by considering extra conditions on the problem instances. Our first step would be specializing the games to win-lose games. Recall that in win-lose games, the given payoffs are in $\{0,1\}$. The following lemma shows that this restriction makes the problem easy.

Lemma 3.3 In a win-lose game G=(A, B) every pure strategy that is weakly dominated by a mixed strategy is also weakly dominated by a pure strategy.

Proof: Consider a mixed strategy $\mathbf{x} = (x_1, x_2, \dots, x_m)$ that dominates a pure strategy i (without loss of generality, both of course, of the row player). Clearly, for any strategy j of the column player where $a_{ij} = 0$, the expected payoff of playing the mixed strategy in the

column j is at least 0. Therefore, we only need to consider columns j where $a_{ij}=1$. Let j_0 be a first column where $a_{ij_0}=1$. Because $\mathbf x$ dominates the strategy i there is a row (strategy) i in the mixed strategy i where i where i we claim that row i weakly dominates row i. We just need to show that i for any column i where i where i however, if i however, if i hen for the i-th column we have i have i

Lemma 3.4 MINIMUM MIXED DOMINATING STRATEGY SET is in **P** (that is, it can be decided in polynomial time) if it is limited to win-lose games.

Proof: By Lemma 3.3, if a pure strategy i is dominated by a mixed strategy \mathbf{x} , then there exits a pure strategy i' that dominates i. Therefore, the problem reduces to the problem of finding a pure strategy that dominates i. This can be done in polynomial time in the size of the game. \square Our first effort for specializing the problem makes it an easy problem (class \mathbf{P}). Therefore, instead of limiting the payoffs, we will work on limiting the number of non-zero entries in each row and each column of the payoff matrix of the row player. The MINIMUM MIXED DOMINATING STRATEGY SET problem remains \mathbf{NP} -complete even on r-sparse games with $r \geq 3$ (Conitzer & Sandholm 2005, Garey & Johnson 1979).

Theorem 3.5 MINIMUM MIXED DOMINATING STRATEGY SET problem for r-sparse games (when considering r as the parameter) is in the class FPT.

Proof: Consider an r-sparse instance of MINIMUM MIXED DOMINATING STRATEGY SET. Without loss of generality we can assume the last row of the first player is the strategy to be dominated by a mixture of another k strategies. Because of Proposition 3.2, finding a mixed strategy that weakly dominates the distinguished strategy reduces to the problem of determining the support of the mixed strategy. Consider the following procedure.

- **Step 1:** We remove (in polynomial time) all columns where the last row has a zero payoff. Because, all payoffs are at least zero in each column, any mixed strategy that dominates those columns with positive entries of the distinguished strategy also does so where the distinguished strategy has zeros. As the game is *r*-sparse, this step reduces the size of the payoff matrix of the row player to a matrix with *r* columns.
- **Step 2:** If there is a column where all entries in that column are less than the last entry in the column, then the instance is a no-instance.
- **Step 3:** Now remove all rows that are made completely of zeros. Because there are at most r entries different than zero in each column, the matrix now has at most r^2 rows. We can test exhaustively all subsets of rows of size k of the first r^2-1 rows for domination of the now r^2 -th row. If none of the tests results in domination, we have a no-instance, otherwise we have a yes-instance and a certificate of the domination.

The only step that is not polynomial is the exhaustive verification at the end; however, this is polynomial in r as there are $\binom{r^2-1}{k} = O(r^{2k})$ such subsets. This problem can be solved in f(r)poly(n) because $k < r^2$.

4 ITERATED WEAK DOMINANCE (IWD)

As discussed earlier, iterated elimination of strictly dominated strategies is conceptually straightforward in a sense that regardless of the elimination order the same set of strategies will be identified, and all Nash equilibria of the original game will be contained in this set. However, this process becomes a bit trickier with the iterated elimination of weakly dominated strategies. In this case, the elimination order does make a difference, that is, the set of strategies that survive iterated elimination can differ depending on the order in which dominated strategies are eliminated. Therefore, the problem such as deciding whether a strategy can be eliminated in a path of iterated weakly dominated absorbed more attention. IT-ERATED WEAK DOMINANCE is a NP-complete problem (Conitzer & Sandholm 2005) even in games with payoffs in $\{(0,0),(0,1),(1,0)\}$ (Brandt et al. 2009). Here, we show its hardness in terms of parameterized complexity theory.

Theorem 4.1 *The* IWD STRATEGY ELIMINATION *problem is W*[2]-hard.

We prove this by providing a parameterized reduction from SET COVER. Therefore, consider an instance of SET COVER. That is, we are given a set $S = \{1, 2, \ldots, n\}$ and a family $\mathcal F$ of proper subsets S_1, \ldots, S_r that cover S (that is, $S_i \subset S$, for $i=1,\ldots,r$ and $S=\bigcup_{i=1}^r S_i$). The question is whether there is a sub-family of k or fewer sets in $\mathcal F$ that also covers S.

Our proof constructs, a game $\mathcal{G}=(A,B)$ and the question, whether the last row of the matrix A can be eliminated by iterated weak domination in k+1 or fewer steps. Because k is the parameter of the SET COVER instance and k'=k+1 is the parameter of the ITERATED WEAK DOMINANCE (IWD) this would be a parameterized reduction.

We start by describing the payoff matrices of the game $\mathcal{G}=(A,B)$. The number of rows of the matrices is $|\mathcal{F}|+1=r+1$. The number of columns is r+n+1.

We first describe the payoff matrix A of the row player. The last row of A will be

$$a_{r+1,j} = \begin{cases} 1, & j < n+r+1, \\ 0, & \text{otherwise.} \end{cases}$$

That is, this row has a 1 everywhere except for the last column.

The last column of A has a similar form.

$$a_{i,n+r+1} = \begin{cases} 1, & i < r+1, \\ 0, & \text{otherwise.} \end{cases}$$

That is, this column has a 1 everywhere except for the last row.

Now, the first block of r columns and r rows of A have a diagonal full with the value 0 and the value 1 everywhere else. We let the following entries of A defined by

$$a_{i,j} = \left\{ \begin{array}{ll} 1, & i \leq r \text{ and } j \leq r \text{ and } i \neq j, \\ 0, & i \leq r \text{ and } j \leq r \text{ and } i = j. \end{array} \right.$$

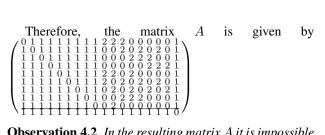
Finally, after the r-th column, the i-row has the characteristic vector of the set S_i scaled by k.

$$a_{i,j} = \left\{ \begin{array}{ll} k, & j-r \in S_i \ i \leq r \ \text{and} \ r+1 \leq j \leq r+n, \\ 0, & j-r \in S \setminus S_i, i \leq r \ \text{and} \ r+1 \leq j \leq r+n. \end{array} \right.$$

We illustrate this construction with an example. Consider the set $S = \{1, 2, 3, 4, 5, 6, 7, 8\}$, and the parame-

ter k=2. The family \mathcal{F} is defined as follows

$$\begin{split} S_1 &= \{1,2,3\}, \quad S_2 = \{3,5,7\}, \quad S_3 = \{4,5,6\}, \\ S_4 &= \{6,7,8\}, \quad S_5 = \{1,2,4\}, \quad S_6 = \{1,3,5,7\}, \\ S_7 &= \{2,4,6,8\}, \quad S_8 = \{3,4,5\}, \quad \text{and } S_9 = \{2\}. \end{split}$$



Observation 4.2 *In the resulting matrix* A *it is impossible to perform a row elimination to eliminate the* r + 1*-th row.*

Any convex combination of strategies in $\{1, 2, \ldots, r\}$ would add to less than one in one column in $\{1, 2, \ldots, r\}$. Thus, there would be one column blocking such elimination.

Observation 4.3 Consider a yes-instance of the SET COVER problem, where I is the set of indexes in $\{1,2,\ldots,r\}$ such that $|I| \leq k$ and $S \subseteq \bigcup_{i \in I} S_i$. Removing the columns in I from A results in a configuration where the linear combination of rows in I with probability 1/|I| eliminate row r+1 in one step.

To confirm this observation first note that any convex combination of rows in $\{1, 2, \ldots, r\}$ produces domination in the r+1-th column, and thus in particular the rows in I.

Now we show that the removal of the columns in I causes no longer a blockage. First, consider a column $j \leq r$. Since r is not in I, when we consider the convex combination of rows in I, that combination will add to a payoff of 1, which is equal to the value in row r+1 and column j.

Finally, consider a column j with r < j < r+n+1. Because I is the set of indexes of a cover, all entries in the rows indexed by I have value k in column j. Therefore the linear combination with uniform probability 1/|I| on the rows with index I will have at least one entry with weight $k/|I| \ge 1$ since $|I| \le k$.

To continue, we now need to describe the payoff matrix B for the column player. This matrix is made of two blocks. The first block is the first r columns, while the second block is the last n+1 columns. All values are 0 for the first block and all values are 1 for the second block.

$$B = (0_{r+1 \times r} | 1_{r+1 \times n+1}).$$

Observation 4.4 The only columns that can be eliminated by a column elimination are one of the first r columns.

This observation follows trivially from the structure of B, since the only dominations are strict dominations from a column in the later block of columns full of the value 1 to a column in the first r columns full of the value 0.

Observation 4.5 A row elimination cannot happen in matrix A until a set $I \subseteq \{1, 2, ..., r\}$ of columns is eliminated by column eliminations, and the set I defines a cover of S

We know the process of elimination must start with a column elimination. Because of the structure of the first r columns of A, the only row elimination possible after some columns eliminations must be a convex combination of a subset of indexes of the already eliminated indexes.

However, this would be a possible row elimination only if the linear combination also implies a set cover because of the structure of the next n columns of matrix A.

Now clearly if there is a path of length k+1 (or less) that eliminates row r+1 in matrix A it must consist of k (or less) column eliminations defining k (or less) indexes of the covering sets, and the last elimination is the corresponding row elimination with uniform weight on the same indexes. This completes the proof.

5 Conclusion

There are many interesting decision problems regarding the notion of dominant strategies (Gilboa et al. 1993, Conitzer & Sandholm 2005). We showed that one of those problems is parameterized hard problem (e.g. MINIMUM MIXED DOMINATING STRATEGY SET). Furthermore, we showed that special cases of those problems are in **P** or **FPT**. However, many other problems are still open. For example, the paper "On the Complexity of Iterated Weak Dominance in Constant-Sum Games" (Brandt et al. 2009) indicates that the ITERATED WEAK DOMINANCE problem is in P for Constant-Sum games, but still NPcomplete for win-lose games. In fact, the paper is more precise. If of each pair of actions (i, j), where i is for the row player and j is for the column player, the corresponding entries in \hat{A} and \hat{B} are only (1,0),(0,1),(0,0) the problem remains **NP**-complete. Disallowing (0,0) makes the problem restricted to constant-sum games and thus becomes a problem in P. However, our parameterized hardness proof uses other entries different from $\{0,1\}$. We do not know the fixed-parameter complexity of ITERATED WEAK DOMINANCE in win-lose games or in the more restricted class of games where only (1,0), (0,1), (0,0) are allowed.

References

- Brandt, F., Brill, M., Fischer, F. & Harrenstein, P. (2009), On the complexity of iterated weak dominance in constant-sum games, *in* 'SAGT-09, Proceedings of the 2nd International Symposium on Algorithmic Game Theory', Springer-Verlag, Berlin, Heidelberg, pp. 287–298.
- Conitzer, V. & Sandholm, T. (2005), Complexity of (iterated) dominance, *in* J. Riedl, M. J. Kearns & M. K. Reiter, eds, 'EC-05, Proceedings of the 6th ACM Conference on Electronic Commerce', ACM, Vancouver, BC, Canada, pp. 88–97.
- Downey, R. & Fellows, M. R. (1998), *Parameterized Complexity*, Monographs in Computer Science, Springer, New York.
- Garey, M. R. & Johnson, D. S. (1979), Computers and Intractability: A Guide to the Theory of NP-Completeness, W. H. Freeman & Co., New York, USA.
- Gilboa, I., Kalai, E. & Zemel, E. (1993), 'The complexity of eliminating dominated strategies', *Mathematics of Operations Research* **18**(3), 553–565.