Complexity Results for Some Classes of Strategic Games

Dissertation an der Fakultät für Mathematik, Informatik und Statistik der Ludwig-Maximilians-Universität München

vorgelegt von Felix Fischer

am 11.03.2009

Complexity Results for Some Classes of Strategic Games

Dissertation an der Fakultät für Mathematik, Informatik und Statistik der Ludwig-Maximilians-Universität München

vorgelegt von Felix Fischer

am 11.03.2009

Betreuer:	Dr. Felix Brandt
	Ludwig-Maximilians-Universität München
Erster Berichterstatter:	Prof. Martin Hofmann, Ph.D.
	Ludwig-Maximilians-Universität München
Zweiter Berichterstatter:	Prof. Lane A. Hemaspaandra, Ph.D.
	University of Rochester, NY, USA
Datum des Rigorosums:	03.07.2009

to you

Contents

At	ostrac	t	xi
Zu	samn	nenfassung	xiii
Ac	know	ledgements	xv
1	Intro	oduction	1
2	Gam	les, Solutions, and Complexity	5
	2.1	Strategic Games	5
	2.2	Solution Concepts	7
	2.3	Elements of Complexity Theory	10
	2.4	A Few Words on Encodings	20
3	State	e of the Art and Our Contribution	23
4	Ran	king Games	29
	4.1	An Introductory Example	30
	4.2	Related Work	31
	4.3	The Model	32
	4.4	Games With Non-Pure Equilibria	35
	4.5	Solving Ranking Games	37
	4.6	Comparative Ratios	46
	4.7	Discussion	55
5	Ano	nymous Games	57
	5.1	Related Work	58
	5.2	The Model	60
	5.3	Pure Nash Equilibria	63
	5.4	Iterated Weak Dominance	77
	5.5	Discussion	98

viii

6	Graj	phical Games	101				
	6.1	Related Work	102				
	6.2	The Model	103				
	6.3	3 A Tight Hardness Result for Pure Equilibria					
	6.4	Pure Equilibria of Graphical Games with Anonymity	111				
	6.5	Interlude: Satisfiability in the Presence of a Matching	123				
	6.6	Mixed Equilibria	124				
	6.7	Discussion	125				
7 Qu	Qua	si-Strict Equilibria	127				
	7.1	Related Work	128				
	7.2	Preliminaries	129				
	7.3	Two-Player Games	129				
	7.4	A Hardness Result for Multi-Player Games	131				
	7.5	Discussion	136				
8	Shap	pley's Saddles	137				
	8.1	Related Work	138				
	8.2	Preliminaries	138				
	8.3	Strict Saddles	140				
	8.4	Weak Saddles of Confrontation Games	141				
	8.5	A Hardness Result for Weak Saddles	146				
	8.6	Discussion	148				
Re	eferen	ces	149				
Le	bensl	auf	161				

List of Figures

2.1	Alice, Bob, or Charlie?	7
2.2	The Matching Pennies game	9
4.1	The game of Figure 2.1 as a single winner game	31
4.2	A ranking game form	34
4.3	A ranking game associated with the ranking game form of Figure 4.2	34
4.4	Four-player ranking game in which all equilibria are pure	37
4.5	Mapping from binary two-player games to three-player single-loser games $\ .$	39
4.6	Iterated weak dominance solvability in two-player ranking games	43
4.7	Three-player ranking game used in the proof of Theorem 4.7	44
4.8	Three-player ranking game used in the proof of Theorem 4.10	48
4.9	Three-player single-winner game used in the proof of Theorem 4.11	48
4.10	Three-player ranking game used in the proof of Theorem 4.11	49
4.11	Three-player ranking game used in the proof of Theorem 4.14	52
4.12	Four-player ranking game used in the proof of Theorem 4.14	53
4.13	Three-player ranking game $\Gamma_{\!5}$ used in the proof of Theorem 4.15 \ldots .	54
4.14	Dual linear program for computing a correlated equilibrium	55
5.1	Inclusion relationships between anonymous, symmetric, self-anonymous,	
	and self-symmetric games	61
5.2	Relationships between the payoffs of anonymous, symmetric, self-anony-	
	mous, and self-symmetric games	62
5.3	Anonymous game with a unique, non-symmetric Nash equilibrium	64
5.4	Matching problem for the game of Figure 5.3	65
5.5	Game used in the proof of Theorem 5.3	66
5.6	Integer flow network for the game of Figure 5.3	76
5.7	A matrix and a sequence of eliminations	80
5.8	Matrix Y used in the proof of Lemma 5.18	81
5.9	Overall structure of the graph used in the proof of Theorem 5.22	87
5.10	Variable gadget used in the proof of Theorem 5.22	88
5.11	Clause gadget used in the proof of Theorem 5.22	89
5.12	Gadget to consume remaining labels, used in the proof of Theorem 5.22	90

х	

5.13 5.14	Labeled graph for the matrix elimination instance of Figure 5.7 Payoffs of a player in a self-anonymous game with three players and three	. 91
	actions	. 96
5.15	Payoff structure of a player of the self-anonymous game used in the proof	07
	of Theorem 5.27	. 97
6.1	Payoffs for input, positive literal, and negative literal players, used in the	
	proof of Theorem 6.3	106
6.2	Payoffs for AND and OR players, used in the proof of Theorem 6.3	. 106
6.3	Payoffs for NAND players, used in the proof of Theorem 6.11	. 112
6.4	Output gadget, used in the proof of Theorem 6.11	. 114
6.5	Equality gadget, used in the proof of Theorem 6.12	. 115
6.6	NAND gadget, used in the proof of Theorem 6.12	116
6.7	Neighborhood graph of a graphical game with seven players, corresponding	
	to the hypergraph given by the lines of the Fano plane \ldots \ldots \ldots	. 117
6.8	Graphical game with eight players and neighborhoods of size four, used in	
	the proof of Theorem 6.17	120
6.9	NOR gadget, used in the proof of Theorem 6.17	123
7.1	Single-winner game, repeated from Figure 4.1	. 128
7.2	Linear programs for computing minimax strategies in zero-sum games	130
7.3	Linear program for computing quasi-strict equilibria in zero-sum games	131
7.4	Linear program for computing quasi-strict equilibria in symmetric zero-sum	
	games	132
7.5	Somebody has to do the dishes	132
7.6	Payoff structure of a symmetric game with two actions	133
7.7	Three-player game used in the proof of Theorem 7.5	. 135
8.1	Strict and weak saddles of a zero-sum game	. 139
8.2	Symmetric zero-sum game with multiple weak saddles	. 140
8.3	Payoff structure of a $4k\times 4k$ symmetric zero-sum game with at least 5^k	
	weak saddles	. 146

Abstract

Game theory is a branch of applied mathematics studying the interaction of self-interested entities, so-called agents. Its central objects of study are games, mathematical models of real-world interaction, and solution concepts that single out certain outcomes of a game that are meaningful in some way. The solutions thus produced can then be viewed both from a descriptive and from a normative perspective. The rise of the Internet as a computational platform where a substantial part of today's strategic interaction takes place has spurred additional interest in game theory as an analytical tool, and has brought it to the attention of a wider audience in computer science.

An important aspect of real-world decision-making, and one that has received only little attention in the early days of game theory, is that agents may be subject to resource constraints. The young field of algorithmic game theory has set out to address this shortcoming using techniques from computer science, and in particular from computational complexity theory. One of the defining problems of algorithmic game theory concerns the computation of solution concepts. Finding a Nash equilibrium, for example, i.e., an outcome where no single agent can gain by changing his strategy, was considered one of the most important problems on the boundary of P, the complexity class commonly associated with efficient computation, until it was recently shown complete for the class PPAD. This rather negative result for general games has not settled the question, however, but immediately raises several new ones: First, can Nash equilibria be approximated, i.e., is it possible to efficiently find a solution such that the potential gain from a unilateral deviation is small? Second, are there interesting classes of games that do allow for an exact solution to be computed efficiently? Third, are there alternative solution concepts that are computationally tractable, and how does the value of solutions selected by these concepts compare to those selected by established solution concepts?

The work reported in this thesis is part of the effort to answer the latter two questions. We study the complexity of well-known solution concepts, like Nash equilibrium and iterated dominance, in various classes of games that are both natural and practically relevant: ranking games, where outcomes are rankings of the players; anonymous games, where players do not distinguish between the other players in the game; and graphical games, where the well-being of any particular player depends only on the actions of a small group other players. In ranking games, we further compare the payoffs obtainable in Nash equilibrium outcomes with those of alternative solution concepts that are easy to compute. We finally study, in general games, solution concepts that try to remedy some of the shortcomings associated with Nash equilibrium, like the need for randomization to achieve a stable outcome.

xii

Zusammenfassung

Die Spieltheorie ist ein Teilgebiet der angewandten Mathematik, das sich mit der Interaktion eigennütziger Akteure, so genannter Agenten, beschäftigt. Sie untersucht dazu Spiele, mathematische Modelle in der realen Welt auftretender Interaktion, und Lösungskonzepte, die bedeutsame Ergebnisse eines Spieles hervorheben. Die so erhaltenen Lösungen können dann sowohl aus deskriptiver als auch aus normativer Sicht betrachtet werden. Der Aufstieg des Internet zu einer Umgebung, die für einen erheblichen Teil heutiger strategischer Interaktion verantwortlich ist, hat das Interesse an Spieltheorie als analytischem Werkzeug weiter vorangetrieben und ihr insbesondere zu einem erhöhten Bekanntheitsgrad in der Informatik verholfen.

Einem wichtigen Aspekt der Entscheidungsfindung in der realen Welt wurde in den Gründungstagen der Spieltheorie nur geringe Aufmerksamkeit zu Teil. Er betrifft die Tatsache, dass Agenten üblicherweise Ressourcenbeschränkungen unterliegen. Das junge Gebiet der algorithmischen Spieltheorie hat nun damit begonnen, sich diesem Defizit mittels Techniken der Informatik, und insbesondere der Komplexitätstheorie, zu widmen. Ein zentrales Problem der algorithmischen Spieltheorie stellt die Berechnung von Lösungskonzepten dar. Das Finden eines Nash-Gleichgewichts, d.h. eines Ergebnisses, in dem kein Agent durch Änderung seiner eigenen Strategie eine Verbesserung erreichen kann, galt beispielsweise als eines der wichtigsten Probleme an der Grenze der üblicherweise mit effizienter Berechnung verbundenen Komplexitätsklasse P, bis kürzlich seine Vollständigkeit für die Klasse PPAD gezeigt wurde. Dieses eher als negativ einzuschätzende Resultat in Bezug auf allgemeine Spiele hat die Frage an sich jedoch keineswegs vollständig beantwortet, sondern wirft umgehend neue Fragen auf: Können Nash-Gleichgewichte angenähert werden, d.h., kann effizient eine Lösung berechnet werden, die den möglichen Gewinn durch einseitige Abweichung gering hält? Gibt es interessante Teilklassen von Spielen, die die effiziente Berechnung exakter Lösungen erlauben? Existieren schließlich alternative, effizient berechenbare Lösungskonzepte, und wie verhält sich ihr Nutzen zu dem etablierter Lösungskonzepte.

Diese Arbeit beschäftigt mit den beiden letzteren Fragen. Wir untersuchen dazu die Komplexität bekannter Lösungskonzepte, wie Nash-Gleichgewicht und iterierte Dominanz, in verschiedenen natürlichen und praktisch relevanten Klassen von Spielen: Ranglistenspielen, in denen jedes Ergebnis eine Rangliste der Spieler ist; anonymen Spielen, in denen die Spieler nicht zwischen anderen Spielern unterscheiden; und graphischen Spielen, bei denen das Wohlergehen eines bestimmten Spielers nur von einem kleinen Teil der anderen Spieler abhängt. In Ranglistenspielen vergleichen wir außerdem den Nutzen von Ergebnissen im Nash-Gleichgewicht mit denen alternativer, effizient berechenbarer Lösungskonzepte. Schließlich betrachten wir in allgemeinen Spielen Lösungskonzepte, die einige mit Nash-Gleichgewichten verbundene Schwächen zu beheben suchen, wie etwa die Notwendigkeit von Randomisierung zum Erreichen eines stabilen Ergebnisses.

Acknowledgements

To my supervisor Felix Brandt, for allowing me to roam freely and talk impudently, for knowing an answer to every question and a question to every answer. I take pride in being his first Ph.D. student, there will no doubt be many more. To Lane Hemaspaandra and Martin Hofmann, for agreeing to act as referees for this thesis. To my coauthors and collaborators, for teaching me everything I know about research. Obviously, this thesis would not have been written without them. To Markus Brill, Hermann Gruber, and Paul Harrenstein for proofreading parts of this thesis and providing valuable feedback. To Jeff Rosenschein, for inviting me to work in his group at the Hebrew University of Jerusalem and take part in the family Seder. To my friends in and outside academia, for the fun we had. To my family, for supporting me in every imaginable way.

Chapter 1

Introduction

One does not have to be a follower of Machiavelli, nor a pessimist, to acknowledge that the world is pervaded by conflict. Conflicts arise almost automatically whenever a situation may lead to different outcomes depending on the choices of several individuals, or *agents*, who disagree about some aspect of these outcomes. As a consequence, the study of conflicts among a group of agents is at the core of many academic disciplines within the social sciences, like sociology, economics, or political science. Ultimately, this study is driven by two questions. From an individual's point of view, what is the best way to act in a given situation in order to achieve one's goals? From the point of view of society as a whole, how can the negative effects of conflicts be alleviated, in order to achieve outcomes that benefit all.

In 1944, von Neumann and Morgenstern broke ground to a rigorous mathematical treatment of the study of conflict with their work "Theory of Games and Economic Behavior." Game theory studies mathematical models, called games, that abstract from conflict situations in the real world and focus on certain aspects that seem worth studying. Consider a situation in which agents move sequentially, taking turns in a particular way that may or may not depend on previous choices or include an element of chance. One way to represent such a situation by a game of strategy, commonly referred to as the strategic form or normal form of a game, starts from a set of *players*, and a set of *strate*gies for each of the players. Each of these strategies provides a player with a complete plan of action for any eventuality that might arise in the course of the game. A strategy profile, consisting of exactly one strategy for each player, thus completely determines the outcome of the game and leads to one of several possible outcomes. The force that drives agents' behavior in the real world is modeled by having players entertain preferences over outcomes. Rational behavior is then characterized in terms of solution concepts that single out certain strategy profiles. A famous example is the solution concept of Nash equilibrium, which requires the strategies of the different players to be best responses to each other, such that no agent can achieve a preferred outcome by unilaterally changing his own strategy. In general, a solution concept can be viewed as *prescribing* a certain behavior by each agent in a given situation, but also as *describing* the outcome that will arise from the interaction of rational agents in this situation.

An important technique in the analysis of games is their classification according to natural parameters like the number of players and actions, but also according to the structure of players' preferences. Well-known classes, which also contain many games relevant in practice, are those of (two-player) zero-sum games, in which the interests of two players are diametrically opposed, or anonymous games, in which no distinction is made between different other players.

Like any area of applied mathematics, game theory walks the thin red line that is the appropriate level of abstraction: abstract enough to be handled in a rigorous way, and general enough such that the results thus obtained are sufficiently interesting and relevant to the real world. An important aspect of real-world decision-making, and one that has received only little attention in the "early days" of game theory, is that decision-makers may be subject to resource constraints. Game theory avoids this issue by assuming perfect rationality, i.e., has each agent choose an action that given a certain state of knowledge leads to his most preferred outcome. The importance of the issue of resource bounded reasoning, however, has by no means escaped the attention of game theorists. Nobel laureate Robert Aumann for example expressed the following opinion in an interview with van Damme (1998, pp. 201–202):

It is important to have an applicable model. It sounds a little like the man who had lost his wallet and was looking under the lamppost for it. His friend asked him: Why do you look only under the lamppost? And he answered: That's because there is light there, otherwise I wouldn't be able to see anything. It sounds crazy, but when you look at it more closely it is really important. If you have a theory that somehow makes a lot of sense, but is not calculable, not possible to work with, then what's the good of it? As we were saying, there is no "truth" out there; you have to have a theory that you can work with in applications, be they theoretical or empirical. ...

... My own viewpoint is that inter alia, a solution concept must be calculable, otherwise you are not going to use it.

Calculability as used by Aumann appears to be grounded in an informal notion of convenience experienced by humans working with a given solution concept. It is thus not immediately clear how it should be treated in a rigorous way. Enter another field that was greatly influenced by work of von Neumann (1945): computer science. With merely a year between two key publications in either field, both of which were authored or coauthored by von Neumann, it took another half century until *algorithmic game theory* set out to reveal deep connections between game theory and computer science (see, e.g., Nisan et al., 2007). One of these connections became obvious by the rise of the Internet as a computational platform and spurred the interest of computer scientists in game theory as a framework to analyze interaction between self-interested entities. Another one concerns Aumann's calculability, and involves the areas of algorithms and complexity theory. Complexity theory classifies problems according to the resources some idealized computational device requires to solve them. In the case of the Turing machine, the resulting notion of complexity is particularly meaningful: "Despite its weak and clumsy appearance, the Turing machine can simulate arbitrary algorithms with inconsequential loss of efficiency" (Papadimitriou, 1994a, p. 19).

Papadimitriou (2007) recasts Aumann's statement in the framework of computational complexity theory, and further argues that this statement still applies if one subscribes to a purely descriptive or analytical view of solution concepts:

But why should we be interested in the issue of computational complexity in connection to Nash equilibria? After all, a Nash equilibrium is above all a conceptual tool, a prediction about rational strategic behavior by agents in situations of conflict—a context that is devoid of computation.

We believe that this matter of computational complexity is one of central importance here, and indeed that the algorithmic point of view has much to contribute to the debate of economists about solution concepts. The reason is simple: If an equilibrium concept is not efficiently computable, much of its credibility as a prediction of the behavior of rational agents is lost—after all, there is no clear reason why a group of agents cannot be simulated by a machine. Efficient computability is an important modeling prerequisite for solution concepts. (Papadimitriou, 2007, pp. 29–30)

For the very reasons outlined above, the computational complexity of game-theoretic solution concepts has come under increased scrutiny. The work reported on in this thesis is part of the endeavor to study algorithmic aspects of game-theoretic solutions. We concentrate on the very general class of games in normal form, and various natural subclasses. In the following chapter, we formally define normal-form games along with several wellknown solution concepts, and lay out certain elements of the framework of computational complexity theory. In Chapter 3, we then take a closer look at the current state of the art and outline the contribution of this thesis. Since many problems have turned out to be hard for general games, researchers have considered various restricted classes of games. This is also the approach we take for the main part of this thesis. In Chapters 4 through 6, we study the computational complexity of various game-theoretic solution concepts in four natural classes of normal-form games. A more detailed account of the significance of each individual class and relevant existing work will be given in the respective chapter. Finally, in Chapters 7 and 8, we consider two solution concepts that are less well-known and try to address two shortcomings of Nash equilibrium: the potential indifference between actions that are played and actions that are not played, and the need for randomness in the choice of actions.

Chapter 2

Games, Solutions, and Complexity

In this chapter we review relevant concepts from game theory and computational complexity theory. While parts of this chapter also provide a high-level overview of the objects this thesis is concerned with, and the techniques used to analyze them, its main purpose is to lay out in detail the formal framework for the results that appear in later chapters. It may therefore be advisable to skip some of the technical details for the time being, and return to this chapter later for reference. For additional details we refer to the textbooks of Myerson (1991) and Osborne and Rubinstein (1994) on game theory, and of Papadimitriou (1994a), Goldreich (2008), and Vollmer (1999) on complexity theory.

2.1 Strategic Games

This thesis is concerned with finite games in normal form. Such a game is given by a finite set of players, and a finite non-empty set of actions for each player. Players move simultaneously to select an action profile, containing exactly one action for each player. This selection leads to a unique outcome, and each player entertains preferences over the set of possible outcomes. A standard assumption, which we also follow in this thesis, restricts attention to von Neumann-Morgenstern preferences over lotteries over outcomes (von Neumann and Morgenstern, 1944). Under this assumption, the preferences of a player can be represented by a real-valued payoff function from the set of action profiles into the reals. Two games are equivalent if there exist bijections between their respective sets of players and actions, and if the corresponding payoff functions can be obtained from each other via positive affine transformations. All solution concepts we consider are invariant under such transformations. We arrive at the following definition (e.g., Myerson, 1991).

DEFINITION 2.1 (normal-form game). A (normal-form) game is given by a tuple $\Gamma = (N, (A_i)_{i \in N}, (p_i)_{i \in N})$ where N is a finite set of *players*, and for each player $i \in N$, A_i is a nonempty and finite set of *actions* available to i and $p_i : (X_{i \in N}A_i) \to \mathbb{R}$ is a function mapping each action profile, i.e., each combination of actions, to a real-valued *payoff* for i.

We write $A_N = X_{i \in N} A_i$ for the set of action profiles and n = |N| for the number of players in a game. Subscripts will generally be used to identify the player to which an action belongs, superscripts to index the actions of a particular player. For example, we write a_i for a typical action of player i, and a_i^j for the jth action of player i. In the case of games with few players, or when we do not explicitly distinguish between specific players, we also use lower case roman letters a^j , b^j , etc., for the players' actions. An action profile $(a_i)_{i \in N} \in A_N$ we abbreviate by a_N .

A necessary condition for studying the computational properties of games is that these games have a finite representation. We therefore restrict our attention to games whose payoffs are *rational* numbers, and simply refer to these as "games" throughout the thesis. We further call a game *binary* if $p_i(a_N) \in \{0, 1\}$ for all $i \in N$ and $a_N \in A_N$. A two-player game $(\{1, 2\}, (A_1, A_2), (p_1, p_2))$ is alternatively called a bimatrix game, because it can be represented by two matrices M_1 and M_2 with rows and columns indexed by A_1 and A_2 , respectively, and $M_i(a_1, a_2) = p_i(a_1, a_2)$ for $i \in \{1, 2\}$ and all $a_1 \in A_1$, $a_2 \in A_2$. A two-player game satisfying $p_1(a, b) = -p_2(a, b)$ for all $(a, b) \in A_1 \times A_2$ is called *zero-sum game* or matrix game, and can be represented by a single matrix M containing the payoffs for the first player. Since all solution concepts considered in this thesis are invariant under positive affine transformations, the results about zero-sum games in fact apply to the larger class of *constant-sum* games, in which the payoffs of the two players always sum up to the same constant. For games with more than two players, this property is far less interesting, as we can always add an extra player who "absorbs" the payoffs of the others (von Neumann and Morgenstern, 1947).

The concept of an action profile can be generalized to that of a mixed strategy profile by letting players randomize over their actions. We have $S_i = \Delta(A_i)$ denote the set of probability distributions over player i's actions, the mixed strategies available to player i, and $S_N = X_{i \in N} S_i$ the set of mixed strategy profiles. Analogously to action profiles, we abbreviate a strategy profile $(s_i)_{i \in N} \in S_N$ by s_N . In the following, A_{-i} and S_{-i} respectively denote the set of action and strategy profiles for all players but i. Accordingly, we write $a_{-i} \in A_{-i}$ for the vector of all actions in a_N but a_i , and $s_{-i} \in S_{-i}$ for the vector of all strategies in s_N but s_i . We further denote by $s_i(a_i)$ and $s_N(a_i)$ the probability player i assigns to action a_i in strategy s_i or strategy profile s_N . The pure strategy s_i such that $s_i(a_i) = 1$ we identify with a_i whenever this causes no confusion. Moreover, we use (s_{-i}, t_i) to refer to the strategy profile obtained from s_N by replacing s_i by t_i . Payoff functions naturally extend to mixed strategy profiles, and we will frequently write $p_i(s_N) = \sum_{a_N \in A_N} p_i(a_N) (\prod_{i \in N} s_i(a_i))$ for the expected payoff of player i, and $p(s_N) =$ $\sum_{i \in N} p_i(s_N)$ for the social welfare under strategy profile $s_N \in S_N$. For better readability we usually avoid double parentheses and write, e.g., $p(s_{-i}, t_i)$ instead of $p((s_{-i}, t_i))$.

To illustrate these concepts, consider a situation in which Alice, Bob, and Charlie are to designate one of them as the winner. They do so by raising their hand or not, simultaneously and independently of one another. Alice wins if the number of hands raised, including her own, is odd, whereas Bob is victorious if this number equals two.

c ¹		c ²			
	b1	b^2		b1	b^2
a ¹	(0,0,1)	(1,0,0)		(1,0,0)	(0,1,0)
a ²	(1,0,0)	(0,1,0)		(0,1,0)	(1,0,0)

Figure 2.1: Alice, Bob, or Charlie? Alice chooses row a^1 or a^2 , Bob chooses column b^1 or b^2 , and Charlie chooses matrix c^1 or c^2 . Outcomes are denoted by a vector where the ith component is the payoff to player i.

Should nobody raise their hand, Charlie wins. The normal form of this game is shown in Figure 2.1. Player 1, Alice, chooses between rows of the table, labeled a^1 and a^2 . Action a^1 corresponds to her *not* raising her hand, whereas a^2 corresponds to her raising her hand. Similarly, player 2, Bob, chooses between the left or right column, labeled b^1 and b^2 , and player 3, Charlie, between the left or right matrix, labeled c^1 and c^2 . Outcomes are denoted as vectors of payoffs, the ith component corresponding to the payoff of player i. The fact that a player wins or loses is represented by a payoff of one or zero, respectively. For example, the top right entry in the left matrix corresponds to the action profile (a^1, b^2, c^1) where only Bob raises his hand, which in turn causes Alice to win.

2.2 Solution Concepts

Now that we have found a way to formalize the conflict between Alice, Bob, and Charlie, how should they play in order to be successful? Game theory tries to answer this question in a general way by providing a number of *solution concepts*. On a normative interpretation, solution concepts identify reasonable, desirable, or otherwise significant strategy profiles in games.

Perhaps the most cautious way for a player to proceed is to ensure a certain minimum payoff even if all other players were to conspire against him.

DEFINITION 2.2 (maximin strategy and security level). Let $\Gamma = (N, (A_i)_{i \in N}, (p_i)_{i \in N})$ be a normal-form game. A strategy $s_i^* \in S_i$ is called a *maximin strategy* for player $i \in N$ if

$$s_i^* \in \underset{s_i \in S_i}{\operatorname{argmax}} \min_{s_{-i} \in S_{-i}} p_i(s_i, s_{-i}).$$

The value $v_i = \max_{s_i \in S_i} \min_{s_{-i} \in S_{-i}} p_i(s_i, s_{-i})$ is called the *security level* of player i.

Given a particular game Γ , we write $\nu_i(\Gamma)$ for the security level of player i in Γ . In the game of Figure 2.1, Alice can guarantee a payoff of at least 1/2 by uniformly randomizing over her actions, i.e., by raising her hand with probability 1/2. We leave it to the reader to verify that this is indeed her security level, and that the security level for both Bob and Charlie is zero.

A different way to identify desirable strategy profiles is to rule out those that are *not* desirable. An action of a particular player in a game is said to be *weakly dominated* if there exists a strategy guaranteeing him at least the same payoff for any profile of actions of the other players, and strictly more payoff for some such action profile. A dominated action may be discarded for the simple reason that the player will never face a situation where he would benefit from using this action. The removal of one or more dominated actions from the game may render additional actions dominated, and the solution concept of *iterated dominance* works by removing a dominated action and applying the same reasoning to the reduced game.

DEFINITION 2.3 (iterated dominance). Let $\Gamma = (N, (A_i)_{i \in N}, (p_i)_{i \in N})$ be a game. An action $d_i \in A_i$ is said to be *(weakly)* dominated by strategy $s_i \in S_i$ if for all $b \in X_{i \in N}A_i$, $p_i(b_{-i}, d_i) \leq \sum_{a_i \in A_i} s_i(a_i)p_i(b_{-i}, a_i)$ and for at least one $\hat{b} \in X_{i \in N}A_i$, $p_i(\hat{b}_{-i}, d_i) < \sum_{a_i \in A_i} s_i(a_i)p_i(\hat{b}_{-i}, a_i)$.

An elimination sequence of Γ is a finite sequence of actions in $\cup_{i \in N} A_i$. For a particular elimination sequence $d = (d^1, d^2, \dots, d^k)$ denote by $\Gamma(d)$ the induced subgame where the actions in d have been removed, i.e., $\Gamma(d) = (A'_1, A'_2, u|_{A'_1 \times A'_2})$ where $A'_1 = A_1 \setminus \{d^1, d^2, \dots, d^k\}$ and $A'_2 = A_2 \setminus \{d^1, d^2, \dots, d^k\}$. Then, an elimination sequence $d = (d^1, d^2, \dots, d^m)$ of Γ is called *valid* if either it is the empty sequence, or if $(d^1, d^2, \dots, d^{m-1})$ is valid in Γ and d^m is weakly dominated in $\Gamma(d^1, d^2, \dots, d^{m-1})$.

An action $a \in \bigcup_{i \in N} A_i$ is called *eliminable* if there exists a valid elimination sequence d such that a is weakly dominated in $\Gamma(d)$. Game Γ is called *solvable* if it is possible to obtain a game where only one action remains for each player, i.e., if there is some valid elimination sequence d such that $\Gamma(d) = (N, (A'_i)_{i \in N}, (p'_i)_{i \in N})$ with $|A'_i| = 1$ for all $i \in N$.

Again consider the game of Figure 2.1. Charlie never wins by raising her hand, but sometimes wins by not doing so, such that c^1 dominates c^2 . Assuming that Charlie never raises her hand, b^2 dominates b^1 . Finally assuming that Bob will always raise his hand, a^1 dominates a^2 , and the only remaining action profile is the one where Bob alone raises his hand.

In general, the result of iterated weak dominance elimination depends on the order in which actions are removed, since the elimination of an action may render actions of another player undominated (e.g., Apt, 2004). This is in contrast to iterated *strict* dominance, which requires the inequality to be strict for every action profile of the other players. We consider two problems concerning iterated dominance in this thesis. Iterated dominance solvability (IDS) asks whether for a given game Γ there exists a sequence of eliminations of length $\sum_{i \in \mathbb{N}} (|A_i| - 1)$, i.e., one that leaves only one action for each player. Iterated dominance eliminability (IDE) is given an action $a_i \in A_i$ of some player $i \in \mathbb{N}$ and asks whether it is possible to eliminate a_i . Our results are often fairly robust as to the particular way these problems are defined. For example, results about IDE can easily be adapted to the problem of deciding whether *some* action of a particular player can be eliminated.

	b ¹	b ²	
a ¹	(1,0)	(0,1)	
\mathfrak{a}^2	(0,1)	(1,0)	

Figure 2.2: The Matching Pennies game, a constant-sum game without pure equilibria. Each of two players turns a penny to heads or tails. The first player wins if both coins show the same side, otherwise the second player wins.

A restricted variant of (iterated) dominance can be obtained by requiring that the dominating strategy s_i be pure. We will frequently exploit that the two variants are equivalent, which obviously holds for games with two actions, but also for games with only two different payoffs (Conitzer and Sandholm, 2005a).¹ Unless explicitly stated otherwise, results hold for dominance by pure strategies and for dominance by mixed strategies. An alternative definition of iterated dominance allows for the deletion of a set of dominated actions in each step (e.g., Apt, 2004). A different notion of solvability merely requires the remaining action profiles to yield a unique payoff to each of the players (e.g., Moulin, 1979). We note, but do not show explicitly, that all hardness and tractability results extend to these definitions as well.

One of the best-known solution concept for strategic games is Nash equilibrium (Nash, 1951). Nash equilibrium requires that the strategy of each player is a *best response* to the other players' strategies, such that no player could increase his payoff by *unilaterally* deviating and playing another strategy.

DEFINITION 2.4 (Nash equilibrium). A strategy profile $s_N^* \in S_N$ is called *Nash equilibrium* if for each player $i \in N$ and every strategy $s_i \in S_i$,

$$p_{i}(s_{N}^{*}) \geq p_{i}(s_{-i}^{*}, s_{i})$$

A Nash equilibrium is called *pure* if it is a pure strategy profile.

An equilibrium of the game in Figure 2.1 is for example attained when Alice and Charlie do not raise their hands, and Bob raises his hand with probability at least one half. The game thus possesses infinitely many equilibria. We leave it to the reader to verify that the only pure equilibrium is the action profile where Bob alone raises his hand.

Nash (1951) has shown that every normal-form game possesses at least one equilibrium. Since the proof is not constructive, it makes sense to consider the problem of *finding* an equilibrium of a given game. Pure Nash equilibria, on the other hand, are not guaranteed to exist, as is illustrated by the well-known Matching Pennies game depicted in Figure 2.2. If they do exist, however, they have two distinct advantages over mixed ones. For one, requiring randomization in order to reach a stable outcome has been criticized for

¹The game in Figure 2.1 clearly satisfies both of these properties.

various reasons. In multi-player games, where action probabilities in equilibrium can be irrational numbers, randomization is particularly questionable. Secondly, pure equilibria as computational objects are usually much smaller in size than mixed ones. We will thus also consider the problem of deciding, for a given game, whether this game possesses at least one pure Nash equilibrium.

Some additional solution concepts will be introduced later in the thesis. Correlated equilibrium generalizes Nash equilibrium by assuming the existence of a device or trusted third party that selects actions according to some joint probability distribution, and informs each player only about his own action (Aumann, 1974). In Chapter 4 we compare the quality of Nash and correlated equilibrium by requiring that every best response is played with positive probability (Harsanyi, 1973). In Chapter 7 we analyze the complexity of quasi-strict equilibrium in general strategic games and in some classes of games studied in earlier chapters, and also use it to shed some light on certain peculiarities of Nash equilibrium in the setting of Chapter 4. Finally, in Chapter 8, we consider a class of ordinal set-valued solution concepts due to Shapley (1964). These solution concepts, called saddles by Shapley, replace the notion of stability underlying Nash equilibrium by a more elementary one that is based on dominance, thereby eliminating the need for randomization as a prerequisite for the existence of a stable outcome.

2.3 Elements of Complexity Theory

The reasoning of real-world agents, both human and artificial, is often restricted by bounds on resources, like the time available for a thought process or the capacity of their memory. It is thus natural to study the resource requirements of game-theoretic solution concepts, for the obvious reason that solutions that cannot be found in practice are of very limited value. This holds irrespective of the fact whether a solution concept is to be used in an analytical or purely descriptive way. Computational complexity theory provides a rigorous mathematical framework to address this type question, and we introduce the necessary concepts in this section.

Complexity theory assigns problems to different *complexity classes*, each of which is characterized by several parameters: the underlying computational model, a computational paradigm, a resource, and an upper bound on this resource. The computational *model* describes the basic operations that can be used in a computation. A prominent example are Turing machines, which provide an abstract and idealized view of today's personal computers but are in fact able to compute any function that one would intuitively consider computable with only inconsequential loss of efficiency. Another example are Boolean circuits, which formalize the type of parallel computation performed by integrated circuits. The computational *paradigm* determines in which way the computation is performed. In a *deterministic* Turing machine, every intermediate state of a computation has exactly one followup state. A *nondeterministic* machine, on the other hand, may investigate several followup states of each state at once, and then use the result of one of the branches as the overall result. This latter mode can alternatively be interpreted as verifying a given solution to a problem. The assumption that a solution was already known is of course unrealistic, but nondeterminism has nevertheless proven very useful in analyzing computational problems. For a Boolean circuit, the paradigm is given by the type of the gates the circuit is composed of. Finally, a complexity class is characterized by a particular *resource*, like the *time* or *space* required for the computation, or the *depth* and overall *number of gates* of a Boolean circuit, as well as an *upper bound* on this resource as a function on the size of the problem instance.

Unfortunately, the current state of complexity theory in many cases does not allow for the *separation* of complexity classes, i.e., for a distinction of problems that can or cannot be solved under certain resource constraints. Quite often, however, like in the famous case of the classes P and NP, there is fairly strong evidence that two complexity classes are indeed distinct. What can be done using current techniques is to identify the hardest problems in each class, i.e., those not contained in a smaller class should the two be distinct. Hardness is established via *reductions* that transform instances of an arbitrary problem in a particular class into those of the problem in question. It is easy to see that the reductions we employ in this thesis compose, so a reduction from a problem that is itself hard for a particular class effectively shows that the problem in question cannot be easier to solve than any problem in that class.

When introducing the necessary concepts from complexity theory we restrict our attention to functions whose input and output are finite strings of bits. Definitions and results then carry over to more general functions by observing that their input and output can be encoded as bitstrings. Functions with several arguments, for example, can easily be obtained from the one-argument case by introducing a new symbol, say "o", using this symbol to separate the different arguments, and then encoding each of the three values 0, 1, and \circ by a pair of bits. Some issues related to encodings of games, and our interpretation of these issues, will briefly be discussed in Section 2.4. In most cases, however, it will be clear that an encoding with the desired properties exists, and we will avoid dealing with the details of any particular encoding in these cases.

Let us define the basic concepts more formally. In the context of this thesis, an *algorithm* for computing a function f will consist of a finite set of instructions describing how f(x) can be obtained for an arbitrary input $x \in \{0, 1\}^*$. The algorithm is allowed to use a scratchpad to write down intermediate results and, finally, the output. Each instruction starts by reading a bit of the input and a bit from the scratchpad. Based on the values that have been read, it then writes a bit to the scratchpad, and either halts or chooses the next instruction be to executed. Thus, while there is only a finite number of instructions, each of them may be executed an arbitrary number of times depending on the input. The Turing machine formalizes this idea.

DEFINITION 2.5 (Turing machine). A (k-tape) Turing machine is given by a tuple $M = (Q, \Sigma, \delta, b, q_s, q_h)$, where Q is a finite set of states, Σ is a finite alphabet, $\delta : Q \times \Sigma^k \to$

 $Q \times \Sigma^{k-1} \times \{L, R\}^k$ is a transition function, $b \in \Sigma$ is a specific blank symbol, and $q_s \in Q$ and $q_h \in Q$ are start and halting states, respectively.

The Turing machine's equivalent to the scratchpad are k tapes of infinitely many cells, each of which contains a symbol in Σ . By convention, the last tape is designated the *output tape* that will eventually bear the result of the computation. For each of the tapes, a *tape head* determines the current position. A *configuration* consists of a state, the content of each tape, and the position of the heads. The initial configuration is the one with state q_s , the input x in the first |x| cells of the first tape, and heads on the leftmost cell of each tape. The (infinitely many) remaining cells to the right of the input are filled with blanks.

If at some point the Turing machine is in state $q \in Q$, the symbol at the position of the ith head is σ_i , and $\delta(q, \sigma_1, \sigma_2, \ldots, \sigma_k) = (q', \sigma'_2, \ldots, \sigma'_k, z_1, \ldots, z_k)$, then at the next step the entry σ_i on the ith tape for $i \ge 2$ will have been replaced by σ'_i , the machine will be in state q', and the ith head will have moved one cell to the left if $z_i = L$ and if this cell exists, and one cell to the right if $z_i = R$. By convention the first tape is assumed to be read-only. The transition function δ is further assumed never to leave the halting state q_h once it has entered it, and not to modify the content of the tapes while in q_h . Entering state q_h can thus be interpreted as *halting*.

We are now ready to define what it means for a Turing machine to compute a function under resource constraints. The running time of a Turing machine will be the number of steps before it halts.

DEFINITION 2.6 (running time). Let $f : \{0, 1\}^* \to \{0, 1\}^*$ and $T : \mathbb{N} \to \mathbb{N}$ be two functions. Then, f can be computed in T-time if there exists a Turing machine M with the following property: for every $x \in \{0, 1\}^*$, if M is started in the initial configuration with input x, then after at most $\max(1, T(|x|))$ steps it halts with f(x) written on its output tape.

Bounds on the space used by a Turing machine can be defined in a similar way. Since we will specifically be interested in computations that require less space than is needed to store the input, we exclude the read-only input tape.

DEFINITION 2.7 (space bound). Let $f : \{0, 1\}^* \to \{0, 1\}^*$ and $S : \mathbb{N} \to \mathbb{N}$ be two functions. Then, f can be computed using S-space if there exists a Turing machine M with the following property: for every $x \in \{0, 1\}^*$, if M is started in the initial configuration with input x, then it halts with f(x) written on its output tape after a finite number of steps, and the number of cells of tapes 2 to k that differ from b at some intermediate step is at most $\max(1, S(|x|))$.

The exact details of Definition 2.5, like the number of tape symbols or the number of tapes, seem rather arbitrary. It is therefore worth noting that their effect on the time and space needed to compute a function will not be significant for the problems considered in this thesis. In particular, the complexity classes P and L defined below are very robust

to such modifications. Similar statements also apply to the much stronger properties of modern computers, like random access to the tape cells.

Let $R \subseteq \{0,1\}^* \times \{0,1\}^*$ be a relation. We say that R is *polynomial-time recognizable* if its characteristic function, i.e., the function $r : \{0,1\}^* \times \{0,1\}^* \to \{0,1\}$ such that r(x,y) = 1if and only if $(x,y) \in R$, can be computed in polynomial time. We call R *polynomially balanced* if there exists a polynomial function $p : \mathbb{N} \to \mathbb{N}$ such that $(x,y) \in R$ implies $|y| \leq p(|x|)$.

Associated with a relation R are three different problems. The decision problem asks, for a given instance $x \in \{0, 1\}^*$, whether there exists a solution $y \in \{0, 1\}^*$ such that $(x, y) \in R$. The search problem is to find a solution, i.e., an element $y \in \{0, 1\}^*$ satisfying $(x, y) \in R$. Finally, the counting problem asks for the number $|\{y \in \{0, 1\}^* : (x, y) \in R\}|$ of solutions for x. An example for a decision problem relevant in the context of this thesis is that of deciding, for a given normal-form game x, whether this game has a pure Nash equilibrium y, i.e., a vector of strategies that are mutual best responses to each other. To solve the corresponding search or counting problem one would have to do more, namely compute such a vector whenever one exists, or count the number of different vectors satisfying the property.

2.3.1 Decision Problems

Let us focus on decision problems for a moment. The decision problem associated with a polynomially balanced relation $R \subseteq \{0,1\}^* \times \{0,1\}^*$ can alternatively be looked at in terms of the *language* $L_R = \{x \in \{0,1\}^* : (x,y) \in R \text{ for some } y \in \{0,1\}^*\}$. In the following we say that a Turing machine *decides* a language L if it computes its characteristic function $f_L : \{0,1\}^* \to \{0,1\}$ such that $f_L(x) = 1$ if and only if $x \in L$. Let us define the class P of languages that can be decided in polynomial time, which is often used synonymously with efficient solvability.

DEFINITION 2.8 (the class P). For a function $T : \mathbb{N} \to \mathbb{N}$, let DTIME(T) be the set of all languages that can be decided in $c \cdot T$ -time for some constant c > 0. Then, $P = \bigcup_{k \ge 1} DTIME(n^k)$.

We proceed to define the class NP of decision problems that can be *verified* efficiently. While this is exactly the class of problems associated with polynomial-time recognizable and polynomially balanced relations, the name NP, short for *nondeterministic* polynomial time, derives from the way the class has traditionally been defined. A *nondeterministic Turing machine* differs from the Turing machine of Definition 2.5 in that δ is no longer a function mapping a configuration to a follow-up configuration, but a relation between successive configurations. A nondeterministic Turing machine is said to decide a language L if for each $x \in \{0, 1\}^*$, the following holds if and only if $x \in L$: there exists a sequence of configurations, connected by δ , that begins with the initial configuration for input x and ends in the halting state with 1 written on the output tape. Running time and space requirements are defined analogously to the deterministic case.

DEFINITION 2.9 (the class NP). For a function $T : \mathbb{N} \to \mathbb{N}$, let NTIME(T) be the set of all languages that can be decided by a nondeterministic Turing machine in $c \cdot T$ -time for some constant c > 0. Then, $NP = \bigcup_{c \ge 1} NTIME(n^c)$.

The relationship to polynomial-time recognizable relations yields an alternative characterization: a language L is in NP if there exists a polynomial function $p : \mathbb{N} \to \mathbb{N}$ and a (deterministic) Turing machine M such that for all $x \in \{0, 1\}^*$, $x \in L$ if and only if there exists a *certificate* $y \in \{0, 1\}^{p(|x|)}$ such that M accepts (x, y).

The relative complexity of different decision problems can be captured in terms of *reductions*. Intuitively, a reduction from one problem to another transforms every instance of the former into an equivalent instance of the latter, where equivalence means that both of them yield the same decision. For this transformation to preserve the complexity of the original problem, the reduction should of course have less power than is required to actually solve the original problem. For comparing problems in NP, the type of reduction most commonly used is the one that can itself be computed in (deterministic) polynomial time.

DEFINITION 2.10 (polynomial time reduction, NP-hardness). A language $P \subseteq \{0, 1\}^*$ is called *polynomial-time (many-one) reducible* to a language $Q \subseteq \{0, 1\}^*$, denoted $P \leq_p Q$, if there exists a function $f : \{0, 1\}^* \to \{0, 1\}^*$ computable in polynomial time such that for every $x \in \{0, 1\}^*$, $x \in P$ if and only if $f(x) \in Q$. A language Q is called *NP-hard* if for every language P in NP, $P \leq_p Q$.

A problem will be called *complete* for a particular class if it is both hard for and contained in this class. It is easy to see that the relation \leq_p and all other reducibility relations defined below are transitive, and that membership of a hard problem from one class in a smaller class implies that the two classes coincide. The existence of complete problems for particular classes is less obvious, but holds for all but one of the classes considered in this thesis.

Let us now turn to space-bounded computation, and in particular to the class L of decision problems that require only logarithmic space. This class is highly relevant for problems in large open systems like the Internet, where the input is often too large to be stored locally.

DEFINITION 2.11 (the class L). For a function $S : \mathbb{N} \to \mathbb{N}$, let SPACE(S) be the set of all languages that can be decided using $c \cdot S$ -space for some constant c > 0. Then, $L = SPACE(\log n)$.

The class NL of problems with solutions verifiable in logarithmic space is obtained by again considering nondeterministic Turing machines. An equivalent characterization of NL in terms of certificates exists, but requires that each bit of the certificate is read only once. DEFINITION 2.12 (the class NL). For a function $S : \mathbb{N} \to \mathbb{N}$, let NSPACE(S) be the set of all languages that can be decided by a nondeterministic Turing machine using $c \cdot S$ -space for some constant c > 0. Then, NL = NSPACE(log n).

An appropriate type of reduction for NL is one that itself requires only logarithmic space. Since the length of the output of such a function can be at most logarithmic in the length of its input, the following definition uses functions for which any single bit can be computed in logarithmic space.

DEFINITION 2.13 (log-space reduction). A function $f: \{0,1\}^* \to \{0,1\}^*$ such that $f(x) \leq |x|^c$ for some c > 0 and every $x \in \{0,1\}^*$ is called *implicitly log-space computable* if the languages $L_f = \{(x,i) : f(x)_i = 1\}$ and $L'_f = \{(x,i) : i \leq |f(x)|\}$ are in L. A language $P \subseteq \{0,1\}^*$ is called *log-space reducible* to a language $Q \subseteq \{0,1\}^*$, denoted $P \leq_{\ell} Q$, if there exists an implicitly log-space computable function $f: \{0,1\}^* \to \{0,1\}^*$ such that for all $x \in \{0,1\}^*$, $x \in P$ if and only if $f(x) \in Q$.

In some cases it is interesting to study decision problems that are the *complements* of languages in a specific complexity class, and the prefix "co" is commonly used to denote the resulting class. The class coNP, for example, can informally be described as the class of problems for which non-existence of a solution can be verified efficiently. It is obvious that deterministic complexity classes are closed under complementation. Immerman (1988) and Szelepcsenyi (1988) show that this also holds for nondeterministic space complexity classes, and for NL and coNL in particular.

2.3.2 Search Problems

We might ask what we have lost by looking only at decisions problems. Define FNP as the class of search problems associated with polynomial-time recognizable and polynomially balanced relations, and FP as the subclass of these problems that are solvable in polynomial time. Call the search problem of a relation R *self-reducible* if it can be reduced, using an appropriate type of reduction, to the corresponding decision problem of R. It turns out that self-reducibility holds for a large class of natural problems, and in particular for any search problem such that the corresponding decision problem is NP-complete. This directly implies that FP=FNP if and only if P=NP, and means that in many cases it is indeed enough to consider only the decision version of a problem.

An interesting subclass of FNP for which this strong relationship seems to break down is obtained by considering search problems in which every instance is guaranteed to have a solution. An example relevant in the context of this thesis is the problem of finding a Nash equilibrium of a given normal-form game, the existence of which has been shown by Nash (1951). Call TFNP, for *total* functions in NP, the class of search problems associated with polynomial-time recognizable and polynomially balanced relations $R \subseteq \{0, 1\}^* \times \{0, 1\}^*$ such that for every $x \in \{0, 1\}^*$, there exists $y \in \{0, 1\}^*$ with $(x, y) \in R$ (Megiddo and Papadimitriou, 1991). We henceforth write R to denote both the relation and the corresponding search problem. Unfortunately the mathematical lemmas that ensure existence of solutions to problems in TFNP are very diverse, and TFNP is likely not to possess any complete problems. It therefore makes sense to study subclasses of TFNP corresponding to the different lemmas. These classes are most conveniently defined via complete problems, so we begin by introducing a notion of reducibility among search problems. The appropriate type of reduction basically consists of a homomorphism between two relations, together with a second function witnessing that this homomorphism indeed preserves the structure of the solutions. As before, reductions will also be used to define hardness for a class of search problems: a search problem R will be called hard for a particular class if every problem in that class reduces to R.

DEFINITION 2.14 (reducibility between search problems). A search problem $P \subseteq \{0, 1\}^* \times \{0, 1\}^*$ is called *polynomial-time (many-one) reducible* to a search problem $Q \subseteq \{0, 1\}^* \times \{0, 1\}^*$, denoted $P \leq_p Q$, if there exist two functions $f : \{0, 1\}^* \rightarrow \{0, 1\}^*$ and $g : \{0, 1\}^* \times \{0, 1\}^* \rightarrow \{0, 1\}^*$ computable in polynomial time, such that for every $x \in \{0, 1\}^*$ and for every $y \in \{0, 1\}^*$ such that $(f(x), y) \in B$, it also holds that $(x, g(x, y)) \in R$.

We consider two particular subclasses of TFNP in this thesis, the class PPAD of polynomial search problems guaranteed to have a solution by a directed parity argument, and the class PLS of polynomial local search problems. Underlying both of these classes is a large graph, which can in turn be seen as describing a generic exhaustive search algorithm for solving problems in the respective class. The graphs underlying the two complete problems will be graphs of functions from the set of all bitstrings of a certain length to itself. For this, let $F \subseteq \{f : \{0,1\}^n \rightarrow \{0,1\}^n : n \in \mathbb{N}\}$ be a set of functions, and consider some encoding of the members of F by elements of $\{0,1\}^*$. Denoting by $f_x : \{0,1\}^n \rightarrow \{0,1\}^n$ the function with encoding $x \in \{0,1\}^*$, we require that |x| is polynomial in n. We further assume that for any n > 0, the set F contains all functions $f : \{0,1\}^n \rightarrow \{0,1\}^n$ computable by a polynomial-size Boolean circuit as introduced in Definition 2.19 below, and that the argument length of a function can be determined from its encoding in polynomial time.

In the case of PPAD (Papadimitriou, 1994b), the underlying graph is the graph of a partial injective function, computable in polynomial time, whose range is strictly contained in its codomain, i.e., a graph where the in- and outdegree of every vertex is bounded by one and some vertex has indegree zero. Such a function can for example be defined via a pair of functions s and p such that x is mapped to y if and only if p(s(x)) = y. A distinguished vertex with indegree zero is provided explicitly, and the set of solutions of the search problem is defined as the set of all vertices, apart from the distinguished vertex, whose in- or outdegree is zero.

DEFINITION 2.15 (the class PPAD). Let $X \subseteq \{0,1\}^*$ be the set of encodings of functions $f_x : \{0,1\}^{2n} \to \{0,1\}^{2n}$ satisfying the following condition: there exist two functions $p_x : \{0,1\}^n \to \{0,1\}^n$ and $s_x : \{0,1\}^n \to \{0,1\}^n$ such that for all $y, z \in \{0,1\}^n$, $f_x(y \circ z) = p_x(y) \circ s_x(z)$, where \circ denotes concatenation of bitstrings. Let $R \subseteq \{0,1\}^* \times \{0,1\}^*$ be

the relation such that $(x, y) \in R$ if and only if $x \in X$, $s_x(p_x(0^{|y|})) \neq 0^{|y|}$, and either $p(s(y)) \neq y$, or $s(p(y)) \neq y$ and $y \neq 0^{|y|}$. Then, PPAD is the class of all search problems P such that $P \leq_p R$.

A local search problem is given by a set $I \subseteq \{0,1\}^*$ of instances and three functions F, N, and c. Function F assigns to each instance $x \in I$ a set $F(x) \subseteq \{0,1\}^*$ of feasible solutions. Function N defines, for each feasible solution $y \in F(x)$, a neighborhood $N(y,x) \subseteq F(x)$. Finally, c assigns an integer cost c(y,x) to each pair of an instance $x \in I$ and a solution $y \in F(x)$. The actual solutions for each instance are those feasible solutions that have optimal cost within their neighborhood, i.e., either minimal or maximal cost depending on the exact definition of the problem. Membership in PLS (Johnson et al., 1988, Schäffer and Yannakakis, 1991) then essentially requires that I is polynomialtime decidable, and that for each instance, all of the following can be done in polynomial time: finding an initial feasible solution, deciding optimality of a feasible solution, and finding a better neighbor of one that is not optimal.

In other words, a local search problem is given by a partial order on the vertex set of an undirected graph, with solutions corresponding to either the minimal or the maximal elements within neighborhoods. The following definition uses a graph on the set of all bitstrings of a certain length. The cost associated with a vertex is given by the interpretation of the corresponding bitstring as a binary number, and two vertices are adjacent if the corresponding bitstrings have Hamming distance one, i.e., differ in exactly one bit.

DEFINITION 2.16 (the class PLS). Let $c : \{0,1\}^* \to \mathbb{N}$ be the function such that for every $y \in \{0,1\}^*$, $c(y) = \sum_{i=1}^{|y|} 2^i y_i$, where y_i denotes the ith bit of y. Further let $R \subseteq \{0,1\}^* \times \{0,1\}^*$ be the relation such that $(x,y) \in R$ if and only if |x| = |y| and $c(f_x(y)) \leq \min\{c(f_x(z)) : z \in \{0,1\}^*, |z| = |y|, \text{ and } |\{i : y_i \neq z_i\}| = 1\}$. Then, PLS is the class of all search problems P such that $P \leq_p R$.

Implicit in the definition of PLS is a standard algorithm that is guaranteed to find a locally optimal solution for a given instance: start with an initial feasible solution, and repeatedly find a neighbor with strictly better cost, breaking ties in some convenient manner. The standard algorithm problem can be phrased as follows: given $x \in I$, find the locally optimal solution output by the standard algorithm on input x. Schäffer and Yannakakis (1991) introduce the notion of a *tight* reduction and show that tight reductions compose and preserve both hardness of the standard algorithm problem and exponential worst-case running time of the standard algorithm.

DEFINITION 2.17 (tight PLS reduction). Let $P, Q \subseteq \{0,1\}^* \times \{0,1\}^*$ be in PLS. Then a reduction (f,g) from P to Q is called *tight* if for any instance x of P there exists a set Y of feasible solutions of f(x) with the following properties:

- (i) Y contains all local optima of f(x), i.e., $Y \supseteq \{y \in \{0,1\}^* : (f(x),y) \in P\}$.
- (ii) For every feasible solution z of x, a solution $y \in Y$ satisfying g(x,y) = z can be found in polynomial time.

(*iii*) Consider a set $Y' = \{y_0, y_1, \dots, y_\ell\}$ of feasible solutions of f(x) such that $Y' \cap Y = \{y_0, y_\ell\}$ and for all $i < \ell$, y_{i+1} is a strictly better neighbor of y_i . Then either $g(x, y_\ell) = g(x, y_0)$, or $g(x, y_\ell)$ is a strictly better neighbor of $g(x, y_0)$.

An interesting feature of all problems in PLS is that they have a fully polynomial-time approximation scheme, i.e., they can be approximated to any factor in time polynomial in the size of the input and in the desired approximation factor (Orlin et al., 2004).

2.3.3 Counting Problems

As for the third type of problem, define #P as the class of counting problems associated with polynomial-time recognizable and polynomially balanced relations. In the context of this thesis, a problem will be called #P-hard if all problems in #P reduce to it via a type of reduction that allows us to efficiently compute the number of solutions of one problem from that of the other. Other notions of #P-hardness that can be found in the literature are those based on polynomial-time many-one and polynomial-time Turing reductions.

DEFINITION 2.18 (reducibility between counting problems, #P-hardness). A counting problem $P \subseteq \{0,1\}^* \times \{0,1\}^*$ is called *polynomial-time reducible* to a counting problem $Q \subseteq \{0,1\}^* \times \{0,1\}^*$, denoted $P \leq_p Q$, if there exist two functions $f : \{0,1\}^* \to \{0,1\}^*$ and $g : \{0,1\}^* \times \mathbb{N} \to \mathbb{N}$ computable in polynomial time, such that for every $x \in \{0,1\}^*$, $|\{y \in \{0,1\}^* : (x,y) \in P\}| = g(x, |\{y \in \{0,1\}^* : (f(x),y) \in Q\}|)$. A counting problem Q is called #P-hard if for every problem P in #P, $P \leq_p Q$.

Hardness for the class #L of counting problems associated with polynomially balanced relations that can be recognized in logarithmic space is defined analogously.

2.3.4 Circuit Complexity

Let us now consider a different computational model, which captures the type of parallel computation characteristic for the lower levels of modern computers, but also for decentralized systems involving many agents: the Boolean circuit. For a directed graph G = (V, E) and a particular vertex $v \in V$, let indeg(v) and outdeg(v) denote the in- and outdegree of v in G, i.e., $indeg(v) = |\{u \in V : (u, v) \in E\}|$ and $outdeg(v) = |\{u \in V : (v, u) \in E\}|$.

DEFINITION 2.19 (Boolean circuit). Let B be a set of Boolean functions. Then, a *Boolean* circuit over B with n inputs and m outputs is a tuple $C = (V, E, \alpha, \beta, \omega)$, where (V, E) is a directed acyclic graph, $\alpha : V \to \mathbb{N}$ is an injective function, $\omega : \{1, \ldots, m\} \to V$, and $\beta : V \to B \cup \{1, \ldots, n\}$ is a function such that the following holds for all $\nu \in V$: if $indeg(\nu) = 0$, then either $\beta(\nu) \in \{1, \ldots, n\}$, or $\beta(\nu) \in B$ is a 0-ary Boolean function; if $indeg(\nu) = k > 0$, then $\beta(\nu) \in B$ is a k-ary Boolean function.

A vertex $v \in V$ is also called a *gate* of C, and indeg(v) and outdeg(v) are respectively referred to as its *fan-in* and *fan-out*. The function α induces an ordering on any subset

18

of V, and β assigns a *type* to each gate, such that each of them either corresponds to one of the n inputs, or to a Boolean function in B with inputs given by the gate's predecessors in the graph. The function ω finally identifies certain vertices that yield the output of the circuit. For a given input, we can inductively assign a unique truth value to each gate of a circuit, and those corresponding to the m outputs in particular, such that the value of the gate coincides with the value of the associated Boolean function, given the truth values of the predecessors of the gate as inputs. This is true because every vertex is reachable from one with fan-in zero, which must either correspond to an input, or to a 0-ary function, i.e., a constant. A circuit with n inputs and m outputs can thus be seen as computing a function $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$. The following definition makes this relationship explicit.

DEFINITION 2.20 (function computed by a circuit). Let $C = (V, E, \alpha, \beta, \omega)$ be a Boolean circuit over B with n inputs and m outputs. For input $x \in \{0, 1\}^n$, let $\varphi_x : V \to \{0, 1\}$ be the unique function such that for all $v \in V$, $\varphi_x(v) = x_i$ if $\beta(v) = i$ for some $i \in \mathbb{N}$, and $\varphi_x(v) = g(\varphi(z_1), \varphi(z_2), \ldots, \varphi(z_k))$ if $\beta(v) = g$ for some $g \in B$, such that $(z_i, v) \in E$ for $1 \leq i \leq k$ and $\alpha(z_i) < \alpha(z_j)$ for $1 \leq i < j \leq k$. Then, C is said to compute a function $f : \{0, 1\}^n \to \{0, 1\}^m$ if for every $x \in \{0, 1\}^n$, $(\varphi_x(\omega(1)), \varphi_x(\omega(2)), \ldots, \varphi_x(\omega(m))) = f(x)$.

Apart from the allowed types of gates and their fan-in, the functions computable by a class of circuits will depend on the size and depth of circuits in the class. The *size* and *depth* of a Boolean circuit $C = (V, E, \alpha, \beta, \omega)$ are respectively defined as size(C) = |E| and $depth(C) = \max\{d \in \mathbb{N} : \text{there exists a path of length } d \text{ in } (V, E)\}.$

By now, a subtle difference between Turing machines and Boolean circuits may have become apparent. Turing machines provide a *uniform* computational model: an algorithm is described by a single machine that works for every input length. To describe algorithms by Boolean circuits, on the other hand, a different circuit has to be given for each input length, and in particular the size of these circuits may grow with the input length. One way to relate the two computational models to each other is to define a family of circuits, one for every input length, and require that this infinite family has a finite description. Such a description can for example be given in terms of a Turing machine that answers queries about the structure of the circuit for a given input length (Ruzzo, 1981). In the context of this thesis, we will consider *log-space uniform* circuit families, i.e., circuit families that can be described by a Turing machine with logarithmic space. Henceforth, when we talk about a complexity class defined by Boolean circuits, we mean the log-space uniform version of this class. We further say that a family of circuits computes a function f if for every input length, the function computed by the respective member of the family coincides with f.

We consider two different circuit complexity classes for Boolean functions. The first one is given by circuits of polynomial size and constant depth with three types of gates corresponding to the logical connectives NOT, AND, and OR, where the latter two types are allowed to have unbounded fan-in. In the following, let B₀ be the set containing functions for the logical connectives NOT, AND, and OR, i.e., B₀ = {¬} \cup { \wedge^k , \vee^k : $k \ge 0$ }, where $\neg : \{0, 1\} \rightarrow \{0, 1\}$ with $\neg(0) = 1$ and $\neg(1) = 0$, and for all $k \ge 0$, $\wedge^k : \{0, 1\}^k \rightarrow \{0, 1\}$ with $\wedge^k(x) = 1$ if and only if $x = 1^k$, and $\vee^k : \{0, 1\}^k \rightarrow \{0, 1\}$ with $\vee^k(x) = 0$ if and only if $x = 0^k$. Observe that in particular, $\wedge^0 = 1$ and $\vee^0 = 0$.

DEFINITION 2.21 (the class AC^0). A Boolean function $f : \{0, 1\}^* \to \{0, 1\}$ is in the class AC^0 if there exists a log-space uniform family of circuits over B_0 with polynomial size and constant depth that computes f.

It is worth noting that inputs can be copied and used multiple times without significantly increasing the size of a circuit, e.g., by using an AND gate and a constant. Thus, the fact that there can only be a single edge between any pair of gates is not a restriction.

The second class of interest is obtained by adding an additional type of gate, which outputs 1 if a majority of its inputs is 1.

DEFINITION 2.22 (the class TC^0). For $k \ge 0$, let g^k be the majority function for input length k, i.e., $g^k : \{0,1\}^k \to \{0,1\}$ with $g^k(x) = 1$ if $|\{i : x_i = 1\}| \ge |\{i : x_i = 0\}|$. A function $f : \{0,1\}^* \to \{0,1\}$ is in the class TC^0 if there exists a log-space uniform family of circuits over $B_0 \cup \{g^k : k \ge 0\}$ with polynomial size and constant depth that computes f.

It is interesting to note that among all the complexity classes defined above, the only known separation is between AC^0 and TC^0 : whether the majority of bits of a bitstring is 1 cannot be decided by a circuit with polynomial size, constant depth, and unbounded fanin, when using only AND, OR, and NOT gates. The same can be shown for other functions, by reducing the majority function to them. The appropriate type of reduction itself must of course not use the majority function. In the following, for a function $f^k : \{0,1\}^k \to \{0,1\}^n$ and for $1 \leq i \leq n$, let $f_i^k : \{0,1\}^k \to \{0,1\}$ be the function computing the ith bit of f^k , i.e., $f_i(x) = y_i$ whenever $f(x) = y_1 y_2 \dots y_n$. For a family $f = (f^k)_{k \geq 1}$ of functions, let $B(f) = \{f_i^k : k \geq 1, i \geq 1\}$. Intuitively, a family f of functions will be considered reducible to another family g if it is possible to construct circuits for members of f if one is allowed to use "oracle" gates that compute members of g at unit cost.

DEFINITION 2.23 (constant-depth reducibility). Let $f = (f^k)_{k \ge 1}$ and $g = (g^k)_{k \ge 1}$ with $f^k : \{0,1\}^k \to \{0,1\}$ and $g^k : \{0,1\}^k \to \{0,1\}$. Then, f is called *constant-depth reducible* to g, denoted $f \leq_{cd} g$, if for every $k \ge 1$, there exists a Boolean circuit $C = (V, E, \alpha, \beta, \omega)$ over $B_0 \cup B(g)$ computing f^k , such that size(C) is polynomial in k and there is no directed path between vertices $u, v \in V$ in (V, E) with $\beta(u) \in B(g)$ and $\beta(v) \in B(g)$.

It is worth noting that the size restriction effectively restricts the use of members of g to those that have input length polynomial in k, when computing f^k .

2.4 A Few Words on Encodings

As we have seen in the previous section, complexity theory measures the complexity of a problem relative to the size of the input defining a particular problem instance. While
nobody would want to argue whether this is a meaningful definition, it may sometimes allow for results that are unsatisfactory from a practical point of view. For example, the existence of a pure Nash equilibrium can trivially be decided in polynomial time for any game that is given explicitly, i.e., lists the payoff for every action profile: simply check for each of them whether it satisfies the equilibrium condition. This does not take into account, however, that the number of different action profiles of a game, and thus the size required for the explicit representation of its payoffs, grows exponentially in a natural parameter, the number of players. More precisely, a general game in normal form with n players and k actions per player comprises $n \cdot k^n$ numbers, which means exponential growth already in the case k = 2. Computational statements over such large objects, and efficient algorithms for problems whose input size is already exponential in a natural parameter are of course somewhat dubious (cf. Papadimitriou and Roughgarden, 2005).

Games emanating from real-world situations, on the other hand, would certainly be expected to possess additional structure that allows them to be played rationally and efficiently at the same time. In fact, it is hard to imagine how the physical world could give rise to a game the payoffs of which cannot be represented in this world. Since the real world is what we are ultimately interested in, we will henceforth restrict ourselves to games that can be represented in space polynomial in their natural parameters, like the number of players or actions. We will try to characterize the complexity of solving these games in terms of their natural parameters, while making as few assumptions as possible about any particular encoding.

Unless explicitly stated otherwise, we assume that the number of players of a game is polynomial in the size of its representation. We further assume that each player can determine efficiently whether a particular action is a best response for a given action profile of the other players, which obviously is both necessary and sufficient for *playing* a game rationally and efficiently at the same time. Tractability results then hold for *any* encoding satisfying these properties. Hardness, on the other hand, is established via *some* encoding which allows efficient and rational play.

Chapter 3

State of the Art and Our Contribution

It is known since the early days of game theory that the security level of a player is the solution to a linear program, and that the two-player case is in fact equivalent to linear programming (Dantzig, 1951). The latter problem in turn is computationally tractable (Khachiyan, 1979). This result also extends to Nash equilibria of two-player *zero-sum* games, because Nash equilibrium strategies and maximin strategies coincide in these games (e.g., Myerson, 1991).

Among the first to study game-theoretic solution concepts in the framework of computational complexity theory, Gilboa and Zemel (1989) show that deciding the existence of a Nash equilibrium possessing one of several natural properties is NP-hard already in twoplayer games. Examples for such properties are that a given action is played with positive probability, or that the equilibrium is not unique. A uniform reduction given by Conitzer and Sandholm (2008) subsumes these results, and also shows that it is already NP-hard to find an equilibrium that is reasonably close to maximizing certain properties, like social welfare or support size. Thus, while deciding the existence of a Nash equilibrium is trivial due to Nash's (1951) existence result, virtually every additional property one might require makes the problem hard. In the case of non-uniqueness of an equilibrium, NP-hardness already holds for binary games (Codenotti and Stefankovic, 2005). Similarly, Gilboa et al. (1993) show NP-completeness for several decision problems concerning the iterated removal of weakly dominated actions, like eliminability of a given action, or solvability of a game in the sense that only one action remains for each player. These results were again strengthened by Conitzer and Sandholm (2005a). Iterated strict dominance, on the other hand, had earlier been shown tractable, and in fact P-complete, by Knuth et al. (1988).

As we have observed earlier, the problem of *finding* a Nash equilibrium of a two-player game is in the class TFNP of search problems a solution of which is guaranteed to exist and can be verified in polynomial time. More precisely, Megiddo and Papadimitriou (1991) show membership in the class PPAD via a reduction to the computation of Brouwer fixed points. Finding an exact lower bound for this problem was for a long time considered a "most important concrete open question on the boundary of P" (Papadimitriou, 2001,

p. 749), and became one of the defining problems of algorithmic game theory. A series of papers recently established that the problem is in fact PPAD-complete in the two-player case (Goldberg and Papadimitriou, 2006, Daskalakis et al., 2006, Chen and Deng, 2006), i.e., as hard as finding a fixed point of any continuous function from the n-dimensional unit ball to itself. This can be seen as fairly strong evidence that the problem cannot be solved efficiently. By a result of Abbott et al. (2005), hardness already holds for *binary* two-player games. In games with more than two players, equilibria might require randomization with irrational probabilities, in which case they cannot be computed exactly. Daskalakis et al. (2009a) show that the problem of finding an ϵ -equilibrium in this case, i.e., a strategy profile where for each player the potential gain from unilateral deviation is at most ϵ , is in PPAD. Etessami and Yannakakis (2007) show that finding an ϵ -approximation of an equilibrium in the three-player case, i.e., a strategy profile that has distance at most ϵ from an equilibrium in each component, is at least as hard as the square-root sum problem, a long standing problem about arithmetic circuits that is not known to be even in NP.

These rather negative results have not settled the question, however, but immediately raise several new ones: First, can an ϵ -equilibrium be computed efficiently? Second, are there *interesting classes* of games that do allow for an exact solution to be computed efficiently? Third, are there *alternative solution concepts* that are computationally tractable, and how does the quality of solutions selected by these concepts compare to those selected by established solution concepts? Regarding the first question, results by Chen et al. (2006, 2007) indicate that there might *not* exist a fully polynomial-time approximation scheme for Nash equilibria, i.e., an algorithm that allows them to be approximated to any factor in time polynomial in the size of the game and in the desired approximation factor. Limited progress has however been made regarding constant factor approximations (see Spirakis, 2008). The work reported in this thesis is part of the effort to answer the latter two questions. In Chapters 4 through 6, we study the complexity of various solution concepts in four natural classes of normal-form games. In Chapters 7 and 8, we then consider two solution concepts that are less well-known, and analyze their computational complexity in both general normal-form games and in restricted classes.

Ranking Games In two-player zero-sum games, any profile of maximin strategies forms a Nash equilibrium, and an equilibrium can therefore be found in polynomial time. It is natural to ask whether there is some property that captures the same degree of competitiveness in the multi-player case and maintains tractability of Nash equilibria. In Chapter 4, we consider *ranking games*, where each outcome is a ranking of the players and each player entertains preferences over ranks, preferring higher ranks over lower ones. Indeed, the outcomes of many strategic situations such as parlor games or competitive economic scenarios are rankings of the participants.¹ We investigate the computational complexity of a variety of common game-theoretic solution concepts in ranking games, and

 $^{^{1}}$ The game in Figure 2.1 on Page 7 describes a special case of such a scenario, where each participant is only interested in being ranked first.

give hardness results for iterated weak dominance and mixed Nash equilibrium when there are more than two players, and for pure Nash equilibrium when the number of players is unbounded but the game is described succinctly. This dashes hope that multi-player ranking games can be solved efficiently, despite their structural restrictions. Based on these findings, we study two alternative solution concepts, maximin strategies and correlated equilibrium, that are known to be efficiently computable even in general games. In particular, we provide matching upper and lower bounds for comparative ratios that measure the quality of outcomes selected by these solution concepts relative to that of Nash equilibrium outcomes: the *price of cautiousness*, the *mediation value*, and the *enforcement value*.

Chapter 4 is based on joint work with Felix Brandt, Paul Harrenstein, and Yoav Shoham (Brandt et al., 2006, 2007a, 2009c).

Anonymous Games Pure Nash equilibria, if they exist, can be found easily by checking the equilibrium condition for each action profile. A problem associated with general games, however, is their massive input complexity in the multi-player case, making it highly questionable that such games could even be *played* efficiently. By contrast, one would certainly expect games in the real world to exhibit additional structure, and to be given in some implicit way that allows efficient play. An interesting question thus concerns the properties of realistic games that allow for such a compact representation. Symmetries are an example for such a property.

Strategic games may exhibit symmetries in a variety of ways. A characteristic feature, enabling the compact representation of games even when the number of players is unbounded, is that players cannot, or need not, distinguish between the other players. In Chapter 5, we investigate the computational complexity of pure Nash equilibrium and iterated weak dominance in four classes of anonymous games obtained by considering two additional properties: *identical payoff functions* for all players and the ability to distinguish oneself from the other players. In contrast to other types of compactly representable multi-player games, the pure equilibrium problem turns out to be tractable in all four classes when only a constant number of actions is available to each player. Identical payoff functions make the difference between TC^{0} -completeness and membership in AC^{0} , while a growing number of actions renders the equilibrium problem NP-hard for three of the classes and PLS-hard for the most restricted class for which the existence of a pure equilibrium is guaranteed. Our results also extend to larger classes of threshold anonymous games where players are unable to determine the exact number of players playing a certain action. On the other hand, we show that deciding whether a game can be solved by means of iterated weak dominance is NP-complete for anonymous games with three actions. For the case of two actions, this problem can be reformulated quite naturally as an elimination problem on a matrix. While enigmatic by itself, the latter turns out to be a special case of matching along paths in a directed graph, which we show to be computationally hard in general, but also use to identify tractable cases of matrix elimination.

We further identify different classes of anonymous games where iterated dominance is in P and NP-complete, respectively.

Chapter 5 is based on joint work with Felix Brandt and Markus Holzer (Brandt et al., 2008a, 2009d).

Graphical Games A different structural element that facilitates efficient play in games with many players is locality. While realistic situations may involve many agents, the weal and woe of any particular agent often depends only on the decisions made by a small part of the overall population: neighbors, business partners, or friends. Graphical games formalize this notion by assigning to each player a subset of the players, his neighborhood, and defining his payoff as a function of the actions of these players. More formally, a graphical game is given by a (directed or undirected) graph on the set of players of a normal-form game, such that the payoff of each player depends only on the actions of his neighbors in this graph. Any graphical game with neighborhood sizes bounded by a constant can be represented using space polynomial in the number of players.

In Chapter 6, we first strengthen a result of Gottlob et al. (2005) concerning hardness of the pure equilibrium problem in graphical games. To be precise, we show that two actions per player, two-bounded neighborhood, and two-valued payoff functions suffice for NP-completeness. This is the best possible result, because deciding the existence of a pure Nash equilibrium becomes trivial in the case of a single action for each player and tractable for one-bounded neighborhood. In fact, we show the latter problem to be NL-complete in general, and thus solvable in deterministic polynomial time. Interestingly, it turns out that the number of actions in a game with one-bounded neighborhood is a sensitive parameter: restricting the number of actions for each player to a constant makes the problem even easier than NL, unless L=NL. In this way, we obtain a nice alternative characterization of the determinism-nondeterminism problem for Turing machines with logarithmic space in terms of the number of actions for games with one-bounded neighborhood.

We then turn to graphical games that additionally satisfy one of four types of anonymity within neighborhoods. We establish that deciding the existence of a pure Nash equilibrium is NP-hard in general for all four types. Using a characterization of games with pure equilibria in terms of even cycles in the neighborhood graph, as well as a connection to a generalized satisfiability problem, we identify tractable subclasses of the games satisfying the most restrictive type of symmetry. Hardness for a different subclass is obtained via a satisfiability problem that remains NP-hard in the presence of a matching, a result that may be of independent interest. Finally, games with symmetries of two of the four types are shown to possess a symmetric *mixed* equilibrium which can be computed in polynomial time. We thus obtain a class of games where the pure equilibrium problem is computationally harder than the mixed equilibrium problem, unless P=NP.

Chapter 6 is based on joint work with Felix Brandt, Markus Holzer, and Stefan Katzenbeisser (Fischer et al., 2006, Brandt et al., 2008b).

Quasi-Strict Equilibria Despite its ubiquity, Nash equilibrium has been criticized for various reasons. Its existence relies on the possibility of randomizing over actions, which in many cases is deemed unsuitable, impractical, or even infeasible. Furthermore, players may be indifferent between actions they do and do not play in equilibrium, thus calling into question the underlying notion of stability. In Chapters 7 and 8, we consider two solution concepts that try to address these two shortcomings. Chapter 7 is devoted to an investigation of the computational properties of quasi-strict equilibrium, an attractive equilibrium refinement proposed by Harsanyi and recently shown to always exist in twoplayer games. Quasi-strict equilibrium strengthens the equilibrium condition by requiring that every best response be played with positive probability. We prove that deciding the existence of a quasi-strict equilibrium in games with more than two players is NP-hard. We further show that unlike Nash equilibria, quasi-strict equilibria in zero-sum games have a unique support, and propose a linear program to compute a quasi-strict equilibrium in these games. Finally, we prove that every symmetric multi-player game where each player has two actions at his disposal possesses an efficiently computable quasi-strict equilibrium which may itself be asymmetric.

Chapter 7 is based on joint work with Felix Brandt (Brandt and Fischer, 2008a).

Shapley's Saddles In work dating back to the early 1950s, Shapley proposed ordinal setvalued solution concepts for zero-sum games that he refers to as *strict and weak saddles*. These concepts are intuitively appealing, they always exist, and are unique in several important classes of games. In Chapter 8, we study of computational aspects of Shapley's saddles and provide polynomial-time algorithms for computing strict saddles in general normal-form games, and weak saddles in a subclass of symmetric zero-sum games. On the other hand, we show that several problems associated with weak saddles are NP-hard already in two-player games, which provides rather strong evidence that they cannot be computed efficiently.

Chapter 8 is based on joint work with Felix Brandt, Markus Brill, and Paul Harrenstein (Brandt and Fischer, 2008b, Brandt et al., 2009a).

Chapter 4 Ranking Games

The situations studied by the theory of games may involve different levels of antagonism. On the one end of the spectrum are games of pure coordination, on the other those in which the players' interests are diametrically opposed. In this chapter, we introduce and study a new class of competitive multi-player games whose outcomes are *rankings* of the players, i.e., orderings representing how well they have done in the game relative to one another. We assume players to weakly prefer a higher rank over a lower one and to be indifferent as to the other players' ranks. Indeed, this type of situation is very common in the real world, with examples such as parlor games, sports competitions, patent races, competitive resource allocation, social choice settings, and other strategic situations where players are merely interested in performing optimal *relative* to their opponents rather than in absolute measures. Formally, ranking games can be defined as normal-form games in which the payoff functions represent the players' von Neumann-Morgenstern preferences over lotteries over rankings.

Apart from their practical relevance, ranking games promise to be interesting also from a computational perspective. Two-player ranking games form a subclass of constant-sum games, such that the computationally tractable maximin solution the solution concept of choice. With more than two players, there no longer is any inclusion relationship between ranking games and constant-sum games. The notion of a ranking, however, is most natural and relevant in multi-player settings, a fact that is much less true for the requirement that the sum of payoffs in all outcomes be zero or constant. Indeed, *any* game can be transformed into a zero-sum game by introducing an additional player, with only one action at his disposal, who absorbs the payoffs of the other players (von Neumann and Morgenstern, 1947).

The maximin solution does not unequivocally extend to general n-player games, or to n-player ranking games. Numerous alternative solution concepts have been proposed to cope with this type of situation, some of which we have introduced in Chapter 2. None of them, however, seems to be as compelling as maximin is for two-player zero-sum games, and computational intractability serves as an additional threat to many of them. One could nevertheless hope that the notion of competitiveness captured by ranking games might cause tractability of solution concepts to carry over to the multi-player case.

We will see in Section 4.5 that unfortunately this does not seem to be the case, and solving ranking games becomes considerably more complicated as soon as more than two players are involved, just as games in which both contrary and common interests prevail. In particular, we give NP-hardness and PPAD-hardness results, respectively, for iterated weak dominance and (mixed) Nash equilibria when there are more than two players, and an NP-hardness result for pure Nash equilibria in games with an unbounded number of players. This dashes hope that multi-player ranking games can be solved efficiently thanks to their structural restrictions. Remarkably, all hardness results hold for arbitrary preferences over ranks, provided they meet the minimal requirements given above. Accordingly, even very restricted subclasses of ranking games such as *singlewinner games*—in which players only care about winning—or *single-loser games*—in which players merely wish not to be ranked last—are computationally hard to solve.

By contrast, maximin strategies as well as correlated equilibria (Aumann, 1974) are known to be computationally easy via linear programming even in general games. Against the potency of these concepts, however, other objections can be brought in. Playing a maximin strategy is extremely defensive and a player may have to forfeit a considerable amount of payoff in order to guarantee his security level. Correlation, on the other hand, may not be feasible in all practical applications, and may fail to provide an improvement of social welfare in restricted classes of games (Moulin and Vial, 1978). In Section 4.6 we thus come to consider the following comparative ratios in an effort to facilitate the quantitative analysis of solution concepts in ranking games: the price of cautiousness, i.e., the ratio between an agent's minimum payoff in a Nash equilibrium and his security level; the mediation value, i.e., the ratio between the social welfare obtainable in the best correlated equilibrium vs. the best Nash equilibrium; and the enforcement value, i.e., the ratio between the highest obtainable social welfare and that of the best correlated equilibrium. Very much like solution concepts themselves, these values can be interpreted descriptively and used as tool for comparing the properties of different games, or they can be viewed from a normative perspective as providing an index of what can be gained or lost by following a more daring rather than a more conservative course of action. Each of the above values obviously equals 1 in the case of two-player ranking games, as these form a subclass of constant-sum games. An interesting question thus concerns bounds on these values for ranking games with more than two players.

4.1 An Introductory Example

To further illustrate the issues addressed in this chapter, recall the game of Figure 2.1, where Alice, Bob, and Charlie are to designate one of them as the winner, by simultaneously and independently raising their hand or not. Alice wins if the number of hands raised, including her own, is odd, whereas Bob is victorious if this number equals two. Should nobody raise their hand, Charlie wins. It is obvious from the description, and can



Figure 4.1: The game of Figure 2.1, involving Alice (player 1), Bob (player 2), and Charlie (player 3). Outcomes are now denoted by the index of the winning player. The dashed square marks the only pure Nash equilibrium. Dotted rectangles mark an equilibrium in which Alice and Charlie randomize uniformly over their respective actions.

also be seen from Figure 2.1, that this game is a ranking game, and in fact a single-winner game. We show the game again in Figure 4.1, now denoting each outcome by the index of the winning player.

What course of action would you recommend to Alice? There is a Nash equilibrium in which Alice raises her hand, another one in which she does not raise her hand, and still another one in which she randomizes uniformly between these two options. In the only pure equilibrium of the game, Alice does not raise her hand. For the latter to occur, Alice would have to believe that Bob will raise his hand and Charlie will not. This assumption is unreasonably strong, however, in that no such beliefs can be derived from the mere description of the game. Moreover, both Bob and Charlie could in the above case deviate from their respective strategies to *any* other strategy without decreasing their chances of winning. After all, they cannot do any worse than losing. On the other hand, by playing her maximin strategy, Alice would guarantee a particular payoff, or winning probability, no matter which actions her opponents choose. Alice's security level in this particular game is 1/2 and can be obtained by randomizing uniformly between both actions. The same expected payoff is achieved in the mixed equilibrium where Alice and Charlie randomize uniformly and Bob invariably raises his hand, indicated by the dotted rectangles in Figure 4.1.

4.2 Related Work

In game theory, several proposals have been made for broader classes of games that maintain some of the desirable properties of two-player zero-sum games. The term strict competitiveness is usually reserved for two-player zero-sum games. Friedman (1983) shows that any convex combination of equilibria in such games is again an equilibrium. In an attempt to characterize games which can always be solved, in the sense that players have a single optimal strategy and the outcome is strictly determined, Aumann (1961) defines *almost strictly competitive* games. A two-player game is almost strictly competitive if a pair of strategies is an equilibrium point, i.e., no player can increase his payoff by unilaterally changing his strategy, if and only if it is a so-called twisted equilibrium point, i.e., no player can decrease the payoff of his opponent. These games permit a set of optimal strategies for each player and a unique value that is obtained whenever a pair of such strategies is played. Moulin and Vial (1978) call a game strategically zero-sum if it is best-response equivalent to a zero-sum game. In the case of two players, and only in this case, one obtains exactly the class of games for which no completely mixed equilibrium can be improved upon by a correlated equilibrium. A game is unilaterally competitive, as defined by Kats and Thisse (1992), if any deviation by a player that (weakly) increases his own payoff must (weakly) decrease the payoffs of all other players. Unilaterally competitive games retain several interesting properties of two-player zero-sum games in the n-player case: all equilibria yield the same payoffs, equilibrium strategies are interchangeable, and the set of equilibria is convex provided that some mild conditions hold. It was later shown by Wolf (1999) that *pure* Nash equilibria of n-player unilaterally competitive games are always profiles of maximin strategies. When there are only two players, all of the above classes contain those of constant-sum games and thus two-player ranking games. Neither is contained in the other in the n-player case. The notion of competitiveness as embodied in ranking games is remotely related to *spitefulness* (Morgan et al., 2003, Brandt et al., 2007b), where agents aim at maximizing their payoff relative to the payoff of all other agents.

Most work on comparative ratios in game theory has been inspired by the literature on the price of anarchy (Koutsoupias and Papadimitriou, 1999, Roughgarden, 2005), i.e., the ratio between the highest obtainable social welfare and that of the best Nash equilibrium. Similar ratios for correlated equilibria were introduced by Ashlagi et al. (2005): the value of mediation, i.e., the ratio between the social welfare obtainable in the best correlated equilibrium and the best Nash equilibrium, and the enforcement value, i.e., the ratio between the highest obtainable social welfare and that of the best correlated equilibrium. It is known that the mediation value of strategically zero-sum games is 1 and that of almost strictly competitive games is greater than 1, showing that correlation can be beneficial even in games of strict antagonism (Raghavan, 2002). To our knowledge, Tennenholtz (2002) was the first to conduct a quantitative comparison of Nash equilibrium payoffs and security levels. This work is inspired by an intriguing example game due to Aumann (1985), in which the only Nash equilibrium yields each player no more than his security level although the equilibrium strategies are different from the maximin strategies. In other words, the equilibrium strategies yield security level payoffs without guaranteeing them.

4.3 The Model

Intuitively, a ranking game is a normal-form game whose outcomes are rankings of the players. A ranking indicates how well each player has done relative to the other players in the game. Formally, a ranking $r = [r_1, r_2, ..., r_n]$ is an ordering of the players in N

in which player r_1 is ranked first, player r_2 ranked second, and so forth, with player r_n ranked last. Obviously, this limits the number of possible outcomes to n! irrespective of the number of actions the players have at their disposal. The set of rankings over a set N of players we denote by R_N .

We assume that all players weakly prefer higher ranks over lower ranks, and strictly prefer being ranked first to being ranked last. Furthermore, each player is assumed to be indifferent as to the ranks of the other players. Even so, a player may prefer to be ranked second for certain to having a fifty-fifty chance of being ranked first or being ranked third, whereas other players may judge quite differently. Accordingly, we have a rank payoff function $p_i: R_N \to \mathbb{R}$ represent player i's von Neumann-Morgenstern preferences over lotteries over R_N . For technical convenience, we normalize the payoffs to the unit interval [0, 1]. Formally, a rank payoff function p_i over R_N satisfies the following three conditions for all rankings $r, r' \in R_N$:

- (i) $p_i(r) \ge p_i(r')$, if $r_k = r'_m = i$ and $k \le m$,
- (*ii*) $p_i(r) = 1$, if $i = r_1$, and
- (*iii*) $p_i(r) = 0$, if $i = r_n$.

It will be convenient to make the relationship between the outcomes of a game and payoffs obtained in these outcomes explicit. To this end, let a game form be a tuple $(N, (A_i)_{i \in N}, \Omega, g)$ where N is a set of players, A_i is a set of actions available to player $i \in N, \Omega$ is a set of outcomes, and $g: X_{i \in N}A_i \to \Omega$ is an outcome function mapping each action profile to an outcome in Ω . A game form then is a ranking game form if the set of outcomes is given by the set of rankings of the players, i.e., if $\Omega = R_N$. We are now in a position to formally define the concept of a ranking game.

DEFINITION 4.1 (ranking game). A normal-form game $\Gamma = (N, (A_i)_{i \in N}, (p_i)_{i \in N})$ is called ranking game if there exists a ranking game form $(N, (A_i)_{i \in N}, R_N, g)$ and rank payoff functions p'_i over R_N such that for all $a \in A_N$ and all $i \in N$, $p_i(a_N) = p'_i(g(a_N))$.

Condition (i) above implies that a player's payoff for a ranking r only depends on the rank assigned to him in r. Accordingly, for $1 \le k \le n$, we have p_i^k denote the *unique* payoff player i obtains in any ranking r in which he is ranked kth. The rank payoff function of player i can then conveniently and compactly be represented by his rank payoff vector $\vec{p}_i = (p_i^1, p_i^2, \dots, p_i^n)$.

In a binary ranking game, a player is completely satisfied up to a certain rank, and not satisfied at all for any lower rank. The expected payoff of a player given a strategy profile can then be taken as his chances of being satisfied. In this case, the use of expected payoffs, and thus randomized strategies, is justified without relying on the von Neumann-Morgenstern axioms (see also Aumann, 1987). An interesting subclass of binary ranking games are so-called *single-winner games*, in which all players are only interested in being ranked first. Formally, a single-winner game is a ranking game in which $\vec{p}_i = (1, 0, ..., 0)$

	с	1	c^2		
	b^1	b^2	b1	b^2	
a ¹	[1,3,2]	[2, 1, 3]	[3, 2, 1]	[3, 1, 2]	
a ²	[2, 3, 1]	[3, 2, 1]	[2, 1, 3]	[1,3,2]	

Figure 4.2: A $2 \times 2 \times 2$ ranking game form. One player chooses rows, another columns, and a third matrices. Each combination of actions results in a ranking. For example, action profile (a^2, b^2, c^2) leads to the row player 1 being ranked first, the matrix player 3 second and the column player 2 third.

	C	1	c^2		
	b1	b ²	b1	b ²	
a ¹	(1,0,1)	$(\frac{1}{2}, 1, 0)$	(0,0,1)	$(\frac{1}{2}, 0, 1)$	
a ²	(0,1,1)	(0,0,1)	$(\frac{1}{2}, 1, 0)$	(1,0,1)	

Figure 4.3: A ranking game associated with the ranking game form of Figure 4.2. The rank payoff for the three players are given by $\vec{p}_1 = (1, \frac{1}{2}, 0), \vec{p}_2 = (1, 0, 0)$ and $\vec{p}_3 = (1, 1, 0)$.

for all $i \in N$. When considering mixed strategies, the expected payoff in a single-winner ranking game equals the probability of winning. Analogous to single-winner games, we can define *single-loser games* as ranking games in which the players' only concern is not to be ranked last, as for instance in a round of musical chairs. Formally, single-loser games are ranking games where $\vec{p}_i = (1, \ldots, 1, 0)$ for each player i. For an example illustrating the definitions of a ranking game form and a ranking game the reader is referred to Figures 4.2 and 4.3, respectively.

At this point, a remark as to the relationship between ranking games and n-player constant-sum games is in order. By virtue of conditions (*ii*) and (*iii*), two-player ranking games constitute a subclass of constant-sum games. If more than two players are involved, however, any such relation with n-person constant-sum games ceases to hold. A strategic game can be converted to a zero-sum game via positive affine transformations only if all outcomes of the game lie on an (n - 1)-dimensional hyperplane in the n-dimensional outcome space. Clearly, there are ranking games, with non-identical rank payoff vectors and more than two players, for which this is *not* the case, like a three-player ranking game with rank payoff vectors $\vec{p}_1 = \vec{p}_2 = (1,0,0)$ and $\vec{p}_3 = (1,1,0)$ that has among its outcomes the rankings [1,2,3], [2,1,3], [3,1,2], and [1,3,2]. As a consequence, ranking games are no subclass of constant-sum games. It is readily appreciated that the opposite inclusion does not hold either.

4.4 Games With Non-Pure Equilibria

We noticed earlier that in the only pure equilibrium of the game in Figure 4.1, two of the three players have no reason whatsoever to stick to their equilibrium strategies. This in fact indicates an inherent weakness of pure Nash equilibrium as a solution concept for ranking games. In any outcome of a ranking game, some player is ranked last and receives his minimum payoff. This means that in *any* pure Nash equilibrium ,some player must be indifferent between *all* of his actions: no matter what he does, he will remain a loser. Obviously, the stability of a solution like that is highly questionable, if it is to be considered a viable solution at all. This is particularly true for single-winner games, where in a pure equilibrium all players but the winner are indifferent over which action to play.

On the other hand, it is very well possible that all the actions in the support of a *mixed* equilibrium yield a strictly higher expected payoff than any action not in the support, mitigating the phenomenon mentioned above. Equilibria satisfying this condition are called *quasi-strict*, and will be treated in more detail in Chapter 7.

For now we concentrate on *non-pure* equilibria, i.e., equilibria in which at least one player randomizes. We conjecture that every single-winner game possesses at least one such equilibrium, and prove this claim for three subclasses.

THEOREM 4.2. The following classes of ranking games always possess at least one non-pure equilibrium:

- (i) two-player ranking games,
- (ii) three-player single-winner games where each player has two actions, and
- (iii) n-player single-winner games where the security level of at least two players is positive.

Proof. Statement (i) follows directly from the fact that every two-player game has a quasistrict equilibrium (Norde, 1999), and the above observation that quasi-strict equilibria of ranking games are never pure. Here we give a simple alternative proof. Assume for contradiction that there is a two-player ranking game that only possesses pure equilibria and consider, without loss of generality, a pure equilibrium s_N^* in which player 1 wins. Since player 2 must be incapable of increasing his payoff by deviating from s_N^* , player 1 has to win no matter which action the second player chooses. As a consequence, the strategies in s_N^* remain in equilibrium even if player 2's strategy is replaced with an arbitrary randomization among his actions.

As for (ii), consider a three-player single winner game with actions $A_1 = \{a^1, a^2\}, A_2 = \{b^1, b^2\}$, and $A_3 = \{c^1, c^2\}$. Assume for contradiction that there are only pure equilibria in the game and consider, without loss of generality, a pure equilibrium $s_N^* = (a^1, b^1, c^1)$ in which player 1 wins. In the following, we say that a pure equilibrium is *semi-strict* if at least one player strictly prefers his equilibrium action over all his other actions, given that

the other players play their equilibrium actions. In single-winner games, this player has to be the winner in the pure equilibrium. We first show that if s_N^* is semi-strict, i.e., if player 1 does not win in action profile (a^2, b^1, c^1) , then there must exist a non-pure equilibrium. For this, consider the strategy profile $s_N^1 = (a^1, s_2^1, c^1)$, where s_2^1 is the uniform mixture of player 2's actions b^1 and b^2 , along with the strategy profile $s_N^2 = (a^1, b^1, s_3^2)$, where s_3^2 is the uniform mixture of actions c^1 and c^2 of player 3. Since player 1 does not win in (a^2, b^1, c^1) , he has no incentive to deviate from either s_N^1 or s_N^2 , even if he wins in (a^2, b^2, c^1) and (a^2, b^1, c^2) . Consequently, player 3 must win in (a^1, b^2, c^2) in order for s_N^1 not to be an equilibrium. Analogously, for s_N^2 not to be an equilibrium, player 2 has to win in the same action profile (a^1, b^2, c^2) , contradicting the assumption that the game is a single-winner game. The existence of a semi-strict pure equilibrium thus implies that of a non-pure equilibrium. Now assume that s_N^* is not semi-strict. When any of the action profiles in $B = \{(a^2, b^1, c^1), (a^1, b^2, c^1), (a^1, b^1, c^2)\}$ is a pure equilibrium, this also yields a non-pure equilibrium because two pure equilibria that only differ by the action of a single player can be combined into infinitely many mixed equilibria. For B not to possess any pure equilibria, there must be (exactly) one player for every profile in B who deviates to a profile in $C = \{(a^2, b^2, c^1), (a^2, b^1, c^2), (a^1, b^2, c^2)\}$, because the game is a single-winner game and because s_N^* is not semi-strict. Moreover, either player 1 or player 2 wins in (a^2, b^2, c^1) , player 2 or player 3 in (a^1, b^2, c^2) , and player 1 or player 3 in (a^2, b^1, c^2) . This implies two facts. First, the action profile $s_N^3 = (a^2, b^2, c^2)$ is a pure equilibrium because no player will deviate from s_N^3 to any profile in C. Second, the player who wins in s_N^3 strictly prefers the equilibrium outcome over the corresponding action profile in C, implying that s_N^3 is semi-strict. The above observation that every semi-strict equilibrium also yields a non-pure equilibrium completes the proof.

As for (iii), recall that the payoff a player obtains in equilibrium must be at least his security level. A positive security level for player i thus rules out all equilibria in which player i receives payoff zero, in particular all pure equilibria in which he does not win. If there are two players with positive security levels, both of them have to win with positive probability in any equilibrium of the game. In single-winner games, this can only be the case in a non-pure equilibrium.

We conjecture that this existence result in fact applies to the entire class of singlewinner games. To see this it does *not* extend to general ranking games, consider a 4-player game in which the first three players have two actions. Payoffs for the case when player 4 plays his first action d¹ are shown in Figure 4.4. We can in fact restrict our attention to this case by assigning payoffs (1, 0, 0, 0) to any action profile where player 4 plays one of his potential other actions, thereby making them strictly dominated. The resulting game is then similar to a 3-player game by van Damme (1983) that does not have any *quasi-strict* equilibria. It is furthermore straightforward to show that the game is a ranking game by virtue of rank payoff vectors $r^1 = r^3 = (1, 0, 0, 0), r^2 = (1, \frac{1}{2}, 0, 0), and r^4 = (1, 1, 1, 0), and$ that the pure equilibria (a¹, a², a³) and (b¹, b², b³) are the only equilibria of this game.

	c ¹			c^2			
	b1	b ²		b1	b^2		
a ¹	(0,0,0,1)	(1,0,0,1)		$(0, \frac{1}{2}, 0, 1)$	(0,0,1,1)	d	
a ²	(0,0,1,1)	(0, 1, 0, 1)		(1,0,0,1)	(0,0,0,1)	u	

Figure 4.4: Four-player ranking game in which all equilibria are pure

It remains open whether there exists a 3-player ranking game with only pure equilibria.

4.5 Solving Ranking Games

The question we will try to answer next is whether the rather restricted payoff structure of ranking games makes it possible to compute instances of common solution concepts more efficiently than in general games. For this reason, we focus on solution concepts that are known to be intractable for general games, namely (mixed) Nash equilibria (Chen and Deng, 2006, Daskalakis et al., 2009a), iterated weak dominance (Conitzer and Sandholm, 2005a), and pure Nash equilibria in games with many players and polynomial-time computable payoff functions (Schoenebeck and Vadhan, 2006). Graphical games, in which pure Nash equilibria are also known to be intractable (Gottlob et al., 2005), are of very limited use for representing ranking games. If two players are not connected by the neighborhood relation, either directly or via a common player in their neighborhood, then their payoffs are completely independent from each other. For a single-winner game with the reasonable restriction that every player wins in at least one outcome, this implies that there must be one designated player who alone decides which player wins the game. Similar properties hold for arbitrary ranking games. For iterated strict dominance (Conitzer and Sandholm, 2005a) or correlated equilibria (Papadimitriou, 2005) efficient algorithms exist even for general games. There thus is no further need to consider these solution concepts here. When in the following we refer to the hardness of a game we mean NP-hardness or PPAD-hardness of solving the game using a particular solution concept.

4.5.1 Mixed Nash Equilibria

Let us first consider Nash equilibria of games with a bounded number of players. Twoplayer ranking games only allow outcomes (1,0) and (0,1) and thus constitute a subclass of constant-sum games. Nash equilibria of constant-sum games can be found by linear programming (Dantzig, 1951), for which there is a polynomial time algorithm (Khachiyan, 1979).

To prove hardness for the case with more than two players, it suffices to show that three-player ranking games are at least as hard to solve as general two-player games. To appreciate this, observe that any n-player ranking game can be turned into an (n + 1)player ranking game by adding a player who has only one action at his disposal and who is invariably ranked last, keeping relative rankings of the other players intact. Nash equilibria of the (n + 1)-player game then naturally correspond to Nash equilibria of the n-player game. A key concept in our proof is that of a Nash homomorphism, a notion introduced by Abbott et al. (2005). We generalize their definition to games with more than two players.

DEFINITION 4.3 (Nash homomorphism). A Nash homomorphism is a mapping h from a set of games into a set of games, such that there exists a polynomial-time computable function f that, when given a game Γ and an equilibrium s_N^* of $h(\Gamma)$, returns an equilibrium $f(s_N^*)$ of Γ .

Obviously, the composition of two Nash homomorphisms is again a Nash homomorphism. Furthermore, any sequence of polynomially many Nash homomorphisms that maps some class of games to another class of games provides us with a polynomial-time reduction from the problem of finding Nash equilibria in the former class to finding Nash equilibria in the latter. Any efficient, i.e., polynomial-time, algorithm for the latter directly leads to an efficient algorithm for the former. On the other hand, hardness of the latter implies hardness of the former.

A very simple example of a Nash homomorphism is the one that scales the payoff of each player by means of a positive affine transformation. It is well-known that Nash equilibria are invariant under this kind of mapping, and f can be taken to be the identity. We will now combine this Nash homomorphism with a more sophisticated function, which maps payoff profiles of a two-player binary game to corresponding three-player subgames with two actions for each player, and obtain Nash homomorphisms from two-player games to three-player ranking games for all possible rank payoff profiles.

LEMMA 4.4. For every rank payoff profile, there exists a Nash homomorphism from the set of two-player games to the set of three-player ranking games.

Proof. Abbott et al. (2005) have shown that there is a Nash homomorphism from twoplayer games to binary two-player games. Since a composition of Nash homomorphisms is again a Nash homomorphism, we only need to provide a homomorphism from binary two-player games to three-player ranking games. Furthermore, outcome (1, 1) is Paretodominant and therefore constitutes a pure Nash equilibrium in any binary game, since no player can benefit from deviating. Instances containing such an outcome are easy to solve and need not be considered in our mapping.

In the following, we denote by $(1, p_i^2, 0)$ the rank payoff vector of player i, and by [i, j, k] the outcome where player i is ranked first, j is ranked second, and k is ranked last. First of all, consider ranking games where $p_i^2 < 1$ for some player $i \in N$, i.e., the class of all ranking games except single-loser games.

Outcome	С	Constant-sur outcome	n		Ranking	subgame	
(0,0)	\longmapsto	$(\frac{1}{2}, \frac{1}{2}, 1)$	\longmapsto	(1,0,1) (0,1,1)	(0, 1, 1) (1, 0, 1)	(0, 1, 1) (1, 0, 1)	(1,0,1) (0,1,1)
(1,0)	\longmapsto	$(1, \frac{1}{2}, \frac{1}{2})$	\longmapsto	(1,0,1) (1,1,0)	(1,1,0) (1,0,1)	(1,1,0) (1,0,1)	(1,0,1) (1,1,0)
(0,1)	\longmapsto	$(\frac{1}{2}, 1, \frac{1}{2})$	\longmapsto	(0,1,1) (1,1,0)	(1,1,0) (0,1,1)	(1, 1, 0) (0, 1, 1)	(0,1,1) (1,1,0)

Figure 4.5: Mapping from binary two-player games to three-player single-loser games

Without loss of generality let i = 1. Then, a Nash homomorphism from binary twoplayer games to the aforementioned class of games can be obtained by first transforming the payoffs according to

$$(x_1, x_2) \longmapsto ((1 - p_1^2)x_1 + p_1^2, x_2)$$

and then adding a third player who only has a single action and whose payoff is chosen such that the resulting game is a ranking game (but is otherwise irrelevant). We obtain the following mapping, which is obviously a Nash homomorphism:

(0, 0)	\longmapsto	$(p_1^2, 0)$	\longmapsto	[3, 1, 2]
(1, 0)	\longmapsto	(1,0)	\longmapsto	[1, 3, 2]
(0, 1)	\longmapsto	$(p_1^2, 1)$	\longmapsto	[2, 1, 3].

Interestingly, three-player single-loser games with only one action for some player $i\in N$ are easy to solve because either

- there is an outcome in which i is ranked last and the other two players both receive their maximum payoff of 1, i.e., a Pareto-dominant outcome, or
- i is *not* ranked last in any outcome, such that the payoffs of the other two players always sum up to 1 and the game is equivalent to a two-player zero-sum game.

As soon as the third player is able to choose between *two* different actions, however, binary games can again be mapped to single-loser games. For this, consider the mapping from binary two-player games to three-player single-loser games shown in Figure 4.5. As a first step, binary two-player games are mapped to three-player constant-sum games according to

$$(\mathbf{x}_1, \mathbf{x}_2) \longmapsto \left(\frac{1}{2}(\mathbf{x}_1 + 1), \frac{1}{2}(\mathbf{x}_2 + 1), 1 - \frac{1}{2}(\mathbf{x}_1 + \mathbf{x}_2)\right).$$

The first two players and their respective sets of actions are the same as in the original game, the third player only has one action c. It is again obvious that this constitutes a Nash homomorphism. Next, outcomes of the three-player constant-sum game are replaced by three-player single-loser subgames. Let Γ be a binary game, and denote by Γ' and Γ'' the three-player constant-sum game and the three-player single-loser game, respectively, obtained by applying the two steps of the mapping in Figure 4.5 to Γ . We further write p'_i , and p''_i for the payoff function of player i in Γ' and Γ'' , and a^1_i and a^2_i for the two actions of player i in Γ'' .

The second part of the mapping in Figure 4.5 is chosen such that for all strategy profiles s_N , all players i and all actions $a_i \in A_i$ in Γ' we have

$$\frac{1}{2}p_{i}''(a_{i}^{1},s_{-i}) + \frac{1}{2}p_{i}''(a_{i}^{2},s_{-i}) = p_{i}'(a_{i},f(s_{N})_{-i}),$$
(4.1)

where for each strategy profile s_N of Γ'' , $f(s_N)$ is the strategy profile in Γ' such that for each player $i \in \{1, 2, 3\}$ and each action $a_i \in A_i$,

$$f(s_N)(a_i) = s_i(a_i^1) + s_i(a_i^2).$$

An important property of this construction is that each player can guarantee his payoff in Γ'' , for any strategy profile of the other players, to be at least as high as his payoff under the corresponding strategy profile in Γ' , by distributing the weight on a_i uniformly on a_i^1 and a_i^2 .

Let s_N^* be a Nash equilibrium in Γ'' . We first prove that for every player $i \in \{1, 2, 3\}$ and each action a_i of player i in Γ' ,

$$s_{i}^{*}(a_{i}^{1})p_{i}^{\prime\prime}(a_{i}^{1},s_{-i}^{*}) + s_{i}^{*}(a_{i}^{2})p_{i}^{\prime\prime}(a_{i}^{2},s_{-i}^{*}) = (f(s_{N}^{*})(a_{i}))p_{i}^{\prime}(a_{i},f(s_{N}^{*})_{-i}).$$
(4.2)

Recall that we write $s_N(a_i)$ for the probability of action a_i in strategy profile s_N , so $f(s_N^*)(a_i)$ is the probability with which a_i is played in strategy profile $f(s_N^*)$ of Γ' . The above equation thus states that the expected joint payoff from a_i^1 and a_i^2 in equilibrium s_N^* equals that from a_i under the corresponding strategy profile $f(s_N^*)$ of Γ' . To see this, first assume for contradiction that for some player i and some action $a_i \in A_i$,

$$s_{i}^{*}(a_{i}^{1})p_{i}^{\prime\prime}(a_{i}^{1},s_{-i}^{*}) + s_{i}^{*}(a_{i}^{2})p_{i}^{\prime\prime}(a_{i}^{2},s_{-i}^{*}) < (f(s_{N}^{*})(a_{i}))p_{i}^{\prime}(a_{i},f(s_{N}^{*})_{-i}) < (f(s_{N}^{*})(a_{i}))p_{i}^{\prime\prime}(a_{i},f(s_{N}^{*})_{-i}) < (f(s_{N}^{*})(a_{i}))p_{i}^{\prime\prime}(a_{i}) < (f(s_{N}^{*})(a_{i}))p_{i}^{\prime\prime}(a_{i}))$$

i.e., that the expected joint payoff from a_i^1 and a_i^2 in Γ'' is strictly smaller than the expected payoff from a_i in Γ' . Define s_i to be the strategy of player i in Γ'' such that $s_i(a_i^1) = s_i(a_i^2) = \frac{1}{2}(s_i^*(a_i^1) + s_i^*(a_i^2))$ and $s_i(a_i') = s_i^*(a_i')$ for all actions $a_i' \in A_i$

distinct from a_i^1 and a_i^2 . It then holds that

$$\begin{split} s_{i}^{*}(a_{i}^{1})p_{i}^{\prime\prime}(a_{i}^{1},s_{-i}^{*}) + s_{i}^{*}(a_{i}^{2})p_{i}^{\prime\prime}(a_{i}^{2},s_{-i}^{*}) \\ &< (f(s_{N}^{*})(a_{i}))p_{i}^{\prime}(a_{i},f(s_{N}^{*})_{-i}) \\ &= (s_{i}^{*}(a_{i}^{1}) + s_{i}^{*}(a_{i}^{2}))p_{i}^{\prime}(a_{i},f(s_{N}^{*})_{-i}) \\ &= (s_{i}^{*}(a_{i}^{1}) + s_{i}^{*}(a_{i}^{2}))(\frac{1}{2}p_{i}^{\prime\prime}(a_{i}^{1},s_{-i}^{*}) + \frac{1}{2}p_{i}^{\prime\prime}(a_{i}^{2},s_{-i}^{*})) \\ &= \frac{1}{2}(s_{i}^{*}(a_{i}^{1}) + s_{i}^{*}(a_{i}^{2}))p_{i}^{\prime\prime}(a_{i}^{1},s_{-i}^{*}) + \frac{1}{2}(s_{i}^{*}(a_{i}^{1}) + s_{i}^{*}(a_{i}^{2}))p_{i}^{\prime\prime}(a_{i}^{2},s_{-i}^{*}) \\ &= s_{i}(a_{i}^{1})p_{i}^{\prime\prime}(a_{i}^{1},s_{-i}^{*}) + s_{i}(a_{i}^{2})p_{i}^{\prime\prime}(a_{i}^{2},s_{-i}^{*}). \end{split}$$

The second and last step respectively follow from the definition of f and s_i . The third step follows from (4.1). We conclude that player i obtains a higher payoff by playing s_i instead of s_i^* , contradicting the assumption that s_N^* is a Nash equilibrium. In particular we have shown that for all $i \in N$ and every $a_i \in A_i$,

$$s_{i}^{*}(a_{i}^{1})p_{i}^{\prime\prime}(a_{i}^{1},s_{-i}^{*}) + s_{i}^{*}(a_{i}^{2})p_{i}^{\prime\prime}(a_{i}^{2},s_{-i}^{*}) \ge (f(s_{N}^{*})(a_{i}))p_{i}^{\prime}(a_{i},f(s_{N}^{*})_{-i}).$$
(4.3)

Now assume, again for contradiction, that for some player i and some action $a_i \in A_i$,

$$s_{i}^{*}(a_{i}^{1})p_{i}^{\prime\prime}(a_{i}^{1},s_{-i}^{*}) + s_{i}^{*}(a_{i}^{2})p_{i}^{\prime\prime}(a_{i}^{2},s_{-i}^{*}) > (f(s_{N}^{*})(a_{i}))p_{i}^{\prime}(a_{i},f(s_{N}^{*})_{-i}),$$

i.e., that the expected joint payoff to i from a_i^1 and a_i^2 in Γ'' is strictly greater under s_N^* than the expected payoff from a_i in Γ' . It follows from (4.3) that the expected payoff player i receives from *any* action under $f(s_N^*)$ cannot be greater than the expected joint payoff from the corresponding pair of actions under s_N^* , and thus $p_i''(s_N^*) > p_i'(f(s_N^*))$. Since Γ' and Γ'' are both constant-sum games, there exists some player $j \neq i$ who receives strictly less payoff under s_N^* in Γ'' than under $f(s_N^*)$ in Γ' . In particular, there has to be an action $a_j \in A_j$ such that

$$s_{j}^{*}(a_{j}^{1})p_{j}^{\prime\prime}(a_{j}^{1},s_{-j}^{*}) + s_{j}^{*}(a_{j}^{2})p_{j}^{\prime\prime}(a_{j}^{2},s_{-j}^{*}) < (f(s_{N}^{*})(a_{j}))p_{j}^{\prime}(a_{j},f(s_{N}^{*})_{-j}),$$

contradicting (4.3).

We are now ready to prove that the mapping in Figure 4.5 is indeed a Nash homomorphism. To this end, let s_N^* be a Nash equilibrium of Γ'' , and assume for a contradiction that $f(s_N^*)$ is not a Nash equilibrium of Γ' . Then there has to be a player i and some action $a_i \in A_i$ such that $p'_i(a_i, f(s_N^*)_{-i}) > p'_i(f(s_N^*))$. Define s_i to be the strategy of player i in Γ'' such that $s_i(a_i^1) = s_i(a_i^2) = \frac{1}{2}$. Then, by (4.1), $p''_i(s_i, s_{-i}) = p'_i(a_i, f(s_N^*)_{-i})$. It further follows from (4.2) that for every player j, $p''_j(s_N^*) = p'_j(f(s_N^*))$, and for j = i in particular. Thus,

$$p_i''(s_N^*) = p_i'(f(s_N^*)) < p_i'(a_i, f(s_N^*)_{-i}) = p_i''(s_i', s_{-i}^*),$$

contradicting the assumption that s_N^* is a Nash equilibrium in Γ'' .

The ground has now been cleared for the main result of this section.

THEOREM 4.5. Computing a Nash equilibrium of a ranking game with more than two players is PPAD-hard for any rank payoff profile. If there are only two players, equilibria can be found in polynomial time.

Proof. According to Lemma 4.4, ranking games with more than two players are at least as hard to solve as general two-player games. We already know that solving general games is PPAD-hard in the two-player case (Chen and Deng, 2006).

Two-player ranking games, on the other hand, form a subclass of two-player constantsum games, in which Nash equilibria can be found efficiently via linear programming. \Box

4.5.2 Iterated Weak Dominance

We now turn to iterated weak dominance. If there are only two players, the problem of deciding whether a ranking game can be solved via iterated weak dominance is tractable.

THEOREM 4.6. For two-player ranking games, iterated weak dominance solvability can be decided in polynomial time.

Proof. First we recall that if an action in a binary game is weakly dominated by a mixed strategy, it is also dominated by a pure strategy (Conitzer and Sandholm, 2005a). Accordingly, we only have to consider dominance by pure strategies. Now consider a path of iterated weak dominance that ends in a single action profile (a_1^*, a_2^*) . Without loss of generality we may assume that player 1, i.e., the row player, is the winner in this profile. This implies that player 1 wins in (a_1^*, a_2) for any $a_2 \in A_2$, i.e., in the entire row. For contradiction, assume the opposite and consider the particular action a_2^1 such that player 2 wins in (a_1^*, a_2^1) and a_2^1 is eliminated last on the path that solves the game. It is easy to see that such an elimination could only have taken place via another action a_2^2 such that player 2 also wins in (a_1^*, a_2^2) , contradicting the assumption that a_2^1 is eliminated last. We now claim that a ranking game with two players is solvable by iterated weak dominance if and only if there exists a unique action a_1^* of player 1 by which he always wins, and an action a_2^* of player 2 by which he wins for a strictly maximal set of actions of player 1. More precisely, the latter property means that there exists a set of actions of player 1 against which player 2 always wins when playing a_2^* and loses in at least one case for every other action he might play. This is illustrated in Figure 4.6, and can be verified efficiently by ordering the aforementioned sets of actions of player 1 according to strict inclusion. If the ordering does not have a maximal element, the game cannot be solved by means of iterated weak dominance. If it does, we can use a_1^* to eliminate all actions $a_1 \in A_1$ such that player 2 does not win in (a_1, a_2^*) , whereupon a_2^* can eliminate all other actions of player 2, until finally a_1^* eliminates player 1's remaining actions and solves the game.¹

¹Since two-player ranking games are a subclass of constant-sum games, weak dominance and *nice weak*



Figure 4.6: Iterated weak dominance solvability in two-player ranking games

As soon as there are three players, iterated weak dominance solvability becomes hard to decide.

THEOREM 4.7. For ranking games with more than two players, and for any rank payoff profile, deciding iterated weak dominance solvability is NP-complete.

Proof. Membership in NP is immediate. We can simply guess a sequence of eliminations and then verify in polynomial time that this sequence is valid and solves the game.

For hardness, we first reduce eliminability in binary two-player games, which asks whether there exits a sequence of eliminations that contains a given action and has recently been shown to be NP-hard (Conitzer and Sandholm, 2005a), to the same problem in ranking games. A game Γ of the former class is mapped to a ranking game Γ' as follows:

- Γ' features the two players of Γ , denoted by 1 and 2, and an additional player 3.
- Players 1 and 2 have the same actions as in Γ , player 3 has two actions c^1 and c^2 .
- Payoffs of Γ are mapped to rankings of Γ' according to



In the following, we write p and p' for the payoff functions of Γ and Γ' , respectively.

First observe that we can restrict our attention to dominance by pure strategies. This property holds for binary games by Lemma 1 of Conitzer and Sandholm (2005a), and thus

dominance (Marx and Swinkels, 1997) coincide, making iterated weak dominance order independent up to payoff-equivalent action profiles. This fact is mirrored by Figure 4.6, since there cannot be a row of 1s and a column of 2s in the same matrix.

		c^1			c^2	
	b1		\mathfrak{b}^k	b1		\mathfrak{b}^k
a^1	$\left[\cdot,2,\cdot\right]$		[·,2, ·]	$[\cdot, 1, \cdot]$		$[\cdot, 1, \cdot]$
÷	÷	·	÷		·	÷
a ^m	$[\cdot, 2, \cdot]$		$[\cdot, 2, \cdot]$	$[\cdot, 1, \cdot]$		$[\cdot, 1, \cdot]$
a^{m+1}	[3, 2, 1]		[3, 2, 1]	[2, 1, 3]		[2, 1, 3]

Figure 4.7: Three-player ranking game Γ' used in the proof of Theorem 4.7

also for actions of player 3, who receives a payoff of either 0 or 1 in any outcome. For players 1 and 2 we can essentially apply the same argument, because each of them can obtain only two different payoffs for any fixed action profile of the remaining two players.

We now claim that irrespective of the rank payoffs $p_i = (1, p_i^2, 0)$, and for any subsets of the actions of players 1 and 2, a particular action of these players is dominated in the restriction of Γ' to these subsets if and only if the corresponding action is dominated in the restriction of Γ to the same subsets. To see this, observe that if player 3 plays c^1 , then for any action profile $(a_1, a_2) \in A_1 \times A_2$, player 1 receives the same payoff he would receive for the corresponding action profile in Γ , i.e., $p'_1(a_1, a_2, c^1) = p_1(a_1, a_2)$, whereas player 2 receives a payoff of p_2^2 . If on the other hand player 3 plays c^2 , then player 1 obtains a payoff of p_1^2 , and the payoff of player 2 for any action profile $(a_1, a_2) \in A_1 \times A_2$ is the same as that for the corresponding profile in Γ , i.e., $p'_2(a_1, a_2, c^2) = p_2(a_1, a_2)$. Moreover, the implication from left to right still holds if one of the actions of player 3 is removed, because this leaves one of players 1 and 2 indifferent between all of his remaining actions but does not have any effect on dominance between sequences of eliminations in Γ and Γ' , which in turn implies NP-hardness of deciding whether a particular action of a ranking game with at least three players can be eliminated.

It also follows from the above that Γ can be solved by iterated weak dominance if Γ' can. The implication in the other direction does not hold, however, because it may not always be possible to eliminate an action of player 3. To this end, assume without loss of generality that some player of Γ' has at least two actions, and that this player is player 1. Otherwise both Γ and Γ' are trivially solvable. We augment Γ' by adding to the action set $A_1 = \{a^1, a^2, \ldots, a^m\}$ of player 1 an additional action a^{m+1} such that for every action b^j of player 2, $g(a^{m+1}, b^j, c^1) = [3, 2, 1]$ and $g(a^{m+1}, b^j, c^2) = [2, 1, 3]$. The structure of the resulting game is shown in Figure 4.7.

It is easily verified that the above arguments about Γ' still apply, because player 1 never receives a higher payoff from a^{m+1} than from any other action, and player 2 is indifferent between all of his actions when player 1 plays a^{m+1} . Now assume that Γ

can be solved. Without loss of generality we may assume that (a^1, b^1) is the remaining action profile. Clearly, for Γ to be solvable, player 1 must be ranked first in some outcome of Γ' , and it must hold that $p_1(a^1, b^1) = 1$ or $p_2(a^1, b^1) = 1$. We distinguish two cases. If $p_1(a^1, b^1) = p_2(a^1, b^1) = 1$, then Γ' can be solved by performing the eliminations that lead to the solution of Γ , followed by the elimination of c^2 and a^{m+1} . Otherwise we can start by eliminating a^{m+1} , which is dominated by the action for which player 1 is sometimes ranked first, and proceed with the eliminations that solve Γ . In the two action profiles that then remain, player 3 is ranked first and last, respectively, and he can eliminate one of his actions to solve Γ' .

4.5.3 Pure Nash Equilibria in Games with Many Players

When a game is given explicitly, pure Nash equilibria can be found efficiently by simply checking for every action profile whether it satisfies the equilibrium condition. As the number of players increases, however, the number of profiles to check grows exponentially, just as the explicit representation of the game. An interesting question is whether pure equilibria can be computed efficiently given a representation of a game that only uses space polynomial in n.

We proceed to show that for ranking games this is most likely not the case. More precisely, we show NP-completeness of deciding whether there is a pure Nash equilibrium in ranking games with *efficiently computable outcome functions*, which is one of the most general representations of multi-player games one might think of. It should be noted that in contrast to Theorems 4.5 and 4.7, we now fix the number of actions and let the number of players grow.

THEOREM 4.8. For ranking games with an unbounded number of players and a polynomial-time computable outcome function, and for any rank payoff profile, deciding the existence of a pure Nash equilibrium is NP-complete. Hardness holds even for games with two actions.

Proof. Since we can check in polynomial time whether a particular player strictly prefers one rank over another, *membership* in NP is immediate. We can guess an action profile a_N and verify in polynomial time whether a_N is a Nash equilibrium. For the latter, we check for each player $i \in N$ and for each action $a'_i \in A_i$ whether $p_i(a_{-i}, a'_i) \leq p_i(a_N)$.

For *hardness*, recall that circuit satisfiability (CSAT), i.e., deciding whether for a given Boolean circuit C there exists an input such that the output is *true*, is NP-complete (e.g., Papadimitriou, 1994a). Given a particular Boolean circuit C with m inputs, we construct a game Γ with players $N = \{1, 2, ..., m\} \cup \{x, y\}$, actions $A_i = \{0, 1\}$ for all $i \in N$, and payoff functions as follows:

 The outcome function of Γ is computed by a Boolean circuit that takes m + 2 bits of input (a₁, a₂,..., a_m, a_x, a_y), corresponding to the actions of the players in N, and computes two bits of output $o = (o_1, o_2)$, given by $o_1 = C(a_1, a_2, ..., a_m)$ and $o_2 = (o_1 \lor (a_x \text{ XOR } a_y))$.

• The possible outputs of the circuit are identified with permutations of the players in N such that the permutation π_{00} corresponding to o = (0,0) and the permutation π_{11} corresponding to o = (1,1) both rank x first and y last, the permutation π_{01} corresponding to o = (0,1) ranks y first and x last, and all other players are ranked in the same order in all three permutations. It should be noted that no matter how permutations are actually encoded as strings of binary values, the encoding of the above permutations can always be computed using a polynomial number of gates.

We now claim that, for arbitrary rank payoffs, Γ has a pure Nash equilibrium if and only if C has a satisfying assignment. This can be seen as follows:

- If (a₁, a₂,..., a_m) is a satisfying assignment of C, only a player in {1,2,...,m} could possibly change the outcome of the game by changing his action. However, these players are ranked in the same order in all the possible outcomes, so none of them can get a higher payoff by doing so. Thus, every action profile (a₁, a₂,..., a_m, a_x, a_y) such that (a₁, a₂,..., a_m) is a satisfying assignment of C is a Nash equilibrium.
- If in turn (a₁, a₂,..., a_m) is not a satisfying assignment of C, both x and y are able to switch between outcomes π₀₀ and π₀₁ by changing their own action. Since further every player strictly prefers being ranked first over being ranked last, x strictly prefers outcome π₀₀ over π₀₁, while y strictly prefers π₀₁ over π₀₀. Thus, (a₁, a₂,..., a_m, a_x, a_y) cannot be a Nash equilibrium, since either x or y could play a different action to get a higher payoff.

4.6 Comparative Ratios

Despite its conceptual elegance and simplicity, Nash equilibrium has been criticized on various grounds (see, e.g., Luce and Raiffa, 1957, for a discussion). It might require randomization on the part of the players, and it is unclear how this randomization should be carried out reliably. In the common case when there are multiple equilibria, it is unclear which one should be selected. Coalitions might benefit from a joint deviation. There might not exist an efficient algorithm for finding Nash equilibria. Finally, players may be utterly indifferent among equilibrium and non-equilibrium strategies, which we saw is the case particularly often in ranking games. We therefore devote the remainder of this chapter to a comparison between Nash equilibrium and two alternative solution concepts: maximin strategies and Aumann's correlated equilibrium. In particular, we study how much a player can lose by playing a maximin strategy instead of a Nash equilibrium, and how much society can benefit from correlation.

4.6 · Comparative Ratios

4.6.1 The Price of Cautiousness

A compelling question is how much worse off a player can be when if he were to revert to his most defensive course of action—his maximin strategy—instead of hoping for an equilibrium outcome. This difference in payoff can be represented by a numerical value which we refer to as the *price of cautiousness*. In the following, let \mathcal{G} denote the class of all normal-form games, and for $\Gamma \in \mathcal{G}$, let $N(\Gamma)$ be the set of Nash equilibria of Γ . Further recall that $v_i(\Gamma)$ denotes player i's security level in game Γ .

DEFINITION 4.9 (price of cautiousness). Let Γ be a normal-form game with non-negative payoffs, $i \in N$ a player such that $v_i(\Gamma) > 0$. The *price of cautiousness* for player i in Γ is defined as

$$PC_{i}(\Gamma) = \frac{\min\{p_{i}(s_{N}) : s_{N} \in N(\Gamma)\}}{\nu_{i}(\Gamma)}.$$

We further write $PC_i(\mathcal{C}) = \sup_{\Gamma \in \mathcal{C}} PC_i(\Gamma)$, where $\mathcal{C} \subseteq \mathcal{G}$ can be any class of games involving player i. In other words, the price of cautiousness of a player is the ratio between his minimum payoff in a Nash equilibrium and his security level. It thus captures the worst-case loss the player may incur by playing his maximin strategy instead of a Nash equilibrium.² For a player whose security level equals his minimum payoff of zero, *every* strategy is a maximin strategy. Since we are mainly interested in a comparison of normative solution concepts, we will thus only consider games where the security level of at least one player is positive.

As we have already mentioned, Nash equilibrium and minimax strategies coincide in two-player ranking games by virtue of the Minimax Theorem of von Neumann (1928), so the price of cautiousness equals 1 for these games. In general ranking games, on the other hand, the price of cautiousness is unbounded.

THEOREM 4.10. Let \mathcal{R} be the class of ranking games with more than two players that involve player i. Then, the price of cautiousness is unbounded, i.e., $PC_i(\mathcal{R}) = \infty$, even if \mathcal{R} only contains games without weakly dominated actions.

Proof. Consider the game Γ_1 of Figure 4.8, which is a ranking game for rank payoff vectors $\vec{p}_1 = (1, \epsilon, 0)$, $\vec{p}_2 = (1, 0, 0)$, and $\vec{p}_3 = (1, 1, 0)$, and rankings [2, 3, 1], [1, 3, 2], [1, 2, 3], [2, 1, 3], and [3, 1, 2]. It is easily verified that none of the actions of Γ_1 is weakly dominated and that $\nu_1(\Gamma_1) = \epsilon$. Let further $s_N = (s_1, s_2, c^1)$ be the strategy profile where s_1 and s_2 are uniform mixtures of a^1 and a^2 , and of b^1 and b^2 , respectively. It is easily verified that s_N is a Nash equilibrium of Γ_1 , and we will argue that it is in fact the only one. For this, consider the possible strategies of player 3. If player 3 plays c^1 , the game reduces to the well-known Matching Pennies game for players 1 and 2, the only Nash equilibrium being the one described above. If on the other hand player 3 plays c^2 ,

 $^{^{2}}$ In our context, the choice of whether to use the worst or the best equilibrium when defining the price of cautiousness is merely a matter of taste. All results in this section still hold when the best equilibrium is used instead of the worst one.

	С	1	c^2		
	b^1	b ²	b^1	b ²	
a ¹	(0,1,1)	(1,0,0)	(e, 1, 0)	$(\varepsilon, 0, 1)$	
a ²	(1,0,1)	(0,1,1)	$(\epsilon, 1, 0)$	$(\varepsilon, 0, 1)$	

Figure 4.8: Three-player ranking game Γ_1 used in the proof of Theorem 4.10

	С	1	с	2
	\mathfrak{b}^1	b^2	\mathfrak{b}^1	b^2
a ¹	2	1	3	1
\mathfrak{a}^2	1	2	1	1

Figure 4.9: Three-player single-winner game used in the proof of Theorem 4.11. Dotted boxes mark all Nash equilibria, one player may mix arbitrarily in boxes that span two outcomes.

action b^1 strictly dominates b^2 . If b^1 is played, however, player 3 will deviate to c^1 to get a higher payoff. Finally, if player 3 randomizes between actions c^1 and c^2 , the payoff obtained from both of these actions must be the same. This can only be the case if either player 1 plays a^1 and player 2 randomizes between b^1 and b^2 , or if player 1 plays a^2 and player 2 plays b^2 . In the former case, player 2 will deviate to b^1 . In the latter case, player 1 will deviate to a^1 . Since the payoff of player 1 in the above equilibrium is 1/2, we have $PC(\Gamma_1) = 1/(2\varepsilon) \to \infty$ for $\varepsilon \to 0$.

We proceed to show that, due to their structural limitations, the price of cautiousness in *binary* ranking games is bounded from above by the number of actions of the respective player. We also derive a matching lower bound.

THEOREM 4.11. Let \mathcal{R}_b be the class of binary ranking games with more than two players involving a player i with exactly k actions. Then, $PC_i(\mathcal{R}_b) = k$, even if \mathcal{R}_b only contains single-winner games or games without weakly dominated actions.

Proof. By definition, the price of cautiousness takes its maximum for maximum payoff in a Nash equilibrium, which is bounded by 1 in a ranking game, and minimum security level. We require the security level to be strictly positive, so for every opponent action profile $s_{-i} \in S_{-i}$ there is some action $a_i \in A_i$ such that $p_i(a_i, s_{-i}) > 0$, i.e., $p_i(a_i, s_{-i}) =$ 1. It is then easily verified that player i can ensure a security level of 1/k by uniform randomization over his k actions. This results in a price of cautiousness of at most k.

For a matching lower bound, consider the single-winner game depicted in Figure 4.9. We will argue that all Nash equilibria of this game are mixtures of the action profiles

	с	1	c ²		
	b ¹	b ²	b1	b ²	
a ¹	(0,1,1)	(1,0,0)	(0,1,0)	(1,0,0)	
a ²	(1,0,0)	(0,1,0)	(1,0,1)	(1,0,1)	

Figure 4.10: Three-player ranking game Γ_2 used in the proof of Theorem 4.11

 (a^2, b^1, c^2) , (a^2, b^2, c^2) and (a^1, b^2, c^2) . Each of these equilibria yields payoff 1 for player 1, twice as much as his security level of 1/2. To appreciate this, consider the possible strategies for player 3. If player 3 plays c^1 , the game reduces to the well-known Matching Pennies game for players 1 and 2, in which they will randomize uniformly over both of their actions. In this case, player 3 will deviate to c^2 . If player 3 plays c^2 , we immediately obtain the equilibria described above. Finally, if player 3 randomizes between actions c^1 and c^2 , the payoff obtained from both of these actions should be the same. This can only be the case if either player 1 plays a^2 and player 2 randomizes between b^1 and b^2 , or if player 1 randomizes between a^1 and a^2 and player 2 plays b^2 . In the former case, player 2 will play b^2 , causing player 1 to deviate to a^1 . In the latter case, player 1 will play a^1 , causing player 2 to deviate to b^1 .

The above construction can be generalized to k > 2 by virtue of a single-winner game with actions $A_1 = \{a^1, a^2, \ldots, a^k\}, A_2 = \{b^1, b^2, \ldots, b^k\}$, and $A_3 = \{c^1, c^2\}$, and payoffs

$$p(a^{i}, b^{j}, c^{\ell}) = \begin{cases} (0, 1, 0) & \text{if } \ell = 1 \text{ and } i \neq k - j + 1 \\ (0, 0, 1) & \text{if } \ell = 2 \text{ and } i = j = 1 \\ (1, 0, 0) & \text{otherwise.} \end{cases}$$

It is easily verified that the security level of player 1 in this game is 1/k, while, by similar arguments as above, his payoff in every Nash equilibrium equals 1. This shows tightness of the upper bound of k on the price of cautiousness for single-winner games.

Now consider the game Γ_2 of Figure 4.10, which is a ranking game for rank payoff vectors $\vec{p}_1 = \vec{p}_2 = (1,0,0)$ and $\vec{p}_3 = (1,1,0)$, and rankings [2,3,1], [1,2,3], [2,1,3], and [1,3,2]. It is easily verified that none of the actions of Γ_2 is weakly dominated and that $v_1(\Gamma_2) = 1/2$. On the other hand, we will argue that all Nash equilibria of Γ_2 are mixtures of action profiles (a^2, b^1, c^2) and (a^2, b^2, c^2) , corresponding to a payoff of 1 for player 1. To see this, we again look at the possible strategies for player 3. If player 3 plays c^1 , players 1 and 2 will again randomize uniformly over both of their actions, causing player 3 to deviate to c^2 . If player 3 plays c^2 , we immediately obtain the equilibria described above. Finally, assume that player 3 randomizes between actions c^1 and c^2 , and let α denote the probability with which player 1 plays a^1 . Again, player 3 must be indifferent between c^1 and c^2 , which can only hold for $1/2 \leq \alpha \leq 1$. In this case, however, player 2 will deviate to b^1 .

This construction can be generalized to k > 2 by virtue of a game with actions $A_1 = \{a^1, a^2, \ldots, a^k\}, A_2 = \{b^1, b^2, \ldots, b^k\}$, and $A_3 = \{c^1, c^2\}$, and payoffs

$$p(a^{i}, b^{j}, c^{\ell}) = \begin{cases} (0, 1, 1) & \text{if } i = j = \ell = 1\\ (1, 0, 0) & \text{if } \ell = 1 \text{ and } i = k - j + 1\\ & \text{or } \ell = 2, i = 1 \text{ and } j > 1\\ (1, 0, 1) & \text{if } \ell = 2 \text{ and } j > 2\\ (0, 1, 0) & \text{otherwise.} \end{cases}$$

Again, it is easily verified that player 1 has a security level of 1/k, while his payoff is 1 in every Nash equilibrium by similar arguments as above. Thus, the upper bound of k for the price of cautiousness is tight as well for binary ranking games without weakly dominated actions.

Informally, the previous theorem states that the payoff a player with k actions can obtain in Nash equilibrium can be at most k times his security level.

4.6.2 The Value of Correlation

Nash equilibrium is based on the assumption that players select their actions *independently* from each other. Aumann (1974) generalizes the notion of a strategy profile by allowing players to coordinate their actions by means of a device or agent that randomly selects one of several action profiles and recommends the actions of this profile to the respective players. More formally, a *correlated strategy* $\mu \in \Delta(A_N)$ is a probability distribution over the set of action profiles. The corresponding equilibrium concept is then defined as follows.

DEFINITION 4.12 (correlated equilibrium). A correlated strategy $\mu \in \Delta(A_N)$ is called a *correlated equilibrium* if for all $i \in N$ and all $a_i^*, a_i \in A_i$,

$$\sum_{a_{-i}\in A_{-i}}\mu(a_{-i},a_i^*)(p_i(a_{-i},a_i^*)-p_i(a_{-i},a_i)) \ge 0.$$

In other words, a correlated equilibrium of a game is a probability distribution μ over the set of action profiles, such that, if a particular action profile $a_N^* \in A_N$ is chosen according to this distribution, and every player $i \in N$ is only informed about his own action a_i^* , it is optimal in expectation for every player $i \in N$ to play a_i^* , given that he only knows the conditional distribution over values of a_{-i}^* . Correlated equilibrium makes stronger assumptions than Nash equilibrium in that it assumes the existence of a trusted third party who can recommend behavior, but cannot enforce it. Using cryptographic means, this requirement can essentially be reduced to the ability to carry out a distributed computation among the players (Dodis et al., 2000).

4.6 · Comparative Ratios

It can easily be seen that every Nash equilibrium naturally corresponds to a correlated equilibrium. Nash's existence result thus carries over to correlated equilibria. Again consider the game of Figure 4.1 on Page 31. The correlated strategy that assigns probability 1/4 each to action profiles (a^1, b^1, c^1) , (a^1, b^2, c^1) , (a^2, b^1, c^1) , and (a^2, b^1, c^2) is a correlated equilibrium, with an expected payoff of 1/2 for player 1 and 1/4 for players 2 and 3. In this particular case, the correlated equilibrium is a convex combination of Nash equilibria, and correlation can be achieved by means of a publicly observable random variable. Quite surprisingly, Aumann (1974) has shown that in general the (expected) social welfare of a correlated equilibrium may exceed that of every Nash equilibrium, and that correlated equilibrium payoffs may in fact be outside the convex hull of the Nash equilibrium payoffs. This is of course not possible if social welfare is identical in all outcomes, as is the case in our example.

We will now turn to the question whether, and by which amount, social welfare in a ranking game can be improved by allowing players to correlate their actions. Just as the payoff of a player in any Nash equilibrium is at least his security level, social welfare in the best correlated equilibrium is at least as high as social welfare in the best Nash equilibrium. In order to quantify the value of correlation in strategic games with nonnegative payoffs, Ashlagi et al. (2005) introduce the *mediation value* of a game as the ratio between the maximum social welfare in a correlated versus that in a Nash equilibrium, and the *enforcement value* as the ratio between the maximum social welfare in any outcome versus that in a correlated equilibrium. Whenever social welfare, i.e., the sum of all players' payoffs, is used as a measure of global satisfaction, one implicitly assumes the inter-agent comparability of payoffs. While this assumption is controversial, social welfare is nevertheless commonly used in the definitions of comparative ratios such as the price of anarchy (Koutsoupias and Papadimitriou, 1999). For $\Gamma \in \mathcal{G}$ and $X \subseteq \Delta(A_N)$, let $C(\Gamma)$ denote the set of correlated equilibria of Γ and let $v_X(\Gamma) = \max\{p(s_N) : s_N \in X\}$. Recall that $N(\Gamma)$ denotes the set of Nash equilibria of game Γ .

DEFINITION 4.13 (mediation value, enforcement value). Let Γ be a normal-form game with non-negative payoffs. Then, the *mediation value* $MV(\Gamma)$ and the *enforcement value* $EV(\Gamma)$ of Γ are defined as

$$MV(\Gamma) = rac{
u_{C(\Gamma)}(\Gamma)}{
u_{N(\Gamma)}(\Gamma)}$$
 and $EV(\Gamma) = rac{
u_{S_N}(\Gamma)}{
u_{C(\Gamma)}(\Gamma)}$

If both numerator and denominator are 0 for one of the values, the respective value is defined to be 1. If only the denominator is 0, the value is defined to be ∞ . For any class $\mathcal{C} \subseteq \mathcal{G}$ of games, we further write $MV(\mathcal{C}) = \sup_{\Gamma \in \mathcal{C}} MV(\Gamma)$ and $EV(\mathcal{C}) = \sup_{\Gamma \in \mathcal{C}} EV(\Gamma)$.

Ashlagi et al. (2005) have shown that both the mediation value and the enforcement value cannot be bounded for games with an arbitrary payoff structure, as soon as there are more than two players, or some player has more than two actions. This holds even if payoffs are normalized to the interval [0, 1]. Ranking games also satisfy this normalization criterion, and here social welfare is also strictly positive for every outcome of the game.

	c ¹		c^2			c ³	
	b^1	b^2	b^1	b ²		b ¹	b ²
a ¹	(1,1,0)	(1,0,0)	(0,1,1)	(0,1,0)		(1,0,0)	(0,0,1)
\mathfrak{a}^2	(0,1,0)	(0,1,1)	(1,0,0)	(1, 1, 0)		(0,0,1)	(1,0,0)

Figure 4.11: Three-player ranking game Γ_3 used in the proof of Theorem 4.14

Ranking games with identical rank payoff vectors for all players, i.e., ones where $p_i^k = p_j^k$ for all $i, j \in N$ and $1 \leq k \leq n$, are constant-sum games. Hence, social welfare is the same in every outcome so that both the mediation value and the enforcement value are 1. This in particular concerns all ranking games with two players. In general, social welfare in an arbitrary outcome of a ranking game is bounded by n - 1 from above and by 1 from below. Since the Nash and correlated equilibrium payoffs must lie in the convex hull of the feasible payoffs of the game, we obtain trivial lower and upper bounds of 1 and n - 1, respectively, on both the mediation and the enforcement value. It turns out that the upper bound of n - 1 is tight for both the mediation value and the enforcement value.

THEOREM 4.14. Let \Re' be the class of ranking games with n > 2 players, such that in games with only three players at least one player has more than two actions. Then, $MV(\Re') = n - 1$.

Proof. It suffices to show that for any of the above cases there is a ranking game with mediation value n-1. For n=3, consider the game Γ_3 of Figure 4.11, which is a ranking game for rank payoff vectors $\vec{p}_1 = \vec{p}_3 = (1,0,0)$ and $\vec{p}_2 = (1,1,0)$. First of all, we will argue that every Nash equilibrium of this game has social welfare 1, by showing that there are no Nash equilibria where c^1 or c^2 are played with positive probability. Assume for contradiction that s_N^* is such an equilibrium. The strategy played by player 3 in s_N^* must either be (i) c^1 or c^2 as a pure strategy, (ii) a mixture of c^1 and c^3 or between c^2 and c^3 , or (iii) a mixture where both c^1 and c^2 are played with positive probability. If player 3 plays a pure strategy, the game reduces to a two-player game for players 1 and 2. In the case of c^1 , this game has the unique equilibrium (a^1, b^1) , which in turn causes player 3 to deviate to c^2 . In the case of c^2 , the unique equilibrium is (a^2, b^2) , causing player 3 to deviate to c^1 . Now assume that player 3 mixes between c^1 and c^3 , and let α and β denote the probabilities with which players 1 and 2 play a^1 and b^1 , respectively. Since player 3's payoff from c^1 and c^3 must be the same in such an equilibrium, we must either have $\alpha = \beta = 1$, in which case player 3 will deviate to c^2 , or $0 \le \alpha \le 1/2$ and $0 \le \beta \le 1/2$, causing player 2 to deviate to b^1 . Analogously, if player 3 mixes between c^2 and c^3 , we must either have $\alpha = \beta = 0$, in which case player 3 will deviate to c^1 , or $1/2 \leq \alpha \leq 1$ and $1/2 \leq \beta \leq 1$, causing player 2 to deviate to b². Finally, if both c¹ and c² are played with positive probability, we must have $\alpha + \beta = 1$ for player 3 to get an identical payoff of $\alpha\beta \leq 1/4$ from both c¹ and c². In this case, however, player 3 can deviate to c³ for a

	с	1	с	2	
	b ¹	b ²	b ¹	b ²	
a ¹	(1,1,0,1)	(1,0,0,0)	(0, 1, 1, 1)	(0,1,0,0)	d1
\mathfrak{a}^2	(0,1,0,0)	(0, 1, 1, 1)	(1,0,0,0)	(1,1,0,1)	u
a ¹	(0,0,0,1)	(0,0,0,1)	(0,0,0,1)	(0,0,0,1)	d2
\mathfrak{a}^2	(0,0,0,1)	(0,0,0,1)	(0,0,0,1)	(0, 0, 0, 1)	u

Figure 4.12: Four-player ranking game Γ_4 used in the proof of Theorem 4.14

strictly greater payoff of $1 - 2\alpha\beta$. Thus, a strategy profile s_N^* as described above cannot exist.

Now let μ^* be the correlated strategy where action profiles (a^1, b^1, c^1) , (a^2, b^2, c^1) , (a^1, b^1, c^2) , and (a^2, b^2, c^2) are played with probability 1/4 each. This correlation can for example be achieved by tossing two coins independently. Players 1 and 2 observe the first coin toss and play a^1 and b^1 if the coin falls on heads, and a^2 and b^2 otherwise. Player 3 observes the second coin toss and plays c^1 if the coin falls on heads and c^2 otherwise. The expected payoff for player 2 under μ^* is 1, so he cannot gain by changing his action. If player 1 observes heads, he knows that player 2 will play b^1 , and that player 3 will play c^1 and c^2 with probability 1/2 each. He is thus indifferent between a^1 and a^2 . Player 3 knows that players 1 and 2 will play (a^1, b^1) and (a^2, b^2) with probability 1/2 each, so he is indifferent between c^1 and c^2 and strictly prefers both of them to c^3 . Hence, none of the players has an incentive to deviate, μ^* is a correlated equilibrium. Moreover, the social welfare under μ^* is 2, and thus $MV(\Gamma_3) = 2$.

Now consider the four-player game Γ_4 of Figure 4.12, which is a ranking game for rank payoffs $\vec{p}_1 = \vec{p}_3 = (1,0,0,0)$, $\vec{p}_2 = (1,1,0,0)$, and $\vec{p}_4 = (1,1,1,0)$, and rankings [1,2,4,3], [1,3,2,4], [3,2,4,1], [2,3,1,4], and [4,1,2,3]. It is easily verified that none of the action profiles with social welfare 2 is a Nash equilibrium. Furthermore, player 4 strictly prefers action d^2 over d^1 as soon as one of the remaining action profiles for players 1 to 3, i.e., those in the upper half of the game where the social welfare is 1, is played with positive probability. Hence, d^1 is not played with positive probability in any Nash equilibrium of Γ_4 , and every Nash equilibrium of Γ_4 has social welfare 1. In turn, consider the correlated strategy μ^* where actions profiles (a^1, b^1, c^1, d^1) , (a^2, b^2, c^1, d^1) , (a^1, b^1, c^2, d^1) , and (a^2, b^2, c^2, d^1) are played with probability 1/4 each. It is easily verified that none of the players can increase his payoff by unilaterally deviating from μ^* . Hence, μ^* is a correlated equilibrium with social welfare 3, and $MV(\Gamma_4) = 3$.

For n > 4, we can restrict our attention to games where the additional players only have a single action. We return to the game Γ_4 of Figure 4.12 and transform it into a

	с	1	c ²		
	b^1	b^2	b1	b ²	
a ¹	(1,0,0)	$(0, 1, \varepsilon)$	(0,0,1)	$(1,0,\varepsilon)$	
a ²	$(0, 1, \epsilon)$	(1,0,1)	$(1,0,\epsilon)$	$(1,0,\varepsilon)$	

Figure 4.13: Three-player ranking game Γ_5 used in the proof of Theorem 4.15

game Γ_4^n with n > 4 players by assigning to players 5, 6, ..., n a payoff of 1 in the four action profiles (a^1, b^1, c^1, d^1) , (a^2, b^2, c^1, d^1) , (a^1, b^1, c^2, d^1) , and (a^2, b^2, c^2, d^1) that constitute the correlated equilibrium with maximum social welfare, and a payoff of zero in all other action profiles. Since the additional players cannot influence the outcome of the game, this construction does not affect the equilibria of the game. To see that the resulting game is a ranking game, consider the rank payoff vectors $\vec{p}_1 = \vec{p}_3 = (1, 0, 0, ..., 0)$, $\vec{p}_2 = (1, 1, 0, ..., 0)$, $r_m^k = 1$ if $k \leq m - 1$ and 0 otherwise, for $m \geq 4$. It is easily verified that we can retain the original payoffs of players 1 to 4 and at the same time assign a payoff of 0 or 1, respectively, to players 5 to n by ranking the latter according to their index and placing either no other players or exactly one other player behind them in the overall ranking. More precisely, Γ_4^n is a ranking game by virtue of the above rank payoffs and rankings [1, 2, 4, 5, ..., n, 3], [1, 3, 2, 4, 5, ..., n], [3, 2, 4, 5, ..., n, 1], [2, 3, 1, 4, 5, ..., n], and [4, 1, 2, 3, 5, ..., n]. Furthermore, $MV(\Gamma_4^n) = n - 1$.

THEOREM 4.15. Let \mathfrak{R} be the class of ranking games with n > 2 players. Then, $EV(\mathfrak{R}) = n - 1$, even if \mathfrak{R} only contains games without weakly dominated actions.

Proof. It suffices to show that for any $n \ge 3$ there is a ranking game with enforcement value n - 1 in which no action is weakly dominated. Consider the ranking game Γ_5 of Figure 4.13, which is a ranking game by virtue of rank payoff vectors $\vec{p}_1 = (1, 1, 0)$, $\vec{p}_2 = (1, 0, 0)$, and $\vec{p}_3 = (1, \epsilon, 0)$ and rankings [1, 2, 3], [2, 3, 1], [3, 1, 2], [3, 2, 1], and [1, 3, 2]. Obviously, all of the actions of Γ_5 are undominated and $\nu_{S_N}(\Gamma_5) = 2$. It remains to be shown that the social welfare in any correlated equilibrium of Γ_5 is at most $(1 + \epsilon)$, such that $\nu_{C(\Gamma_5)}(\Gamma_5) \rightarrow 1$ and $EV(\Gamma_5) \rightarrow 2$ for $\epsilon \rightarrow 0$.

Finding a correlated equilibrium that maximizes social welfare constitutes a *linear* programming problem constrained by the inequalities of Definition 4.12 and the probability constraints $\sum_{\alpha_N \in A_N} \mu(\alpha_N) = 1$ and $\mu(\alpha_N) \ge 0$ for all $\alpha_N \in A_N$. Feasibility of this problem is a direct consequence of Nash's existence theorem. Boundedness follows from boundedness of the quantity being maximized. To derive an upper bound for social welfare in a correlated equilibrium of Γ_5 , we will transform the above linear program into its dual. Since the primal is feasible and bounded, the primal and the dual will have the same optimal value, in our case the maximum social welfare in a correlated equilibrium. The latter constitutes a minimization problem and finding a feasible solution with objective value ν shows that the optimal value cannot be greater than ν . Since

```
\begin{array}{ll} \begin{array}{l} \mbox{minimize } \nu \\ \mbox{subject to} \\ \mbox{-}x_1 + y_1 + z_1 + \nu & \geqslant 1, \\ \mbox{}x_2 - y_1 + \nu & \geqslant 1 + \epsilon, \\ \mbox{}x_1 - y_2 + \nu & \geqslant 1 + \epsilon, \\ \mbox{-}x_2 + y_2 + (\epsilon - 1)z_1 + \nu & \geqslant 2, \\ \mbox{}x_1 - z_2 + \nu & \geqslant 1, \\ \mbox{-}x_2 + \nu & \geqslant 1 + \epsilon, \\ \mbox{}\nu & \geqslant 1 + \epsilon, \\ \mbox{}\nu & \geqslant 1 + \epsilon, \\ \mbox{}(1 - \epsilon)z_2 + \nu & \geqslant 1 + \epsilon, \\ \mbox{}x_1 \geqslant 0, \ x_2 \geqslant 0, \ y_1 \geqslant 0, \ y_2 \geqslant 0, \ z_1 \geqslant 0, \ \text{and} \ z_2 \geqslant 0. \end{array}
```

Figure 4.14: Dual linear program for computing a correlated equilibrium of Γ_5 , used in the proof of Theorem 4.15

there are three players with two actions each, the primal has six constraints of the form $\sum_{a_{-i}\in A_{-i}} \mu(a_{-i}, a_i^*)(p_i(a_{-i}, a_i^*) - p_i(a_{-i}, a_i)) \ge 0$. For $j \in \{1, 2\}$, let x_j , y_j , and z_j denote the variable of the dual associated with the constraint for the jth action of player 1, 2, and 3, respectively. Furthermore, let ν denote the variable of the dual associated with constraint $\sum_{a_N\in A_N} \mu(a_N) = 1$ of the primal. The dual is given in Figure 4.14.

Now let $x_2 = y_1 = z_2 = 0$, $x_1 = y_2 = (\epsilon - 1)^2/\epsilon$, $z_1 = (1 - 2\epsilon)/\epsilon$, and $\nu = 1 + \epsilon$, and observe that for every $\epsilon > 0$, this is a feasible solution with objective value $1 + \epsilon$. However, the objective value of any feasible solution to the dual is an upper bound for that of the optimal solution, which in turn equals $\nu_{C(\Gamma_5)}(\Gamma_5)$.

The above construction can easily be generalized to games Γ_5^n with n > 4 by adding additional players that receive payoff 1 in action profile a_N if $a_1 = a^2$, $a_2 = b^2$, and $a_3 = c^1$, and payoff 0 otherwise. This can for example be achieved by means of rank payoff vectors $\vec{p}_1 = (1, 0, ..., 0)$, $\vec{p}_2 = (1, 1, 0, ..., 0)$, $\vec{p}_3 = (1, \epsilon, 0, ..., 0)$, and $\vec{p}_m^k = 1$ if $k \leq m - 1$ and 0 otherwise for $m \geq 4$. By the same arguments as in the proof of Theorem 4.14, this does not affect the maximum social welfare achievable in a correlated equilibrium. It is thus easily verified that $EV(\Gamma_5^{k_1 \times \cdots \times k_4}) \rightarrow n - 1$ for $\epsilon \rightarrow 0$.

4.7 Discussion

In this chapter we proposed a new class of normal-form games, so-called ranking games, which model settings in which players are merely interested in outperforming their opponents. Arguably this class of games is very natural and relevant for many realistic scenarios, and provides a meaningful generalization of the property of strict competitiveness attributed to constant-sum games in the two-player case. Despite the structural simplicity of ranking games, however, various solution concepts turned out to be just as hard to compute as in general normal-form games. In particular we obtained hardness results for mixed Nash equilibria and iterated weak dominance in games with more than two players and pure Nash equilibria in games with an unbounded number of players. As a consequence, the mentioned solution concepts appear to be of limited use in large instances of ranking games that do not possess additional structure. This underlines the importance of alternative, efficiently computable, solution concepts for ranking games such as maximin strategies or correlated equilibrium.

Based on these findings, we then quantified and bounded ratios comparing different solution concepts in ranking games. It turned out that playing one's maximin strategy in binary ranking games with only few actions might be a prudent choice, not only because this strategy guarantees a certain payoff even when playing against irrational opponents, but also because of the limited price of cautiousness and the inherent weakness of Nash equilibria in ranking games.

We also investigated the relationship between Nash equilibria and correlated equilibria. While correlation can never *decrease* social welfare, it is an important question which scenarios permit an *increase*. This question is particularly relevant for scenarios that would intuitively be considered to be highly competitive. To this end, we showed that in ranking games with many players and asymmetric preferences over ranks, i.e., with non-identical rank payoff vectors, overall satisfaction can be improved substantially by allowing players to correlate their actions.
Chapter 5

Anonymous Games

A major obstacle when considering normal-form games with an unbounded number of players is the exponential size of the explicit representation of the payoffs, an issue we already touched upon in Section 4.5.3 when we investigated the computational complexity of pure Nash equilibria in ranking games. Realistic situations, however, often possess additional structure, and allow an agent to make rational decisions without reasoning explicitly about all the possible ways the other agents may behave. A particular property often found in situations with many agents is that the agents are in some sense similar to each other, such that a particular agent does not, or need not, distinguish between the other agents.

In this chapter, we consider four classes of *anonymous games* and study the complexity of pure Nash equilibria and iterated dominance in these games. More formally, an anonymous game is characterized by the fact that players do not distinguish between other players in the game, i.e., their payoff only depends on the numbers of other players playing the different actions, but not on their identities. Anonymous games constitute a very natural class of multi-player games which is also highly relevant in practice (cf. Daskalakis and Papadimitriou, 2007). A *symmetric* game additionally has identical payoff functions for all players. Two more classes, called self-anonymous and self-symmetric games in this thesis, are obtained by assuming that a player does not distinguish himself from the other players.

It turns out that in all four classes pure Nash equilibria can be found efficiently if only a constant number of actions is available to each player. Moreover, identical payoff functions for all players further reduce the computational complexity of pure equilibria, an effect that is nullified as soon as there are two different payoff functions. The fact that a player cannot, or does not, distinguish himself from the other players, does not seem to offer any computational advantage. Finally, computing pure equilibria becomes intractable in all four classes of anonymous games when the number of actions grows in the number of players.

In the second part of the chapter we turn to iterated weak dominance. We show that iterated dominance solvability is NP-hard for symmetric games with a growing number of actions, and tractable for symmetric games with a constant number of actions. The only case that remains is that of anonymous games with a constant number of actions. When restricted to two actions, it can be reformulated as a natural elimination problem on matrices. The complexity of this problem remains open, but it turns out to be related to a matching problem on paths of a directed graph. The latter problem, which may be of independent interest, is intractable in general but allows us to obtain efficient algorithms for restricted versions of matrix elimination. We finally use the matching formulation to show NP-hardness of iterated dominance in anonymous games with three actions.

5.1 Related Work

Symmetries in games have been investigated since the earliest days of game theory. Von Neumann (1928) and von Neumann and Morgenstern (1947) were the first to consider symmetries of *cooperative* games, calling a game in characteristic form symmetric if the value of a coalition depends only on its size. In the context of normal-form games, a game is usually called symmetric if the payoff functions of all players are *identical* and *symmetric* in the other players' actions (von Neumann and Morgenstern, 1947, Luce and Raiffa, 1957). For two-player normal-form games, this restriction corresponds to a skew-symmetric payoff matrix (e.g., Borel, 1921, Gale et al., 1950).

Most early research on symmetries in games has concentrated on these symmetric games. One of the reasons for this may have been the strong focus of the early research in non-cooperative game theory on two-player games, where anonymity as defined in this chapter does not impose any restrictions. Gale et al. (1950) provide a (polynomialtime) reduction from arbitrary two-player games to symmetric two-player games which preserves equilibria. The recent PPAD-completeness result of Chen and Deng (2006) thus also applies to symmetric games with two players and a large number of actions. An early result by Nash (1951) shows that there always exists an equilibrium that "respects" all symmetries of a game, which in symmetric games implies the existence of a symmetric equilibrium, i.e., one where all players play the same (mixed) strategy. Papadimitriou and Roughgarden (2005) capitalize on this existence result and show that a Nash equilibrium of a symmetric game with n players and k actions can be computed in polynomial time if k = $O(\log n / \log \log n)$. While their tractability results for correlated equilibrium (Aumann, 1974) do not rely on identical payoff functions and hence apply to anonymous games as well, this is not the case for the results about Nash equilibria. The aforementioned existence of symmetric Nash equilibria neither extends to pure equilibria, nor does it hold for anonymous games. For example, Figure 5.3 on Page 64 shows an anonymous game without any symmetric equilibria.

Computational aspects of the larger class of anonymous games have recently come under increased scrutiny due to their importance in modeling large anonymous environments like the Internet. PPAD-hardness of the Nash equilibrium problem in general games has led to an increased interest in approximate equilibria, and anonymous games with a constant number of actions have recently been shown to admit a polynomial-time approximation scheme (Daskalakis and Papadimitriou, 2008).

Obviously, deciding the existence of a *pure* equilibrium is easy if the number of candidates for such an equilibrium, i.e., the number of action profiles, is polynomial in the size of the game. This is certainly the case for the explicit representation of a game as a multi-dimensional table of payoffs, but no longer holds if the game is represented succinctly. For example, deciding the existence of a pure equilibrium is NP-complete for games in graphical normal form (Gottlob et al., 2005), which we consider in Chapter 6, and games in circuit form (Schoenebeck and Vadhan, 2006). Quite a few classes of games are related to anonymity in that they exploit some form of independence among certain actions or players playing these actions. In congestion games (Rosenthal, 1973), the available actions consist of sets of resources, and the payoff depends on the number of other players that have played the same action and selected the same resources. Congestion games always have a pure equilibrium (Rosenthal, 1973), and finding one is PLS-complete even for symmetric congestion games, but in P in the symmetric network case (Fabrikant et al., 2004). For singleton (or simple) congestion games, where only a single resource can be selected, there is a polynomial-time algorithm for finding a social-welfare-maximizing equilibrium (Ieong et al., 2005). In local-effect games (Leyton-Brown and Tennenholtz, 2003), the payoff from an action may also depend on (a function of) the number of agents playing "neighboring" actions. Dunkel and Schulz (2008) give hardness results for the pure equilibrium problem in several classes of congestion and local-effect games. Unlike congestion games and local-effect games, action-graph games (Bhat and Leyton-Brown, 2004) can encode arbitrary payoff functions. For action-graph games of bounded degree, expected payoffs and the Jacobian of the payoff function can be computed in polynomial time. The latter forms the practical bottleneck step of the algorithm of Govindan and Wilson (2003) for finding Nash equilibria, but the algorithm may still take exponentially many steps to converge even for bounded degree. In fact, the pure equilibrium problem is NP-complete for symmetric action-graph games with bounded degree, but becomes tractable if the treewidth is bounded (Jiang and Leyton-Brown, 2007). In general actiongraph games, the pure equilibrium problem is NP-complete even if the action-graph is a bounded-degree tree (Daskalakis et al., 2009b).

Deciding whether a general game can be solved by iterated weak dominance is NPcomplete already for games with two players and two different payoffs (Gilboa et al., 1993, Conitzer and Sandholm, 2005a), even when restricted to dominance by pure strategies. The corresponding problem for *strict* dominance, which requires the dominating strategy to be strictly better under any circumstance, can be solved in polynomial time (e.g., Conitzer and Sandholm, 2005a). Knuth et al. (1988) provide an improved algorithm for the case of two players and dominance by pure strategies, and show that computing the reduced game in this case is P-complete. Apart from the results given in Section 4.5.2, and recent work by Brandt et al. (2009b), we are not aware of any complexity results for iterated dominance in restricted classes of games.

5.2 The Model

Symmetry as a property of a mathematical object refers to its invariance under a certain type of transformation. Symmetries of games usually mean invariance of the payoffs under automorphisms of the set of action profiles induced by some group of permutations of the set of players. Since such an automorphism preserves the number of players that play a particular action, a characteristic feature of symmetries in games is the inability to distinguish between other players. The most general class of games with this property will be called *anonymous*. Four different classes of games are obtained by considering two additional characteristics: *identical payoff functions* for all players and the ability to *distinguish oneself* from the other players. The games obtained by adding the former property will be called *symmetric*, and presence of the latter will be indicated by the prefix "*self*". For the obvious reason, we will henceforth talk about games where the set of actions is the same for all players, and write $A = A_1 = A_2 = \cdots = A_n$ and k = |A|, respectively, to denote this set and its cardinality.

Let Γ be such a game. For any permutation $\pi: N \to N$ of the set of players, let $\pi': A^N \to A^N$ be the permutation of the set of action profiles such that $\pi'((a^1, a^2, \dots, a^n)) = (a^{\pi(1)}, a^{\pi(2)}, \dots, a^{\pi(n)})$. Then, Γ is anonymous if $p_i(a_N) = p_i(\pi'(a_N))$ for all $a_N \in A^N$, $i \in N$ and all π with $\pi(i) = i$. Similarly, Γ is symmetric if $p_i(a_N) = p_j(\pi'(a_N))$ for all $a_N \in A^N$, $i, j \in N$ and all π with $\pi(j) = i$. Finally, Γ is self-anonymous if $p_i(a_N) = p_i(\pi'(a_N))$ for all $a_N \in A^N$, $i \in N$, and self-symmetric if $p_i(a_N) = p_j(\pi'(a_N))$ for all $a_N \in A^N$, $i, j \in N$. Since π' is an automorphism of the set of action profiles that preserves the number of players who play a particular action, an intuitive way to describe anonymous games is in terms of equivalence classes of the automorphism group of π' , using a notion introduced by Parikh (1966) in the context of context-free languages. Given a set A of actions, the commutative image of an action profile $a_N \in A^N$ is given by $\#(a_N) = (\#(a, a_N))_{a \in A}$ where $\#(a, a_N) = |\{i \in N : a_i = a\}|$. In other words, $\#(a, a_N)$ denotes the number of players playing action a in action profile a_N , and $\#(a_N)$ is the vector of these numbers for all the different actions. This definition naturally extends to action profiles for subsets of the players.

DEFINITION 5.1 (anonymity). Let $\Gamma = (N, (A_i)_{i \in N}, (p_i)_{i \in N})$ be a normal-form game, A a set of actions such that $A_i = A$ for all $i \in N$. Γ is called

- anonymous if $p_i(a_N) = p_i(a'_N)$ for all $i \in N$ and all $a_N, a'_N \in A^N$ with $a_i = a'_i$ and $\#(a_{-i}) = \#(a'_{-i})$,
- symmetric if $p_i(a_N) = p_j(a'_N)$ for all $i, j \in N$ and all $a_N, a'_N \in A^N$ with $a_i = a'_j$ and $\#(a_{-i}) = \#(a'_{-i})$,
- self-anonymous if $p_i(a_N) = p_i(a'_N)$ for all $i \in N$ and all $a_N, a'_N \in A^N$ with $\#(a_N) = \#(a'_N)$, and



Figure 5.1: Inclusion relationships between anonymous, symmetric, self-anonymous, and self-symmetric games

• self-symmetric if $p_i(a_N) = p_j(a'_N)$ for all $i, j \in N$ and all $a_N, a'_N \in A^N$ with $\#(a_N) = \#(a'_N)$.

When talking about anonymous games, we write $p_i(a_i, x_{-i})$ to denote the payoff of player i under any action profile a_N with $\#(a_{-i}) = x_{-i}$. For self-anonymous games, $p_i(x)$ is used to denote the payoff of player i under any profile a_N with $\#(a_N) = x$. It is easily verified that the class of self-symmetric games equals the intersection of symmetric and self-anonymous games, which in turn are both strictly contained in the class of anonymous games. An illustration of these inclusions is shown in Figure 5.1. Figure 5.2 illustrates the different payoff structures for n = 3 and k = 2.

In terms of the above characterization, a game is anonymous if the payoff $p_i(a_N)$ of player $i \in N$ in action profile a_N depends, besides his own action a_i , only on the number $\#(a, a_{-i})$ of other players playing each of the actions $a \in A$, but not on who plays them. If two players exchange actions, all other players' payoffs remain the same. For two-player games, anonymity does not impose any restrictions, because action sets of equal size can simply be achieved by adding dummy actions. A game is symmetric if it is anonymous and if the payoff function is the same for all players. Hence, if two players exchange actions, their payoffs are also exchanged, while all other players' payoffs remain the same. Many well-known games like the Prisoner's Dilemma, Rock-Paper-Scissors, or Chicken are examples of symmetric (two-player) games. Simple congestion games (leong et al., 2005) are an example for the multi-player case. In a self-anonymous game the payoff of each player depends only on the number $\#(a, a_N)$ of players playing each of the actions $a \in A$, including the player himself. If two players exchange actions, the payoffs of all players remain the same. Matching Pennies is a self-anonymous two-player game, and voting with identical weights can be seen as an example for the multi-player case. Finally, in a self-symmetric game the payoff is always the same for all players and stays the same if two players exchange actions. Self-symmetric games thus are a special case of so-called *common payoff* games, in which all players get the same payoff. Obviously such games always have a pure equilibrium, namely an action profile with maximum payoff. Other games guaranteed to possess a pure equilibrium, and the complexity of finding an equilibrium in these games, have been investigated by Fabrikant et al. (2004).

Гı·	(\cdot, \cdot, \cdot)	(a, g, c)	(a, b, \cdot)	(\cdot, e, f)
11.	(\cdot, b, c)	(d, e, \cdot)	(d, \cdot, f)	(\cdot,\cdot,\cdot)
Гa·	(a, a, a)	(b, c, b)	(b, b, c)	(e, d, d)
12.	(c, b, b)	(d, d, e)	(d, e, d)	(f, f, f)
Га	(\cdot, \cdot, \cdot)	(a, b, c)	(a, b, c)	(d, e, f)
13.	(a, b, c)	(d, e, f)	(d, e, f)	(\cdot, \cdot, \cdot)
Γı	(a, a, a)	(b, b, b)	(b, b, b)	(c, c, c)
•4•	(b, b, b)	(c, c, c)	(c, c, c)	(d, d, d)

Figure 5.2: Relationships between the payoffs of anonymous (Γ_1), symmetric (Γ_2), selfanonymous (Γ_3), and self-symmetric (Γ_4) games for n = 3 and k = 2. Players 1, 2, and 3 simultaneously choose rows, columns, and tables, respectively, and obtain payoffs according to the vector in the resulting cell. Each lower case letter stands for a payoff value, dots denote arbitrary payoff values. As an example for the separation of the different classes, Γ_1 is not symmetric if $a \neq c$ and not self-anonymous if $b \neq g$. Γ_2 is not selfanonymous if $b \neq c$. Γ_3 is not self-symmetric if $a \neq c$.

There are $\binom{n+k-1}{k-1}$ distributions of n players among k actions. Since these are exactly the equivalence classes of the set of action profiles for n-1 players under the commutative image, an anonymous game can be represented using at most $n \cdot k \cdot \binom{n+k-2}{k-1}$ numbers. In the following, we call the *explicit representation* of an anonymous game the one that simply lists the payoffs for each of the above equivalence classes, and note that the explicit representation requires space polynomial in n if and only if k is bounded by a constant. On the other hand, its size becomes super-polynomial in n even for the slightest growth of k. Nevertheless, space polynomial in n may still suffice to encode certain subclasses of anonymous games with a larger number of actions if we use an implicit representation of the payoff functions like a Boolean circuit. It is easy to see that for games with a constant number of actions, any encoding of a game that has size at least linear in the number of players and satisfies the basic assumptions of rational and efficient play made throughout the thesis, is equivalent to its explicit representation under polynomial-time reductions.

Interestingly, the ability to distinguish oneself from the other players does not increase the complexity of the pure equilibrium and iterated dominance problems when players only have two actions.

5.3 · Pure Nash Equilibria

LEMMA 5.2. There exists a constant-depth reduction from anonymous games with two actions to self-anonymous games with two actions that preserves pure Nash equilibria, dominance by pure strategies, and identical payoff functions.

Proof. Let $\Gamma = (N, (\{a^1, a^2\})_{i \in N}, (p_i)_{i \in N})$ be an anonymous game. We can define a new game $\Gamma' = (N, (\{a^1, a^2\})_{i \in N}, (p'_i)_{i \in N})$ such that for all $i \in N$ and for all $x \in \{0, 1, ..., n-1\}$,

- 1. $p'_i((x, n-x)) > p'_i((x+1, n-x-1))$ if and only if $p_i(a^1, (x, n-x-1)) > p_i(a^2, (x, n-x-1))$,
- 2. $p'_i((x, n-x)) < p'_i((x+1, n-x-1))$ if and only if $p_i(a^1, (x, n-x-1)) < p_i(a^2, (x, n-x-1))$, and
- 3. $p'_i((x, n-x)) = p'_i((x+1, n-x-1))$ if and only if $p_i(a^1, (x, n-x-1)) = p_i(a^2, (x, n-x-1))$.

Depending on the payoff structure of Γ , it may be necessary to use up to n different payoffs in Γ' , even when in Γ there are only two. It is now easily verified that Γ' is self-anonymous in general, and self-symmetric if the original game Γ is symmetric.

It should be noted that the above construction cannot in general be extended to games where players have more than two actions, because it may lead to cyclic preference relations. The symmetric two-player game Rock-Paper-Scissors is an example for a game that cannot be mapped to a corresponding self-symmetric game using the above technique.

5.3 Pure Nash Equilibria

For general games, simply checking the equilibrium condition for each action profile takes time polynomial in the size of their explicit representation. Using a succinct representation for games where the size of the explicit representation grows exponentially in the number of players, which is the case for k = 2 already, quickly renders the problem NP-hard (Fischer et al., 2006, Schoenebeck and Vadhan, 2006). On the other hand, the polynomial size even of the explicit representation for anonymous games with a constant number of actions might suggest that finding pure equilibria is easy by a similar argument as above. This reasoning is flawed, however, since a single entry in the payoff table corresponds to an exponential number of action profiles, and it is very well possible that only a single one of them is an equilibrium while all others are not. The anonymous game given in Figure 5.3 illustrates this fact.

5.3.1 Games with a Constant Number of Actions

We begin by investigating games with a constant number of actions. Obviously, solving a game cannot be easier than playing it optimally given that the opponents' actions are known. The most interesting upper bounds for the former problem will thus be obtained

(0,1,1)	(0,0,1)	(0,1,0)	(0,0,0)
(1,1,1)	(0, 0, 0)	(0,1,0)	(1,0,1)

Figure 5.3: Anonymous game with a unique, non-symmetric Nash equilibrium at the action profile with payoff (1, 1, 1). Players 1, 2, and 3 choose rows, columns, and tables, respectively. Outcomes are denoted as a vector of payoffs for the three players. Action profiles with the same commutative image as the equilibrium are shaded.

when the latter problem is easy. We therefore assume throughout this section that for any action profile of his opponents, a player can compute the payoff of a particular action in AC^0 , i.e., by evaluating a Boolean circuit with constant depth and bounded fan-in. This particularly holds if payoffs are given explicitly. It will further be obvious from the proofs that for payoff functions that are harder to compute, the complexity of the pure equilibrium problem exactly matches that of computing the payoff function.

As we have noted earlier, the potential hardness of finding pure equilibria in games with succinct representation stems from the fact that the number of action profiles that are candidates for being an equilibrium is exponential in the size of the representation of the game. While anonymous games do satisfy this property, the pure equilibrium problem nevertheless turns out to be tractable. The following theorem concerns games where the number of players is polynomial in the size of the representation.

THEOREM 5.3. Deciding whether an anonymous or self-anonymous game with a constant number of actions has a pure Nash equilibrium is TC^0 -complete under constant-depth reducibility. Hardness holds even for games with three different payoffs and two different payoff functions.

Proof. For membership in TC^0 , we propose an algorithm that decides whether there exists a pure Nash equilibrium with a given commutative image. The theorem then follows by observing that the number of different commutative images is polynomial in the number of players if the number of actions is constant.

Let $\Gamma = (N, (A_i)_{i \in N}, (p_i)_{i \in N})$ be an anonymous game, $A = \{a^1, a^2, \dots, a^k\}$ a set of actions such that $A_i = A$ for all $i \in N$. Given the commutative image $x = (x^1, x^2, \dots, x^k)$ for some action profile of Γ , call an action $a^{\ell} \in A$ a *potential best response* for player i in x if $x^{\ell} > 0$ and

$$p_{i}(a^{\ell}, x^{-\ell}) \geqslant p_{i}(a^{m}, x^{-\ell}) \qquad \qquad \text{for all } a^{m} \in A, \tag{5.1}$$

where $x^{-\ell} = (x^1, \dots, x^{\ell-1}, x^{\ell} - 1, x^{\ell+1}, \dots, x^k).$

Fix a particular commutative image $x = (x^1, x^2, \dots, x^k)$, and define a bipartite graph G = (V, E) such that

$$\begin{split} &V = V_1 \cup V_2, \quad V_1 = N, \quad V_2 = \{ (a^j, \ell) : a^j \in A, 1 \leqslant \ell \leqslant x^j \}, \quad \text{and} \\ &E = \{ (\mathfrak{i}, (a^j, \ell)) : a^j \text{ is a potential best response for } \mathfrak{i} \text{ under } x \}. \end{split}$$



Figure 5.4: Matching problem for the game of Figure 5.3 (left) and representation of the same problem by a graph with a constant number of vertices (right), as used in the proof of Theorem 5.3.

In other words, the two sides of G respectively correspond to the players and actions of Γ , with action multiplicities according to x. Edges connect each player to his potential best responses. The graph corresponding to the shaded action profiles in Figure 5.3 is shown on the left of Figure 5.4.

It is now readily appreciated that a pure equilibrium of Γ with commutative image x directly corresponds to a perfect matching of G, and vice versa. Furthermore, by Hall's Theorem, G has a perfect matching if and only if $|v(V')| \ge |V'|$ for all $V' \subseteq V_1$, where $v(V') = \{v \in V_2 : (u, v) \in E, u \in V_1\}$ is the neighborhood of vertex set V' (e.g., Bollobás, 1998).

Observe that Hall's condition cannot be verified efficiently in general. We will argue, however, that this can in fact be done for G, by considering a new graph obtained from G which possesses only a constant number of vertices. More precisely, we want to show that Hall's condition can be verified by a threshold circuit with unbounded fan-in, constant depth, and a polynomial number of gates. From the description given below it is easy to see that the constructed circuit is log-space uniform.

Assume without loss of generality that for all $\nu \in V_1$, $\nu(\nu) \neq \emptyset$, and define an equivalence relation $\sim \subseteq V \times V$ such that for all $\nu, \nu' \in V$, $\nu \sim \nu'$ if and only if $\nu(\nu) = \nu(\nu')$. By construction of G, and since both the number of actions and the number of possible subsets of actions are constant, the set V/\sim of equivalence classes has constant size, and $V/\sim = (V_1/\sim) \cup (V_2/\sim)$. Each element of V_1/\sim corresponds to the set of players having a particular set of actions as their potential best responses in x. Each element of V_2/\sim corresponds to an action in A. The neighborhood function ν can naturally be extended to equivalence classes by letting for each $U \in V_1/\sim$, $\nu(U) = \{U' \in V_2/\sim : \nu \in \nu(u) \text{ for some } u \in U, \nu \in U'\}$. This yields a bipartite graph with vertex set V/\sim , the graph corresponding to the game in Figure 5.3 is shown on the right of Figure 5.4. It is now easily verified that G has a perfect matching, and Γ a pure equilibrium, if and only if for every $Y \subseteq V_1/\sim$, $\sum_{X \in Y} |\nu(X)| \ge |Y|$.

Since V_1/\sim has only a constant number of subsets, we can construct a constant depth threshold circuit which uses sub-circuits UNARY-COUNT and UNARY-TO-BINARY as described by Chandra et al. (1984) to sum over elements of the equivalence classes, and COMPARISON sub-circuits to verify the inequalities. The former is easily realized with

5 · Anonymous Games



Figure 5.5: Game Γ used in the proof of Theorem 5.3

the help of MAJORITY gates. It thus remains to be shown that for any $X \in V_1/\sim$, |X|and |v(X)| can be computed in TC^0 . For this, recall that a particular element of V_1/\sim corresponds to the set of players that have a certain set of actions as their set of best responses in x. To compute the number of such players we first construct a circuit of constant depth that uses COMPARISON sub-circuits to check whether (5.1) is satisfied for a fixed commutative image x, a particular player $i \in N$, and a particular action $a \in A$. To check whether $C \subseteq A$ is the set of best responses for player i under x, we simply combine the outputs of the above circuits for all actions $a \in A$ into a single AND gate, negating the outputs of circuits for actions $a \notin C$. The desired number of players is then obtained by adding up the outputs of these gates for all players $i \in N$, again using UNARY-COUNT sub-circuits. On the other hand, |v(X)| corresponds to the number of players bound to play an action from a certain subset in every action profile with commutative image x, and can easily be obtained by summing over the respective elements of x.

For hardness, we reduce the problem of deciding whether exactly ℓ bits of a string of m bits are 1 to that of deciding the existence of a pure equilibrium in a self-anonymous game. Hardness of the former problem is immediate from that of MAJORITY (e.g., Chandra et al., 1984). For a particular m-bit string b, define a game Γ with m+2 players of two different types 0 and 1 and actions $A = \{0, 1\}$. The ith player of Γ is of type 0 or 1 if the ith bit of b is 0 or 1, respectively. Player m + 1 is of type 0, player m + 2 is of type 1. The payoffs p_0 and p_1 for the two types are given in Figure 5.5, the column labeled j specifies the payoff when exactly j players, including the player himself, play action 1. It is easily verified that this yields a constant-depth reduction.

We now claim that Γ possesses a pure equilibrium if and only if exactly ℓ bits of b are 1, and observe the following:

- An action profile a_N cannot be an equilibrium of Γ if #(1, a_N) ≠ ℓ + 1. In this case, the players of one of the two types get a higher payoff at both #(1, a_N) 1 and #(1, a_N) + 1, or at one of these in case #(1, a_N) = 0 or #(1, a_N) = m + 2. By construction there exists at least one player of each type, so there always is a player who can change his action to get a higher payoff.
- If there are l + 1 players of type 1, the action profile where all players of type 0

5.3 · Pure Nash Equilibria

play action 0 and all players of type 1 play action 1 is an equilibrium. None of the players of type 0 can gain by changing his action to 1, and none of them can change his action to 0, because all of them already play 0. A symmetric condition holds for players of type 1.

In turn, if the number of players of type 1 does not equal l+1, an action profile a_N with #(1, a_N) = l+1 cannot be an equilibrium. In this case, there exists (i) a player of type 0 playing action 1 in a_N, or (ii) a player of type 1 playing 0. This player can change his action to get a higher payoff.

Hence, a pure equilibrium exists if and only if there are $\ell + 1$ players of type 1, i.e., if and only if b has ℓ 1-bits.

In contrast to anonymous games, if a_N is a pure equilibrium of a symmetric game, so are all a'_N satisfying $\#(a'_N) = \#(a_N)$. This is due to the fact that the payoff functions of all players, and thus the situation of all players playing the same action $a \in A$, is identical, as would be the situation of any other player exchanging actions with someone playing a. We exploit this property to show that deciding the existence of a pure equilibrium in symmetric games with a constant number of actions is strictly easier than for anonymous and self-anonymous games.

THEOREM 5.4. The problem of deciding whether a symmetric game with a constant number of actions has a pure Nash equilibrium is in AC^0 .

Proof. Like with anonymous games, an action profile a_N is an equilibrium of a symmetric game if and only if, for all $i \in N$, a_i is a best response to $\#(a_{-i})$, i.e., if

$$p_{i}(a_{i}, \#(a_{-i})) \ge p_{i}(a'_{i}, \#(a_{-i})) \qquad \text{for all } a'_{i} \in A.$$
(5.2)

For a particular player $i \in N$ and for constant k, checking this inequality requires only a constant number of comparisons and can be done using a circuit of constant depth and polynomial size (e.g., Chandra et al., 1984). When it comes to checking (5.2) for the different players, the observation about action profiles with identical commutative images affords us a considerable computational advantage as compared to, say, anonymous or self-anonymous games. More precisely, we only have to check if (5.2) is satisfied for a player *playing a certain action*, of which there are at most k. Again, this can be done using a circuit of constant depth and polynomial size if k is a constant.

Finally, to decide whether game Γ has a pure equilibrium, we have to verify (5.2) for the different values of $\#(a_N)$ for $a_N \in A^N$. If k is constant, there are only polynomially many of these, so the complete check requires only polynomial size and constant depth. \Box

The reasoning in the proof of Theorem 5.4 also provides a nice illustration of the fact that every symmetric game with two actions possesses a pure equilibrium, as shown by Cheng et al. (2004). In the case of two actions, p_i depends only on player i's action, 0

or 1, and on the number of other players playing action 1. A pure equilibrium exists if for some m neither the players playing 0, who see m players playing 1, nor the players playing 1, who see m - 1 other players playing 1, have an incentive to deviate, i.e., $p_i(0,m) \ge p_i(1,m)$ and $p_i(1,m-1) \ge p_i(0,m-1)$. For m = 0 and m = n, one of the conditions is trivially satisfied, because there are no players playing 1 or 0, respectively. It is now straightforward to show that at least one m satisfying both conditions must exist. Alternatively, the existence of pure equilibria in symmetric games with two actions can also be obtained as an immediate consequence of Lemma 5.2: we can transform every symmetric game with two actions into a self-symmetric game with the same set of pure equilibria, and every game in the latter class is guaranteed to have at least one pure equilibrium.

As stated earlier, self-symmetric games always possess a pure equilibrium, namely an action profile with maximum payoff for every player. We proceed to show that such an action profile, with the additional property that it maximizes social welfare, can be found in AC^{0} .

THEOREM 5.5. The problem of finding a social-welfare-maximizing pure Nash equilibrium of a self-symmetric game with a constant number of actions is in AC^0 .

Proof. Since self-symmetric games belong to the class of common payoff games, any action profile with maximum payoff for one player automatically is a social-welfare-maximizing equilibrium, and Pareto dominates any other strategy profile. Finding such an equilibrium is in turn equivalent to finding the maximum of $\binom{n+k-2}{k-1}$ integers. The exact number is irrelevant as long as it is polynomial in the size of the input, which is certainly the case if k is bounded by a constant. Chandra et al. (1984) have shown that the maximum of m binary numbers of m bits each can be computed by an unbounded fan-in, constant-depth Boolean circuit of size polynomial in m. Since m is of course polynomial in the size of the input, the size of this circuit is as well.

5.3.2 Games with a Growing Number of Actions

The proofs we have seen so far in this chapter could exploit the fact that for constant k the explicit representation of an anonymous game is equivalent, under the appropriate type of reduction, to any kind of payoff function computable in a particular complexity class inside P. This need no longer be the case if k is unbounded, because then the size of the explicit representation grows exponentially in n. Such games may of course still admit a polynomial-size representation, for example if payoff functions are encoded by a Boolean circuit. We will now show that deciding the existence of a pure equilibrium in anonymous, symmetric, and self-anonymous games becomes NP-hard if the number of actions grows in n. For self-symmetric games, which always have a pure equilibrium, the associated search problem will be shown to be PLS-hard. In particular, we show NP-completeness and PLS-completeness, respectively, for games that have a polynomial

5.3 · Pure Nash Equilibria

number of players—like those covered by Theorems 5.3 and 5.4—and a number of actions that grows linearly in the number of players. It will be obvious from the proofs that hardness for the respective classes also holds for games with an exponential number of players and logarithmic growth of the number of actions. The corresponding case with a constant number of actions, on the other hand, remains open.

If the number of actions in a game is large enough, they can in principle be used to distinguish the players playing them. We will exploit this fact and prove the following theorems by reductions from satisfiability of a Boolean circuit. While as a matter of fact we establish hardness via a particular encoding of a game, they nevertheless provide interesting insights into the influence of restricted classes of payoff functions on the complexity of solving a game. After all it is far from obvious that hardness results for general games extend to anonymous and symmetric games.

Recall that circuit satisfiability (CSAT), i.e., the problem of deciding whether a Boolean circuit has a satisfying assignment, is NP-complete (e.g., Papadimitriou, 1994a). We provide a reduction from CSAT to the problem of deciding the existence of a pure equilibrium in a special class of games. For a particular circuit \mathbb{C} with inputs $M = \{1, 2, \ldots, m\}$, we define a game Γ with at least m players and actions $A = \{a^{j0}, a^{j1} : j \in M\} \cup \{b\}$. An action profile a_N of Γ where $\#(a^{j0}, a_N) + \#(a^{j1}, a_N) = 1$ for all $j \in M$, i.e., one where exactly one action of each pair a^{j0} , a^{j1} is played, directly corresponds to an assignment c of \mathbb{C} , the jth bit of this assignment being 1 if and only if a^{j1} is played. Observe that in this case the auxiliary action b has to be played by exactly n - m players. We can thus distinguish the action profiles of Γ corresponding to a satisfying assignment of \mathbb{C} from those corresponding to a non-satisfying assignment and those not corresponding to an assignment at all.

THEOREM 5.6. Deciding whether a self-anonymous game has a pure Nash equilibrium is NP-complete, even if the number of actions is linear in the number of players and there is only a constant number of different payoffs.

Proof. Membership in NP is straightforward. Since the number of players is polynomial, we can simply guess an action profile and verify that it satisfies the equilibrium condition.

For hardness, we reduce satisfiability of a Boolean circuit C with a set $M = \{1, 2, ..., m\}$ of inputs to the existence of a pure equilibrium in a game Γ with $n \ge m$ players, actions $A = \{a^{j0}, a^{j1} : j \in M\} \cup \{b\}$, and payoff functions p_i as follows:

- If action profile a_N corresponds to a *satisfying* assignment of C, we let $p_i(a_N) = 1$ for all $i \in N$.
- Otherwise we let

 $-p_1(a_N) = 1$ and $p_2(a_N) = 0$ if $\#(b, a_N)$ is even,

- $(-p_1(a_N) = 0 \text{ and } p_2(a_N) = 1 \text{ if } \#(b, a_N) \text{ is odd, and}$
- $-p_i(a_N) = 1$ for all $i \in N \setminus \{1, 2\}$.

We observe the following:

- In all of the above cases, the payoff of player i only depends on the number of players playing certain actions. If two players exchange actions, the payoff to all players remains the same. Hence, Γ is self-anonymous.
- Clearly, every action profile a_N corresponding to a satisfying assignment of C is an equilibrium, because in this case all players receive the maximum payoff of 1.
- For an action profile a_N not corresponding to a satisfying assignment of C, either player 1 or player 2 receives a payoff of 0. Furthermore, by choosing his own action to be either b or some other action, this player can determine the parity of the number of players playing b. Changing the parity strictly increases the player's payoff. This means that a_N cannot be an equilibrium.

We have hence established a direct correspondence between satisfying assignments of \mathcal{C} and pure equilibria of Γ . The transformation from \mathcal{C} to Γ essentially works by writing down Boolean circuits that compute p_i . Observing that this can be done in time polynomial in the size of \mathcal{C} if $n \leq 2^k$ completes the proof.

As the reader may have noticed, the construction used in this proof has two distinguished players play *Matching Pennies* for any assignment that does not satisfy the Boolean circuit. Not only is this game a well-known example for a game that does not possess a pure equilibrium, it is also self-anonymous on its own. On the other hand, it is easily verified that the payoffs in this game do depend on the identities of the players, i.e., that the game is not symmetric. We will have to avail of a different construction for the symmetric case.

THEOREM 5.7. Deciding whether a symmetric game has a pure Nash equilibrium is NP-complete, even if the number of actions is linear in the number of players and there is only a constant number of different payoffs.

Proof. Membership in NP is again straightforward.

For hardness, we provide a reduction from CSAT, mapping a circuit \mathcal{C} with inputs $M = \{1, 2, ..., m\}$ to a game Γ with $n \ge m$ players, actions $A = \{a^{j0}, a^{j1} : j \in M\} \cup \{b\}$, and payoff functions p_i as follows:

- If $#(b, a_N) = n m$, we let
 - $-p_i(a_N) = 2$ if a_N corresponds to a satisfying assignment or if $a_i = a^{j1}$ for some $j \in M$, $\#(a^{j0}, a_N) > 0$, and $\#(a^{j1}, a_N) > 0$,
 - $-p_i(a_N)=1 \text{ if } a_i=a^{j0} \text{ for some } j\in M, \ \#(a^{j0},a_N)>0, \text{ and } \ \#(a^{j1},a_N)=0, \\ \text{ and }$
 - $p_i(a_N) = 0$ otherwise.

- If $\#(b, a_N) < n m$, we let $p_i(a_N) = 1$ if $a_i = b$, and $p_i(a_N) = 0$ otherwise.
- If $\#(b, a_N) > n m$, we let $p_i(a_N) = 0$ if $a_i = b$, and $p_i(a_N) = 1$ otherwise.

We observe the following:

- For all of the above cases, the payoff of player i only depends on his own action and on the number of players playing certain other actions. If two players exchange actions, their payoffs are also exchanged. Hence, Γ is symmetric.
- Clearly, any action profile corresponding to a satisfying assignment of C is an equilibrium, because in this case all players receive the maximum payoff of 2.
- If on the other hand a_N does not correspond to a satisfying assignment, we have one of three different cases, in none of which a_N is an equilibrium:
 - If $\#(b, a_N) < n m$ or $\#(b, a_N) > n m + 1$, then there exists a player that receives payoff 0 and can change his action to receive payoff 1.
 - If $\#(b, a_N) = n m$ and $\#(a^{j0}, a_N) = 1$ for all $j \in M$, player i can change to any $a^{\ell 1}$ such that $a_i \neq a^{\ell 0}$ to increase his payoff from 1 to 2.
 - Otherwise, there has to be some player $i \in N$ who gets payoff 0, and, by the pigeonhole principle, some $j \in M$ such that $\#(a^{j0}, a_{-i}) = \#(a^{j1}, a_{-i}) = 0$. Then, player i can change to a^{j0} to get a higher payoff.

Again, there is a direct correspondence between pure equilibria of Γ and satisfying assignments of \mathcal{C} . The transformation from \mathcal{C} to Γ essentially works by writing down Boolean circuits that compute p_i . Observing that this can be done in time polynomial in the size of \mathcal{C} if $n \leq 2^k$ completes the proof.

By each of the previous two theorems and by the inclusion relationships between the different classes of games, we also have the following.

COROLLARY 5.8. Deciding whether an anonymous game has a pure Nash equilibrium is NP-complete, even if the number of actions is linear in the number of players and there is only a constant number of different payoffs.

Since the proofs of Theorems 5.6 and 5.7 work by mapping satisfying assignments of a Boolean circuit to *a certain number* of pure equilibria of a strategic game, we can show that counting the number of pure equilibria in the above classes of games is hard.

COROLLARY 5.9. For anonymous, symmetric, and self-anonymous games, counting the number of pure Nash equilibria is #P-hard, even if the number of actions is linear in the number of players and there is only a constant number of different payoffs. *Proof.* Recall that in the proof of Theorem 5.6, actions of the game Γ are identified with inputs of the Boolean circuit C. As a direct consequence of anonymity or symmetry, it does not matter which player plays a particular action to assign a value to the corresponding gate. Every satisfying assignment of C thus corresponds to n! equilibria of Γ , so the number of satisfying assignments can be determined by counting the number of pure equilibria, of which there are at most $2^n n!$, and dividing this number by n!. Division of two m-bit binary numbers can be done using a circuit with bounded fan-in and depth $O(\log m)$ (Beame et al., 1986). For $m = \log(2^n n!) = O(n^2)$, we have depth $O(\log n^2) = O(\log n)$. We have thus found a reduction of the problem #SAT of counting the number of satisfying assignments of C, which is #P-complete (e.g., Papadimitriou, 1994a), to the problem of counting the pure equilibria of Γ . The same line of reasoning applies to the proof of Theorem 5.7. Analogously to Corollary 5.8, #P-hardness extends to anonymous games.

As we have already outlined above, every self-symmetric game possesses a pure equilibrium. Theorem 5.5 states that finding even a social-welfare-maximizing equilibrium is very easy as long as the number of actions is bounded by a constant. If now the number of actions is growing but polynomial in the size of the input, we can start at an arbitrary action profile and check in polynomial time whether some player can change his action to increase the (common) payoff. If this is not the case, we have found an equilibrium. Otherwise, we can repeat the process for the new profile, resulting in a procedure called *best-response dynamics* in game theory. Since the payoff strictly increases in each step, we are guaranteed to find an equilibrium in polynomial time if the number of different payoffs is polynomial. Conversely we will show that, given a self-symmetric game with a growing number of actions and an exponential number of different payoffs, finding a pure equilibrium is PLS-complete, i.e., at least as hard as finding a *locally optimal* solution to an NP-hard optimization problem. The proof of the following theorem works along similar lines as those of Theorems 5.6 and 5.7 to give a reduction from the PLS-complete problem FLIP.

THEOREM 5.10. The problem of finding a pure Nash equilibrium in a self-symmetric game is PLS-complete, even if the number of actions is linear in the number of players.

Proof. Neighborhood among action profiles is given by a single player changing his action. Since the number of players and actions is polynomial in the input size, and since the payoff function is computable in polynomial time, *membership* in PLS is immediate.

For hardness, consider a Boolean circuit C with inputs $M = \{1, 2, ..., m\}$ and ℓ outputs. Finding an assignment such that the output interpreted as an ℓ -bit binary number is a local maximum under the FLIP neighborhood, where neighbors are obtained by changing a single input bit, is known to be PLS-complete (Johnson et al., 1988, Schäffer and Yannakakis, 1991). We provide a PLS reduction to the problem of finding a pure equilibrium in a self-symmetric game by mapping a particular circuit C as described above to a game Γ with $n \ge m$ players, actions $A = \{ a^{j0}, a^{j1} : j \in M \}$, and a (common) payoff function p as follows:

- If action profile a_N corresponds to an assignment c of C, let $p(a_N) = n + C(c)$, where C(c) denotes the output of C for input c, interpreted as a binary number.
- Otherwise let p(a_N) = min(#(b, a_N), n-m) + |{j ∈ M : #(a^{j0}, a_N) + #(a^{j1}, a_N) > 0}|. That is, the payoff is at most n − 1 and decreases in the minimum number of players that would have to change their action in order to make a_N correspond to an assignment of C.

We observe the following:

- Obviously, Γ is a common payoff game. Since p is invariant under any permutation of the players in both of the above cases, Γ is self-symmetric.
- If n ≤ 2^k, a Boolean circuit that computes p can be constructed from C in time polynomial in the size of C. Hence, there exists a polynomial time computable function that maps instances of FLIP to instances of the problem under consideration.
- An action profile a_N that does not correspond to an assignment of C cannot be an equilibrium of Γ. In this case, either #(a^{j0}, a_N) + #(a^{j1}, a_N) = 0 for some j ∈ M, or #(a^{j0}, a_N) + #(a^{j1}, a_N) > 1 for some j ∈ M and #(b, a_N) < n-m. Then there exists a player who can increase his payoff (the payoff of all players, actually) by changing his action, to a^{j0} or a^{j1} in the former case and to b in the latter.
- There is a direct correspondence between the FLIP neighborhood of C and a single player changing between a^{j0} and a^{j1} for some j ∈ M. Furthermore, changing to an action profile that does not correspond to an assignment of C will get the player strictly less payoff. Thus, there is a direct correspondence between pure equilibria of Γ and local maxima of C under the FLIP neighborhood. Obviously, the assignment corresponding to an action profile can be computed in polynomial time, if such an assignment exists. The conditions of Definition 2.16 do not require that we map solutions of Γ that are not locally optimal to solutions of C that are not locally optimal. This means that action profiles not corresponding to an assignment can simply be mapped to an arbitrary assignment.

It is easily verified that this satisfies the properties of a PLS reduction. \Box

With some extra work, we can show that the reduction used in the proof of Theorem 5.10 is tight, and draw additional conclusions about the standard algorithm and the standard algorithm problem.

COROLLARY 5.11. The running time of the standard algorithm for finding pure Nash equilibria in self-symmetric games is exponential in the worst case. The standard algorithm problem is NP-hard.

Proof. Johnson et al. (1988) have shown that the standard algorithm for FLIP has an exponential worst-case running time, and the standard algorithm problem is NP-hard. By Lemma 3.3 of Schäffer and Yannakakis (1991) it thus suffices to show that the reduction in the proof of Theorem 5.10 is tight. To this end, choose \mathcal{R} to be the set of action profiles of Γ that correspond to an assignment of \mathcal{C} . Obviously, \mathcal{R} contains all optimal solutions, and a payoff profile corresponding to a particular assignment can be computed in polynomial time. The third condition is trivially satisfied because the measure of any solution inside \mathcal{R} is strictly greater than that of any solution outside of \mathcal{R} .

By a slight modification of the proof of Theorem 5.10, PLS-hardness, exponential worst-case running time of the standard algorithm, and NP-hardness of the standard algorithm problem can also be shown for general, i.e., not necessarily anonymous, common payoff games with k = 2. This fact nicely illustrates the influence of anonymity on the complexity of the pure equilibrium problem.

5.3.3 Threshold Anonymity

We will now extend the tractability results of Section 5.3.1 to games where the players cannot even observe the exact number of players playing a certain action, but only whether this number exceeds certain *thresholds*. Let $\Gamma = (N, (A_i)_{i \in N}, (p_i)_{i \in N})$ be a normal-form game, and A a set of actions such that $A_i = A$ for all $i \in N$. For $T \subseteq \{1, 2, ..., n\}$, let $\sim_T \subseteq A^N \times A^N$ be defined as follows: $a_N \sim_T a'_N$ if for all $a \in A$ and all $x \in T$, $\#(a, a_N) < x$ if and only if $\#(a, a'_N) < x$. The relation \sim_T naturally extends to action profiles for subsets of N. It is then easily verified that for any $T \subseteq \{1, 2, ..., n\}$, \sim_T is an equivalence relation on the set A^M for any $M \subseteq N$. We use \sim_T to generalize Definition 5.1.

DEFINITION 5.12 (threshold anonymity). Let $\Gamma = (N, (A_i)_{i \in N}, (p_i)_{i \in N})$ be a normal-form game, A a set of actions such that $A_i = A$ for all $i \in N$. Let $T \subseteq \{1, 2, ..., n\}$. Γ is called

- T-anonymous if $p_i(a_N) = p_i(a'_N)$ for all $i \in N$ and all $a_N, a'_N \in A^N$ with $a_i = a'_i$ and $a_{-i} \sim_T a'_{-i}$,
- T-symmetric if $p_i(a_N) = p_j(a'_N)$ for all $i, j \in N$ and all $a_N, a'_N \in A^N$ with $a_i = a'_j$ and $a_{-i} \sim_T a'_{-j}$,
- T-self-anonymous if $p_i(a_N) = p_i(a'_N)$ for all $i \in N$ and all $a_N, a'_N \in A^N$ with $a_N \sim_T a'_N$, and
- T-self-symmetric if $p_i(a_N) = p_j(a'_N)$ for all $i, j \in N$ and all $a_N, a'_N \in A^N$ with $a_N \sim_T a'_N$.

For $T = \{1, 2, ..., n\}$, these classes are equivalent to those of Definition 5.1. Moreover, for $T = \{1\}$, we obtain *Boolean anonymity*, where payoffs only depend on the set of actions that are played by at least one player. In general, we call a game *threshold anonymous*, for one of the above classes, if it is T-anonymous for some T and the corresponding class.

5.3 · Pure Nash Equilibria

75

Obviously, the number of payoffs that need to be written down for each player to specify a general T-anonymous game is exactly the number of equivalence classes of \sim_T for action profiles of the other players. A T-anonymous game can be represented using at most $n \cdot k \cdot |A^{n-1}/\sim_T|$ numbers, where X/~ denotes the quotient set of set X by equivalence relation \sim . For Boolean anonymity, the number of equivalence classes equals the number of k-bit binary numbers where at least one bit is 1, i.e., $2^k - 1$. More generally, there cannot be more than $(|T|+1)^k$ equivalence classes if |T| is bounded by a constant, since for every action the number of players playing this action must be between two thresholds. For $T = \{n\}$ there are as few as k + 1. Hence, any T-anonymous game with constant |T| is representable using space polynomial in n if $k = O(\log n)$. It does not matter if the thresholds themselves are constant or not. We are now ready to identify a class of threshold anonymous games for which the pure equilibrium problem is tractable. It should be noted that the proof technique is not limited to this particular class, but in fact applies to the larger class of games for which the kernel of the best response function has polynomial size.

THEOREM 5.13. For threshold anonymous games with $k = O(\log n)$ and a constant number of thresholds, deciding the existence of a pure Nash equilibrium is in P.

Proof. Like in the proof of Theorem 5.3, we provide an algorithm that checks whether there is an equilibrium in a particular equivalence class $X \in A^N/\sim_T$. Since for $k = O(\log n)$ and |T| = O(1), the cardinality of A^N/\sim_T is polynomial in n, it suffices to show that the algorithm requires only polynomial time for every such set. For a particular element $X \in A^N/\sim_T$, the algorithm is again divided into two phases: (i) computing the set of best responses for each player under X, and (ii) checking whether there is a particular action profile $a_N \in X$ where each player plays a best response.

In the first phase, and unlike the case $T = \{1, 2, ..., n\}$ covered by Theorem 5.3, the action a played by player $i \in N$ may or may not yield a different element of $A^{N \setminus \{i\}} / \sim_T$ against which a should be a best response. Instead of just looking for best responses under elements of T^N , we thus look for best responses under those of U^N , where $U = \{u \leq n : u \in T \text{ or } (u-1) \in T\}$. Since the cardinalities of both U^N and of the set of possible best responses is polynomial if |T| = O(1) and $k = O(\log n)$, the first phase requires only polynomial time.

As for the second phase, we show that it can be reduced to deciding the existence of an integer flow with upper and lower bounds in a directed network with $O(2^k)$ vertices. Since this problem is in P if the number of vertices is polynomial (e.g., Garey and Johnson, 1979), observing that 2^k is polynomial in the size of the input if $k = O(\log n)$ completes the proof. Fix $X \in A^N/\sim_T$ and define a directed graph G = (V, E) such that

$$\begin{split} &V = \{s, t, t'\} \cup V_1 \cup V_2, \quad V_1 = 2^A, \quad V_2 = A, \quad \text{and} \\ &E = \{s\} \times V_1 \cup \{(A', a) \in V_1 \times V_2 : a \in A'\} \cup V_2 \times \{t\} \cup \{(t, t')\}. \end{split}$$



Figure 5.6: Integer flow network used in the proof of Theorem 5.13, example for the game of Figure 5.3. Edge e is labeled (u(e), l(e)).

Further define two functions $\ell : E \to \mathbb{N}$ and $u : E \to \mathbb{N}$ such that,

for all $A' \in V_1$, $\ell(s, A') = u(s, A') = |\{i \in N : A' \text{ is the set of best responses of } i \text{ under } X\}|,$ for all $A' \in V_1$ and $a \in A'$, $\ell(A', a) = 0$ and u(A', a) = n,for all $a \in V_2$, $\ell(a, t) = \min_{x \in X} \#(a, x)$ and $u(a, t) = \max_{x \in X} \#(a, x),$ and $\ell(t, t') = u(t, t') = n.$

Figure 5.6 shows the flow network for the game in Figure 5.3. Edge capacities have been computed by checking for each player if his action in the respective (shaded) action profile of Figure 5.3 is a best response. Observe that since this game is not only threshold anonymous but also anonymous, upper and lower bounds coincide for flow leaving vertices in $V_2 = A$.

Obviously every feasible flow from s to t' must have size n. Furthermore, the conditions for flow leaving vertices in V_1 require that there exists an assignment of actions to players such that each player plays a best response, while those for flow leaving edges in V_2 require that the resulting action profile is an element of X. It is thus readily appreciated that a flow from s to t' satisfying lower bounds ℓ and upper bounds u directly corresponds to a Nash equilibrium of Γ , and vice versa.

On the other hand, it is fairly straightforward to modify the games defined in the proofs of Theorems 5.6, 5.7, and 5.10 to obey Boolean thresholds if n = k. We obtain the following corollary.

COROLLARY 5.14. Deciding the existence of a pure Nash equilibrium is NP-hard for threshold anonymous, threshold symmetric, and threshold self-anonymous games, even if thresholds are Boolean, the number of actions is linear in the number of players, and there is only a constant number of different payoffs. For the same classes, counting the number of pure Nash equilibria is #P-hard. For threshold self-symmetric games, finding a pure Nash equilibrium is PLS-hard, even if thresholds are Boolean and the number of actions is linear in the number of players.

Proof. In all constructions, we assume n = m and remove the auxiliary action b. In addition to that, the self-anonymous game used in the proof of Theorem 5.6 is modified by letting players 1 and 2 play Matching Pennies on the parity of the number $|\{j \in M : \#(a^{j0}, a_N) > 0\}|$ of 0-actions that are played by at least one player. It is easily verified that the arguments used to show the correspondence between satisfying assignments of the Boolean circuit and pure equilibria of the respective game still go through. Furthermore, the payoff of a particular player in each of these games only depends on whether certain actions are played by at least one player and, potentially, on the player's own action.

5.4 Iterated Weak Dominance

We now turn to iterated weak dominance, and begin by looking at games with a growing number of actions. Intuitively, a large number of actions nullifies the computational advantage obtained from anonymity by allowing for a distinction of the players by means of the actions they play. This intuition was confirmed in Section 5.3.2, where we saw that the pure equilibrium problem becomes NP-hard or PLS-hard if the number of actions is sufficiently large compared to the number of players. We derive a similar result for iterated dominance solvability (IDS) and eliminability (IDE).

THEOREM 5.15. IDS and IDE are NP-hard for all four classes of anonymous games, even if the number of actions grows logarithmically in the number of players, if only dominance by pure strategies is considered, and if there are only two different payoffs.

Proof. We provide a reduction from CSAT to IDS and IDE for self-symmetric games. Hardness for the other types of anonymity follows by inclusion. For a particular Boolean circuit C with inputs $M = \{1, 2, ..., m\}$, we define a game Γ with $n \ge m$ players and actions $A = \{a^{j0}, a^{j1} : j \in M\} \cup \{a^0, a^1\}$. An action profile a_N of Γ where $\#(a^{j0}, a_N) + \#(a^{j1}, a_N) = 1$ for all $j \in M$, i.e., one where exactly one action of each pair a^{j0} , a^{j1} is played, directly corresponds to an assignment c of C, the jth bit c_j of this assignment being 1 if and only if a^{j1} is played. Observe that in this case the auxiliary actions a^0 and a^1 have to be played by exactly n-m players. We can thus also identify action profiles of Γ that correspond to a *satisfying* assignment of C. Now define the (common) payoff function p by letting $p(a_N) = 1$ if $\#(a^0, a_N) + \#(a^1, a_N) > n - m$, or if a_N corresponds to a satisfying assignment of C and $\#(a^1, a_N) = n - m$. Otherwise, let $p(a_N) = 0$. Since the payoff function is the same for all players, and the payoff only depends on the number of players playing each of the different actions, Γ is self-symmetric. We will further argue that for any $A' \subseteq A$ with $a^1 \in A'$, a^1 dominates every action $a \in A' \setminus \{a^0, a^1\}$ in the restriction of Γ to action set A', and a^1 dominates a^0 in such a game if and only if C has a satisfying assignment. These properties clearly imply that a^0 is eliminable for any player, and Γ is solvable via iterated dominance with action a^1 remaining for each player, if and only if C has a satisfying assignment. Since there are only two different payoffs, we can restrict our attention to dominance by pure strategies (e.g., Conitzer and Sandholm, 2005a).

To see the former property, consider an action profile a_N corresponding to a satisfying assignment of C, and a player $i \in N$ such that $a_i \in \{a^0, a^1\}$. Then, $1 = p(a_{-i}, a^1) > p(a_{-i}, a') = 0$ for any $a' \notin \{a^0, a^1\}$. On the other hand, consider an action profile a_N not corresponding to a satisfying assignment, and some player $i \in N$. Then, for any $a' \notin \{a^0, a^1\}$, $1 = p(a_{-i}, a^1) > p(a_{-i}, a') = 0$ if $\#(a^0, a_{-i}) + \#(a^1, a_{-i}) = n - m$, and $p(a_{-i}, a^1) = p(a_{-i}, a')$ otherwise.

For the latter property, first consider an action profile a_N not corresponding to a satisfying assignment of \mathcal{C} , and some player $i \in N$ such that $a_i \in \{a^0, a^1\}$. Then, $p(a_{-i}, a^1) = p(a_{-i}, a^0)$. On the other hand, consider an action profile a_N corresponding to a satisfying assignment, and some player $i \in N$ such that $\#(a^0, a_{-i}) + \#(a^1, a_{-i}) = n - m - 1$. If $\#(a^1, a_{-i}) < n - m - 1$, then $p(a_{-i}, a^0) = p(a_{-i}, a^1) = 0$. If $\#(a^1, a_{-i}) = n - m - 1$, then $1 = p(a_{-i}, a^1) > p(a_{-i}, a^0) = 0$.

The transformation from \mathcal{C} to Γ essentially works by writing down a Boolean circuit that computes p. Observing that this can be done in time polynomial in the size of \mathcal{C} if $n \leq 2^k$ completes the proof.

In the case of symmetric games, iterated dominance becomes tractable when the number of actions is bounded by a constant.

THEOREM 5.16. For symmetric games with a constant number of actions, IDS and IDE can be decided in polynomial time.

Proof. Since all players have identical payoff functions, a state of iterated dominance elimination can be represented as a vector that counts, for each set $C \subseteq A$, the number of players that have eliminated exactly the actions in C. This vector has constant dimension if the number of actions is constant. The value of each entry is bounded by n, so the number of different vectors is polynomial in n and thus in the size of the game. The elimination process can then be described as a graph that has the above vectors as vertices and a directed edge between two such vectors if the second one can be obtained from the first by adding 1 to some component, and if the action corresponding to this component can indeed be eliminated in the state described by the first vector. For dominance by mixed strategies, this neighborhood relation can be computed in polynomial time via linear programming (Conitzer and Sandholm, 2005a). This reduces the computational problems related to iterated dominance to reachability problems in a directed graph, which in turn can be decided in nondeterministic logarithmic space and thus in polynomial time. For IDS, we need to find a directed path from (the vertex corresponding to) the zero vector

to some vector with sum n(k-1). For IDE, we need to find a path where the respective action is deleted while traversing the final edge.

In light of these two results, only one interesting class remains, namely anonymous games with a constant number of actions. To gain a better understanding of the problem, we restrict ourselves further to games with two actions. It turns out that in this case iterated dominance can be reformulated as a natural elimination problem on matrices. The latter problem will be the topic of the following section.

5.4.1 A Matrix Elimination Problem

Consider an $m \times n$ matrix X with entries from the natural numbers. Call a column c of X increasing for an interval I over the rows of X if the entries in c are monotonically increasing in I, with a strict increase somewhere in this interval. Analogously, call c decreasing for I if its entries are monotonically decreasing in I, with a strict decrease somewhere in this interval. We then say that c is *active* for I if it is either increasing or decreasing for this interval. Now consider a process that starts with X and successively eliminates pairs of a row and a column. Rows will only be eliminated from the top or bottom, such that the remaining rows always form an interval over the rows of X. A column will only be eliminated if it is active for the remaining rows. Elimination of an increasing column is accompanied by elimination of the top row. Similarly, a decreasing column and the bottom row are eliminated at the same time. The process ends when no active columns remain. In this section we study two computational problems. Matrix *elimination* asks whether for a given matrix there exists a sequence of such eliminations of length $\min(m-1,n)$, i.e., one that eliminates all columns of the matrix or all rows but one, depending on the dimensions of the original matrix. Eliminability of a column asks whether a particular column can be eliminated at some point during the elimination process.

More formally, the matrix elimination process can be described by a pair of sequences of equal length, where the first sequence consists of column indices of X and the second sequence of elements of $\{0, 1\}$, corresponding to elimination of the top or bottom row, respectively. The first sequence will contain every column index at most once. The ith element of the second sequence will be 0 or 1, respectively, if the column corresponding to the ith element of the first sequence is increasing or decreasing in the interval described by the number of 0s and 1s in the second sequence up to element i - 1.

Consider for example the sequence of matrices shown in Figure 5.7, obtained by starting with the 5×4 matrix on the left and successively eliminating columns b, a, c, and d. In this particular example, the process ends when all rows and columns of the matrix have been eliminated. Of course, this does not always have to be the case. Again consider the matrix on the left of Figure 5.7, with all entries in the second row from the bottom replaced by 2. It is easy to see that in this case no column will be active after the first elimination step, and elimination cannot continue. Since column b was the only active

	a	b	с	d	a	b	с	d	a	b	с	d		a	b	с	d	
0	1	3	2	1	1		2	1			2	1						
1	0	2	2	1	0		2	1			2	1					1	
2	0	2	3	0	0		3	0			3	0					0	
3	0	2	3	0	0		3	0										
4	3	2	3	0														

Figure 5.7: A matrix and a sequence of eliminations

column in the first place, eliminating just this one column is in fact all that can be done. A related phenomenon can be observed if we instead replace the top entry in the leftmost column by 0, and take a closer look at the matrix obtained after one elimination. While we could continue eliminating at this point, it is already obvious that we will not obtain a sequence of length 4. The reason is that one of the columns not eliminated so far, namely the leftmost one, contains the same value in every row. This column cannot become active anymore, and, as a consequence, will never be eliminated.

Let us define the problem more formally. For a set A, $v \in A^n$, and $a \in A$, denote by $\#(a,v) = |\{\ell \leq n : v_\ell = a\}|$ the *commutative image* of a and v, and write $v_{\dots k} = (c_1, c_2, \dots, c_k)$ for the prefix of v of length $k \leq n$. Further denote $[n] = \{1, 2, \dots, n\}$ and $[n]_0 = \{0, 1, \dots, n\}$.

DEFINITION 5.17 (elimination sequence). Let $X \in \mathbb{N}^{m \times n}$ be a matrix. A column $k \in [n]$ of X is called *increasing* in an interval $[i, j] \subseteq [m]$ if the sequence $x_{ik}, x_{i+1,k}, \ldots, x_{jk}$ is monotonically increasing and $x_{ik} < x_{jk}$, and *decreasing* in $[i, j] \subseteq [m]$ if the sequence $x_{ik}, x_{i+1,k}, \ldots, x_{jk}$ is monotonically decreasing and $x_{ik} > x_{jk}$.

Then, an *elimination sequence* of length k for X is a pair (c, r) such that $c \in [m]^k$, $r \in \{0, 1\}^k$, and for all i, j with $1 \le i < j \le k$, $c_i \ne c_j$ and

- $r_i = 0$ and column c_i is increasing in $[\#(0, r_{...i-1}) + 1, m \#(1, r_{...i-1})]$, or
- $r_i = 1$ and column c_i is decreasing in $[\#(0, r_{...i-1}) + 1, m \#(1, r_{...i-1})]$.

A column will be called *active* in an interval if it is either increasing or decreasing in this interval. What really matters are not the actual matrix entries x_{ij} , but rather the difference between successive entries x_{ij} and $x_{i+1,j}$. A more intuitive way to look at the problem may thus be in terms of a different matrix with the number of rows reduced by one, and entries describing the relative size of x_{ij} and $x_{i+1,j}$, e.g., arrows pointing upward and downward, respectively, depending on whether $x_{ij} > x_{i+1,j}$ or $x_{ij} < x_{i+1,j}$, and empty cells if $x_{ij} = x_{i+1,j}$. According to this representation, a column can be deleted if it contains at least one arrow, and if all arrows in this column point in the same direction. The corresponding row to be deleted is the one at the base of the arrows.

				$n + \left\lfloor \frac{m - (i + j)}{2} ight floor$			n +	$\left\lceil \frac{\mathfrak{m}-(\mathfrak{i}-\mathfrak{i})}{2} \right\rceil$		
	x ₁₁		x _{1n}	0		0	1		1	
	:	۰.	•	÷	·	÷	:	·	÷	
	x _{i1}		x _{in}	0		0	1		1	
	:		•	1		1	0		0	
n	x _{i1}		x _{in}		•				•	$n \perp \left[\frac{m - (i+j)}{m} \right]$
	$x_{i+1,1}$	•••	$x_{i+1,n}$:	••	:	:	••	:	
				1		1	0		0	
	:		:	0		0	1		1	
	$x_{m-j-1,1}$		$\mathbf{x}_{m-j-1,n}$	•	•	•		•.	•	$n + \left \frac{m - (i+j)}{m - (i+j)} \right $
n	$x_{m-j,1}$	•••	x _{m-j,n}	:	•	:	:	••	:	
11	:		•	0		0	1	•••	1	
	$x_{m-j,1}$		$x_{m-j,n}$	1		1	0		0	
	÷	·.	:	÷	·	÷	:	·	÷	
	x _{m1}		x _{mn}	1		1	0		0	

Figure 5.8: Matrix Y used in the proof of Lemma 5.18

We call matrix elimination (ME) the computational problem that asks, for a given matrix $X \in \mathbb{N}^{m \times n}$, whether X has an elimination sequence of length $\min(m-1,n)$. The problem of eliminability of a column (CE) is given $k \in [n]$ and asks whether there exists an elimination sequence (c, r) such that for some i, $c_i = k$. Without restrictions on m and n, ME and CE are equivalent. We prove this statement by showing equivalence to the problem of deciding whether there exists an elimination sequence eliminating certain numbers of rows from the top and bottom of the matrix. Several other questions, like the one for an elimination sequence of a certain length, are equivalent as well.

LEMMA 5.18. CE and ME are equivalent under disjunctive truth-table reductions.

Proof. We provide reductions between both CE and ME and the problem of *matrix elimination up to an interval* (IE): given a matrix X and two numbers k_0 and k_1 , does there exist an elimination sequence (c, r) of X such that $\#(0, r) = k_0$ and $\#(1, r) = k_1$?

To reduce ME to IE, observe that X is a "yes" instance of ME if and only if X and some interval of length $\max(1, m - n)$ form a "yes" instance of IE. Analogously, to reduce CE to IE, X and $i \in [n]$ form a "yes" instance of CE if there is an interval over the rows of X in which column i is active and which together with X forms a "yes" instance of IE.

For a reduction from IE to either ME or CE, let $X \in \mathbb{N}^{m \times n}$ and consider the $(m + 2n) \times (3n + m - (i + j))$ matrix Y shown in Figure 5.8. We claim that a column with

index greater than n, and the entire matrix, can be eliminated if and only if X has an elimination sequence (c, r) satisfying #(0, r) = i and #(1, r) = j.

For the direction from left to right, assume that (c, r) is an elimination sequence for X as above and define (c', r') by

$$\begin{split} c_k' &= \begin{cases} c_k & \text{if } 1 \leqslant k \leqslant i+j, \\ n+k-(i+j) & \text{if } i+j < k \leqslant m+2n, \text{ and} \\ \end{cases} \\ r_k' &= \begin{cases} r_k & \text{if } 1 \leqslant k \leqslant i+j, \\ 1 & \text{if } i+j < k \leqslant n + \lceil \frac{m+(i+j)}{2} \rceil, \\ 0 & \text{if } n + \lceil \frac{m+(i+j)}{2} \rceil < k \leqslant m+2n. \end{cases} \end{split}$$

It is easily verified that (c', r') is an elimination sequence of length m + 2n for Y, i.e., one that eliminates Y entirely.

For the direction from left to right, consider an elimination sequence (c', r') of length m + 2n for Y. Define ℓ to be the smallest index k for which $c'_k > n$, and let $I = [\#(0, r_{\dots \ell-1}) + 1, m - \#(1, r_{\dots \ell-1})]$. Clearly, $\ell > i + m - j$. Now define a sequence c that contains the first i elements c'_k of c' for which $r'_k = 0$, and the first j elements c'_k for which $r'_k = 1$, in the same order in which they appear in c'. Define r to be sequence of corresponding elements of r'. Then, (c, r) is an elimination sequence for Y, because the set of active columns is the same for I and [i, m - j], and also for all intervals in between. Furthermore, c only contains columns with index at most n. Thus (c, r) is also an elimination sequence for X, and the number of rows eliminated from the top and bottom is exactly as required.

We finally observe that the above arguments about CE still apply to the problem of eliminability of a column in a given direction (CED), where in addition to $k \in [n]$ we are given $d \in \{0, 1\}$ and ask for an elimination sequence (c, r) such that for some i, $c_i = k$ and $r_i = d$.

When restricted to the case m > n, CE is at least as hard than ME in the sense that the latter can be reduced to the former while there is no obvious reduction in the other direction. In general, the case of ME where m > n appears easier than the one where $m \leq n$. In the former, *every* column has to appear somewhere in the elimination sequence, while in the latter the set of columns effectively needs to be partitioned into two sets of sizes m and n - m, respectively, of columns to be deleted and columns to be discarded right away.

It will not have gone unnoticed that elimination for a matrix X is closely related to iterated dominance in an anonymous game with two actions 0 and 1 where the payoff of player j when exactly i-1 players play action 1 is given by matrix entry x_{ij} . Given actions for the other players, player j can choose between two adjacent entries of column j, so one of his two actions is dominated by the other one if the column is increasing or decreasing, respectively. Eliminating one of two actions effectively removes a player from the game, whereas elimination of the top or bottom row of the matrix mirrors the fact that at the same time, the number of players who can still choose between both of their actions is reduced by one. The following result formally establishes this relationship.

LEMMA 5.19. IDS and IDE in anonymous games with two actions are equivalent under disjunctive truth-table reductions to ME and CE, respectively, restricted to instances with m = n + 1.

Proof. By Lemma 5.2, an anonymous game with two actions can be transformed into a self-anonymous game while preserving dominance by pure strategies. Since by a result of Conitzer and Sandholm (2005a) dominance by a mixed strategy implies dominance by a pure strategy when there are only two different payoffs, it suffices to prove the equivalences for self-anonymous games. We further recall that CE is equivalent under disjunctive truth-table reductions to the problem CED where a direction for the elimination is given as well. We show equivalence of IDS to ME and of IDE to CED, both under polynomial-time reductions.

Consider a self-anonymous game $\Gamma = (N, (A_i)_{i \in N}, (p_i)_{i \in N})$ such that for all $i \in N$, $A_i = \{0, 1\}$, and assume without loss of generality that for all $i \in N$ and all $a_N \in \{0, 1\}^N$, $p_i(a_N) \in \mathbb{N}$. Since in games with two actions it suffices to consider dominance by pure strategies, we can otherwise construct a game with payoffs from the natural numbers that is equivalent with respect to iterated dominance. Now write down the payoffs of Γ in an $(|N|+1) \times |N|$ matrix X such that the jth column contains the payoffs of player $j \in N$ for the different numbers of players playing action 1, i.e., $x_{ij} = p_j(a_N)$ where $\#(1, a_N) = i$. Then, the jth column of X is increasing in an interval $[k_0, k_1]$ if and only if action 1 dominates action 0 for player j given that at least $k_0 - 1$ and at most $k_1 - 1$ other players play action 1. Analogously, the jth column is decreasing in such an interval if action 0 dominates action 1 under the same conditions. If player j eliminates action 0 or 1, respectively, this decreases the number of players that can still play the respective action, corresponding to the deletion of the top or bottom row of X, respectively. Furthermore, since every player has only two actions, the corresponding column of the matrix can be ignored as soon as one of them has been deleted. Observing that the above does not impose any restrictions on the resulting matrix apart from its dimensions, equivalence of the corresponding problems follows.

A natural way of obtaining restricted versions of ME is to consider special classes of matrices, like matrices with entries in $\{0, 1\}$ or with a bounded number of maximal intervals in which a particular column is increasing or decreasing. One such restriction is to require that all columns are increasing or decreasing in [1, m]. It is not too hard to show that this makes the problem tractable irrespective of the dimensions of the matrix. We will formally state this result in the next section and prove it as a corollary of a more general result. Unfortunately, tractability of this restricted case does not tell us a lot about the complexity of ME in general. The latter obviously becomes almost trivial if the order of elimination for the columns is known, i.e., if we ask for a specific vector $c \in [n]^k$ whether there exists a vector $r \in \{0, 1\}^k$ such that (c, r) is an elimination sequence. This observation directly implies membership in NP. More interestingly, deciding whether for a given $r \in \{0, 1\}^k$ there exist $c \in [n]^k$ such that (c, r) is an elimination sequence is also tractable. The reason is the specific "life cycle" of a column. Consider a matrix X, two intervals I, $J \subseteq [m]$ over the rows of X such that $J \subseteq I$, and a column $c \in [n]$ that is active in both I and J. Then, c must also be active for any interval K such that $J \subseteq K \subseteq I$, and c must either be increasing for all three intervals, or decreasing for all three intervals. Thus, r determines for every $i \in [k]$ a set of possible values for c_i , and leaves us with a matching problem in a bipartite graph with edges in $[n] \times [k]$. A simple greedy algorithm is sufficient to solve this problem in polynomial time. Closer inspection reveals that it can in fact be decomposed into two independent matching problems on convex bipartite graphs, for which the best known upper bound is NC² (Glover, 1967). As we will see in the following section, yet another way to make the problem tractable is to provide a set of k pairs (c_j, r_j) that have to appear in corresponding places in the sequences of rows and columns, while leaving open the ordering of these pairs.

But what if nothing about c and r is known? While we can only eliminate the top or bottom row of the matrix in each step, this still amounts to an exponential number of possible sequences. The best upper bound currently known for matching in convex bipartite graphs does not allow us to construct an algorithm that determines r nondeterministically and computes a matching on the fly. We can nevertheless use the above reasoning to recast the problem in the more general framework of *matching on paths*. For this, we will identify intervals and pairs of intervals over the rows of X by vertices and edges of a directed graph G, and will then label each edge (I, J) for two intervals I and J by the identifiers of the columns of X that take I to J. An elimination sequence of length k for X then corresponds to a path of length k in G which starts at the vertex corresponding to the interval [1, m], such that there exists a matching of size k between the edges on this path and the columns of X. In particular, by fixing a particular path, we obtain the bipartite matching problem described above. A more detailed discussion of this problem is the topic of the following section. We first study the problem on its own, and return to matrix elimination toward the end of the section.

5.4.2 Matched Paths

Let us define the matching problem described above more formally. This problem generalizes the well-studied class of matching problems between two disjoint sets, or bipartite matching problems, by requiring that the elements of one of the two sets form a certain sub-structure of a combinatorial structure. This problem is particularly interesting from a computational perspective if identifying the underlying combinatorial structure can be done in polynomial time, as for paths like in our case, or for spanning trees.

DEFINITION 5.20 (matching, matched path). Let X be a set, Σ an alphabet, and $\sigma: X \to 2^{\Sigma}$ a labeling function assigning sets of labels to elements of X. Then, a *matching* of σ is a

total function $f: X \to \Sigma$ such that for all $x, y \in X$, $f(x) \in \sigma(x)$ and $f(y) \neq f(x)$ if $y \neq x$.

Let G = (V, E) be a directed graph, Σ an alphabet, and $\sigma : E \to 2^{\Sigma}$ a labeling function for edges of G. Then, a *matched path* of length k in G is a sequence e_1, e_2, \ldots, e_k such that

- for all i, $1 \leq i < k$, $e_i \in E$ and there exist $u, v, w \in V$ such that $e_i = (u, v)$ and $e_{i+1} = (v, w)$, and
- the restriction of σ to $\{e_i : 1 \leq i \leq k\}$ has a matching.

We call matched path (MP) the following computational problem: given the explicit representation of a directed graph G with corresponding labeling function σ and an integer k, does there exist a matched path of length k in G? Variants of this problem can be obtained by asking for a matching that contains a certain set of labels, or a matched path between a particular pair of vertices. These variants also have an interesting interpretation in terms of sequencing with resources and multi-dimensional constraints on the utilization of these resources: every resource can be used in certain states corresponding to vertices of a directed graph, and their use causes transitions between states. The goal then is to find a sequence that uses a specific set or a certain number of resources, or one that reaches a certain state.

In the context of this thesis, we are particularly interested in instances of MP corresponding to instances of ME. We will see later on that the graphs of such instances are layered grid graphs (e.g., Allender et al., 2006), and that the labeling function satisfies a certain convexity property. But let us first look at the general problem. Greenlaw et al. (1995) consider the related *labeled graph accessibility problem*, which, given a directed graph G with a single label attached to each edge, asks whether there exists a path such that the concatenation of the labels along the path is a member of a context free language L given as part of the input. This problem is P-complete in general and LOGCFL-complete if G is acyclic. A matching, however, corresponds to a partial permutation of the members of the alphabet, and Ellul et al. (2004) have shown that the number of nonterminal symbols of any context-free grammar in Chomsky normal form for the permutation language over Σ grows faster than any polynomial in the size of Σ . It should thus not come as a surprise if the problem becomes harder when we ask for a matching. Indeed, MP bears some resemblance to the NP-complete problem forbidden pairs of finding a path in a directed or undirected graph if certain pairs of nodes or edges may not be used together (Gabow et al., 1976). Instead of trying to reduce forbidden pairs to MP, however, we show NP-hardness of a restricted version of MP using a slightly more complicated construction. We will then be able to build on this construction in Section 5.4.3.

In the following we restrict our attention to the case where G is a layered grid graph.

DEFINITION 5.21 (layered grid graph). A directed graph G = (V, E) is an $m \times n$ grid graph if $V = [m]_0 \times [n]_0$. An edge $(u, v) \in E$ is called *south edge* if u = (i, j) and v = (i + 1, j) for some i and j, and *east edge* if u = (i, j) and v = (i, j + 1) for some i and j. A grid graph is called *layered* if it contains only south and east edges.

THEOREM 5.22. MP is NP-complete. Hardness holds even if G is a layered grid graph, if $|\sigma(e)| = 1$ for every $e \in E$, and $|\{e \in E : \lambda \in \sigma(e)\}| \leq 2$ for every $\lambda \in \Sigma$.

Proof. Membership in NP is immediate. We can simply guess a sequence of edges of the required length as well as an assignment of labels to these edges, and verify in polynomial time that we have in fact obtained a path and a matching on this path.

For hardness, we provide a reduction from the NP-complete problem balanced onein-three 3SAT (B3SAT) (Parberry, 1991) to MP with the above restrictions. A B3SAT instance over a set U of variables is given by a set $C \subseteq U^3$ of clauses of length three such that every variable occurs in exactly three clauses, i.e., $|\{(x_1, x_2, x_3) \in C : x_i = x\}| = 3$ for all $x \in U$. An instance is called satisfiable if there exists an assignment to the variables such that exactly one element of each clause is true, i.e., a set $V' \subseteq V$ such that $|\{i : x_i \in V'\}| = 1$ for all $(x_1, x_2, x_3) \in C$. It is easily verified that |U| = |C| for every instance of B3SAT, and |U| = 3|V'| for every assignment V' satisfying C. In particular, satisfiable instances must have |U| divisible by three.

Given a particular B3SAT instance C, we construct an MP instance consisting of a complete layered grid graph G = (V, E) and a labeling function $\sigma : E \to 2^{\Sigma}$ such that a path between two designated nodes s and t of G has a matching if and only if C is satisfiable. For the moment, we will put aside the restrictions that all sets of labels are singletons and every label occurs on at most two different edges. This allows us to prove hardness for a labeled graph with special structure, which will then also be used in the proof of Theorem 5.27. At the end of the current proof, we will see that the construction can easily be modified to meet the above requirements for σ .

Now let $m = |\mathbf{U}|$, and define G as a complete $6m \times 5m$ layered grid graph. Figure 5.9 illustrates the overall structure of the labeling function σ . From s to t, G is composed of gadgets for each of the variables of C, gadgets for the clauses, and a final path of 2m east edges. We write s_i and t_i , $1 \leq i \leq 2m$, for the initial and final node of the ith of these gadgets. Before we take a closer look at both types of gadgets, let us define the set Σ of labels available for labeling edges of G. For every variable x_i of C, $1 \leq i \leq m$, we have six labels λ_{ij} , $1 \leq j \leq 6$, appearing on east edges only. Labels λ_{ij}^{ν} , $j \in \{1, 2\}$, and λ_{ij}^{c} , $j \in \{1, 2, 3\}$, on the other hand, are exclusive to south edges. The labeling function σ is defined in such a way that labels on east edges appear on *every* south edge in the respective rows of the grid, and labels on south edges appear on *every* south edge in the respective columns. Furthermore, for each label, there are at most two sets of subsequent rows or columns where this label appears. Intuitively, the gadget for variable x_i lies at the intersection of columns carrying labels λ_{ij}^{ν} and rows carrying labels λ_{ij} and rows carrying labels λ_{ij} and rows carrying labels that appear in c_i .

Figures 5.10 and 5.11 illustrate the gadgets for variables and clauses of C. The labeling



Figure 5.9: Overall structure of the layered grid graph G used in the proof of Theorem 5.22

function is defined using the following subsets of Σ :

$$\begin{array}{lll} \Sigma_{i} &= \{\lambda_{i1}, \lambda_{i2}, \lambda_{i3}\} \\ \Sigma_{i}^{-} &= \{\lambda_{i4}, \lambda_{i5}, \lambda_{i6}\} \\ \Sigma^{-} &= \bigcup_{1\leqslant i\leqslant m} \Sigma_{i}^{-} \\ \Sigma_{0}^{\flat} &= \{\lambda_{11}^{\flat}\} \\ \Sigma_{i}^{\flat} &= \{\lambda_{i,1}^{\flat}, \lambda_{i,2}^{\flat}, \lambda_{i+1,1}^{\flat}\} & \text{for } 1\leqslant i\leqslant m-1 \\ \Sigma_{m}^{\flat} &= \Sigma_{0}^{c} = \{\lambda_{m1}^{\flat}, \lambda_{m2}^{\flat}, \lambda_{11}^{c}, \lambda_{12}^{c}\} \\ \Sigma_{i}^{c} &= \{\lambda_{i1}^{c}, \lambda_{i2}^{c}, \lambda_{i3}^{c}, \lambda_{i+1,1}^{c}, \lambda_{i+1,2}^{c}\} & \text{for } 1\leqslant i\leqslant m-1 \\ \Sigma_{m}^{c} &= \{\lambda_{m1}^{c}, \lambda_{m2}^{c}, \lambda_{m3}^{c}\} \end{array}$$

Labels in Σ_i and Σ_i^- correspond to a positive and negative assignment of the ith variable, respectively. Sets Σ_i^{ν} and Σ_i^c contain auxiliary labels for the ith variable gadget and the ith clause gadget. Note that while labels in sets Σ_i^- are marked as "negative," assigning them to edges in the variable gadget actually sets variable x_i to *true*, because the selection of labels from the corresponding set Σ_i will have to take place in the respective clause gadgets. Returning to Figure 5.9, the final path of east edges from s_{2m+1} to t has length 2m, and each of the edges carries all "negative" variable labels λ_{ij} for $1 \leq i \leq m$ and $j \in \{4, 5, 6\}$. It is readily appreciated that G and σ can be constructed from C in polynomial time.

Two properties of G and σ will be useful in the following. First, every path from s



Figure 5.10: Gadget for variable x_i used in the proof of Theorem 5.22

to t traverses exactly 6m east edges and 5m south edges, which equals the overall number of labels for both directions. Secondly, a matched path from s to t must traverse every edge (t_i, s_{i+1}) for $1 \leq 1 \leq 2m$. To see this, assume for contradiction that there is an edge $(v, v') \neq (t_i, s_{i+1})$ on the path such that t_i is reachable from v but not from v'. If v is to the west from t_i , i.e., (v, v') is a south edge, then the number of south edges on the path up to v' exceeds the number of labels available for these. If v is to the north from t_i , i.e., (v, v') is an east edge, then the number of labels for south edges that do not appear on any edge reachable from v' exceeds the number of south edges on the path to v'. In both cases, the number of edges differs from the number of labels available for these edges, and the path cannot have a matching.

Now assume that there exists a satisfying assignment for C. We construct a path from s to t via all s_i and t_i for $1 \le i \le 2m$, as well as a matching for this path of size $|\Sigma|$. For vertices s_i and t_i with $1 \le i \le m$, i.e., the gadget for variable x_i , we select the path labeled with elements of Σ_i if $x_i = true$, and the path labeled with elements of Σ_i^- otherwise. For nodes s_i and t_i with $m < i \le 2m$, i.e., the gadget for clause c_i , we select the (unique) path labeled with λ_{jk} for some k such that $x_j = true$. In both cases, we arbitrarily assign one of the available labels to each edge. By this, "positive" labels $\lambda \in \Sigma_i$ corresponding to variable x_i are assigned to edges in clause and variable gadgets, respectively, depending on whether or not $x_i = true$. Every "positive" label is used exactly once on the path from s to t_{2m} , and none of the "negative" labels to true, and since, by construction of G, 2m of the "negative" labels are not assigned to any edge on a labeled path from s to t_{2m} , arbitrarily assigning these labels to the edges on the path from t_{2m} to t yields a matching for the path from s to t.

Conversely assume that there is a matched path from s to t. As observed above, this path must traverse s_i and t_i for all $1 \le i \le 2m$. Furthermore, by construction of G, the "positive" labels for a particular variable x_i either all have to be assigned to



Figure 5.11: Gadget for clause $c_i = (x_j \lor x_k \lor x_\ell)$ used in the proof of Theorem 5.22. o_{ji} denotes the number of times variable x_i occurs in clauses up to and including c_i .

edges in the gadget for x_i , or to edges in the gadgets for the clauses where x_i appears, but not both. It is then easily verified that setting a variable to *true* if and only if the corresponding "positive" labels are assigned to edges in clause gadgets yields a satisfying assignment. Thus, there is some path from s to t in G that has a matching if and only if C is satisfiable.

It remains to be shown that the above construction can be simplified such that every edge can be labeled with exactly one label and every label appears on at most two different edges. For this, we first remove all edges that cannot be part of a path from s that has a matching, i.e., those that are not part of any gadget. Then, for every set of labels defined above, the number of edges labeled with this set within a particular gadget equals the cardinality of the set, and we can assign a different singleton to each of these edges. The path from s_{2m+1} to t requires some additional attention. We know that, at the time we have found a path from s to s_{2m+1} that does not use any of the labels more than once, exactly 2m labels in Σ^- have not yet been assigned to an edge, but we do not know which. To ensure that the remaining labels can be chosen in an arbitrary order, we replace the path starting at s_{2m+1} by 2m/3 additional gadgets of the form shown in Figure 5.12, which use $2m^2/3$ additional labels λ_{ji}^e for $1 \leq j \leq m$ and $1 \leq i \leq 2m/3$. It is easily verified that the modified labeling function satisfies the desired constraints.

Let us now return to matrix elimination. In light of Theorem 5.22, an efficient algorithm for ME would have to exploit additional structure of MP instances induced by instances of ME. It turns out that this structure is indeed quite restricted in that edges carrying a particular label satisfy a "directed" convexity condition: if a particular label λ appears on two edges e = (u, v) and e' = (u', v'), then λ must appear on *all* south edges

5 · Anonymous Games



Figure 5.12: Gadget to consume remaining labels, used in the proof of Theorem 5.22

or on all east edges that lie on a path from u to v', but not both. In particular, if there is such a path, it cannot be that one of e and e' is a south edge and the other is an east edge. This fact is illustrated in Figure 5.13, which shows the labeled graph for the ME instance of Figure 5.7, as well as a matched path corresponding to an elimination sequence of maximum length.

Let us formally define the above property, along with a second property which requires the set of edges carrying a particular label to form a weakly connected subgraph of G. We henceforth concentrate on *complete* layered grid graphs, i.e., ones that contain all south and all east edges.

DEFINITION 5.23 (directed convexity, connectedness). Let G = (V, E) be a complete layered grid graph. A labeling function $\sigma : E \to 2^{\Sigma}$ for G is called directed convex if for every label $\lambda \in \Sigma$ and for every set of three edges e_1, e_2, e_3 with $e_i = (u_i, v_i)$, such that u_2 is reachable from u_1 , u_3 is reachable from u_2 , and $\lambda \in \sigma(e_1) \cap \sigma(e_3)$, it holds that e_1 and e_3 have the same direction and $\lambda \in \sigma(e_2)$ if and only if e_2 has the same direction as well. A labeling function σ is called *connected* if for every $\lambda \in \Sigma$ and every pair of edges $e_1, e_2 \in E$ such that $\lambda \in \sigma(e_1) \cap \sigma(e_2)$ there exists $(u, v) \in E$ such that $\lambda \in \sigma(u, v)$ and both e_1 and e_2 are reachable from u.

It is not too hard to see that instances corresponding to ME have a directed convex labeling function. Connectedness is related to a restricted version of ME which we term *matrix elimination with given directions* (MED): given a matrix X, a labeling function σ , and a total function $d : [n] \rightarrow \{0, 1\}$, does there exist an elimination sequence (c, r) with directions given by d, i.e., one such that for all i, j satisfying d(i) = j there is some $\ell \in \mathbb{N}$ for which $c_i = i$ and $r_i = j$. We will see shortly that this problem has a natural interpretation in terms of iterated dominance.



Figure 5.13: Labeled graph for the matrix elimination instance of Figure 5.7. A matched path and its matching are shown in bold.

LEMMA 5.24. ME is polynomial time reducible to MP restricted to layered grid graphs and directed convex labeling functions. MED is polynomial time equivalent to MP restricted to layered grid graphs and directed convex and connected labeling functions.

Proof. Consider the following reduction from ME to MP. For a matrix $X \in \mathbb{N}^{m \times n}$, define a layered grid graph G = (V, E) with $V = [m]_0 \times [n]_0$ and a labeling function $\sigma : E \mapsto 2^{[n]}$ such that for all $\lambda \in [n]$, $\lambda \in \sigma(e)$ if for some $i, j \in \mathbb{N}$, e = ((i, j), (i + 1, j)) and column ℓ of X is increasing in [i + 1, m - j], or e = ((i, j), (i, j + 1)) and column ℓ of X is decreasing in [i + 1, m - j]. Now consider $k \in \mathbb{N}$, $c \in [n]^k$, and $r \in \{0, 1\}^k$. Let $p = e_1, e_2, \ldots, e_k$ be a path in G such that $e_1 = ((0, 0), v)$ for some $v \in V$, and e_i is a south edge if and only if $r_i = 0$. Further define a function $f : E \to [n]$ by letting $f(e_i) = c_i$ for all $i \leq k$. It is not too hard to see that (c, r) is an elimination sequence of X if and only if f is a matching for the restriction of σ to the edges on p.

For directed convexity of σ , consider $e_1, e_2, e_3 \in E$ with $e_i = (u_i, v_i)$, such that u_2 is reachable from u_1 and u_3 is reachable from u_2 . For $\ell = 1, 2, 3$, define an interval $I_{\ell} = [i + 1, m - j]$ for $i, j \in \mathbb{N}$ such that $e_{\ell} = ((i, j), v)$ for some $v \in V$. Further consider $\lambda \in \sigma(e_1) \cap \sigma(e_3)$. By definition of σ , column λ of X must be active in both I_1 and I_3 . Since $I_3 \subseteq I_1$, λ must either be increasing in both of them, or decreasing in both of them. Since $I_3 \subseteq I_2$ and $I_2 \subseteq I_1$, the same must also be true for I_2 .

For MED, consider a total function $d: [n] \to \{0, 1\}$ and define $\sigma': E \mapsto 2^{[n]}$ such that for all $e \in E$ and $\lambda \in [n]$, $\lambda \in \sigma'(e)$ if $\lambda \in \sigma(e)$ and if either e is a south edge and $f(\lambda) = 0$ or e is an east edge and $f(\lambda) = 1$. It is not hard to see that σ' is directed convex and connected. On the other hand consider a layered grid graph G = (V, E) and a directed convex and connected labeling function σ . Then, for every $\lambda \in \Sigma$, there exists a unique pair of vertices $u, v \in V$ such that $\lambda \in \sigma(e)$ for exactly those south edges or exactly those east edges e that are reachable from \mathfrak{u} but not from ν . It is now possible to define a matrix X with a column for λ that is active exactly in every interval I such that $I \subseteq [\mathfrak{i}, \mathfrak{j}]$ and $I \cap [\mathfrak{i}', \mathfrak{j}'] \neq \emptyset$, and increasing if $f(\lambda) = 0$ and decreasing if $f(\lambda) = 1$. By the same reasoning as above, elimination sequences of X correspond to matched paths of G and σ with initial vertex (0, 0).

Label a in the instance of Figure 5.13 serves as an example that the labeling function of an instance of MP corresponding to one of ME does not have to be connected, and it even appears on both east edges and south edges. On the other hand, MP can be solved in polynomial time when restricted to instances that do satisfy connectedness in addition to directed convexity. This also means that we can decide in polynomial time whether there exists an elimination sequence with a specific direction of elimination for every column of a matrix.

THEOREM 5.25. Let G be a layered grid graph, σ a directed convex and connected labeling function for G. Then MP for G, σ and $k = |\Sigma|$ is in P.

Proof. It suffices to show how to decide whether there exists a matched path from s = (0,0) to a particular vertex $t = (k_s, k_e)$ such that $k_s + k_e = k$. Different values for t can then be checked sequentially.

Given a path p from some vertex $v_1 \in V$ to t, we define two labeling functions $\sigma_s^p : [k_s] \to \Sigma$ and $\sigma_e^p : [k_e] \to \Sigma$, one for south edges and one for east edges of paths from s to t. We will argue that a pair of matchings for σ_s^p and σ_e^p can easily be combined into a matching for p, while nonexistence of a matching for either of the two implies that a large set of paths in G cannot be *matched* paths. The latter will ultimately provide us with a succinct certificate that a particular pair of a graph G and a labeling function σ does *not* have a matched path of length k.

More formally, consider a complete layered grid graph G = (V, E) and a labeling function $\sigma : E \to 2^{\Sigma}$. For a path $p = e_{\ell}, e_{\ell+1}, \ldots, e_k$, define σ_s^p and σ_e^p such that for every $\lambda \in \Sigma$,

$$\begin{split} \lambda \in \sigma^p_s(i) & \text{ if there exists a path } e'_1, e'_2, \dots, e'_k \text{ with } e'_i = e_i \text{ for all } i \geqslant \ell, \\ & \text{ and } j \in [k], \, i' \in [k_e] \text{ such that} \\ & e_j = ((i-1,i'),(i,i')) \text{ and } \lambda \in \sigma(e_j), \text{ and} \\ \lambda \in \sigma^p_e(i) & \text{ if there exists a path } e'_1, e'_2, \dots, e'_k \text{ with } e'_i = e_i \text{ for all } i \geqslant \ell, \\ & \text{ and } j \in [k], \, i' \in [k_s] \text{ such that} \\ & e_j = ((i',i-1),(i',i)) \text{ and } \lambda \in \sigma(e_j). \end{split}$$

In other words, σ_s^p and σ_e^p provide an "optimistic" version of the matching problems obtained by restricting σ to a path in G that contains p as a sub-path, by allowing a certain label to be matched to the ith south edge or east edge of these paths, respectively, if it appears on the ith edge in the respective direction of *some* such path. It follows from
directed convexity and connectedness of σ that for every path p, σ_s^p and σ_e^p are convex functions and $\{a \in \sigma_s^p(i) : i \in [k_s]\} \cap \{a \in \sigma_e^p(i) : i \in [k_e]\} = \emptyset$. We can further assume without loss of generality that for every path p, σ_s^p and σ_e^p have images of size k_s and k_e , respectively.

Now let p be a particular path from s to t. By definition, there is a one-to-one correspondence between σ_s^p and the restriction of σ to south edges of p, and also between σ_e^p and the restriction of σ to east edges of p. Any pair of matchings for σ_s^p and σ_e^p thus directly corresponds to a matching for the restriction of σ to p, and existence of the former implies that p is a matched path.

On the other hand, consider a path $p = e_{\ell}, e_{\ell+1}, \ldots, e_k$ and an edge $e \in E$ such that e = (u, v) and $e_{\ell} = (v, w)$ for some $u, v, w \in V$. Denote $p' = e, e_{\ell}, e_{\ell+1}, \ldots, e_n$, and assume that both σ_s^p and σ_e^p have a matching while $\sigma_s^{p'}$ or $\sigma_e^{p'}$ does not. First consider the case where e is an east edge, and where the function that does not have a matching is $\sigma_e^{p'}$. Let $i, j \in \mathbb{N}$ such that u = (i, j). By definition, σ_e^p and $\sigma_e^{p'}$ only differ with respect to labels λ such that $\lambda \in \sigma_e^p(j')$ if and only if j' < j. Since $\sigma_e^{p'}$ is a convex function that does not have a matching, and since the image of $\sigma_e^{p'}$ has size k_e , there has to be some interval in $[k_e]$ the size of which is strictly larger, and some interval the size of which is strictly smaller than the number of labels $\sigma_e^{p'}$ assigns exclusively to elements of this interval. Furthermore, every matching f of σ_e^p must satisfy $f(j) \notin \sigma(e)$, since a matching with $f(j) \in \sigma(e)$ would also be a matching for $\sigma_e^{p'}$. This means that there actually must exist an interval I of the second type such that $I \subseteq [1, j - 1]$. Now consider any path p''from a vertex u' south of u to t, i.e., a vertex u' = (i', j) such that i' > i. Clearly, the number of labels appearing exclusively in I cannot be smaller for $\sigma_e^{p''}$ than it is for $\sigma_e^{p'}$. This means that $\sigma_e^{p''}$ does not have a matching, and thus that no matched path of G and σ can traverse u'.

Now assume that the function that does not have a matching is $\sigma_s^{p'}$. This again means that there has to be an interval such that the number of different labels assigned by $\sigma_s^{p'}$ to elements of this interval is strictly smaller than the length of this interval. Since σ_s^p has a matching, since the restrictions of σ_s^p and $\sigma_s^{p'}$ to $[i, k_s]$ are identical, and by convexity of $\sigma_s^{p'}$, there has to exist an interval I with this property such that $I \subseteq [0, i - 1]$. Now consider any path p'' from a vertex u' west of u to t, i.e., a vertex u' = (i, j') such that j' < j. By definition, for any $\lambda \in \Sigma$, $\lambda \in \sigma_e^{p''}$ only if $\lambda \in \sigma_e^{p'}$, such that the number of different labels assigned by $\sigma_s^{p''}$ to elements of I is strictly smaller than |I|. Thus $\sigma_s^{p''}$ does not have a matching, and thus no matched path of G and σ can traverse u'.

If e is a south edge, then by symmetrical arguments either no matched path of G and σ can traverse any vertex north of u, or no such path can traverse any vertex east of u.

Now consider an algorithm which starts at t and tries to iteratively construct a path from s to t by traversing edges of G backwards. Given a path p_{ℓ} of length ℓ from a vertex v_{ℓ} to t, the algorithm selects $p_{\ell+1}$ to be a path of length $\ell + 1$ containing p_{ℓ} as a sub-path such that both $\sigma^{s}p_{\ell+1}$ and $\sigma^{s}p_{\ell+1}$ have a matching. If the algorithm runs for k steps, we obtain a path p_{k} from s to t with this property, i.e., a matched path. Assume on the other hand that for some ℓ no path satisfying the above requirements exists, and denote by P the set of paths obtainable by adding a predecessor of ν_{ℓ} to p_{ℓ} . This set contains one or two paths depending on whether ν_{ℓ} has one or two predecessors. Then, for every $p \in P$, one of σ_s^p and σ_e^p does not have a matching, and from the above reasoning we obtain a set of vertices such that no matched path can traverse any of these vertices. It is easily verified that the union of these sets for the different elements of P always forms a cut that separates s from t, implying that a matched path from s to t cannot exist. \Box

An interesting consequence of Lemma 5.24 and Theorem 5.25 is that we can decide in polynomial time whether an anonymous game with two actions can be solved such that particular actions remain for the different players. Matrix elimination when all columns are active from the beginning is a special case of Theorem 5.25. With some additional work, we can derive a better upper bound.

COROLLARY 5.26. Let $X \in \mathbb{N}^{m \times n}$ be a matrix every column of which is active in [1, m]. Then ME for X is in L.

Proof. Consider the graph G and the labeling function σ corresponding to X. For any path p in G with final vertex t, consider the functions σ_s^p and σ_e^p defined in the proof of Theorem 5.25. Since every column of X is active in [1, m], σ_s^p and σ_e^p are convex. Moreover, for every label $\lambda \in \Sigma$, $\lambda \in \sigma_s^p(1)$ or $\lambda \in \sigma_e^p(1)$. It is not too hard to see that for a path $p = \nu_k, \nu_{k+1}, \ldots, \nu_n$ with $\nu_n = t$, σ_s^p and σ_e^p have a matching if and only if there exists a path p' from $\nu_{k+1} = (i, j)$ to t such that $\sigma_s^{p'}$ and $\sigma_e^{p'}$ have matchings f_s and f_e , and if the number of labels both in $\sigma_s^p(i-1) \setminus (\bigcup_{k \ge i} f_s(k))$ and in $\sigma_s^p(i-1) \setminus (\bigcup_{k \ge j} f_e(k))$ are strictly positive. We can thus construct a path by moving backwards from t, and storing a pointer to the current source of the path, and the numbers of labels that are currently active but have not been assigned to edges between the current source and t. This can clearly be done using only logarithmic space.

The complexity of ME remains open, and additional insights will be necessary to solve this question. The proof of Theorem 5.25 hinges on connectedness of the labeling function. On the other hand, directed convexity of the labeling function corresponding to an ME instance means that we cannot use a construction similar to the one used in the proof of Theorem 5.22 to show NP-hardness of MP. Another interesting question is whether the case m > n is easier than the general case.

5.4.3 Self-Anonymous Games With a Constant Number of Actions

It is natural to ask what happens for games with more than two actions, and whether there still exists a nice interpretation in terms of row and column eliminations in a matrix or matrix-like structure. It turns out there is such an interpretation, but its formulation is rather complicated. Consider a self-anonymous game with n players and k actions for each player. As before, the payoff of a particular player i only depends on the number of players, including himself, that play each of the different actions. For a particular player we thus have payoff values for each tuple (j^1, j^2, \ldots, j^k) with $\sum_{\ell=1}^k j^\ell = n$. These can be represented as entries in a discrete simplex of dimension k-1. When writing down the payoffs of all players, one obtains a structure $X = (x_i^{j^1 \dots j^k})_{i \in N, \sum_{\ell=1}^k j^\ell = n}$ where $x_i^{j^1 j^2 \dots j^k} \in \mathbb{R}$ denotes the payoff of player $i \in N$ if for each ℓ , j^ℓ players play action a^ℓ . This structure has the aforementioned simplices as columns, and resembles a triangular prism in the case k = 3.

Restricting our attention to dominance by pure strategies, action $a^{\ell} \in A$ weakly dominates action $a^m \in A$ for player $i \in N$ if i can never decrease his payoff by playing a^{ℓ} instead of a^m , no matter which actions the other players play, and if the payoff strictly increases for at least one combination of actions played by the other players. This corresponds to the values in the ith column of X being *increasing* from a^m to a^{ℓ} , i.e., weakly increasing with a strict increase at some position (j^1, j^2, \ldots, j^k) . If m players have eliminated action a^{ℓ} , tuples with $j^{\ell} > n - m$ are no longer reachable, corresponding to a cut along the ℓ th 0-face of the simplex. Eliminations of a particular action have the same effect on the payoff simplex of every single player and thus correspond to cuts along the respective edge of the prism in the case k = 3. Given a vector $d = (d^i)_{1 \leq i \leq k}$ with $1 \leq d^i \leq n$, we will write X(d) to denote the structure obtained from X by performing, for each i, d^i eliminations in dimension i, i.e., $X(d) = (x_i^{j^1 j^2 \dots j^k})_{i^i \leq n - d^i}$.

Now, let $\Gamma = (N, (A_i)_{i \in N}, (p_i)_{i \in N})$ be a self-anonymous game, and let X be defined by $x_i^{j^1 j^2 \dots j^k} = p_i(j^1, j^2, \dots, j^k)$ for all $i \in N$ and $j^1, j^2, \dots, j^k \in \mathbb{N}_0$ such that $\sum_{\ell=1}^k j^\ell = n$. Then, Γ is solvable using iterated dominance by pure strategies if there exists a pair (c, r)of sequences $c \in N^{(k-1)n}$ and $r \in A^{(k-1)n}$ such that

- (i) $|\{1 \le i \le (k-1)n : c_i = j\}| = k-1 \text{ for all } j \in N$,
- (*ii*) $c_i = c_j$ and $r_i = r_j$ implies i = j for all $1 \le i, j \le (k-1)n$, and
- (*iii*) for each i, $1 \le i \le (k-1)n$, there exists some $r^* \in A$ such that, for all $j < i, c_j \neq c_i$ or $r_j \neq r^*$, and c_i is increasing from r_i to r^* in $X(r_1, r_2, \ldots, r_{i-1})$.

That is, a game is solvable if there exists a sequence of (k-1)n pairs of a player and an action such that (i) every player deletes exactly k-1 times, (ii) no player deletes the same action twice, and (iii) every action is deleted using some other action that has not itself been deleted.

The left hand side of Figure 5.14 shows the payoffs of a particular player in a selfanonymous game with n = 3 and k = 3. Compared to matrix elimination as introduced in Definition 5.17 and illustrated in Figure 5.7, we notice an interesting shift. Curiously, this shift has nothing to do with the added possibility of dominance by mixed strategies in games with more than two actions. Rather, a particular action $a \in A$ may now be eliminated by either one of several other actions in $A \setminus \{a\}$, and the situations where a a can be eliminated no longer form a convex set. Recalling the proof of Theorem 5.22, our



Figure 5.14: Payoffs of a particular player in a self-anonymous game with three players and three actions. Initially all actions are pairwise undominated. If one of the other players eliminates action 1, action 3 weakly dominates action 1. Action 1 then becomes undominated if some player deletes action 3, and dominated by action 2 if one more player deletes action 3, and some player deletes action 2.

strategy becomes clear: try to construct a layered grid graph and a labeling of its edges for which the existence of a matched path is NP-hard to decide, and which are induced by a self-anonymous game with three actions for each player. It turns out that this is indeed possible.

THEOREM 5.27. IDS and IDE are NP-complete. Hardness holds even for self-anonymous games with three actions and two different payoffs.

Proof. Membership in NP is immediate. We can simply guess a sequence of eliminations and verify that all of them are valid and that they eventually leave only a single action for each player.

For hardness of IDS, recall the construction used in the proof of Theorem 5.22. Given a B3SAT instance C, we constructed an MP instance consisting of a layered grid graph G = (V, E) and a labeling function $\sigma : E \to 2^{\Sigma}$ such that a path between two designated nodes s and t has a matching if and only if C is satisfiable. We will now show that G and σ correspond to iterated dominance in a specific self-anonymous game Γ with k = 3 when only actions 1 and 2 of each player are considered. Observing that, given a matched path from s to t, all players in Γ can also eliminate action 3 at some vertex on the path without affecting the restriction of the labeling function to the remainder of the path effectively reduces B3SAT to IDS.

Given a particular grid graph G, a set Σ of labels, and a labeling function σ as defined in the proof of Theorem 5.22, we construct a game Γ with players $N = \Sigma$ and actions $A = \{1, 2, 3\}$. Action 1 is associated with east edges of G, action 2 is associated with south edges. Now consider a particular label $i \in \Sigma$. By construction of G, there exist two numbers k^1 and k^2 such that i appears exclusively on east edges (south edges, respectively) that can be reached from s by traversing exactly k^1 or k^2 south edges (east edges). In game Γ , this is modeled by a player that can eliminate action 1 (action 2) after exactly k^1 or k^2 players have eliminated action 2 (action 1). Since we only use payoffs 0 and 1, it follows from Lemma 1 of Conitzer and Sandholm (2005a) that an action is dominated by a



Figure 5.15: Payoff structure of a particular player of the self-anonymous game Γ used in the proof of Theorem 5.27. There are two types of players, eliminating action 2 and 1, respectively, actions of the second type are shown in parentheses. The player may eliminate action 2 (action 1, respectively) by action 3 after exactly k¹ players have eliminated action 1 (2), and by action 1 (2) after exactly k² players have eliminated action 1 (2).

mixed strategy if and only if it is dominated by a pure strategy. We can thus concentrate exclusively on dominance by pure strategies.

The payoff structure for players of Γ is shown in Figure 5.15. Clearly, Γ can be constructed from G in polynomial time. In addition to the aforementioned properties regarding the elimination of action 1 or 2, it is easily verified that every player can also eliminate action 3 after k^2 eliminations of action 2 or 1, and that this has no effect whatsoever on the ability of other players to eliminate their actions. In other words, Γ actually induces a three-dimensional grid graph, where each layer in the third dimension is identical to G, and transitions between different layers may take place at vertices where some player has arrived at k^2 . This means, however, that a matched path from s to t corresponds to a sequence of eliminations of actions 1 and 2 in Γ , which can in turn be transformed into a sequence of eliminations that solves Γ by iterated dominance by letting each player eliminate action 3 at a certain well-defined point. On the other hand, the possible future transitions within a particular layer of the three-dimension grid graph do not depend on the layer, i.e., a player may not gain the ability to eliminate actions 1 or 2 by first eliminating action 3. Hence, if there is no matched path from s to t, then some player of Γ will not be able to eliminate either action 1 or action 2, meaning that Γ is not solvable by iterated dominance.

Hardness of IDE can finally be obtained by adding an additional player that can only eliminate once the lowest level of the grid graph has been reached. \Box

5.5 Discussion

In this chapter, we introduced four classes of anonymous games and investigated the computational complexity of pure Nash equilibrium and iterated weak dominance in these classes. We established that the former solution concept is tractable for games with a constant number of actions, but becomes intractable if the number of actions grows at least linearly in the number of players. It is worth noting that, for games with a constant number of actions, the pure equilibrium problem happens to lie in the complexity class NC¹ for all types of anonymity and is thus open to parallel computation. NP-hardness also holds for games with an exponential number of players and logarithmic growth of the number of actions. For games with an exponential number of players in which the number of actions grows sub-logarithmically, the complexity remains open. Iterated dominance, on the other hand, is tractable in symmetric games with any constant number of actions, but NP-hard in anonymous and self-anonymous games with two actions remains open.

In future work, it would be interesting to extend the tractability results to larger classes of games. For example, games with a certain number of player types, where indistinguishability holds only for players of the same type, can be obtained by restricting Definition 5.1 to permutations that map players from a certain subset to players of the same set. Given a game in this class, we can construct an anonymous game with the same set of players and an action set that is the Cartesian product of the original set of actions and the set of player types. By assigning a unique minimum payoff to all actions not corresponding to the type of the respective player, we can ensure that players only play actions corresponding to their type in every equilibrium of the new game, effectively allowing us to distinguish players of different types in the new game. For games with a constant number of players the size of the new game is polynomial in the size of the original game, and the tractability result of Theorem 5.3 carries over immediately. A different notion, such that players of the same type have identical payoff functions, does not seem to provide additional structure. As we have already shown, only two different payoff functions suffice to make the pure equilibrium problem TC^{0} -hard for a constant number of actions and NP-hard for a growing number of actions. More generally, one might investigate games where payoffs are invariant under particular sets of permutations. For example, von Neumann and Morgenstern (1947) regard the number of permutations under which the payoffs of a game are invariant as a measure for the degree of anonymity. The question is in how far the computational complexity of solving a game depends on this degree.

With respect to iterated dominance, the most important open question concerns iterated weak dominance solvability in anonymous games with two actions, and the equivalent problem of matrix elimination. More generally, we looked at a problem concerning matchings on paths in a directed graph. This problem was mainly introduced as a proxy to matrix elimination, but appears to be interesting in their own right, with connections

$5.5 \cdot Discussion$

to ordinary matching problems, sequencing, and planning. It will therefore be worthwhile to investigate versions of this problem with restrictions on the graph structure or labeling function.

Chapter 6

Graphical Games

Another structural element commonly found in real-world interaction, besides the one considered in the previous chapter, is locality. Often a situation involves many agents, but the weal and woe of any particular agent depends only on the decisions made by a select few. *Graphical games* (Kearns et al., 2001) formalize this notion by assigning to each player a subset of the other players, his neighborhood, and defining his payoff as a function of the actions of these players. More formally, a graphical game is given by a (directed or undirected) graph on the set of players of a normal-form game, such that the payoff of each player depends only on the actions of his neighbors in this graph. Any graphical game with neighborhood sizes bounded by a constant can be represented using space polynomial in the number of players.

Gottlob et al. (2005) investigate the complexity of pure Nash equilibria in graphical games, and show that deciding the existence of a pure equilibrium is NP-complete already for a very restricted class, namely that where each player can choose from at most three different actions and his payoff depends on at most three other players. We begin this chapter by strengthening this result to apply to an even more restrictive setting. To be precise, we show that two actions per player, two-bounded neighborhood, and two-valued payoff functions suffice for NP-completeness. This result is tight, because deciding the existence of a pure Nash equilibrium becomes trivial in the case of a single action for each player and tractable for one-bounded neighborhood. In fact, we show the latter problem to be NL-complete in general, and thus solvable in deterministic polynomial time. Interestingly, it turns out that the number of actions in a game with one-bounded neighborhood is a sensitive parameter: restricting the number of actions for each player to a constant makes the problem even easier than NL unless L = NL. In this way, we obtain a nice alternative characterization of the determinism-nondeterminism problem for Turing machines with logarithmic space in terms of the number of actions for games with one-bounded neighborhood.

We then move on to investigate the computational complexity of the pure equilibrium problem in graphical games which additionally satisfy one of four types of anonymity within each neighborhood. Despite these additional restrictions, the question for tractable classes of games is answered mostly in the negative. For three of the four types of anonymity, deciding the existence of a pure equilibrium remains NP-hard for games with two actions, two payoffs, and neighborhoods of size two. Assuming the most restricted type of anonymity, the problem becomes NP-hard when either there are three different payoffs, or neighborhoods of size four. On the other hand, we use interesting connections of the latter class to the even cycle problem in directed graphs and to generalized satisfiability to identify tractable classes of games. One such class for example arises from a situation where each agent is faced with the decision of producing one of two types of complementary goods within a regional neighborhood. In a sense, agents are not only producers but also consumers, and thus happier when both products are available within their neighborhood.

As a corollary, we further exhibit a satisfiability problem that remains NP-hard in the presence of a matching, a result which may be of independent interest. Finally, we show that mixed equilibria in games with two of four types of anonymity can be found in polynomial time if the number of actions grows only slowly in the neighborhood size. Quite interestingly, there exists a class of games where deciding the existence of a pure equilibrium is likely to be harder than finding a mixed one.

6.1 Related Work

The problem of finding (mixed) Nash equilibria in graphical games with neighborhood sizes bounded by three is equivalent to the same problem for general n-player games with $n \ge 4$ (Goldberg and Papadimitriou, 2006), and thus complete for the complexity class PPAD (Daskalakis et al., 2006). It is not surprising that the structure of the neighborhood graph greatly influences the complexity of the equilibrium problem. PPAD-hardness holds even if the underlying graph has constant pathwidth, but becomes tractable for undirected graphs of degree two, i.e., for *paths* (Elkind et al., 2006). All known algorithms for the more general case of *trees* have exponential worst-case running time even on trees with bounded degree and pathwidth two, but equilibria satisfying various fairness criteria can be computed in polynomial time if additionally there are only two actions per player and the best response policy, a data structure representing *all* Nash equilibria of a game, has polynomial size (Elkind et al., 2007).

A different line of research has investigated the problem of deciding the existence of *pure* equilibria. In addition to the above-mentioned hardness result, Gottlob et al. (2005) show tractability of the pure equilibrium problem for certain classes of games with restricted graph structure, in particular for graphs with *bounded treewidth*. The results of Chapter 5, and to some extent also the approximability results of Daskalakis and Papadimitriou (2008), fuel hope that tractability results can be obtained for larger classes of games satisfying some kind of *symmetry*. In this regard, Daskalakis and Papadimitriou (2005) consider games on a d-dimensional undirected torus or grid with payoff functions that are identical for all players and symmetric in the actions of the players in the neighborhood,

a condition we refer to as symmetry. The authors show that deciding the existence of a pure Nash equilibrium in such games is NL-complete when d = 1 and NEXP-complete for $d \ge 2$. In the second part of this chapter, we investigate the pure equilibrium problem in graphical games satisfying, within each neighborhood, one of the four types of anonymity introduced in Chapter 5. This can be seen as a refinement of the work of Gottlob et al. (2005) and of Daskalakis and Papadimitriou (2005).

6.2 The Model

A graphical game is given by a graph on the set of players, such that the payoff of a particular player depends only on his own action, and on the actions of his neighbors in the graph. In the following definition, the underlying graph is directed, corresponding to a neighborhood relation that is not necessarily symmetric.¹

DEFINITION 6.1 (graphical game). Let $\Gamma = (N, (A_i)_{i \in N}, (p_i)_{i \in N})$ be a normal-form game, $\nu : N \to 2^N$. Γ is a graphical game with neighborhood ν if for all $i \in N$ and $a_N, a'_N \in A_N$, $p_i(a_N) = p_i(a'_N)$ whenever $a_{\hat{\nu}(i)} = a'_{\hat{\nu}(i)}$, where $\hat{\nu}(i) = \nu(i) \cup \{i\}$.

A game Γ is said to have k-bounded neighborhoods if there exists $\nu : N \to 2^N$ such that Γ is a graphical game with neighborhood ν and for all $i \in N$, $|\nu(i)| \leq k$.

We assume throughout the chapter that graphical games are encoded by listing the payoffs of each player as a function of the actions of his neighbors. This encoding has polynomial size in the number of players if and only if neighborhood sizes are bounded by a constant.

Symmetry as a property of a mathematical object refers to its invariance under a certain type of transformation. Symmetries of games usually mean invariance of the payoffs under automorphisms of the set of action profiles induced by some group of permutations of the set of players. Anonymous games, for example, as considered in Chapter 5, require the set of available actions to be the same for all players, and the payoff of a particular player to remain the same under any permutation of the elements of an action profile. This imposes constraints on individual payoff functions only, and can therefore directly be applied to graphical games as well. In general, however, it does not make much sense from a computational point of view to consider symmetries of the payoff functions without requiring the neighborhood graph to be symmetric in an appropriate way as well. Consider, for example, the class of all graphical games whose payoff functions are invariant under automorphisms in the automorphism group of the neighborhood graph. While this class of games is very natural, it does not impose meaningful computational restrictions. Indeed, it is not too hard to see that any graphical game can be encoded by a game in the

¹While results can be transferred between graphical games on directed and undirected graphs, tightness of bounds on the size of the respective neighborhoods is obviously lost in the process. Results for directed graphs will in general be more expressive.

above class that has a neighborhood graph with a trivial automorphism group. Hardness results for both pure and mixed equilibria thus carry over immediately.

In general, different types of restrictions on the neighborhood structure will be required for different kinds of symmetries of the payoff functions. We take a slightly different approach by considering properties found in anonymous and symmetric games and studying graphical games that possess these properties. A characteristic feature of symmetries in games is the inability to distinguish between other players. As in Chapter 5, the most general class of games with this property will be called *anonymous*. Four different classes of games are again obtained by considering two additional characteristics: *identical payoff functions* for all players and the ability to *distinguish oneself* from the other players. The games obtained by adding the former property will be called *symmetric*, and presence of the latter will be indicated by the prefix "*self*." For ease of exposition, we assume the set of actions, and possibly the payoff functions, in these games to be the same for *all* players rather than just those with intersecting neighborhoods. The set of actions will be denoted by $A = A_1 = A_2 = \cdots = A_n$.

Again, an intuitive way to describe anonymous games is in terms of equivalence classes of the aforementioned automorphism group. Given a set A of actions, let $\#(a_N)$ denote the *commutative image* of an action profile $a_N \in A^N$, i.e., $\#(a_N) = (\#(a, a_N))_{a \in A}$, where $\#(a, a_N) = |\{i \in N : a_i = a\}|$. This definition naturally extends to action profiles for subsets of the players.

DEFINITION 6.2 (anonymity). Let $\Gamma = (N, (A_i)_{i \in N}, (p_i)_{i \in N})$ be a graphical game with neighborhood ν , A a set of actions such that for all $i \in N$, $A_i = A$. Γ is called

- anonymous if for all $i \in N$ and all $a_N, a'_N \in A^N$, $p_i(a_N) = p_i(a'_N)$ whenever $a_i = a'_i$ and for all $a \in A$, $\#(a, a_{\nu(i)}) = \#(a, a'_{\nu(i)})$;
- symmetric if for all $i, j \in N$ and all $a_N, a'_N \in A^N$, |v(i)| = |v(j)| and $p_i(a_N) = p_j(a'_N)$ whenever $a_i = a'_j$ and for all $a \in A$, $\#(a, a_{v(i)}) = \#(a, a'_{v(j)})$;
- self-anonymous if for all $i \in N$ and all $a_N, a'_N \in A^N$, $p_i(a_N) = p_i(a'_N)$ whenever for all $a \in A$, $\#(a, a_{\hat{v}(i)}) = \#(a, a'_{\hat{v}(i)})$; and
- self-symmetric if for all $i, j \in N$ and all $a_N, a'_N \in A^N$, |v(i)| = |v(j)| and $p_i(a_N) = p_j(a'_N)$ whenever for all $a \in A$, $\#(a, a_{\hat{v}(i)}) = \#(a, a'_{\hat{v}(i)})$.

It should be noted that a graphical game in one of the four classes does not in general belong to the corresponding class of Chapter 5, unless the neighborhood of every player contains all other players. When talking about self-anonymous and self-symmetric games with two actions, we write $\mathbf{p}_i(m) = \mathbf{p}_i(a_N)$ for any action profile a_N with $\#(1, a_{\hat{\mathbf{v}}(i)}) = m$ to denote the payoff of player i when m players in his neighborhood, including himself, play action 1, and $\mathbf{p}_i = (\mathbf{p}_i(m))_{0 \leqslant m \leqslant |\hat{\mathbf{v}}(i)|}$ for the vector of payoffs for the possible values of m.

6.3 A Tight Hardness Result for Pure Equilibria

Gottlob et al. (2005) show that deciding the existence of a pure equilibrium is NP-complete for graphical games with three-bounded neighborhoods and at most three actions for each player. We improve upon this result by showing that hardness holds already in the case of two-bounded neighborhoods, two actions, and two-valued payoff functions. Schoenebeck and Vadhan (2006) independently showed NP-hardness for three-bounded symmetric neighborhoods. Both constructions can easily be adapted to show the respective other result.

THEOREM 6.3. Deciding whether a graphical game has a pure Nash equilibrium is NP-complete. Hardness holds even for games with two-bounded neighborhoods, two actions for each player, and two-valued payoff functions.

Proof. Membership in NP is obvious. We can guess an action profile s and verify in polynomial time that it satisfies the equilibrium condition.

For *hardness*, recall that circuit satisfiability (CSAT), i.e., deciding whether for a given Boolean circuit \mathcal{C} with k inputs and one output there exists an assignment such that \mathcal{C} evaluates to *true*, is NP-complete (e.g., Papadimitriou, 1994a). Assume without loss of generality that \mathcal{C} contains at least one input and one (internal) gate, and that NOT gates only occur at the input layer. For an arbitrary circuit, all NOT gates can be moved to the input layer in polynomial time by successive application of de Morgan's law.

Given a Boolean circuit C, we define a graphical game $\Gamma = (N, (A_i)_{i \in N}, (p_i)_{i \in N})$, and argue that Γ has a pure Nash equilibrium if and only if C is satisfiable. As for the players of Γ , there is one for each input of C, one for each (positive or negative) literal, and one for each gate of types AND and OR. We denote the respective sets of players by N_i , N_x , $N_{\bar{x}}$, N_{\wedge} , and N_{\vee} . The output of C corresponds to a particular player $o \in N_{\wedge} \cup N_{\vee}$. For each input player $i \in N_i$, let $\nu(i) = \emptyset$. For each gate player $i \in N_{\wedge} \cup N_{\vee}$, let $\nu(i)$ be the set of players corresponding to the inputs of the gate—whenever a gate is connected to one of the k inputs or to a NOT-gate, $\nu(i)$ contains the player corresponding to the appropriate positive or negative literal. Finally, for each literal player $i \in N_x \cup N_{\bar{x}}$, let $\nu(i)$ contain the appropriate input player and the output player o. For every player $i \in N$, let $A_i = \{1, 0\}$, where 1 and 0 can be interpreted as truth values. Finally define the payoff functions as follows:

- For input players $i \in N_i$, let $p_i(a_N) = 1$ for every action profile $a_N \in A_N$.
- For positive literal players i ∈ N_x, let p_i(a_N) = 1 if in action profile a_N ∈ A_N, i plays the same action as the input player j ∈ v(i) and the output player o plays 1, or if i plays 1 and o plays 0; p_i(a_N) = 0 otherwise.
- For negative literal players i ∈ N_{x̄}, let p_i(a_N) = 1 if in action profile a_N ∈ A_N, i plays the opposite of the action of the input player j ∈ v(i) and the output player o plays 1, or if i plays 1 and o plays 0; p_i(a_N) = 0 otherwise.

		$\mathfrak{a}_{\mathbf{v}(\mathfrak{i})}$								$\mathfrak{a}_{\mathbf{v}(\mathfrak{i})}$				
				11	01	10	00			11	01	10	00	
1	1	-	1	1	1	0	1	-	1	0	1	1	1	
0	1		0	0	0	1	0		0	1	0	0	0	

Figure 6.1: Payoffs $p_i(a_N)$ for input, positive literal, and negative literal players, used in the proof of Theorem 6.3. The latter two depend on the actions of the corresponding input player and the output player o.

		a_{γ}	'(i)				$\mathfrak{a}_{\mathbf{v}(\mathfrak{i})}$				
	11	01	10	00			11	01	10	00	
1	1	0	0	0	1		1	1	1	0	
0	0	1	1	1	С)	0	0	0	1	

Figure 6.2: Payoffs $p_i(a_N)$ for AND and OR players, used in the proof of Theorem 6.3. Payoffs depend on the actions of players corresponding to the inputs of the gate.

- For AND players i ∈ N∧, let p_i(a_N) = 1 if i and both players j ∈ v(i) play 1, or if i and at least one j ∈ v(i) play 0; p_i(a_N) = 0 otherwise.
- For OR players $i \in N_{\vee}$, let $p_i(a_N) = 1$ if i plays 1 and at least one player $j \in v(i)$ plays 1, or if g and both players $j \in v(i)$ play 0; $p_i(a_N) = 0$ otherwise.

The payoff functions for the different types of players are summarized in Figures 6.1 and 6.2. It is readily appreciated that each player of Γ has at most two neighbors and two different actions, that all payoff functions take at most two different values, and that for a particular Boolean circuit the neighborhood graph and the payoff functions can be constructed in polynomial time.

Now consider a pair of a Boolean circuit \mathcal{C} and the corresponding game Γ . We claim that \mathcal{C} is satisfiable if and only if Γ has a pure Nash equilibrium. For the direction from left to right, assume that \mathcal{C} has a satisfying assignment ϕ , and consider an action profile a_N of Γ where (i) each input player plays according to ϕ , (ii) each literal player correctly reproduces the action of the corresponding input player, i.e., positive literals play the same action as their input, negative ones playing the opposite action, and (iii) each gate player correctly implements the truth function of the respective gate depending on the inputs, i.e., actions of his neighbors. By construction of Γ , and since ϕ is a satisfying assignment, the output player plays action 1 in a_N , and each player receives a payoff of 1. Since 1 is the maximum payoff, a_N is a Nash equilibrium.

For the direction from right to left, we use the following properties of action profiles of Γ :

1. A profile where all literal players play action 1 and the output player o plays 0 cannot

be a pure Nash equilibrium. Since no negations occur above the literal players, any gate player $j \in N_{\wedge} \cup N_{\vee}$ with $\nu(j) \subseteq (N_x \cup N_{\bar{x}})$ who plays 0 could increase his payoff by playing 1. By induction over the structure of the gate this in fact holds for all gate players, and for o in particular. This is a contradiction.

- 2. A profile where o plays 0 cannot be a pure Nash equilibrium. In this case, any literal player not playing 1 could improve his payoff by playing 1, contradicting Property 1.
- 3. A profile where o plays 1 and some literal player $i \in P_x \cup P_{\bar{x}}$ plays an action that does not correctly implement the value of the corresponding input cannot be a Nash equilibrium. In this case, i could change his action to increase his payoff, a contradiction.
- 4. A profile in which some gate player i ∈ P_∧ ∪ P_∨ plays an action that does not correctly implement the Boolean function of the corresponding gate cannot be a Nash equilibrium. In this case, i could change his action to increase his payoff, a contradiction.
- 5. A profile where the input players do not play a satisfying assignment but o plays 1 cannot be a Nash equilibrium. By Property 4, all gate players would have to play according to the Boolean function they implement. In particular, since the assignment is not satisfying, player o would play 0, contradicting Property 2.

By combining Properties 1 to 4, we conclude that every pure Nash equilibrium of Γ corresponds to a satisfying assignment of \mathcal{C} , and thus there is a one-to-one correspondence between satisfying assignments of \mathcal{C} and pure Nash equilibria of Γ .

The proof of Theorem 6.3 also shows the following.

COROLLARY 6.4. Counting the pure Nash equilibria of a graphical game is #P-complete. Hardness holds even for games with two-bounded neighborhoods, two actions for each player, and two-valued payoff functions.

Obviously, deciding the existence of a pure Nash equilibrium is trivial if players only have one action or if payoff functions are single-valued. Hence, the only interesting case that remains concerns games with one-bounded neighborhood.

The interaction among players of a graphical game Γ with neighborhood ν can be represented as a neighborhood graph $G(\Gamma) = (N, E)$, with vertices corresponding to players of N and directed edges to the neighbors of each player, i.e., $(i, j) \in E$ if $j \in \nu(i)$. The neighborhood graph of a game with one-bounded neighborhood is a (directed) *pseudo-forest*, i.e., a graph where every vertex has outdegree at most one. Each component of such a graph can be obtained by taking a rooted tree with edges oriented towards the root, and possibly adding an additional edge from the root to some other vertex. The following lemma states that the existence of a pure Nash equilibrium of a game whose

neighborhood graph is a pseudoforest only depends on players corresponding to vertices on cycles of the neighborhood graph.

LEMMA 6.5. Let $\Gamma = (N, (A_i)_{i \in N}, (p_i)_{i \in N})$ be a graphical game with neighborhood ν such that $G(\Gamma)$ is a pseudoforest. Denote by Γ' the game obtained by restricting Γ to those players whose corresponding vertex in $G(\Gamma)$ lies on directed cycle, i.e., $\Gamma' = (N', (A_i)_{i \in N'}, (p_i)_{i \in N'})$ where $N' = \{i \in N : there is a path from i to i in <math>G(\Gamma)\}$. Then, Γ has a pure Nash equilibrium if and only if Γ' has a pure Nash equilibrium.

Proof. The implication from left to right is trivial. For the direction from right to left, consider a pure equilibrium a'_N of Γ' and define an action profile a_N of Γ as follows. For all $i \in N'$, let $a_i = a'_i$. As long as there exists a player $i \in N$ for which a_i has not yet been defined, find such a player such that a_j is defined for all $j \in v(i)$, and define a_i as an action that maximizes the payoff of player i given $a_{v(i)}$. Obviously this procedure terminates after a finite number of steps. Furthermore, since the subgraph of $G(\Gamma)$ induced by $N \setminus N'$ does not contain a directed cycle, it is readily appreciated that a_N is a Nash equilibrium of Γ .

We are now ready to classify the complexity of the pure Nash equilibrium problem for games with one-bounded neighborhoods.

THEOREM 6.6. Deciding whether a graphical game with one-bounded neighborhoods has a pure Nash equilibrium is NL-complete. Hardness holds even for games with two-valued payoff functions.

Proof. For *membership* in NL, let Γ be a graphical game with one-bounded neighborhood.

By the observation that $G(\Gamma)$ is a pseudoforest, and by Lemma 6.5, it is sufficient to decide whether *every* game induced by the members of a cycle of $G(\Gamma)$ has a pure equilibrium. Since every vertex of $G(\Gamma)$ has outdegree at most one, these cycles can be found in *deterministic* logarithmic space.

For the players on a particular cycle, we can guess an action profile in the following way. We start at an arbitrary deterministically chosen vertex on the cycle and guess an action for the player corresponding to this vertex. We then traverse the cycle backwards, guessing an action for each player and checking whether this action maximizes the player's payoff given the action of the next player. The computation ends when the initially chosen vertex has been reached, accepting if and only if the initially chosen action equals the action guessed in the last step of the traversal. Note that the traversal can be done in deterministic logarithmic space, and that we only need to maintain a constant number of pointers. Thus, the whole computation can be performed by a Turing machine with logarithmic space.

For *hardness*, we reduce the NL-complete reachability problem for directed graphs, i.e., the problem of deciding whether for a given directed graph G = (V, E) and two designated vertices $s, t \in V$ there exists a path from s to t in G (e.g., Papadimitriou, 1994a), to the

pure equilibrium problem for graphical games with one-bounded neighborhoods. Without loss of generality, assume that (i) $V = \{1, 2, ..., n\}$ with $n \ge 4$, s = 1, and t = n, that (ii) every vertex i with $1 \le i \le n - 1$ has outdegree at least one, and that (iii) the only edge leaving n is a self-loop.

Define a graphical game $\Gamma = (N, (A_i)_{i \in N}, (p_i)_{i \in N})$ with players $N = \{1, 2, ..., n - 1\}$ and neighborhoods $\nu(1) = \{n - 1\}$ and $\nu(i) = \{i - 1\}$ for $2 \leq i \leq n - 1$. The set of possible actions is the same for all players and given by $\{a^{ij} : 1 \leq i, j \leq n\}$, where a^{ij} can be interpreted as selection of the (possibly non-existing) edge (i, j) in graph G. For an action profile a_N , define the payoffs as follows:

- The payoff function p_1 of player 1 is such that $p_1(a_N) = 1$ if (i) player 1 plays $a_1 = a^{1k}$ for some $(1, k) \in E$ and player n 1 plays $a_{n-1} = a^{jn}$ for some j, or if (ii) player 1 plays $a_1 = a^{nn}$ and player n 1 plays $a_{n-1} = a^{jk}$ such that $k \neq n$; $p_1(a_N) = 0$ otherwise.
- For players i ∈ N with 2 ≤ t ≤ n − 1, the payoff function p_i closely resembles the transition matrix of G. More precisely, p_i(a_N) = 1 if i plays a_i = a^{jk} for some (j,k) ∈ E and i − 1 plays a_{i-1} = a^{lj} for some l; p_i(a_N) = 0 otherwise.

We claim that there exists a path from vertex 1 to vertex n in G if and only if Γ has a pure equilibrium.

For the direction from left to right, assume that there exists a path $1 = v_1, v_2, \ldots, v_n = n$ of length n in G. A shorter path from 1 to n can be extended to this length by virtue of the self-loop at n. Now consider an action profile a_N of Γ where $a_i = a^{v_i v_{i+1}}$ for each $i \in N$. It is easily verified that in this case, each player receives the maximum payoff of 1, such that a_N is a Nash equilibrium.

For the direction from right to left, we need to show that a Nash equilibrium of Γ yields a path connecting vertices 1 and n in G. We exploit the following properties of action profiles of Γ :

- An action profile where some player i with 2 ≤ i ≤ n 1 plays a^{kl} and player i 1 plays a^{ij} for some j ≠ k cannot be a Nash equilibrium. In this case, player i would obtain a payoff of 0, and by construction there exists an alternative action a^{kl} with (k, l) ∈ E with payoff 1, because every vertex of G has outdegree at least one. This is a contradiction.
- 2. An action profile where some player i with $1 \le i \le n-2$ plays a^{kn} for some $1 \le k \le n$ and some player j with $i+1 \le j \le n-1$ does not play a^{nn} cannot be a Nash equilibrium. Assume without loss of generality that j is the smallest such number. Then player j could increase his payoff by playing a^{nn} , a contradiction.
- 3. An action profile where player n-1 plays $a^{k\ell}$ with $\ell \neq n$ cannot be a Nash equilibrium. In this case the best response for player 1 would be to play a^{nn} , contradicting Property 2.

4. An action profile where player 1 plays a^{ij} with i ≠ 1 cannot be a Nash equilibrium. We distinguish two different cases. If player n - 1 played a^{kn} for some k, player 1 could increase his payoff by playing a^{1j}, a contradiction. If instead player n - 1 played a^{kl} for some l ≠ n, player 1 could improve his payoff by playing aⁿⁿ, contradicting Property 2.

By combining these properties, we conclude that the only action profiles that are pure Nash equilibria of Γ are those where (i) player 1 plays an action a^{1j} with $(1,j) \in E$, (ii) each pair of players i - 1 and i with $2 \leq i \leq n - 2$ play actions a^{jk} and $a^{k\ell}$, respectively, for some $1 \leq j, k, \ell \leq n$, with $(k, \ell) \in E$, and (iii) player n - 1 plays an action a^{kn} with $(k, n) \in E$. There thus is a one-to-one correspondence of paths between vertices 1 and n in G and pure Nash equilibria of Γ . We further observe that for a particular graph G, the payoff functions can be constructed in logarithmic space and take two different values, and that each player has at most one neighbor.

An immediate consequence of the above proof is given next.

COROLLARY 6.7. Counting the pure Nash equilibria of a game with one-bounded neighborhoods is #L-complete.

In the hardness part of the proof of Theorem 6.6, the number of actions grows linearly in the number of vertices of graph G. For games with one-bounded neighborhoods where the number of actions grows only slowly, the pure equilibrium problem turns out to be L-complete.

THEOREM 6.8. Deciding whether a graphical game with one-bounded neighborhoods and at most $\log \log n + O(1)$ actions has a pure Nash equilibrium is L-complete under constant-depth reducibility. Hardness holds even for games with two actions and two-valued payoff functions.

Proof. The proof for *membership* follows similar lines as the corresponding part of the proof for Theorem 6.6. However, deterministic logarithmic space suffices to decide the existence of a pure equilibrium in a graphical game whose neighborhood graph is a cycle. When traversing the cycle backwards, we can write down all best responses of a particular player within the space bound, since the overall number of actions is small.

Hardness can be shown by a straightforward reduction from the L-complete reachability problem for directed graphs with outdegree one (Jones, 1975). Let G = (V, E) be a graph with this property and two designated vertices $s, t \in V$. Since outgoing edges of t have no influence on the reachability of t from s, we can assume without loss of generality that $(t, s) \in E$. Now define a game Γ that has G as its neighborhood graph, i.e., $\Gamma = (N, (A_i)_{i \in N}, (p_i)_{i \in N})$ where N = V and $j \in v(i)$ if $(i, j) \in E$. Let $A_i = \{0, 1\}$, and define the payoff functions as follows:

• For player t, let $p_t(a_N) = 1$ for any action profile $a_N \in A_N$ with $a_t \neq a_s$, and $p_t(a_N) = 0$ otherwise.

6.4 · Pure Equilibria of Graphical Games with Anonymity

• For any player $i \neq t$, let $p_i(a_N) = 1$ if $a_i = a_j$, $p_i(a_N) = 0$ otherwise, where $j \in v(i)$.

It is easily verified that this reduction can be computed in constant depth. We further claim that Γ does *not* have a pure equilibrium if and only if y is reachable from x in G. The statement of the theorem then follows by recalling that deterministic logarithmic space is closed under complementation.

First assume that t is reachable from s in G. Then s and t lie on a cycle in the neighborhood graph of Γ , and by definition of the payoff functions p_i for players $i \neq t$, all players on this cycle have to play the same action in every pure equilibrium. Then, by definition of p_t , player t can increase his payoff by changing his action, a contradiction.

Now assume that yt is not reachable from s, and partition V into a set V_1 of vertices from which t is reachable, including t itself, and a set V_2 of vertices from which t is not reachable. Consider the action profile a_N where all players in V_1 play action 0, and all players in V_2 play action 1. Since $s \in V_2$, and since (t, s) is the only edge originating from t, it is readily appreciated that $p_i(a_N) = 1$ for every player $i \in N$, such that a_N is a Nash equilibrium.

By combining Theorems 6.6 and 6.8, we obtain an alternative characterization of the determinism-nondeterminism problem for Turing machines with logarithmic space.

COROLLARY 6.9. The following statements are equivalent:

1. L = NL.

- 2. The existence of a pure Nash equilibrium in a graphical game with one-bounded neighborhoods can be decided in deterministic logarithmic space.
- 3. For every graphical game Γ with one-bounded neighborhoods, we can construct in deterministic logarithmic space a graphical game Γ' with one-bounded neighborhoods and O(log log n) actions such that Γ' has a pure Nash equilibrium if and only if Γ has a pure Nash equilibrium.

6.4 Pure Equilibria of Graphical Games with Anonymity

Let us now consider games with bounded neighborhoods that additionally satisfy anonymity within each neighborhood. For neighborhoods of size one, anonymity does not impose any restrictions. We further know from Theorem 6.6 that the pure equilibrium problem for such games can be decided in polynomial time. On the other hand, it is not hard to strengthen Theorem 6.3 to apply to anonymous graphical games as well. While the game used in the proof of that theorem is not anonymous, an anonymous game can be obtained by adapting the construction of Schoenebeck and Vadhan (2006) mentioned earlier to directed neighborhood graphs. Like the one in the proof of Theorem 6.3, this

$\#(1, \mathfrak{a}_{\mathbf{v}(\mathfrak{i})})$												
	0	1	2		$\#(1,\mathfrak{a}_{\hat{\mathbf{v}}(\mathfrak{i})})$	0	1	2	3			
0	0	0	1		$p_i(a_N)$	0	1	2	0			
1	1	1	0									

Figure 6.3: Payoffs $p_i(a_N)$ for NAND players in the symmetric and the self-symmetric case, used in the proof of Theorem 6.11. Columns correspond to the different values of the commutative image of a_N with respect to v(i) or $\hat{v}(i)$. In the symmetric case, rows correspond to the two actions of player i.

construction models gates of a Boolean circuit by players of a graphical game, but then has two additional players play a game with or without a pure equilibrium, depending on the output of the circuit. By observing that all payoff functions required for this construction are anonymous, we have the following.

THEOREM 6.10 (Schoenebeck and Vadhan, 2006). Deciding whether a graphical game has a pure Nash equilibrium is NP-complete, even if every player has only two neighbors, two actions, and two different payoffs, and when restricted to anonymous games.

6.4.1 Symmetry and Self-Symmetry

We now turn to the more restrictive class of *symmetric* graphical games, where the payoff functions of all players are identical. The proof of the following theorem again uses a construction similar to the one of Schoenebeck and Vadhan (2006). The main difficulty is to model the two building blocks, a Boolean circuit and games with and without pure equilibria, using only a single payoff function.

THEOREM 6.11. Deciding whether a graphical game has a pure Nash equilibrium is NP-complete, even if every player has only two actions, and when restricted to symmetric games with two different payoffs or to self-symmetric games with three different payoffs.

Proof. Membership in NP is straightforward. We can simply guess an action profile and verify that the action of each player is a best response to the actions of the players in his neighborhood.

For hardness, we provide a reduction from CSAT. For a set N of players with appropriately defined neighborhoods ν , let $\Gamma(N) = (N, \{0, 1\}^N, (p_i)_{i \in N})$ be a graphical game with payoffs satisfying symmetry or self-symmetry as given in Figure 6.3.² We observe the following properties:

²Also recall that every anonymous or symmetric game with two actions per player can respectively be reduced to a self-anonymous or self-symmetric game, while preserving pure equilibria.

- 1. Let N be a set of players, |N| = 3, and $\hat{v}(i) = N$ for all $i \in N$. Then, an action profile a_N of $\Gamma(N)$ is a pure equilibrium if and only if $\#(1, a_N) = 2$. In particular, for every $i \in N$, there exists a pure equilibrium where player i plays action 0 and a pure equilibrium where he plays action 1.
- 2. Let N and N' be two disjoint sets of players with neighborhoods such that for all $i \in N$, $\nu(i) \subseteq N$, and for all $i \in N'$, $\nu(i) \subseteq N'$. Then, $a_{N \cup N'}$ is a pure equilibrium of $\Gamma(N \cup N')$ if and only if a_N and $a_{N'}$ are pure equilibria of $\Gamma(N)$ and $\Gamma(N')$, respectively.
- Let N be a set of players such that Γ(N) has a pure equilibrium, and consider two players 1,2 ∈ N. Further consider an additional player 3 ∉ N with v(3) = {1,2}. Then the game Γ(N ∪ {3}) has a pure equilibrium, and in every pure equilibrium a_{N∪{c}} of Γ(N ∪ {c}), a₃ = 0 if a₁ = a₂ = 1, and a₃ = 1 otherwise. In other words, such an action profile always satisfies a₃ = a₁ NAND a₂.
- 4. Let N be a set of players and consider a particular player 1 ∈ N. Further consider five additional players 2, 3, 4, 5, 6 ∉ N with neighborhoods according to Figure 6.4, and denote N' = N ∪ {2, 3, 4, 5, 6}. Then Γ(N') has a pure equilibrium if and only if Γ(N) has a pure equilibrium a_N with a₁ = 0. For the direction from right to left, assume that Γ(N) has a pure equilibrium a_N with a₁ = 0. For the direction from right to left, assume that Γ(N) has a pure equilibrium a_N where a₁ = 0, and extend it to an action profile a_{N'} for Γ(N') by letting a₂ = 0 and a₃ = a₄ = a₅ = a₆ = 1. It is easily verified that a_{N'} is a Nash equilibrium of Γ(N'). For the direction from left to right, consider an action profile a_{N'} for Γ(N') where a₁ = 1. If a₂ = 0, then action 1 is the unique best response for players 4 and 5, after which action 0 is the unique best response for players 3 and 6. In this case, player 2 could change his action to get a higher payoff. If a₂ = 1, then the unique best response for players 3 and 6 is 1. Again, player 2 could change his action to get a higher payoff.
- 5. Let $N_1 = \{1, 2, 3\}$ be an instance of N in Property 1, and N_2 an instance of N' in Property 4 with $N = \{1\}$. Let N be any set of players such that $\Gamma(N)$ has a pure equilibrium, let $4 \in N$, and denote $N' = N_1 \cup N_2 \cup N$. Further consider an additional player $5 \notin N'$ with $\nu(5) = \{2, 4\}$. Then, $\Gamma(N' \cup \{5\})$ has a pure equilibrium, and in every pure equilibrium $a_{N'\cup\{5\}}$ of $\Gamma(N' \cup \{5\})$, $a_5 = 1 - a_4$. To see this, observe that by Property 1 exactly two players in N_1 must play action 1, which, by Property 4, have to be players 2 and 3. Since ϕ NAND true $= \neg \phi$, the claim follows from Property 3.

Now consider an instance \mathcal{C} of CSAT, and assume without loss of generality that \mathcal{C} consists exclusively of NAND gates and that no variable appears more than once as the input to the same gate. The latter assumption can be made by Property 5. We construct a game $\Gamma = \Gamma(N)$ as follows. For every input of \mathcal{C} we augment N by three players according to



Figure 6.4: Output gadget, used in the proof of Theorem 6.11. All players have payoffs as in Figure 6.3. Player x must play action 0 in every pure equilibrium of the game.

Property 1. We then inductively define Γ by adding, for a gate with inputs corresponding to players $1, 2 \in \mathbb{N}$, a player 3 as described in Property 3. Finally, we construct a player according to Property 5 who plays the opposite action as the one corresponding to the output of \mathcal{C} , and identify this player with player 1 in a new instance of \mathbb{N}' in Property 4. It is now easily verified that a pure equilibrium of Γ corresponds to a computation of \mathcal{C} which outputs *true*, and that such an equilibrium exists if and only if \mathcal{C} has a satisfying assignment.

6.4.2 Self-Anonymity and Two Different Payoffs

Since self-symmetry implies self-anonymity, Theorem 6.11 also implies NP-hardness in the self-anonymous case. The result is not tight, however, in that three different payoffs are required for hardness. A natural question to ask is what happens for self-anonymous games with only two different payoffs. In this section we prove a tight result for the most restricted version of self-anonymity, i.e., the case with only two different payoff *functions*.

The problem with anonymity and the construction used in the proof of Theorem 6.11 is that two different payoffs are not enough to make a player care about his own action irrespective of the actions played by his neighbors. With four different values for $\#(1, a_{\hat{v}(i)})$, there will either be an equilibrium where all players play the same action, or a situation where a player is indifferent between both of his actions. When we want to use games to compute a function, such indifference is clearly undesirable. The key idea that will enable us to prove the following theorem is to isolate pure equilibria that are themselves symmetric in the actions of a subset of the players, i.e., equilibria in which these players all play the same action. To enforce that two particular players play the same action in every equilibrium, we will add two additional players, each of which observes the other as well as one of the original players. Depending on the actions of the original players, the new players will either play a game with a unique pure equilibrium, or a game that is prototypical both for self-anonymous games and for games without pure equilibria, namely Matching Pennies.

THEOREM 6.12. Deciding whether a graphical game has a pure Nash equilibrium is NP-complete, even if every player has only two neighbors, two actions, and two



Figure 6.5: Equality gadget, used in the proof of Theorem 6.12. A pure equilibrium exists if and only if players 1 and 2 play the same action.

different payoffs, and when restricted to self-anonymous games with two different payoff functions.

Proof. Membership in NP is again straightforward.

For hardness, we again give a reduction from CSAT. Let $\Gamma(N) = (N, \{0, 1\}^N, (p_i)_{i \in N})$ denote a graphical game for a set N of players with neighborhood ν and payoff functions p_i satisfying self-anonymity. We observe the following properties:

- 1. Let N be a set of players, $1, 2 \in N$, and consider two additional players $3, 4 \notin N$ with neighborhoods and payoffs according to Figure 6.5. We claim that $\Gamma(N \cup \{3, 4\})$ has a pure equilibrium if and only if $\Gamma(N)$ has a pure equilibrium a_N where $a_1 = a_2$. For the direction from right to left, assume that $\Gamma(N)$ has such an equilibrium, and extend it to an action profile $a_{N'}$ for $\Gamma(N')$ by letting $a_3 = 0$ and $a_4 = 1$. It is easily verified that under this action profile players 3 and 4 both receive the maximum payoff of 1, such that the equilibrium condition is trivially satisfied. For the direction from left to right, assume that one of players 3 or 4 observes action 0 being played by player 1 or player 2, while the other one observes action 1. Then players 3 and 4 effectively play the well-known Matching Pennies game. More precisely, the player observing 0 receives a payoff of 1 if and only if $\#(1, a_{\{3,4\}})$ is odd, while the same is true for the player observing 1 if and only if this number is even. Since both players can change between the two outcomes by changing their own action, there cannot be a pure equilibrium.
- 2. Let $N = \{1, 2, 3\}$ with $\nu(i) = N$ for all $i \in N$, and let p_i be defined according to Figure 6.5. It is then easily verified that a_N with $a_1 = a_2 = a_3 = 1$ or with $a_1 = a_2 = 0$ and $a_3 = 1$ is an equilibrium of $\Gamma(N)$. In particular, there exists an equilibrium where player 1 plays 0, and one where player 1 plays 1.
- 3. Let N and N' be two disjoint sets of players with neighborhoods such that for all $i \in N$, $\nu(i) \subseteq N$, and for all $i \in N'$, $\nu(i) \subseteq N'$. Again, $a_{N \cup N'}$ is a pure equilibrium of $\Gamma(N \cup N')$ if and only if a_N and $a_{N'}$ are pure equilibria of $\Gamma(N)$ and $\Gamma(N')$, respectively.
- 4. Let $N = \{1, 2, 3\}$ with neighborhoods and payoffs as in Property 2, and assume by Property 1 that every pure equilibrium a_N of $\Gamma(N)$ is symmetric, i.e., satisfies



Figure 6.6: NAND gadget, used in the proof of Theorem 6.12. The construction of Figure 6.5 is used to ensure that players connected by "=" play the same action in every pure equilibrium.

 $a_1 = a_2 = a_3$. Then, a_N with $a_1 = a_2 = a_3 = 1$ is the unique pure equilibrium of $\Gamma(N)$. Clearly, a_N is an equilibrium of $\Gamma(N)$, since all players receive the maximum payoff of 1. In the only other symmetric action profile, all players play action 0 and receive a payoff of 0, and each of them could change his action to 1 to receive a higher payoff.

- 5. Let N be a set of players such that Γ(N) has a pure equilibrium, let 1,2 ∈ N, and consider three additional players 3,4,5 ∉ N with neighborhoods and payoffs according to Figure 6.6. Then, Γ(N ∪ {3,4,5}) has a pure equilibrium, and for every pure equilibrium a_{N∪{3,4,5}</sub> of Γ(N ∪ {3,4,5}), a₃ = 0 if a₁ = a₂ = 1, and a₃ = 1 otherwise. To see this, observe that players 3, 4, and 5 get the maximum payoff of 1 under any action profile satisfying a₃ = a₄ = a₅ = 1 and #(1, a_{{1,2}) ≤ 1, or a₃ = a₄ = a₅ = 0 and a₁ = a₂ = 1. On the other hand, by Property 1, an action profile cannot be an equilibrium unless a₃ = a₄ = a₅. If a₁ = a₂ = a₅ = 0 or a₁ = a₂ = a₅ = 1, then player 5 can change his action to receive a higher payoff. If otherwise a₁ ≠ a₂ and a₃ = a₄ = 0, then there exists a player i ∈ {3,4} such that #(1, a_{γ(i)}) = 0, who could change his action to get a higher payoff.
- 6. Let N be a set of players, 4 ∈ N. Let N' = {1,2,3} with neighborhoods as in Property 4, N" = {5,6} with v(5) = {1,6} and v(6) = {4,5}. Then, Γ(N ∪ N' ∪ N") has a pure equilibrium if and only if Γ(N) has a pure equilibrium a_N with a₄ = 1. Clearly, an action profile that is not an equilibrium of Γ(N) cannot be extended to an equilibrium of Γ(N ∪ N' ∪ N"). Conversely assume that a_{N∪N'∪N"} is an equilibrium of Γ(N ∪ N' ∪ N"). Then, by Property 4, a₁ = 1. Furthermore, by Property 1, a₁ = a₄, and thus a₄ = 1.

Now consider an instance \mathcal{C} of CSAT, and assume without loss of generality that \mathcal{C} consists exclusively of NAND gates. Since ϕ NAND true = $\neg \phi$, and using Property 4, we can further assume that no variable appears more than once as an input to the same gate. We construct a game $\Gamma = \Gamma(N)$ as follows: For every input of \mathcal{C} , we add three players according to Property 2. For every gate of \mathcal{C} with inputs corresponding to players 1, $2 \in N$, we add three players according to Property 5. Finally, we add five players according to Property 6, such that player 4 is the one that corresponds to the output of \mathcal{C} . It is now



Figure 6.7: Neighborhood graph of a graphical game with seven players (left), corresponding to the three-uniform square hypergraph given by the lines of the Fano plane (right). A directed edge from vertex i to vertex j of the neighborhood graph denotes that $j \in v(i)$.

readily appreciated that Γ has a pure equilibrium if and only if C is satisfiable. \Box

6.4.3 Self-Symmetry and Two Different Payoffs

Let us return to self-symmetric graphical games. Self-symmetric games as studied in Chapter 5 always possess a pure Nash equilibrium due to the fact that they are commonpayoff games. This existence result does not hold for self-symmetric graphical games, even when there are only two different payoffs. In particular, there exists a seven-player game in the latter class that does not have a pure equilibrium, and in which each player has exactly two actions and two neighbors. It will be instructive to view a graphical game as a hypergraph, with each vertex corresponding to a player and each edge to the set of players in the neighborhood of one particular player *including* the player himself. Corresponding to the set of games with neighborhoods of size m is the set of (m+1)-uniform hypergraphs that possess a matching in the sense of Seymour (1974), i.e., a bijection from the set of vertices to the set of edges that maps every vertex to an edge containing it. It is not hard to see that a self-symmetric game with two actions and payoffs $\mathbf{p}_i = (0, 1, 1, 0)$ for all $i \in N$ has a pure Nash equilibrium if and only if the corresponding hypergraph is vertex two-colorable. Given a two-coloring, every player observes either one or two players in his neighborhood, including himself, who play action 1, and thus obtains the maximum payoff of 1. If on the other hand there is no two-coloring, then there is at least one player for every action profile who plays the same action as all of his neighbors and can deviate to obtain a higher payoff. Figure 6.7 shows the neighborhood of a graphical game with seven players and two neighbors for each player. This graph induces the 3-uniform square hypergraph corresponding to the lines of the Fano plane, which in turn cannot be twocolored (e.g., Seymour, 1974). We leave it to the avid reader to verify that there is no game with the above properties and less than seven players.

An interesting property of the neighborhood graph on the left of Figure 6.7 is that it does not have any cycles of even length. We will begin our investigation of the pure equilibrium problem in self-symmetric games by generalizing this observation to games with arbitrary neighborhoods and $\mathbf{p}_i = (0, 1, 1, \dots, 1, 0)$ for all $i \in N$. The following lemma characterizes games with pure equilibria in the above subclass in terms of cycles in the neighborhood graph. Seymour (1974) provides a similar characterization of the minimal uniform square hypergraphs that do not have a two-coloring.

LEMMA 6.13. Let $\Gamma = (N, (A_i)_{i \in N}, (p_i)_{i \in N})$ be a self-symmetric graphical game with $|A_i| = 2$ and $\mathbf{p}_i = (0, 1, 1, \dots, 1, 0)$ for all $i \in N$. Then, Γ has a pure Nash equilibrium if and only if for all $i \in N$, there exists $j \in N$ reachable from i such that j lies on a cycle of even length.

Proof. For the implication from left to right, assume that there exists a pure equilibrium, i.e., a two-coloring $c : N \to \{0, 1\}$ of the neighborhood graph such that the neighborhood of every player contains some player playing action 0 and some player playing action 1. Now consider an arbitrary player $v_1 \in N$. Using the above property of c, we can construct a path $v_1, v_2, \ldots, v_{|N|+1}, v_i \in N$, such that for all $i, 1 \leq i \leq |N|, c(v_i) = 1 - c(v_{i+1})$. By the pigeonhole principle, there must exist $i, j, 1 \leq i < j \leq |N| + 1$, such that $v_i = v_j$ and for all $j', i < j' < j, v_{j'} \neq v_i$. Then, $v_i, v_{i+1}, \ldots, v_j$ is a cycle of even length reachable from v_1 .

For the implication from right to left, let $N' \subseteq N$ be a set of players such that for every $i \in N$ there exists a directed path to some $j \in N'$, and such that N' induces a set of vertex-disjoint cycles of even length. We construct a two-coloring $c : N \to \{0, 1\}$, corresponding to an assignment of actions to players, as follows. First color the members of N' such that for all $i \in N'$ and $j \in v(i) \cap N'$, c(i) = 1 - c(j). While there are uncolored vertices left, find $i, j \in N$ such that $j \in v(i)$, i is uncolored, and j is colored. Such a pair of vertices must always exist, since for every member of N there is a directed path to some member of N', and thus to a vertex that has already been colored. Color i such that c(i) = 1 - c(j). It is now easily verified that at any given time, and for all $i \in N$ that have already been colored, there exist $j, j' \in \hat{v}(i)$ with c(j) = 0 and c(j') = 1. If all vertices have been colored, then every neighborhood will contain at least one player playing action 0, and at least one player playing action 1. The corresponding action profile is a pure Nash equilibrium.

Thomassen (1985) has shown that for every natural number m, there exists a directed graph without even cycles where every vertex has outdegree m. On the other hand, it is easy to construct graphs that do have even cycles. Together with Lemma 6.13, we thus have that the pure equilibrium problem for the considered class of games is nontrivial.

COROLLARY 6.14. For every $m \in \mathbb{N}$, m > 0, there exist self-symmetric graphical games Γ and Γ' with two actions and |v(i)| = m and $\mathbf{p}_i = (0, 1, 1, \dots, 1, 0)$ for all $i \in N$, such that Γ has a pure Nash equilibrium and Γ' does not.

We are now ready to identify several classes of graphical games where the existence of a pure equilibrium can be decided in polynomial time. THEOREM 6.15. Let Γ be a self-symmetric graphical game with payoffs p_i . The pure equilibrium problem for Γ can be decided in polynomial time if one of the following properties holds:

- (i) for all $i \in N$, $p_i(0) \ge p_i(1)$ or for all $i \in N$, $p_i(|\hat{\nu}(i)|) \ge p_i(|\hat{\nu}(i)|-1)$;
- (ii) for all $i \in N$ and all m with $1 \leq m \leq |v(i)|$, $\mathbf{p}_i(m-1) > \mathbf{p}_i(m)$ and $\mathbf{p}_i(m+1) > \mathbf{p}_i(m)$, or $\mathbf{p}_i(m-1) < \mathbf{p}_i(m)$ and $\mathbf{p}_i(m+1) < \mathbf{p}_i(m)$;
- (iii) for all $i \in N$ and all m with $1 \leq m < |v(i)|$, $p_i(m) = p_i(m+1)$.

Proof. It is easy to see that a game Γ satisfying (i) possesses a pure equilibrium a_N such that $\#(0, a_N) = 0$ or $\#(1, a_N) = 0$.

For a game Γ satisfying (*ii*), we observe that in every equilibrium a_N , $p_i(a_N) = 1$ for all $i \in N$. The pure equilibrium problem for Γ thus corresponds to a variant of generalized satisfiability, with clauses induced by neighborhoods of Γ . The constraints associated with this particular variant require an odd number of variables in each clause to be set to true, and can be written as a system of linear equations over GF(2). Tractability of the pure equilibrium problem for Γ then follows from Theorem 2.1 of Schaefer (1978).

Finally, a game satisfying (*iii*) but not (*i*) can be transformed into a best response equivalent one that satisfies the conditions of Lemma 6.13. We further claim that we can check in polynomial time whether for every $i \in N$, there exists $j \in N$ on a cycle of even length and reachable from i. For a particular $i \in N$, this problem is equivalent to checking whether the subgraph induced by the vertices reachable from i contains an even cycle. The latter problem has long been open, but was recently shown to be solvable in polynomial time (Robertson et al., 1999).

It is easily verified that every self-symmetric graphical game Γ with two different payoffs and neighborhoods of size two or three can be transformed into a game Γ' with the same set of players and the same neighborhoods, such that Γ and Γ' have the same set of pure equilibria and Γ' satisfies one of the conditions of Theorem 6.15. We thus have the following.

COROLLARY 6.16. The problem of deciding whether a self-symmetric graphical game with two different payoffs and three-bounded neighborhoods has a pure equilibrium is in P.

6.4.4 Self-Symmetry and Larger Neighborhoods

The question that remains is whether the pure equilibrium problem can be solved in polynomial time for all self-symmetric graphical games with two payoffs, or whether there is some bound on the neighborhood size where this problem again becomes hard. We will show in this section that the latter is true, and that the correct bound is indeed four, as suggested by Corollary 6.16.



Figure 6.8: Neighborhood graph and payoffs of a graphical game with eight players and neighborhoods of size four, used in the proof of Theorem 6.17. The neighborhood graph satisfies rotational symmetry, the neighborhood of player 1 is highlighted.

We will essentially use the same tools as in Section 6.4.2, but extract the necessary complexity from only a single payoff function. The additional insight required for this extraction is that "constant" players, i.e., players who play the same action in every pure equilibrium of a game, can be used to prune a larger payoff table and effectively obtain different payoff functions for smaller neighborhoods that can then be used to proceed with the original proof. Constructing such players will prove a rather difficult task in its own right.

THEOREM 6.17. Deciding whether a self-symmetric graphical game with two different payoffs has a pure Nash equilibrium is NP-complete, even if every player has exactly four neighbors.

Proof. Membership in NP is straightforward. We can simply guess an action for each player and then verify that no player can increase his payoff by playing a different action instead.

For *hardness*, we again give a reduction from CSAT to the problem at hand. The central idea of this proof will be to guarantee that some players in a neighborhood only play certain well-defined actions in equilibrium. By this, the original payoff table is effectively "pruned" to a smaller one that can then be used, like in earlier proofs, to model the behavior of gates in a Boolean circuit.

As a first step, we will show how to construct "constants," i.e., players who play the same action, 0 or 1, in *every* equilibrium of a game. To achieve this, we will construct a set of four players, such that in every equilibrium two of these players play action 0 and two of them play action 1. A player observing these four players can determine if the number of players in his neighborhood, including himself, who play action 1 is two or three. Clearly, such a player will play action 1 in every equilibrium. By a similar argument, a player who observes four players who play action 1 in every equilibrium will himself play action 0 in every equilibrium.

Consider the graphical game Γ with eight players and neighborhoods of size four given

in Figure 6.8. We will argue that the pure equilibria of this game are exactly those action profiles where two players, with distance two on the outer cycle, play action 0. More formally, an action profile a_N is an equilibrium if and only if there is one pair of distinct players $i, j \in N$ such that $a_i = a_j = 0$, and it holds that $i - j = 2 \pmod{4}$.

For the direction from right to left, we can exploit rotational symmetry of the neighborhood graph and assume without loss of generality that $a_1 = a_3 = 0$. The statement then follows by checking that the equilibrium condition is satisfied for all players. For the direction from left to right, we exploit the following properties of the neighborhood graph:

- 1. For any distinct $i, j \in N$ with $i-j \neq 2 \pmod{4}$, there exist a player $\ell \in N$ such that $\hat{v}(\ell) \cap \{i, j\} = \emptyset$. Assuming without loss of generality that i = 1, the property follows from a case analysis.
- 2. For any $N' \subseteq N$, |N'| = 3, there exists a player $i \in N$ such that $N' \subseteq \hat{\nu}(i)$. Due to rotational symmetry, we can assume without loss of generality that $1 \in N'$. The property then follows by a straightforward if somewhat tedious case analysis.
- For any N' ⊆ N, |N'| = 3, there exists a player i ∈ N such that |N' ∩ v̂(i)| = 2. Showing this property is again straightforward by assuming without loss of generality that 1 ∈ N' and showing that for any pair of distinct players, there exists a player i ∈ N such that either v̂(i) contains player 1 and exactly one element of the pair, or both elements of the pair but not player 1.
- For any N' ⊆ N, |N'| = 4, there exists a player i ∈ N such that |N' ∩ v(i)| = 3. To show this property, we can again assume without loss of generality that 1 ∈ N', and distinguish neighborhoods that contain player 1 from neighborhoods that do not. The analysis is again straightforward.

Now consider an equilibrium a_N of Γ , and observe that due to the structure of the payoffs, it must be the case that $p_i(a_N) = 1$ for all $i \in N$. If $\#(0, a_N) < 2$ or $\#(1, a_N) < 2$, then there exists a player $i \in N$ such that $\#(0, a_{\hat{v}(i)}) = 0$ or $\#(1, a_{\hat{v}(i)}) = 0$, contradicting the assumption that a_N is an equilibrium. If $\#(0, a_N) = 2$, assume without loss of generality that $a_1 = 0$, and consider $i \in N \setminus \{1, 3, 7\}$ such that $a_i = 0$. Then, by Property 1, there exists a player $j \in N$ such that $\#(0, a_{\hat{v}(j)}) = 0$, which again leads to a contradiction. If $\#(0, a_N) = 3$, then by Property 2 there must exist a player $i \in N$ such that $\#(0, a_{\hat{v}(i)}) = 3$ and thus $\#(1, a_{\hat{v}(i)}) = 2$, again contradicting the assumption that a_N is an equilibrium. By Property 4, the same holds if $\#(0, a_N) = 4$. If $\#(0, a_N) = 5$ and thus $\#(1, a_{\hat{v}(i)}) = 2$, a contradiction. The same trivially holds if $\#(1, a_N) = 2$.

Now we augment Γ by a set $\{9, 10, \ldots, 13\}$ of five additional players such that

$$\nu(\mathfrak{i}) = \begin{cases} \{1, 3, 5, 7\} & \text{if } \mathfrak{i} \in \{9, 10\} \\ \{2, 4, 6, 8\} & \text{if } \mathfrak{i} \in \{11, 12\} \\ \{9, 10, 11, 12\} & \text{if } \mathfrak{i} = 13. \end{cases}$$

By construction of the original game with eight players, every pure equilibrium has either two or four players in the common neighborhood of players 9 and 10 play action 1. Furthermore, if players 9 and 10 observe two players who play action 1, then players 11 and 12 will observe four players who play action 1, and *vice versa*. As a consequence, either players 9 and 10 will play action 0, and players 11 and 12 will play action 1, or the other way round. In any case, exactly two players in the neighborhood of player 13 will play action 1 in every equilibrium of the augmented game, and player 13 himself will therefore play action 1.

In the following, we denote by $0_1, 0_2, 0_3$ three players who play action 0 in every equilibrium, and by $1_1, 1_2$ two players that constantly play action 1. Using these players to prune the payoff table, we will then proceed to design games that simulate Boolean circuits. We want these games to satisfy self-symmetry, and the payoff of all players will therefore be determined by the table already used above and shown in Figure 6.8. As for the inputs of the circuit, it is easily verified that a game with players N, |N| = 5, such that for all $i \in N$, $\hat{v}(i) = N$, has pure equilibria a_N and a'_N such that for an arbitrary $i \in N$, $a_i = 0$ and $a'_i = 1$.

As before, we will now construct a subgame that simulates a functionally complete Boolean gate, in this case NOR, and a subgame that has a pure equilibrium if and only if a particular player plays action 1. For a set N of players with appropriately defined neighborhoods ν , let $\Gamma(N) = (N, \{0, 1\}^N, (p_i)_{i \in N})$ be a graphical game with payoff functions p_i satisfying self-symmetry as in Figure 6.8. We observe the following properties:

- 1. Let N and N' be two disjoint sets of players with neighborhoods such that for all $i \in N$, $\nu(i) \subseteq N$, and for all $i \in N'$, $\nu(i) \subseteq N'$. Again, $a_{N \cup N'}$ is a pure equilibrium of $\Gamma(N \cup N')$ if and only if a_N and $a_{N'}$ are pure equilibria of $\Gamma(N)$ and $\Gamma(N')$, respectively.
- 2. Let N be a set of players such that $\Gamma(N)$ has a pure equilibrium, let $1, 2 \in N$, and consider two additional players $3, 4 \notin N$ with $\nu(3) = \{0_1, 0_2, 1, 4\}$, and $\nu(4) = \{0_1, 0_2, 2, 3\}$. Then every pure equilibrium $a_{N \cup \{3,4\}}$ of $\Gamma(N \cup \{3,4\})$ satisfies $a_1 = a_2$.
- 3. Identifying player 2 with 1_1 in the previous construction, we have that $\Gamma(N \cup \{3,4\})$ has a pure equilibrium if and only if $a_1 = 1$ in some pure equilibrium a_N of $\Gamma(N)$.
- 4. Let N be a set of players such that $\Gamma(N)$ has a pure equilibrium, let $1, 2 \in N$, and consider two additional players $3, 4 \notin N$ with neighborhoods given by $\nu(3) = \{0_1, 0_2, 0_3, 4\}$ and $\nu(4) = \{0_1, 0_2, 1, 2\}$. Then $\Gamma(N \cup \{3, 4\})$ has a pure equilibrium, and every pure equilibrium $a_{N \cup \{3, 4\}}$ of $\Gamma(N \cup \{3, 4\})$ satisfies $a_3 = 1$ whenever $a_1 = a_2 = 0$, and $a_3 = 0$ whenever $a_1 \neq a_2$. For every pure equilibrium $a_{N \cup \{3, 4\}}$ with $a_1 = a_2 = 1$, there exists a pure equilibrium $a'_{N \cup \{3, 4\}}$ such that $a_3 \neq a'_3$, and $a_i = a'_i$ for all $i \in N$.
- 5. Consider an additional player $5 \notin N \cup \{3,4\}$, and let $\nu(5) = \{1_1, 1_2, 1, 2\}$. Then $\Gamma(N \cup \{3,4,5\})$ has a pure equilibrium, and every pure equilibrium $a_{N \cup \{3,4,5\}}$ of $\Gamma(N \cup \{3,4,5\})$



Figure 6.9: NOR gadget, used in the proof of Theorem 6.17. Payoffs are identical to those in Figure 6.8. A construction analogous to the one shown in Figure 6.5 is used to ensure that players 3 and 5 play the same action in every pure equilibrium.

 $\{x, y, z\}$ satisfies $a_5 = 1$ whenever $a_1 = a_2 = 0$, and $a_5 = 0$ whenever $a_1 = a_2 = 1$. For every pure equilibrium $a_{N \cup \{3,4,5\}}$ with $a_1 \neq a_2$, there exists a pure equilibrium $a'_{N \cup \{3,4,5\}}$ such that $a_5 \neq a'_5$, and $a_i = a'_i$ for all $i \in N$.

6. By Property 2, we can assume that every equilibrium $a_{N\cup\{3,4,5\}}$ of $\Gamma(N\cup\{3,4,5\})$ satisfies $a_3 = a_5$, and thus that $a_3 = 1$ if and only if $a_1 = a_2 = 0$.

Properties 4 through 6 are illustrated in Figure 6.9.

Now consider an instance C of CSAT, and assume without loss of generality that C consist exclusively of NOR gates and that no variable appears more than once as an input to the same gate. The latter assumption can be made since ϕ NOR *false* = $\neg \phi$, and since there exists a self-symmetric game and a player in this game who plays action 0 in every pure equilibrium. As before, we construct a game Γ by simulating every gate of C according to Property 6 and identifying the player that corresponds to the output of the circuit with player 1 in Property 3. It is now readily appreciated that Γ has a pure equilibrium if and only if C is satisfiable.

Observing that in the constructions used in the proofs of Theorems 6.11, 6.12, and 6.17 there is a one-to-one correspondence between satisfying assignments of a Boolean circuit and pure equilibria of a game, we have that counting the number of pure equilibria in the respective games is as hard as computing the permanent of a matrix.

COROLLARY 6.18. For graphical games with neighborhoods of size two, counting the number of pure Nash equilibria is #P-hard, even when restricted to symmetric graphical games with two different payoffs, to self-anonymous graphical games with two different payoffs and two different payoff functions, or to self-symmetric graphical games with three different payoffs. The same holds for self-symmetric graphical games with neighborhoods of size four and two different payoffs.

6.5 Interlude: Satisfiability in the Presence of a Matching

The analysis at the end of the previous section allows us to derive a corollary that may be of independent interest. Schaefer (1978) completely characterizes which variants of the generalized satisfiability problem are in P and which are NP-complete. Some of the variants become tractable if there exists a matching, i.e., a bijection from variables to clauses that maps every variable to a clause it appears in. In the case of not-all-equal 3SAT, for example, this follows from equivalence with two-colorability of three-uniform hypergraphs and from the work of Robertson et al. (1999). On the other hand, the proof of Theorem 6.17 identifies a variant that is NP-complete and remains so in the presence of a matching. We thus have the following.

COROLLARY 6.19. Generalized satisfiability is NP-complete, even when restricted to instances that have a matching and clauses of size five.

We leave a complete characterization for future work. While the proof techniques developed in this chapter will certainly be useful in this respect, it should be noted that the equivalence between generalized satisfiability and the pure equilibrium problem covered by Theorem 6.17 may fail to hold for instances of the latter where $p_i(a_N) = p_i(a'_N) = 0$ for action profiles a_N , a'_N such that $\#(1, a_{\hat{v}(i)}) = \#(1, a'_{\hat{v}(i)}) + 1$. For example, it would not be possible to show hardness of one-in-three 3SAT (Schaefer, 1978) using the same approach.

6.6 Mixed Equilibria

Let us now briefly look at the problem of finding a mixed equilibrium. The following theorem states that this problem is tractable in symmetric graphical games if the number of actions grows slowly in the neighborhood size. The proof relies on the fact that such games always have a symmetric equilibrium.

THEOREM 6.20. Let $\Gamma = (N, A^N, (p_i)_{i \in N})$ be a symmetric graphical game such that for all $i \in N$, $|A| = O(\log |v(i)| / \log \log |v(i)|)$. Then, a Nash equilibrium of Γ can be computed in polynomial time.

Proof. We show that Γ possesses a symmetric equilibrium, i.e., one where all players play the same (mixed) strategy, and that this equilibrium can be computed efficiently. For this, choose an arbitrary player $i \in N$ and construct a game $\Gamma_i = (N_i, A^{N_i}, (p_{i,j})_{j \in N})$ with players $N_i = \hat{\nu}(i)$, and payoff functions $p_{i,j}$ such that for all $j \in N_i$ and for action profiles $a_N \in A^N$ and $a'_{N_i} \in A^{N_i}$, $p_{i,j}(a'_{N_i}) = p_i(a_N)$ if $a'_j = a_i$ and for all $a \in A$, $\#(a, a'_{N_i}) = \#(a, a_{\nu(i)})$.

Since Γ is a symmetric graphical game, it is easily verified that Γ_i is a symmetric game, and must therefore possess a symmetric equilibrium, i.e., one where all the players in N_i play the same strategy. By a result of Papadimitriou and Roughgarden (2005), one such equilibrium s'_{N_i} can be computed in polynomial time if $|A| = O(\log |N'| / \log \log |N'|)$. Moreover, due to the symmetry of Γ , all the games Γ_i for $i \in N$ are isomorphic, and thus s'_{N_i} is an equilibrium in each of them.

6.7 · Discussion

Now define a strategy profile s_N of Γ by letting, for each $j \in N$, $s_j = s'_i$, and assume for contradiction that s_N is *not* an equilibrium. Then there exists a player $j \in N$ and some strategy $t_j \in \Delta(A)$ for this player such that $p_j(s_{-i}, t_i) > p_i(s_N)$. Further, by definition of $p_{i,j}$, $p_{i,i}(s'_{-i}, t) > p_{i,i}(s'_{N_i})$, contradicting the assumption that s'_{N_i} is an equilibrium of Γ_i .

This result applies in particular to the case where both the number of actions and the neighborhood size are bounded. Since the pure equilibrium problem in symmetric graphical games is NP-complete even in the case of two actions, we have identified a class of games where computing a mixed equilibrium is computationally easier than deciding the existence of a pure one, unless P=NP. A different class of games with the same property is implicit in Theorem 3.4 of Daskalakis and Papadimitriou (2005). It should be noted, on the other hand, that the existence of a symmetric equilibrium does not in general extend to games that are not anonymous or in which players have different payoff functions.

6.7 Discussion

In this chapter we have completely characterized the complexity of deciding the existence of a pure Nash equilibrium in games with bounded neighborhoods. This problem is NPcomplete in games with neighborhoods of size two, two actions, and two-valued payoff functions. For neighborhoods of size one it is NL-complete in general, and L-complete if additionally the number of actions grows only very slowly. Some additional cases become tractable for games that further satisfy the most restrictive type of anonymity considered in Chapter 5 within each neighborhood.

For the other types of anonymity, two neighbors again suffice for NP-hardness. While the construction used in the proof of Theorem 6.17 can be generalized to arbitrary neighborhoods of even size, it is unclear what happens for odd-sized neighborhoods. The extreme case when the neighborhood of every player consists of all other players yields ordinary anonymous games, in which the pure equilibrium problem is in P when the number of actions is bounded. It remains open at which neighborhood size the transition between membership in P and NP-hardness occurs. Another open problem concerns the complexity of the mixed equilibrium problem in anonymous graphical games. A promising direction for proving hardness would be to make the construction of Goldberg and Papadimitriou (2006) anonymous. Finally, as suggested in Section 6.5, it would be interesting to study the complexity of generalized satisfiability problems in the presence of matchings.

Chapter 7

Quasi-Strict Equilibria

Criticism directed at Nash equilibrium has been a recurring theme in previous chapters. In the remaining two chapters of the thesis we consider two more solution concepts, each of them trying to address a particular shortcoming.

Consider again the single-winner game introduced in Chapter 2, in which Alice, Bob, and Charlie select a winner among themselves using a protocol in which each of them has to raise their hand or not. The game is repeated in Figure 7.1. As we have already noted in Chapter 4, the only pure equilibrium of this game, in which Bob raises his hand while Alice and Charlie do not, is particularly weak. Both Bob and Charlie could deviate from their respective strategies to *any* other strategy without decreasing their chances of winning. After all, they cannot do any worse than losing. A similar property in fact applies to all pure equilibria of ranking games: there exists at least one player, namely the one ranked last in the equilibrium outcome, who receives his minimum payoff regardless of his choice of action.

To alleviate the effects of phenomena like this, Harsanyi (1973) suggests to impose the additional restriction that *every* best response of a player be played with positive probability. A Nash equilibrium satisfying this requirement is called a *quasi-strict equilibrium*.¹ Quasi-strict equilibrium is a *refinement* of Nash equilibrium in the sense that the set of quasi-strict equilibria of every game forms a subset of the set of Nash equilibria. This may also be beneficial with respect to another weakness of Nash equilibrium, its potential multiplicity. It turns out that the second equilibrium of the game in Figure 7.1, in which Bob raises his hand while Alice and Charlie randomize uniformly over their respective actions, is quasi-strict. Interestingly, Charlie plays a weakly dominated action with positive probability in this equilibrium.

Quasi-strict equilibria always exist in two-player games (Norde, 1999), but may fail to do so in games with more than two players. The game in Figure 4.9 on Page 48 shows that in fact, quasi-strict equilibria can already fail to exist in the very restricted class

¹Harsanyi originally referred to quasi-strict equilibrium as "quasi-strong". However, this term has been dropped to distinguish the concept from Aumann's (1959) strong equilibrium.



Figure 7.1: Single-winner game involving Alice (player 1), Bob (player 2), and Charlie (player 3), repeated from Figure 4.1. The dashed square marks the only pure equilibrium, dotted rectangles mark an equilibrium in which Alice and Charlie randomize uniformly over their respective actions. The latter is the unique quasi-strict equilibrium of the game.

of single-winner games.² In this chapter, we study the existence and the computational properties of quasi-strict equilibrium in zero-sum games, general normal-form games, and certain classes of anonymous and symmetric games. Section 7.3 focuses on two-player games, and it shown that quasi-strict equilibria, unlike Nash equilibria, have a unique support. We further give linear programs that characterize the quasi-strict equilibria in non-symmetric and symmetric zero-sum games. In Section 7.4 we turn to games with more than two players. We first identify new classes of games where a quasi-strict equilibrium is guaranteed to exist, and can in fact be found efficiently. An interesting example of such a class are symmetric games where every player has two actions. We then show that deciding the existence of a quasi-strict equilibrium in games with more than two players is NP-hard in general. This is in contrast to the two-player case, where the decision problem is trivial due to the existence result by Norde (1999).

7.1 Related Work

To address various drawbacks of Nash equilibrium, a number of concepts that single out particularly reasonable Nash equilibria—so-called equilibrium refinements—have been proposed over the years. A result by Norde et al. (1996) has however cast doubt upon this strand of research. The authors show that Nash equilibrium can be completely characterized by *payoff maximization in one-player games, consistency*,³ and *existence*. As a consequence, all common equilibrium refinements either violate consistency or existence. In particular, all refinements that are guaranteed to exist such as *perfect, proper*, and *persistent* equilibria suffer from inconsistency while consistent refinements such as

²There are only few examples in the literature for games without quasi-strict equilibria, essentially one example by van Damme (1983) and another one by Cubitt and Sugden (1994). For this reason, the game depicted in Figure 4.9 might be of independent interest.

³Consistency as introduced by Peleg and Tijs (1996) is defined as follows. Let s be a solution of game Γ and let Γ' be a reduced game where a subset of players are assumed to invariably play the strategies prescribed by s. A solution concept is *consistent* if the solution s' in which all of the remaining players still play according to s is a solution of Γ' .
7.2 · Preliminaries

quasi-strict, strong, and *coalition-proof* equilibria may fail to exist. Since consistency is a very intuitive and appealing property, its failure may be considered more severe than possible non-existence.

Harsanyi's quasi-strict equilibrium, which refines the Nash equilibrium concept by requiring that every action in the support yields strictly more payoff than actions not in the support, has been shown to always exist in generic n-player games, and thus in "almost every" game (Harsanyi, 1973), and in two-player games (Norde, 1999). Squires (1998) has further shown that quasi-strict equilibrium satisfies a set of axioms due to Cubitt and Sugden (1994), and is therefore very attractive from an axiomatic perspective. This result can be interpreted in such a way that the existence of a quasi-strict equilibrium is sufficient to justify the assumption of common knowledge of rationality. In fact, Quesada (2002) even poses the question whether the existence of quasi-strict equilibrium is sufficient for *any* reasonable justification theory. Finally, isolated quasi-strict equilibria satisfy almost all desirable properties defined in the refinements literature. They are essential, strongly stable, regular, proper, and strictly perfect (e.g., Jansen, 1987, van Damme, 1983, 1991). Using the framework of Peleg and Tijs (1996) and Norde et al. (1996), quasi-strict equilibrium can be characterized by consistency and *strict* payoff maximization in one-player games.

7.2 Preliminaries

The idea behind Nash equilibrium, as introduced in Definition 2.4 on Page 9, is that no player can increase his payoff by unilaterally changing his strategy. It follows directly from the definition that every action played with positive probability yields the same expected payoff. Quasi-strict equilibrium strengthens the equilibrium condition, and thus refines the Nash equilibrium concept, by requiring that the actions played with positive probability are *exactly* those maximizing a player's expected payoff, such that equilibrium actions must yield *strictly* more payoff than actions played with probability zero.

DEFINITION 7.1 (quasi-strict equilibrium). Let $\Gamma = (N, (A_i)_{i \in N}, (p_i)_{i \in N})$ be a normalform game. A Nash equilibrium $s_N^* \in S_N$ is called *quasi-strict equilibrium* if for all $i \in N$ and all $a_i, a'_i \in A_i$ with $s_i^*(a_i) > 0$ and $s_i^*(a'_i) = 0$, $p_i(s_{-i}^*, a_i) > p_i(s_{-i}^*, a'_i)$.

A pure equilibrium satisfying this property is usually called *strict*.

7.3 Two-Player Games

Using an elaborate construction based on Brouwer's fixed point theorem, Norde (1999) shows that quasi-strict equilibria always exists in two-player games. Since every quasi-strict equilibrium is also a Nash equilibrium, the problem of finding a quasi-strict equilibrium is intractable unless P=PPAD (Chen and Deng, 2006). The same is true for symmetric two-player games, because the symmetrization of Gale et al. (1950) preserves

Figure 7.2: Primal and dual linear programs for computing minimax strategies in zero-sum games

quasi-strictness (Jurg et al., 1992). For the restricted class of zero-sum games, however, quasi-strict equilibria, like Nash equilibria, can be found efficiently by linear programming. In contrast to Nash equilibria, quasi-strict equilibria have a unique support in these games.

THEOREM 7.2. In two-player zero-sum games, quasi-strict equilibria have a unique support, and one of them can be found in polynomial time.

Proof. It is known from the work of Jansen (1981) that every two-player game with a convex equilibrium set, and thus every two-player zero-sum game, possesses a quasi-strict equilibrium. We first establish that the support of a quasi-strict equilibrium must contain every action that is played with positive probability in some equilibrium of the game.

Consider a quasi-strict equilibrium (s_1, s_2) with payoff v for player 1, and assume for contradiction that there exists an action $a_1 \in A_1$ played with positive probability in some Nash equilibrium such that $s_1(a_1) = 0$. Then, by Definition 7.1, $p_1(a_1, s_2) < v$. Since a_1 is in the support of some Nash equilibrium, and by the exchangeability of equilibrium strategies in zero-sum games, however, it holds that $p_1(a_1, s_2) = v$, a contradiction.

In order to compute quasi-strict equilibria, consider the two standard linear programs for finding the minimax strategies for player 1 and 2, respectively, given in Figure 7.2 (e.g., Luce and Raiffa, 1957). It is well-known from the minimax theorem (von Neumann, 1928), and also follows from LP duality, that the value v of the game is identical and unique in both cases. We can thus construct a linear feasibility program that computes equilibrium strategies for both players by simply merging the sets of constraints and omitting the minimization and maximization objectives.

Now, quasi-strict equilibrium requires that action a yields strictly more payoff than action a' of the same player if and only if a is in the support and a' is not. For a zero-

maximize ϵ			
subject to			
$\sum_{\mathfrak{a}_2 \in \mathcal{A}_2} s_2(\mathfrak{a}_2) p(\mathfrak{a}_1, \mathfrak{a}_2)$	\leqslant	ν	$\forall a_1 \in A_1$
$s_2(a_2)$	\geq	0	$\forall a_2 \in A_2$
$\sum_{a_2 \in A_2} s_2(a_2)$	=	1	
$s_1(a_1) + v - \sum_{a_2 \in A_2} s_2(a_2) p(a_1, a_2) - \epsilon$	\geq	0	$\forall a_1 \in A_1$
$\sum_{a_1\in A_1} s_1(a_1) p(a_1, a_2)$	\geq	ν	$\forall a_2 \in A_2$
$s_1(a_1)$	\geq	0	$\forall a_1 \in A_1$
$\sum_{a_1 \in A_1} s_1(a_1)$	=	1	
$s_2(a_2) + v - \sum_{a_1 \in A_1} s_1(a_1) p(a_1, a_2) - \epsilon$	\geq	0	$\forall a_2 \in A_2$

Figure 7.3: Linear program for computing quasi-strict equilibria in zero-sum games

sum game with value v this can be ensured by requiring that for every action $a_1 \in A_1$ of player 1, $s_1(a_1) + v > \sum_{a_2 \in A_2} s_2(a_2) p(a_1, a_2)$. If a_1 is not in the support, i.e., $s_1(a_1) = 0$, then it must yield strictly less payoff than the value of the game. If, on the other hand a_1 is in the support, i.e., $s_1(a_1) > 0$, these constraints do not impose any restrictions given that the strategy profile s_N is indeed an equilibrium with value v. The latter is ensured by the remaining constraints. We finally add another variable ϵ to be maximized to the right hand side of the above inequality, to turn the strict inequality into a weak one. Due to the existence of at least one quasi-strict equilibrium, we are guaranteed to find a solution with positive ϵ . The resulting linear program is given in Figure 7.3.

We proceed by showing that every symmetric zero-sum game has a *symmetric* quasistrict equilibrium. This result stands in contrast to Theorem 7.4 in Section 7.4, which shows that the same need *not* be the case for symmetric two-player games in general.

THEOREM 7.3. Every symmetric two-player zero-sum game has a symmetric quasistrict equilibrium which can be found in polynomial time.

Proof. By Theorem 7.2, the support of any quasi-strict equilibrium contains all actions that are played with positive probability in some equilibrium. Clearly these actions coincide for both players in symmetric games, and any minimax strategy using these actions constitutes a symmetric equilibrium. Since both players can enforce their minimax value using the same strategy in a symmetric zero-sum game, the value of the game has to be zero. Using this information, the linear program in Figure 7.3 can be simplified significantly. The resulting linear program is shown in Figure 7.4.

7.4 A Hardness Result for Multi-Player Games

In games with three or more players, the existence of a quasi-strict equilibrium is no longer guaranteed. However, there are very few examples in the literature for games

```
\begin{array}{lll} \mbox{maximize $\varepsilon$} \\ \mbox{subject to} \\ & \sum_{a_2 \in A_2} s(a_2) \ p(a_1, a_2) & \leqslant & 0 & \forall a_1 \in A_1 \\ s(a_2) & \geqslant & 0 & \forall a_2 \in A_2 \\ & \sum_{a_2 \in A_2} s(a_2) & = & 1 & \\ s(a_1) - \sum_{a_2 \in A_2} s(a_2) \ p(a_1, a_2) - \varepsilon & \geqslant & 0 & \forall a_1 \in A_1 \end{array}
```

Figure 7.4: Linear program for computing quasi-strict equilibria in symmetric zero-sum games

	c ¹			c^2		
	b ¹ b ²			b^1	b^2	
a ¹	(1,1,0)	(0,1,1)		(0,1,1)	(1,0,1)	
a ²	(0,1,1)	(1,0,1)		(1,0,1)	(1, 1, 0)	

Figure 7.5: Somebody has to do the dishes, via a self-anonymous single-loser game. Players 1, 2, and 3 choose rows, columns, and matrices, respectively. In the only Nash equilibrium of the game player 1 plays his second action, player 2 plays his first action, and player 3 randomizes uniformly over both of his actions. The game does not have a quasi-strict equilibrium.

without quasi-strict equilibria.⁴ An important question is of course which natural classes of games always possess a quasi-strict equilibrium. We have seen in Chapter 4 that this is not the case for the class of *single-winner* games, which require that all outcomes are permutations of the payoff vector (1, 0, ..., 0).

In the following, we look at anonymous and symmetric games. It turns out that self-anonymous games, and thus also anonymous games, need *not* possess a quasi-strict equilibrium. For this, consider the following protocol used by Alice, Bob, and Charlie to select one among them to do the dishes. Each of them decides to raise their hand or not, simultaneously and independently of the others. Alice loses, and has to do the dishes, if exactly one player raises his hand, Bob loses if exactly two players raise their hands, and Charlie loses if either all or no players raise their hand. The resulting anonymous single-loser game, depicted in Figure 7.5, exhibits some peculiar phenomena, some of which may be attributed to the absence of quasi-strict equilibria. For example, the security level of all players is 1/2, and the expected payoff in the *only* Nash equilibrium, which has Alice raise her hand and Charlie randomize with equal probability, is (1/2, 1/2, 1). The minimax

⁴To our knowledge, there are three examples in addition to the one we will give below (van Damme, 1983, Kojima et al., 1985, Cubitt and Sugden, 1994). All of them involve three players with two actions each, which clearly is optimal.

 $p_{m-1,0}$	p _{m0}		p _{m'0}	$p_{m'+1,0}$	
\wedge	II	•••	II	\vee	
 ₽m−1,1	p _{m1}		p _{m'1}	$p_{m'+1,1}$	

Figure 7.6: Payoff structure of a symmetric game with two actions

strategies of Alice and Bob, however, are different from their equilibrium strategies, i.e., they can *guarantee* their equilibrium payoff by *not* playing their respective equilibrium strategies.⁵ Furthermore, the unique equilibrium is not quasi-strict, i.e., Alice and Bob could as well play any other action without jeopardizing their payoff.

For symmetric and self-symmetric games, on the other hand, the picture is different. Self-symmetric games form a subclass of common-payoff games, where the payoff of all players is identical in every outcome. Starting from an outcome with maximum payoff, a quasi-strict equilibrium can be found by iteratively adding other actions to the support by which some player, and thus all players, also obtain the maximum payoff.

It follows from a theorem by Nash (1951) that every symmetric game has a symmetric Nash equilibrium, i.e., a Nash equilibrium where all players play the same strategy. We can use this result to show that quasi-strict equilibria are guaranteed to exist in symmetric games with two actions for each player. Perhaps surprisingly, all quasi-strict equilibria of such a game may themselves be asymmetric.

THEOREM 7.4. Every symmetric game with two actions for each player has a quasistrict equilibrium. Such an equilibrium can be found in polynomial time.

Proof. Let $\Gamma = (N, \{0, 1\}^N, (p_i)_{i \in N})$ be a symmetric game. For an action profile a_N and an action $a \in \{0, 1\}$, denote by $\#(a, a_N) = |\{i \in N : a_i = a\}|$ the number of players who play a in a_N . It follows from Definition 5.1 that there exist 2n numbers $p_{m\ell} \in \mathbb{R}$ for $0 \leq m \leq n-1$ and $\ell \in \{0, 1\}$, such that for all $i \in N$, $p_i(a_N) = p_{m\ell}$ whenever $a_i = \ell$ and $\#(1, a_{-i}) = m$. We can further assume without loss of generality that $p_{00} = p_{01}$ and $p_{n-1,0} \geq p_{n-1,1}$, and that $p_{m0} \neq p_{m1}$ for some m. To see this, recall that Γ , being a symmetric game, must possess a symmetric equilibrium s_N (Nash, 1951), which we can assume to be the pure strategy profile where all players play action 0 with probability 1. If instead all players played both of their actions with positive probability, s_N would itself be quasi-strict. Now, if one of the former two equations was not satisfied, then one of the two symmetric pure strategy profiles would be a quasi-strict equilibrium. If the latter condition would not to hold, there would exist a quasi-strict equilibrium where all players randomize between their actions.

It is now easily verified that there must exist m and m', $0 \le m \le m' \le n-1$, such that (i) $p_{m''0} = p_{m''1}$ for all m'' with $m \le m'' \le m'$, (ii) either m = 0 or $p_{m-1,0} < m'$

⁵Similar phenomena were also observed by Aumann (1985).

 $p_{m-1,1}$, and (iii) either m' = n - 1 or $p_{m'+1,0} > p_{m'+1,1}$. This situation is illustrated in Figure 7.6. We further claim that any strategy profile s_N where n - m' - 1 players play action 0, m players play action 1, and the remaining m' - m + 1 players randomize between both of their actions, is a quasi-strict equilibrium of Γ . First consider any player $i \in N$ who plays action 0 with probability 1. Then for any action profile a_N played with positive probability in s_N , $m \leq \#(1, a_{-i}) \leq m + (m' - m + 1) = m' + 1$. As a consequence, player i weakly prefers action 0 over action 1 in any such action profile, and the preference is strict for at least one such action profile. Now consider any player $i \in N$ who plays action 1 with probability 1. Then for any action profile a_N played with positive probability in s_N , $m - 1 \leq \#(1, a_{-i}) \leq (m - 1) + (m' - m + 1) = m'$. Player i weakly prefers action 1 over action 0 in any such action profile, and the preference is strict for at least one such action profile. Finally consider any player $i \in N$ who plays both actions with positive probability. It then holds for any action profile a_N played with positive probability in s_N that $m \leq \#(1, a_{-i}) \leq m + (m' - m) = m'$, and player i is indifferent between actions 0 and action 1 in any such action profile.

Since a symmetric equilibrium of a symmetric game with a constant number of actions can be found in polynomial time (Papadimitriou and Roughgarden, 2005), and since the proof of the first part of the theorem is constructive, the second part follows immediately.

We leave it as an open problem whether all symmetric games contain a quasi-strict equilibrium. If the symmetrization procedure due to Gale et al. (1950) can be extended to multi-player games while still preserving quasi-strictness, a counterexample could be constructed from one of the known games without quasi-strict equilibria. Of course, in light of Theorem 7.4, the number of actions per player in such a counter-example has to be greater than two, and may very well be substantially greater than that.

We conclude this chapter by showing that the existence of a quasi-strict equilibrium is NP-hard to decide in general.

THEOREM 7.5. Deciding whether a game in normal form possesses a quasi-strict equilibrium is NP-hard, even if there are only three players and a constant number of payoffs.

Proof. For hardness, we give a reduction from the NP-complete problem CLIQUE (e.g., Papadimitriou, 1994a) reminiscent of a construction used by McLennan and Tourky (2005) to show NP-hardness of various problems related to Nash equilibria in two-player games. Given an undirected graph G = (V, E) and a positive integer $k \leq |E|$, CLIQUE asks whether G contains a clique of size at least k, i.e., a subset $C \subseteq V$ such that $|C| \geq k$ and for all distinct $v, w \in C$, $(v, w) \in E$. Given a particular CLIQUE instance ((V, E), k)with $V = \{1, 2, ..., m\}$, we construct a game Γ with three players, actions $A_1 = \{a^i : i \in V\} \cup \{a^0\}, A_2 = \{b^i : i \in V\} \cup \{b^0\}$ and $A_3 = \{c^1, c^2\}$. The payoff functions p_i are defined as follows, and illustrated in Figure 7.7. If player 3 plays c^1 and players 1 and 2



Figure 7.7: Three-player game Γ used in the proof of Theorem 7.5

play a^i and b^j , respectively, for $i, j \in V$, payoffs are given by a matrix $(m_{ij})_{i,j\in V}$ defined according to G, and by the identity matrix $(e_{ij})_{i,j\in V}$. More precisely,

$$\mathfrak{m}_{\mathfrak{i}\mathfrak{j}} = \begin{cases} 1 & \text{if } (\mathfrak{i},\mathfrak{j}) \in \mathsf{E}, \\ 0 & \text{if } \mathfrak{i} = \mathfrak{j}, \\ -1 & \text{otherwise,} \end{cases}$$

and

$$e_{ij} = egin{cases} 1 & ext{if } i=j, \ 0 & ext{otherwise.} \end{cases}$$

If player 3 instead plays c^2 , he obtains a payoff of K = (2k - 3)/2k. We claim that Γ possesses a quasi-strict equilibrium if and only if there exists a clique of size at least k in G.

For the direction from right to left, assume that there exists a maximal clique $C \subseteq V$ with $|C| \ge k$, and consider the strategy profile s_N with $s_1(c^1) = 1$, and $s_1(a^i) = s_2(b^i) = 1/|C|$ if $i \in C$ and $s_1(a^i) = s_2(b^i) = 0$ otherwise. By construction of Γ , and for all $i \in V \cup \{0\}$, $p_2(s_{-2}, b^i) < p_2(s_N)$ whenever $s_1(a^i) = 0$. Furthermore, by maximality of C, $p_1(s_{-1}, a^i) < p_1(s_N)$ for all $i \notin C$. Finally, $p_3(s_N) = (k-1)/k > (2k-3)/2k = p_3(s_{-3}, c^2)$. Thus, s_N is a quasi-strict equilibrium of Γ .

For the direction from left to right, consider a quasi-strict equilibrium s_N of Γ , and assume for contradiction that G does not have a clique of size at least k. In equilibrium, for all b, b' $\in A_2$, we must have $p_2(s_{-2}, b) = p_2(s_{-2}, b')$ whenever $s_2(b) > 0$ and $s_2(b') > 0$, and thus, for all a, a' $\in A_1$, $s_1(a) = s_1(a')$ whenever $s_1(a) > 0$ and $s_1(a') > 0$. As a consequence, for s_N to be quasi-strict, $s_2(b^i) > 0$ whenever $s_1(a^i) > 0$, for all $i \in V \cup \{0\}$. First consider the case when $s_3(c^1) > 0$. If $s_1(a^0) = s_2(b^0) = 1$, s_N would not be quasistrict for player 1. If on the other hand $s_1(a^i) > 0$ or $s_2(b^i) > 0$ for some $i \in V$, then there would have to be a set $C \subseteq V$ with $|C| \ge k$, such that for all $i \in V$ with $s_1(a^i) > 0$ and all $j \in C \setminus \{i\}$, $p_1(a^i, b^j, c^1) = 1$. By construction of Γ , C would be a clique of size at least k in G, a contradiction. Now consider the case when $s_3(c^2) = 1$. If $s_1(a^0) = 1$ or $s_2(b^0) = 1$, s_N would not be quasi-strict for player 3. If, on the other hand, $s_1(a^i) > 0$ and $s_2(b^j) > 0$ for some $i, j \in V$, then player 1 could change his action to a^0 to get a higher payoff, a contradiction.

It follows that the problem of *finding* a quasi-strict equilibrium in games with more than two players is NP-hard, under polynomial-time Turing reductions. It is unlikely that the same is true for Nash equilibrium.

7.5 Discussion

In this chapter we investigated the computational properties of an attractive equilibrium refinement known as quasi-strict equilibrium. It turned out that quasi-strict equilibria in zero-sum games have a unique support and can be computed efficiently via linear programming. In games with more than two players, deciding the existence of a quasi-strict equilibrium quickly becomes NP-hard.

As pointed out in the beginning of this chapter, classes of games that always admit a quasi-strict equilibrium, such as two-player games, are of vital importance to justify rational play based on elementary assumptions. We specifically looked at anonymous and symmetric games, and found that self-anonymous games, and thus also anonymous games, need not possess a quasi-strict equilibrium, while symmetric games with two actions for each player always do. It is an interesting question whether the latter is still true when there are more than two actions. Other classes of multi-player games where one might hope for the existence of a quasi-strict equilibrium include unilaterally competitive games (Kats and Thisse, 1992), potential games (Monderer and Shapley, 1996), graphical games with bounded neighborhood, and single-winner games where all players have a positive security level.

Chapter 8

Shapley's Saddles

Another, perhaps more severe, problem of Nash equilibrium is that its existence generally requires randomization on behalf of the players. It is not clear if and how players can execute such randomizations in an exact and reliable way, which in turn is essential for the notion of stability underlying Nash equilibrium. The problem becomes more severe in games with more than two players, where randomization with irrational probabilities may be required.

In work dating back to the early 1950s, Shapley proposed ordinal set-valued solution concepts for zero-sum games that he refers to as *saddles* (Shapley, 1953a,b, 1955, 1964). What makes these concepts intuitively appealing is that they are based on the elementary notions of dominance and stability. Call a *generalized saddle point* (GSP) is a tuple of subsets of actions, one for each player, such that every action *not* contained in the GSP is dominated by some action in the GSP, given that the remaining players only play actions from the GSP. A *saddle* then is an inclusion-minimal GSP, i.e., a GSP that contains no other GSP. By varying the underlying notion of dominance, one obtains strict and weak saddles. Shapley (1964) showed that every (two-player) zero-sum game admits a *unique* strict saddle. Duggan and Le Breton (1996) proved that the same is true for the weak saddle in a certain subclass of symmetric zero-sum games.

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. In this chapter, we provide polynomial-time algorithms for computing strict saddles in normal-form games with any number of players, and weak saddles in the subclass of symmetric zero-sum games considered by Duggan and Le Breton (1996). We note, but do not show here, that these results extend to mixed refinements of Shapley's saddles introduced by Duggan and Le Breton (2001). On the other hand, certain problems associated with (pure and mixed) weak saddles in two-player games, such as deciding whether there exists a weak saddle with k actions for some player, are shown to be NP-hard. We restrict our attention to games with a constant number of players, and assume throughout the chapter that these games are given explicitly, i.e., as a multi-dimensional table of payoffs.

8.1 Related Work

Shapley's saddles are based on the notion of dominance, which has also been studied from a computational perspective (e.g., Knuth et al., 1988, Conitzer and Sandholm, 2005a,b, and Chapters 4 and 5 of this thesis). Iterated strict dominance, for example, is one of the few examples of a solution concept known to be efficiently computable in general games. Strict saddles can be seen as a refinement of iterated strict dominance in that all strict saddles of a normal-form game are contained in the subgame that one obtains by iterated elimination of strictly dominated actions.

A related set-valued solution concept are *CURB sets* (Basu and Weibull, 1991). Unlike Shapley's saddles, however, CURB sets are based on randomized strategies. They are therefore not ordinal, and are subject to some of the complications associated with randomized strategies. Every strict saddle represents the support of a CURB set, and thus contains the support of a minimal CURB set. In confrontation games, as defined in Section 8.4, the support of a minimal CURB set and the strict saddle trivially coincide. Moreover, in this particular class of games, the strict mixed saddle is identical to the support of the minimal CURB set when only allowing pure strategies. There appears to be no such relationship between *weak* saddles and CURB sets. Benisch et al. (2006) have recently proposed a polynomial-time algorithm for computing the CURB sets of a two-player game.

8.2 Preliminaries

Existence of a Nash equilibrium is not guaranteed if strategies are required to be pure, as the Matching Pennies game in Figure 2.2 on Page 9 illustrates. Requiring randomization in order to reach a stable outcome, however, is problematic for various reasons. A possible solution is to consider *set-valued* solution concepts that identify, for each player $i \in N$, a subset $X_i \subseteq A_i$ of his actions, such that the vector (X_1, X_2, \ldots, X_n) satisfies some notion of stability. Shapley's saddles generalize strict Nash equilibrium¹ by requiring that for every action $a_i \in A_i \setminus X_i$ of a player $i \in N$ that is *not* included in X_i , there should be some reason for its exclusion, namely an action in X_i that is strictly better than a_i .

To formalize this idea, we need some notation. Henceforth, let $A_N = (A_1, A_2, \ldots, A_n)$. For $X_N = (X_1, X_2, \ldots, X_n)$, write $X_N \subseteq A_N$, and say that X_N is a subset of A_N , if $\emptyset \neq X_i \subseteq A_i$ for all $i \in N$. Further let $X_{-i} = (X_1, \ldots, X_{i-1}, X_{i+1}, \ldots, X_n)$. Finally consider a generalization of the notions of strict and weak dominance to specific sets of actions of the other players. For a player $i \in N$ and two actions $a_i, b_i \in A_i$, say that a_i strictly dominates b_i with respect to $X_{-i} \subseteq A_{-i}$, denoted $a_i \gg_{X_{-i}} b_i$, if $p(a_i, x_{-i}) > p(b_i, x_{-i})$ for all $x_{-i} \in X_{-i}$, and that a_i weakly dominates b_i with respect to X_{-i} .

¹Recall that a pure Nash equilibrium is called *strict* if every player strictly loses when deviating from his equilibrium action.

	b^1	b ²	b ³	
a ¹	(3,-3)	(3, -3)	(4, -4)	
a ²	(2, -2)	(3, -3)	(3, -3)	
a ³	(1,-1)	(2, -2)	(3, -3)	
a ³	(2, -2)	(1,-1)	(5, -5)	

Figure 8.1: Strict and weak saddles of a zero-sum game, respectively indicated by dashed and dotted boxes

denoted $a_i >_{X_{-i}} b_i$, if $p(a_i, x_{-i}) \ge p(b_i, x_{-i})$ for all $x_{-i} \in X_{-i}$, with at least one strict inequality. We are now ready to define strict and weak saddles formally.

DEFINITION 8.1 (strict and weak saddle). Let $\Gamma = (N, (A_i)_{i \in N}, (p_i)_{i \in N})$ be a game, $X_N = (X_1, X_2, \dots, X_n) \subseteq A_N$. Then, X_N is a generalized saddle point (GSP) of Γ if for all $i \in N$ and

for all
$$a_i \in A_i \setminus X_i$$
, there exists $x_i \in X_i$ such that $x_i \gg_{X_{-i}} a_i$. (8.1)

A GSP is called *strict saddle* if it does not contain a GSP as a proper subset.

Similarly, X_N is a *weak generalized saddle point (WGSP)* of Γ if for all $i \in N$ and

for all $a_i \in A_i \setminus X_i$, there exists $x_i \in X_i$ such that $x_i >_{X_{-i}} a_i$. (8.2)

A WGSP is called *weak saddle* if it does not contain a WGSP as a proper subset.

The underlying intuition is that every player $i \in N$ has a distinguished set X_i of actions such that for every action $a_i \notin X_i$, there is some action in X_i that dominates a_i , provided that the other players play only actions from their distinguished sets. Both strict and weak saddle are *ordinal* solution concepts, i.e., they are invariant under order-preserving transformations of the payoff functions. This is in contrast to Nash equilibrium, for which invariance holds only under *positive affine* transformations.

Properties (8.1) and (8.2) are sometimes referred to as *external stability*. Using this terminology, a (W)GSP is a tuple of externally stable sets. Since strict dominance implies weak dominance, every strict saddle is a WGSP and thus contains a weak saddle. Consider for example the two-player zero-sum game shown in Figure 8.1. The pair $X_N = (\{a^1, a^2\}, \{b^1, b^2\})$ is a strict saddle and a WGSP. Since a^1 weakly dominates a^2 with respect to $\{b^1, b^2\}$, and both b^1 and b^2 dominate b^3 with respect to a^1 , the pair $X'_N = (\{a^1\}, \{b^1, b^2\})$ also is a WGSP. Indeed, X'_N does not contain a smaller WGSP, and therefore is a weak saddle. Some reflection reveals that X_N and X'_N are in fact the *unique* strict and weak saddles of the game.

It is easy to see that every game has both a strict and a weak saddle. By definition, the tuple A_N is a GSP. Furthermore every GSP that is not a saddle must contain a GSP

	a ¹	a^2	a ³	a^4
a ¹	(0,0)	(1,-1)	(0, 0)	(0, 0)
a ²	(-1,1)	(0, 0)	(0, 0)	(0, 0)
a ³	(0,0)	(0, 0)	(0, 0)	(1,-1)
a ⁴	(0,0)	(0, 0)	(-1, 1)	(0, 0)

Figure 8.2: Symmetric zero-sum game with multiple weak saddles

that is strictly smaller. Finiteness of A_N implies the existence of a minimal GSP, i.e., a strict saddle. An analogous argument applies to the weak saddle. Strict saddles are *unique* in two-player zero-sum games but not in general games, whereas weak saddles are not even unique in the former class. To see this, consider the symmetric zero-sum game of Figure 8.2. It is easily verified that each of the following pairs is a weak saddle of this game: $(\{a^1, a^2\}, \{a^1, a^2\}), (\{a^3, a^4\}, \{a^3, a^4\}), (\{a^1, a^3\}, \{a^1, a^3\}), (\{a^2, a^3\}, \{a^1, a^4\}), and (\{a^1, a^4\}, \{a^2, a^3\}).$

8.3 Strict Saddles

Shapley (1964) shows that every two-player zero-sum game possesses a unique strict saddle, because the set of GSPs in such games is closed under intersection. He further describes an algorithm, attributed to Harlan Mills, to compute this saddle. The idea behind this algorithm is that given a subset of the saddle, the saddle itself can be computed by iteratively adding actions that are maximal, i.e., not dominated with respect to the current subset of actions of the other player. Shapley further points out that the strict saddle must contain all actions possibly leading to a *minimax* or *maximin* point. Note, however, that being able to find a subset of the saddle is not crucial for the above method. Starting the algorithm from singleton sets of actions, and invoking it for every combination of such sets for the different players, yields a number of candidates for the strict saddle. The strict saddle itself can then easily be identified as the inclusion-minimal set among these candidates. Correctness of this procedure follows from the fact that every candidate set is a GSP, and that the unique strict saddle is contained in every GSP. Furthermore, Mills' iterative procedure is invoked only a polynomial number of times.

In general games, strict saddles are no longer unique. For example, consider the twoplayer *coordination game* in which both players have two actions and receive a payoff of one if they play the same action, and a payoff of zero otherwise. Obviously, this game has two strict saddles, one where both players play their first action, and one where both of them play their second action. From a computational point of view, the existence of multiple strict saddles turns out not to have any serious consequences. THEOREM 8.2. All strict saddles of an n-player game can be computed in polynomial time.

Proof. We generalize Mills' algorithm and show that the inclusion-minimal GSP containing a given input set X_N^0 can be computed in polynomial time. The statement of the theorem then follows by similar arguments as before: we can invoke the algorithm a polynomial number of times, namely for each vector containing for each player a singleton set of actions, and then select the inclusion minimal elements from the resulting set of candidate GSPs.

Let $\Gamma = (N, (A_i)_{i \in N}, (p_i)_{i \in N})$ be a game, $A_N = (A_i)_{i \in N}$, and $X_N^0 \subseteq A_N$, and consider the following algorithm:

- 1. Let $X_N = X_N^0$.
- While there exists a player i ∈ N and an action a_i ∈ A_i \ X_i that is not dominated with respect to X_{-i}, i.e., one for which there is no b_i ∈ A_i with b_i ≫_{X_{-i}} a_i, add a_i to X_i.
- 3. Return X_N .

Clearly, we can check in polynomial time whether a given action is dominated with respect to particular subsets of actions of the other players. Since there is only a polynomial number of actions, and one of them is added to X_N in every iteration, the algorithm runs in polynomial time.

It remains to be shown that the algorithm returns the inclusion minimal GSP containing X_N^0 . To see this, let X_N^{min} be the minimal GSP containing X_N^0 . We show that $X_N \subseteq X_N^{min}$ always holds during the execution of the algorithm. When the algorithm terminates, all elements outside X_N are dominated, meaning that X_N is a GSP. At the beginning of the algorithm, $X_N = X_N^0 \subseteq X_N^{min}$ by definition of X_N^{min} . Now assume that $X_N \subseteq X_N^{min}$ holds at the beginning of a particular iteration, and let A'_i be the set of actions of player i not dominated with respect to X_{-i} , i.e., $A'_i = \{a_i \in A_i \setminus X_i :$ there is no $b_i \in A_i$ such that $b_i \gg_{X_{-i}} a_i\}$. Let $a_i \in A'_i$ be the element added to X_i , and assume for contradiction that $a_i \notin X_i^{min}$. Since X_N^{min} is a GSP, there must exist an action $b_i \in X_i^{min}$ with $b_i \gg_{X_{-i}} a_i$. By the induction hypothesis, $X_{-i} \subseteq X_{-i}^{min}$, which in turn implies $b_i \gg_{X_{-i}} a_i$. This contradicts the assumption that $a_i \in A'_i$.

A similar result can be obtained for the mixed strict saddle, where external stability is defined in terms of dominance by mixed strategies.

8.4 Weak Saddles of Confrontation Games

The computation of *weak* saddles turns out be significantly more complicated than that of strict saddles. Somewhat surprisingly, this is the case even in two-player zero-sum

games. In particular, Mills' algorithm does not easily generalize to weak saddles. When it is invoked on the game of Figure 8.1, for example, and initialized with $X_N^0 = (\{a^1\}, \{b^2\})$, it might add action a^2 in the next step, which is not contained in a minimal WGSP containing X_N^0 . In this section we consider a subclass of *symmetric* zero-sum games that is guaranteed to possess a unique weak saddle. The complexity of the problem for general zero-sum games remains open.

It will be instructive to view a symmetric zero-sum game from a slightly different perspective, namely that of comparison functions as introduced by Dutta and Laslier (1999). Let A be a finite set of *alternatives*. A comparison function on A is a function $g : A \times A \rightarrow \mathbb{R}$ such that for all $x, y \in A$, g(x, y) = -g(y, x). This obviously implies that g(x, x) = 0 for all $x \in X$. For two alternatives $x, y \in A$, the value g(x, y) can be interpreted as the extent to which x is "better" than y.

The subclass of comparison functions g satisfying $g(x, y) \in \{-1, 1\}$ for all $x, y \in A$ with $x \neq y$ corresponds to *tournaments*, i.e., complete and asymmetric² relations, on A. Tournaments for example arise if an odd number of agents with strict preferences conduct a majority vote for each pair of alternatives, or from coalitional games with the property that either a coalition or its complement is winning. A very active subarea of social choice theory is concerned with the study of *tournament solutions* that map each tournament to a nonempty set of "desirable" alternatives (e.g., Laslier, 1997). A particularly appealing tournament solution is the *minimal covering set* of Dutta (1988), which was subsequently generalized to comparison functions by Dutta and Laslier (1999).

An element x is said to cover another element y if it is better according to g, and if for every third element z, x compares at least as well against z as y does.

DEFINITION 8.3 (covering). Let A be a set of alternatives, g a comparison function on A, and $B \subseteq A$. Then, for $x, y \in B$, x covers y in B if g(x, y) > 0 and g(x, z) > g(y, z) for all $z \in B \setminus \{x, y\}$.

We will call the *uncovered set* of a comparison function g on a set B, denoted UC(B,g), the set of alternatives in B not covered by any alternative in B. A covering set then is a set of alternatives that is internally and externally stable, in the appropriate sense, with respect to the covering relation.

DEFINITION 8.4 (covering set). Let A be a set of alternatives, g a comparison function on A, and $B \subseteq A$. Then, B is a covering set for g in A if

- (i) UC(B,g) = B and
- (*ii*) for all $x \in A \setminus B$, $x \notin UC(B \cup \{x\}, g)$.

Dutta and Laslier (1999) show that the intersection of any two externally stable sets is nonempty and itself externally stable, and that the unique inclusion minimal such set also

²A relation $R \subseteq A \times A$ is called asymmetric if for $x, y \in A$ with $x \neq y$, xRy implies that not yRx.

satisfies internal stability. For any set A of alternatives and any comparison function g on A there thus exists a unique minimal covering set MC(A,g).

We will now show that the minimal covering set of a comparison function can be computed efficiently, and then use a result about the relationship between the minimal covering set and the weak saddle to obtain a tractability result for the latter in a restricted class of games. Our algorithm for computing the minimal covering set proceeds in a similar fashion as that for the strict saddle: start with a subset of the minimal covering set and then iteratively add elements outside the current set that are still uncovered. There are two problems, however. For one, it is not clear how to find a subset of the minimal covering set. Secondly, it is important only to add elements that may not be covered in a later iteration, and it is not obvious which elements these should be.

A solution known to yield a subset of the minimal covering set is the essential set of Dutta and Laslier (1999), which generalizes the bipartisan set of Laffond et al. (1993). The essential set makes a connection back to game theory by considering, for a comparison function g on a set A of alternatives, a symmetric zero-sum game ($\{1, 2\}, (A, A), (g, -g)$). The intuition underlying this so-called *tournament game* is that each of two parties proposes an alternative for adoption, and the better of the two, according to the comparison function, gets selected. The essential set itself is defined as the set of actions that are played with positive probability in *some* Nash equilibrium of the tournament game. Since the tournament game is a symmetric zero-sum game, it suffices to consider symmetric equilibria.

DEFINITION 8.5 (essential set). Let A be a set of alternatives, g a comparison function on A. Then, the essential set of g is given by

$$ES(A,g) = \{x \in A : s_1(x) > 0, (s_1,s_1) \text{ is a Nash equilibrium of} \\ \text{the game } (\{1,2\}, (A,A), (g,-g)) \}$$

By Theorem 7.2, the essential set equals the support of the unique *quasi-strict* equilibrium of the tournament game and can be computed in polynomial time by solving a linear program.

The key observation that will allow us to overcome the second problem above is that any element in the *minimal covering set* of the set of elements not covered by the members of a subset of the minimal covering set must itself be in the minimal covering set.

LEMMA 8.6. Let A be a set of alternatives, g a comparison function on A. Further let $B \subseteq MC(A, g)$ and $A' = \bigcup_{a \in A \setminus B} (UC(B \cup \{a\}, g) \cap \{a\})$. Then, $MC(A', g) \subseteq MC(A, g)$.

Proof. Partition A', the set of alternatives not covered by B, into two sets C and C' of elements contained in MC(A, g) and elements not contained in MC(A, g), i.e., C = $A' \cap MC(A, g)$ and C' = $A' \setminus MC(A, g)$. We will show that C is externally stable for A'. Since MC(A', g) must lie in the intersection of all sets that are externally stable for A', this means that $MC(A', g) \subseteq MC(A, g)$.

In the following, we will use an easy consequence of the definition of the covering relation: for two sets X, X' with $X \subseteq X' \subseteq A$ and two alternatives $x, y \in X$, if y covers x in X', then y also covers x in X. We will refer to this property as "covering in subsets".

Let $x \in C'$. Since $x \notin MC(A, g)$, there has to be some $y \in MC(A, g)$ that covers xin $MC(A, g) \cup \{x\}$. It is easy to see that $y \notin B$. Otherwise, since $B \subseteq MC(A, g)$ and by covering in subsets, y would cover x in $B \cup \{x\}$, contradicting the assumption that $x \in A'$. On the other hand, assume that $y \in MC(A, g) \setminus (B \cup C)$. Since y covers x in $MC(A, g) \cup \{x\}$, since $B \subseteq MC(A, g)$, and by covering in subsets, we have that g(y, x) > 0 and for all $w \in B$, $g(y, w) \ge g(x, w)$. Furthermore, since $y \notin A'$, there has to be some $z \in B$ covering y in $B \cup \{y\}$, i.e., one such that g(z, y) > 0 and for all $w \in B$, $g(z, w) \ge g(y, w)$. Combining the two, we get g(z, x) > 0 and for all $w \in B$, g(z, w), i.e., z covers xin $B \cup \{x\}$. This again contradicts the assumption that $x \in A'$. It thus has to be the case that $y \in C$. Since $C \subseteq MC(A, g)$, it follows from covering in subsets that y also covers xin $C \cup \{x\}$. We have shown that for every $x \in C'$, there exists $y \in C$ such that y covers xin $C \cup \{x\}$, i.e., C is externally stable for A'.

Note that this lemma directly yields a recursive algorithm for computing the minimal covering set. Some reflection reveals, however, that this algorithm requires exponentially many steps in the worst case. Since the sets B and A' in the statement of Lemma 8.6 are always disjoint, the lemma also tells us how to find, for every *proper* subset of the minimal covering set, another disjoint and non-empty subset. Again using the insight that a non-empty subset of the minimal covering set can be found efficiently, we finally obtain a polynomial-time algorithm.

THEOREM 8.7. Let A be a set of alternatives, g a comparison function on A. Then, MC(A, g) can be computed in time polynomial in |A|.

Proof. Consider the following simple algorithm:

- 1. Let B = ES(A, g).
- 2. While $A' = \bigcup_{a \in A \setminus B} (UC(B \cup \{a\}, g) \cap \{a\})$ is nonempty, add ES(A', g) to B.
- 3. Return B.

In each iteration of the second step, at least one element is added to set B, so the algorithm is guaranteed to terminate after a linear number of iterations. In each iteration the algorithm first identifies a subset A' of alternatives that are not yet covered by elements of B. This can be done in polynomial time by computing, for each element a outside B, the covering relation for $B \cup \{a\}$, and checking whether a itself is a maximal element of this relation. Then the elements in the essential set of A', which by Theorem 7.2 can be computed in polynomial time, are added to B.

As for correctness, a simple inductive argument shows that $B \subseteq MC(A, g)$ holds at any time. The base case follows directly from the fact that $ES(A, g) \subseteq MC(A, g)$ (Dutta and Laslier, 1999, Theorem 4.3), the induction step from Lemma 8.6. When the algorithm terminates, B is a covering set for A, so we must actually have B = MC(A, g).

Duggan and Le Breton (1996) consider an interesting class of symmetric zero-sum games, and show that every game in this class has a unique weak saddle $X_N = (X_1, X_2)$, which is itself symmetric, i.e., satisfies $X_1 = X_2$, and coincides with the minimal covering set of the corresponding comparison function. Games in this class, which we will term *confrontation games*, are characterized by the fact that the two players get the same payoff if and only if they play the same action. In any other case, one of the players strictly wins and the other strictly loses, and the outcome would be reversed if players were to exchange their actions. Duggan and Le Breton (1996) call this property the *off-diagonal property* in the context of comparison functions.

DEFINITION 8.8 (confrontation game). Let $\Gamma = (N, A^N, (p_i)_{i \in N})$ be a symmetric zero-sum game. Γ is called *confrontation game* if for all $a, a' \in A$, $p_1(a, a') = p_2(a, a') = 0$ if and only if a = a'.

We have the following corollary of Theorem 8.7.

COROLLARY 8.9. The unique weak saddle of a confrontation game can be computed in polynomial time.

In the remainder of this section, we present a family of symmetric zero-sum games that are not confrontation games and have an exponential number of weak saddles. An immediate consequence of this is that the computation of *all* weak saddles of a game requires exponential time in the worst case, even for zero-sum games. For an odd integer $k \ge 1$, define a two-player game Γ_k with action set $A = A^1 \cup A^2 \cup \cdots \cup A^k$ for both players, where $A^j = \{a^{j1}, a^{j2}, a^{j3}, a^{j4}\}$ for $1 \le j \le k$. Let p'_1 denote the payoff function of player 1 in the game of Figure 8.2, and define the payoff function p_1 of player 1 such that for all j_1, j_2 with $1 \le j_1, j_2 \le k$ and all m_1, m_2 with $1 \le m_1, m_2 \le 4$,

$$p_1(a^{j_1m_1}, a^{j_2m_2}) = \begin{cases} p'_1(a^{m_1}, a^{m_2}) & \text{if } j_1 = j_2 \\ (-1)^{j_1+j_2} & \text{if } j_1 < j_2 \\ (-1)^{j_1+j_2+1} & \text{otherwise.} \end{cases}$$

Further let $p_2(s_N) = -p_1(s_N)$ for all $s_N \in A \times A$. The overall structure of this game is shown in Figure 8.3.

Now, consider a pair $(X_1, X_2) \subseteq (A, A)$ such that for all j with $1 \leq j \leq k$, the pair (Y_1, Y_2) with $Y_1 = \{a^m : a^{jm} \in X_1 \cap A^j\}$ and $Y_2 = \{a^m : a^{jm} \in X_2 \cap A^j\}$ is a weak saddle of the game in Figure 8.2. It is not hard to see that (X_1, X_2) is a weak saddle of Γ_k , which means that the total number of weak saddles of Γ_k is at least 5^k .

8 · Shapley's Saddles

	A ¹	A ²	A ³	A^4		A ^k
A ¹	Г	(-1,1)	(1,-1)	(-1,1)		(1,-1)
A ²	(1,-1)	Г	(-1, 1)	(1,-1)		(-1, 1)
A ³	(-1,1)	(1,-1)	Г	(-1, 1)		(1, -1)
A ⁴	(1,-1)	(-1, 1)	(1, -1)	Г		(-1, 1)
	÷				·.	÷
A ^k	(-1,1)	(1,-1)	(-1,1)	(1,-1)		Г

Figure 8.3: Payoff structure of a $4k \times 4k$ symmetric zero-sum game Γ_k for an odd integer k. The game has at least 5^k weak saddles.

8.5 A Hardness Result for Weak Saddles

In this section, we establish a relationship between weak saddles of two-player games and inclusion-maximal cliques of undirected graphs. Our construction is similar to one used by McLennan and Tourky (2005) to show hardness of various problems concerning Nash equilibria, and directly yields a result concerning the computational hardness of weak saddles in general two-player games. This in turn leaves little hope that weak saddles can be found efficiently in general games.

THEOREM 8.10. Deciding whether there exists a weak saddle with a given minimum number of actions, or one with a given average payoff, is NP-hard, even when restricted to two-player games.

Proof. We provide a reduction from the NP-complete problem CLIQUE (e.g., Papadimitriou, 1994a). A *clique* of an undirected graph G is a subset $C \subseteq V$ such that $(i, j) \in E$ for all distinct $i, j \in C$. For a given graph G and $k \in \mathbb{N}$, CLIQUE asks whether G has a clique of size at least k.

Let G = (V, E) be an undirected graph, and define a two-player game Γ_G where both players have V as their set of actions, and payoffs are given by

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

and

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

Obviously Γ_G can be computed from G in polynomial time. We claim that a pair (X_1, X_2) is a weak saddle of Γ_G if and only if $X_1 = X_2$ and X_1 is an inclusion-maximal clique of G.

Recall that a WGSP is a pair of subsets of V that is externally stable for both players. For $v \in V$ and $U \subseteq V$, denote by $p_i(v, U) = (p_i(v, u))_{u \in U}$ the vector of payoffs for player i if he plays v and the other player plays some $u \in U$.

We begin by showing that $(X_1, X_2) \subseteq (V, V)$ is externally stable for player 2 if and only if $\emptyset \neq X_1 \subseteq X_2$. Since this effectively means that player 1 will only play actions he thinks player 2 might play, we will refer to this property as *imitation*. For the direction from left to right, assume that (X_1, X_2) is externally stable for player 2. Obviously, $X_1 \neq \emptyset$. Consider any $x \in X_1$ and assume $x \notin X_2$. Then there exists $x' \in X_2$ with $x' >_{X_1} x$, contradicting the fact that $p_2(x', x) = 0 < 1 = p_2(x, x)$. For the direction from right to left, consider $(X_1, X_2) \subseteq (V, V)$ such that $\emptyset \neq X_1 \subseteq X_2$. We have to show that for all $x \in V \setminus X_2$, there exists $x' \in X_2$ with $x' >_{X_1} x$. Let $x \in V \setminus X_2$. Since $X_1 \subseteq X_2$, it follows that $x \notin X_1$ and thus $p_2(x, X_1) = (0, \ldots, 0)$. Now let $x' \in X_1$. Then $p_2(x', X_1) = (0, \ldots, 0, 1, 0, \ldots, 0)$, with entry 1 at x'. This implies that $x' >_{X_1} x$.

We are now ready to prove that a pair (S_1, S_2) is a weak saddle of Γ_G if and only if $S_1 = S_2$ and S_1 is an inclusion-maximal clique in G. The statement of the theorem then follows immediately. For the direction from right to left, consider a maximal clique C of G. We have to show that (C, C) is a WGSP of Γ_G , i.e., externally stable for both players, and does not contain a WGSP as a proper subset. External stability for player 2 follows directly from imitation. For external stability for player 1, consider any $v \in V \setminus C$. Since C is a maximal clique, there must exist some $x \in C$ with $(v, x) \notin E$ or, equivalently, $p_1(v, x) = -1$. Then, $x >_C v$ because $p_1(x, C) = (1, \ldots, 1, 0, 1, \ldots, 1)$, with entry 0 at x. Now assume for contradiction that there exists a WGSP (X_1, X_2) with $X_1 \subseteq C$ and $X_2 \subseteq C$, such that at least one of the inclusions is strict. By imitation, $X_1 \subseteq X_2$. This in fact means that X_1 must be a strict subset of C, because otherwise $(X_1, X_2) = (C, C)$. Consider some $x \in C \setminus X_1$. Since (X_1, X_2) is a WGSP, there must exist some $x' \in X_1$ with $x' >_{X_2} x$. This is a contradiction, since $x' \in X_2$ and $p_1(x', x') = 0 < 1 = p_1(x, x')$, where the last equality is due to the fact that both x and x' are in the clique C.

For the direction from left to right, let (X_1, X_2) be a weak saddle of Γ_G and observe that $X_1 \subseteq X_2$ by imitation. Further let C be an inclusion-maximal clique in the induced subgraph $G|_{X_1}$ of G with vertex set X_1 . We claim that C is also an inclusion-maximal clique in G. Then, by the above, (C, C) is a weak saddle of Γ_G . Furthermore, $C \subseteq X_1 \subseteq X_2$, and thus $(X_1, X_2) = (C, C)$ by inclusion minimality of (X_1, X_2) . Assume for contradiction that C is *not* an inclusion-maximal clique of G, i.e., that there exists some $v \in V \setminus C$ connected to every vertex in C, such that $p_1(v, C) = (1, \ldots, 1)$. Since (X_1, X_2) is a weak saddle, there exists $x \in X_1$ with $x >_{X_2} v$. In particular, $p_1(x, C) = (1, \ldots, 1)$, implying that $x \notin C$ and that x is connected to all vertices in C. This obviously contradicts maximality of C in $G|_{X_1}$.

8.6 Discussion

In this chapter we have studied computational aspects of Shapley's saddles—ordinal setvalued solution concepts dating back to the early 1950s—by proposing polynomial-time algorithms for computing strict saddles in general normal-form games and pure weak saddles in a subclass of symmetric two-player zero-sum games. The algorithm for the latter class of games is highly non-trivial and relies on linear programs that determine the support of Nash equilibria in certain subgames of the original game. We have also seen that, in general two-player games, natural problems associated with weak saddles, such as deciding the existence of a weak saddle of a certain size or one containing a given action, are NP-hard. Several interesting questions concerning weak saddles, however, remain open. In particular, it is not known whether weak saddles can be computed efficiently in general two-player zero-sum games. Furthermore, the aforementioned NP-completeness results do not imply that finding an arbitrary weak saddle is NP-hard. Finally, gaps remain between the known upper and lower bounds for different problems in two-player games, like membership in some weak saddle or uniqueness of a weak saddle.

All of the above results apply to games with a constant number of players and many actions. It is an interesting question whether strict saddles can still be computed efficiently in certain classes of games that allow for a compact representation when the number of players is growing. Similarly, one might ask for classes of games where weak saddles become tractable. A natural candidate for such a class are anonymous games, studied in Chapter 5.

References

- T. Abbott, D. Kane, and P. Valiant. On the complexity of two-player win-lose games. In Proceedings of the 46th Symposium on Foundations of Computer Science (FOCS), pages 113-122. IEEE Computer Society Press, 2005. [24, 38]
- E. Allender, D. A. Mix Barrington, T. Chakraborty, S. Datta, and S. Roy. Grid graph reachability problems. In Proceedings of the 21st Annual IEEE Conference on Computational Complexity (CCC), pages 299–313, 2006. [85]
- K. R. Apt. Uniform proofs of order independence for various strategy elimination procedures. *Contributions to Theoretical Economics*, 4(1), 2004. [8, 9]
- I. Ashlagi, D. Monderer, and M. Tennenholtz. On the value of correlation. In *Proceedings* of the 21st Annual Conference on Uncertainty in Artificial Intelligence (UAI), pages 34–41. AUAI Press, 2005. [32, 51]
- R. J. Aumann. Acceptable points in general n-person games. In A. W. Tucker and R. D. Luce, editors, *Contributions to the Theory of Games IV*, volume 40 of Annals of Mathematics Studies, pages 287–324. Princeton University Press, 1959. [127]
- R. J. Aumann. Almost strictly competitive games. Journal of the Society of Industrial and Applied Mathematics, 9(4):544-550, 1961. [31]
- R. J. Aumann. Subjectivity and correlation in randomized strategies. Journal of Mathematical Economics, 1:67-96, 1974. [10, 30, 46, 50, 51, 58]
- R. J. Aumann. On the non-transferable utility value: A comment on the Roth-Shafer examples. *Econometrica*, 53(3):667–678, 1985. [32, 133]
- R. J. Aumann. Game theory. In J. Eatwell, M. Milgate, and P. Newman, editors, The New Palgrave: A Dictionary of Economics, volume 2, pages 460–482. MacMillan, 1987. [33]
- K. Basu and J. Weibull. Strategy subsets closed under rational behavior. *Economics Letters*, 36:141-146, 1991. [138]

- P. Beame, S. A. Cook, and H. J. Hoover. Log depth circuits for division and related problems. *SIAM Journal on Computing*, 15(4):994–1003, 1986. [72]
- M. Benisch, G. B. Davis, and T. Sandholm. Algorithms for rationalizability and CURB sets. In Proceedings of the 21st National Conference on Artificial Intelligence (AAAI), pages 598-604. AAAI Press, 2006. [138]
- N. Bhat and K. Leyton-Brown. Computing Nash equilibria of action-graph games. In Proceedings of the 20th Annual Conference on Uncertainty in Artificial Intelligence (UAI), pages 35-42. AUAI Press, 2004. [59]
- B. Bollobás. Modern Graph Theory. Springer-Verlag, 1998. [65]
- E. Borel. La théorie du jeu et les équations intégrales à noyau symétrique. Comptes Rendus de l'Académie des Sciences, 173:1304–1308, 1921. [58]
- F. Brandt and F. Fischer. On the hardness and existence of quasi-strict equilibria. In B. Monien and U.-P. Schroeder, editors, *Proceedings of the 1st International Symposium on Algorithmic Game Theory (SAGT)*, volume 4997 of *Lecture Notes in Computer Science (LNCS)*, pages 291–302. Springer-Verlag, 2008a. [27]
- F. Brandt and F. Fischer. Computing the minimal covering set. Mathematical Social Sciences, 56(2):254-268, 2008b. [27]
- F. Brandt, F. Fischer, and Y. Shoham. On strictly competitive multi-player games. In Y. Gil and R. Mooney, editors, *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI)*, pages 605–612. AAAI Press, 2006. [25]
- F. Brandt, F. Fischer, P. Harrenstein, and Y. Shoham. A game-theoretic analysis of strictly competitive multiagent scenarios. In M. Veloso, editor, *Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1199–1206, 2007a. [25]
- F. Brandt, T. Sandholm, and Y. Shoham. Spiteful bidding in sealed-bid auctions. In M. Veloso, editor, Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI), pages 1207–1214, 2007b. [32]
- F. Brandt, F. Fischer, and M. Holzer. On iterated dominance, matrix elimination, and matched paths. ECCC Report TR08-077, Electronic Colloquium on Computational Complexity (ECCC), 2008a. Working paper. [26]
- F. Brandt, F. Fischer, and M. Holzer. Equilibria of graphical games with symmetries. In C. H. Papadimitriou and S. Zhang, editors, *Proceedings of the 4th International Workshop on Internet and Network Economics (WINE)*, Lecture Notes in Computer Science (LNCS), pages 198–209. Springer-Verlag, 2008b. [26]

- F. Brandt, M. Brill, F. Fischer, and P. Harrenstein. Computational aspects of Shapley's saddles. In K. S. Decker and J. S. Sichman, editors, Proceedings of the 8th International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS), 2009a. [27]
- F. Brandt, M. Brill, F. Fischer, and P. Harrenstein. On the complexity of iterated weak dominance in constant-sum games. In M. Mavronicolas, editor, *Proceedings of the* 2nd International Symposium on Algorithmic Game Theory (SAGT), 2009b. To appear. [59]
- F. Brandt, F. Fischer, P. Harrenstein, and Y. Shoham. Ranking games. Artificial Intelligence, 173(2):221-239, 2009c. [25]
- F. Brandt, F. Fischer, and M. Holzer. Symmetries and the complexity of pure Nash equilibrium. Journal of Computer and System Sciences, 75(3):163-177, 2009d. [26]
- A. K. Chandra, L. Stockmeyer, and U. Vishkin. Constant depth reducibility. SIAM Journal on Computing, 13(2):423-439, 1984. [65, 66, 67, 68]
- X. Chen and X. Deng. Settling the complexity of 2-player Nash-equilibrium. In Proceedings of the 47th Symposium on Foundations of Computer Science (FOCS), pages 261-272. IEEE Press, 2006. [24, 37, 42, 58, 129]
- X. Chen, X. Deng, and S.-H. Teng. Computing Nash equilibria: Approximation and smoothed complexity. In Proceedings of the 47th Symposium on Foundations of Computer Science (FOCS), pages 603-612. IEEE Computer Society Press, 2006. [24]
- X. Chen, S.-H. Teng, and P. Valiant. The approximation complexity of win-lose games. In Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pages 159–168. SIAM, 2007. [24]
- S.-F. Cheng, D. M. Reeves, Y. Vorobeychik, and M. P. Wellman. Notes on equilibria in symmetric games. In *Proceedings of the 6th International Workshop On Game Theoretic And Decision Theoretic Agents (GTDT)*, 2004. [67]
- B. Codenotti and D. Stefankovic. On the computational complexity of Nash equilibria for (0, 1) bimatrix games. *Information Processing Letters*, 94(3):145-150, 2005. [23]
- V. Conitzer and T. Sandholm. Complexity of (iterated) dominance. In Proceedings of the 6th ACM Conference on Electronic Commerce (ACM-EC), pages 88-97. ACM Press, 2005a. [9, 23, 37, 42, 43, 59, 78, 83, 96, 138]
- V. Conitzer and T. Sandholm. A generalized strategy eliminability criterion and computational methods for applying it. In *Proceedings of the 20th National Conference on Artificial Intelligence (AAAI)*, pages 483–488. AAAI Press, 2005b. [138]

- V. Conitzer and T. Sandholm. New complexity results about Nash equilibria. Games and Economic Behavior, 63:621-641, 2008. [23]
- R. Cubitt and R. Sugden. Rationally justifiable play and the theory of non-cooperative games. *Economic Journal*, 104(425):798-803, 1994. [128, 129, 132]
- G. B. Dantzig. A proof of the equivalence of the programming problem and the game problem. In T. C. Koopmans, editor, Activity Analysis of Production and Allocation, pages 330-335. John Wiley & Sons Inc., 1951. [23, 37]
- C. Daskalakis and C. H. Papadimitriou. The complexity of games on highly regular graphs. In Proceedings of the 13th Annual European Symposium on Algorithms (ESA), volume 3669 of Lecture Notes in Computer Science (LNCS), pages 71-82. Springer-Verlag, 2005. [102, 103, 125]
- C. Daskalakis and C. H. Papadimitriou. Computing equilibria in anonymous games. In Proceedings of the 48th Symposium on Foundations of Computer Science (FOCS), pages 83-93. IEEE Computer Society Press, 2007. [57]
- C. Daskalakis and C. H. Papadimitriou. Discretized multinomial distributions and Nash equilibria in anonymous games. In *Proceedings of the 49th Symposium on Founda*tions of Computer Science (FOCS). IEEE Computer Society Press, 2008. [59, 102]
- C. Daskalakis, P. Goldberg, and C. Papadimitriou. The complexity of computing a Nash equilibrium. In Proceedings of the 38th Annual ACM Symposium on the Theory of Computing (STOC), pages 71–78. ACM Press, 2006. [24, 102]
- C. Daskalakis, P. Goldberg, and C. Papadimitriou. The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, 2009a. To appear. [24, 37]
- C. Daskalakis, G. Schoenebeck, G. Valiant, and P. Valiant. On the complexity of Nash equilibria of action-graph games. In *Proceedings of the 20th Annual ACM-SIAM* Symposium on Discrete Algorithms (SODA), 2009b. [59]
- Y. Dodis, S. Halevi, and T. Rabin. A cryptographic solution to a game theoretic problem. In Proceedings of the 20th Annual International Cryptology Conference (CRYPTO), volume 1880 of Lecture Notes in Computer Science (LNCS), pages 112–130. Springer-Verlag, 2000. [50]
- J. Duggan and M. Le Breton. Dutta's minimal covering set and Shapley's saddles. *Journal of Economic Theory*, 70:257–265, 1996. [137, 145]
- J. Duggan and M. Le Breton. Mixed refinements of Shapley's saddles and weak tournaments. Social Choice and Welfare, 18(1):65-78, 2001. [137]

- J. Dunkel and A. S. Schulz. On the complexity of pure-strategy Nash equilibria in congestion and local-effect games. *Mathematics of Operations Research*, 33(4):851-868, 2008. [59]
- B. Dutta. Covering sets and a new Condorcet choice correspondence. Journal of Economic Theory, 44:63-80, 1988. [142]
- B. Dutta and J.-F. Laslier. Comparison functions and choice correspondences. Social Choice and Welfare, 16(4):513-532, 1999. [142, 143, 144]
- E. Elkind, L. A. Goldberg, and P. W. Goldberg. Equilibria in graphical games on trees revisited. In Proceedings of the 7th ACM Conference on Electronic Commerce (ACM-EC), pages 100–109. ACM Press, 2006. [102]
- E. Elkind, L. A. Goldberg, and P. W. Goldberg. Computing good Nash equilibria in graphical games. In *Proceedings of the 8th ACM Conference on Electronic Commerce* (ACM-EC), pages 162–171. ACM Press, 2007. [102]
- K. Ellul, B. Krawetz, J. Shallit, and M.-W. Wang. Regular expressions: New results and open problems. *Journal of Automata, Languages and Combinatorics*, 9(2-3): 233-256, 2004. [85]
- K. Etessami and M. Yannakakis. On the complexity of Nash equilibria and other fixed points. In Proceedings of the 48th Symposium on Foundations of Computer Science (FOCS), pages 113–123. IEEE Computer Society Press, 2007. [24]
- A. Fabrikant, C. H. Papadimitriou, and K. Talwar. The complexity of pure Nash equilibria. In Proceedings of the 36th Annual ACM Symposium on the Theory of Computing (STOC), pages 604-612. ACM Press, 2004. [59, 61]
- F. Fischer, M. Holzer, and S. Katzenbeisser. The influence of neighbourhood and choice on the complexity of finding pure Nash equilibria. *Information Processing Letters*, 99 (6):239-245, 2006. [26, 63]
- J. W. Friedman. On characterizing equilibrium points in two person strictly competitive games. *International Journal of Game Theory*, 12:245–247, 1983. [31]
- H. N. Gabow, S. N. Maheshwari, and L.J. Osterweil. On two problems in the generation of program test paths. *IEEE Transactions on Software Engineering*, 2(3):227-231, 1976. [85]
- D. Gale, H. W. Kuhn, and A. W. Tucker. On symmetric games. In H. W. Kuhn and A. W. Tucker, editors, *Contributions to the Theory of Games*, volume 1, pages 81– 87. Princeton University Press, 1950. [58, 129, 134]

- M. R. Garey and D. S. Johnson. Computers and Intractability: A Guide to the Theory of NP-Completeness. W. H. Freeman, 1979. [75]
- I. Gilboa and E. Zemel. Nash and correlated equilibria: Some complexity considerations. Games and Economic Behavior, 1(1):80–93, 1989. [23]
- I. Gilboa, E. Kalai, and E. Zemel. The complexity of eliminating dominated strategies. Mathematics of Operations Research, 18(3):553-565, 1993. [23, 59]
- F. Glover. Maximum matching in convex bipartite graphs. Naval Research Logistics Quarterly, 14:313-316, 1967. [84]
- P. W. Goldberg and C. H. Papadimitriou. Reducibility among equilibrium problems. In Proceedings of the 38th Annual ACM Symposium on the Theory of Computing (STOC), pages 61-70. ACM Press, 2006. [24, 102, 125]
- O. Goldreich. Computational Complexity: A Conceptual Perspective. Cambridge University Press, 2008. [5]
- G. Gottlob, G. Greco, and F. Scarcello. Pure Nash equilibria: Hard and easy games. Journal of Artificial Intelligence Research, 24:195-220, 2005. [26, 37, 59, 101, 102, 103, 105]
- S. Govindan and R. Wilson. A global Newton method to compute Nash equilibria. *Journal* of *Economic Theory*, 110(1):65–86, 2003. [59]
- R. Greenlaw, H. J. Hoover, and W. L. Ruzzo. Limits to Parallel Computation. Oxford University Press, 1995. [85]
- J. C. Harsanyi. Oddness of the number of equilibrium points: A new proof. International Journal of Game Theory, 2:235-250, 1973. [10, 27, 127, 129]
- S. Ieong, R. McGrew, E. Nudelman, Y. Shoham, and Q. Sun. Fast and compact: A simple class of congestion games. In *Proceedings of the 20th National Conference on Artificial Intelligence (AAAI)*, pages 489–494. AAAI Press, 2005. [59, 61]
- N. Immerman. Nondeterministic space is closed under complementation. SIAM Journal on Computing, 17(5):935-938, 1988. [15]
- M. J. M. Jansen. Regularity and stability of equilibrium points of bimatrix games. *Mathematics of Operations Research*, 6(4):530-550, 1981. [130]
- M. J. M. Jansen. Regular equilibrium points of bimatrix games. OR Spektrum, 9(2): 82-92, 1987. [129]
- A. X. Jiang and K. Leyton-Brown. Computing pure Nash equilibria in symmetric actiongraph games. In Proceedings of the 22nd AAAI Conference on Artificial Intelligence (AAAI), pages 79–85. AAAI Press, 2007. [59]

- D. S. Johnson, C. H. Papadimitriou, and M. Yannakakis. How easy is local search? Journal of Computer and System Sciences, 37:79-100, 1988. [17, 72, 74]
- N. D. Jones. Space-bounded reducibility among combinatorial problems. Journal of Computer and System Sciences, 11:68-85, 1975. [110]
- A. P. Jurg, M. J. M. Jansen, J. A. M. Potters, and S. H. Tijs. A symmetrization for finite two-person games. ZOR – Methods and Models of Operations Research, 36(2): 111–123, 1992. [130]
- A. Kats and J.-F. Thisse. Unilaterally competitive games. International Journal of Game Theory, 21:291-299, 1992. [32, 136]
- M. J. Kearns, M. L. Littman, and S. P. Singh. Graphical models for game theory. In Proceedings of the 17th Annual Conference on Uncertainty in Artificial Intelligence (UAI), pages 253-260. Morgan Kaufmann, 2001. [101]
- L. Khachiyan. A polynomial algorithm in linear programming. Soviet Mathematics Doklady, 20:191–194, 1979. [23, 37]
- D. E. Knuth, C. H. Papadimitriou, and J. N. Tsitsiklis. A note on strategy elimination in bimatrix games. *Operations Research Letters*, 7:103–107, 1988. [23, 59, 138]
- M. Kojima, A. Okada, and S. Shindoh. Strongly stable equilibrium points of N-person noncooperative games. *Mathematics of Operations Research*, 10(4):650-663, 1985.
 [132]
- E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In Proceedings of the 16th International Symposium on Theoretical Aspects of Computer Science (STACS), volume 1563 of Lecture Notes in Computer Science (LNCS), pages 404–413. Springer-Verlag, 1999. [32, 51]
- G. Laffond, J.-F. Laslier, and M. Le Breton. More on the tournament equilibrium set. Mathématiques et sciences humaines, 31(123):37-44, 1993. [143]
- J.-F. Laslier. Tournament Solutions and Majority Voting. Springer-Verlag, 1997. [142]
- K. Leyton-Brown and M. Tennenholtz. Local-effect games. In Proceedings of the 18th International Joint Conference on Artificial Intelligence (IJCAI), pages 772–777, 2003. [59]
- R. D. Luce and H. Raiffa. Games and Decisions: Introduction and Critical Survey. John Wiley & Sons Inc., 1957. [46, 58, 130]
- L. M. Marx and J. M. Swinkels. Order independence for iterated weak dominance. *Games* and *Economic Behavior*, 18:219-245, 1997. [43]

- A. McLennan and R. Tourky. Simple complexity from imitation games. Unpublished manuscript, 2005. [134, 146]
- N. Megiddo and C. H. Papadimitriou. On total functions, existence theorems and computational complexity. *Theoretical Computer Science*, 81:317-324, 1991. [15, 23]
- D. Monderer and L. S. Shapley. Potential games. Games and Economic Behavior, 14 (1):124-143, 1996. [136]
- J. Morgan, K. Steiglitz, and G. Reis. The spite motive and equilibrium behavior in auctions. *Contributions to Economic Analysis & Policy*, 2(1):1102-1127, 2003. [32]
- H. Moulin. Dominance solvable voting schemes. Econometrica, 47:1337-1351, 1979. [9]
- H. Moulin and J.-P. Vial. Strategically zero-sum games: The class of games whose completely mixed equilibria cannot be improved upon. International Journal of Game Theory, 7(3-4):201-221, 1978. [30, 32]
- R. B. Myerson. Game Theory: Analysis of Conflict. Harvard University Press, 1991. [5, 23]
- J. F. Nash. Non-cooperative games. Annals of Mathematics, 54(2):286-295, 1951. [9, 15, 23, 58, 133, 137]
- N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani. Algorithmic Game Theory. Cambridge University Press, 2007. [2]
- H. Norde. Bimatrix games have quasi-strict equilibria. Mathematical Programming, 85: 35-49, 1999. [35, 127, 128, 129]
- H. Norde, J. Potters, H. Reijnierse, and D. Vermeulen. Equilibrium selection and consistency. Games and Economic Behavior, 12(2):219-225, 1996. [128, 129]
- J. B. Orlin, A. P. Punnen, and A. S. Schulz. Approximate local search in combinatorial optimization. *SIAM Journal on Computing*, 33(5):1201-1214, 2004. [18]
- M. J. Osborne and A. Rubinstein. A Course in Game Theory. MIT Press, 1994. [5]
- C. H. Papadimitriou. Computational Complexity. Addison-Wesley, 1994a. [3, 5, 45, 69, 72, 105, 108, 134, 146]
- C. H. Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. *Journal of Computer and System Sciences*, 48(3):498-532, 1994b. [16]
- C. H. Papadimitriou. Algorithms, games, and the internet. In Proceedings of the 33rd Annual ACM Symposium on the Theory of Computing (STOC), pages 749-753. ACM Press, 2001. [23]

- C. H. Papadimitriou. Computing correlated equilibria in multi-player games. In Proceedings of the 37th Annual ACM Symposium on the Theory of Computing (STOC), pages 49-56. ACM Press, 2005. [37]
- C. H. Papadimitriou. The complexity of finding Nash equilibria. In N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*, chapter 2, pages 29-78. Cambridge University Press, 2007. [3]
- C. H. Papadimitriou and T. Roughgarden. Computing equilibria in multi-player games. In Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pages 82–91. SIAM, 2005. [21, 58, 124, 134]
- I. Parberry. On the computational complexity of optimal sorting network verification. In Proceedings of the Conference on Parallel Architectures and Languages Europe (PARLE), volume 505 of Lecture Notes in Computer Science (LNCS), pages 252– 269. Springer-Verlag, 1991. [86]
- R. Parikh. On context-free languages. Journal of the ACM, 13(4):570-581, 1966. [60]
- B. Peleg and S. Tijs. The consistency principle for games in strategic form. International Journal of Game Theory, 25:13-34, 1996. [128, 129]
- A. Quesada. Another impossibility result for normal form games. *Theory and Decision*, 52:73–80, 2002. [129]
- T. E. S. Raghavan. Non-zero-sum two-person games. In R. J. Aumann and S. Hart, editors, Handbook of Game Theory with Economic Applications, volume III, chapter 17, pages 1687-1721. North-Holland, 2002. [32]
- N. Robertson, P. D. Seymour, and R. Thomas. Permanents, Pfaffian orientations, and even directed circuits. Annals of Mathematics, 150:929–975, 1999. [119, 124]
- R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. International Journal of Game Theory, 2:65-67, 1973. [59]
- T. Roughgarden. Selfish Routing and the Price of Anarchy. MIT Press, 2005. [32]
- W. Ruzzo. On uniform circuit complexity. Journal of Computer and System Sciences, 22(3):365–383, 1981. [19]
- T. J. Schaefer. The complexity of satisfiability problems. In *Proceedings of the 10th* Annual ACM Symposium on the Theory of Computing (STOC), pages 216–226. ACM Press, 1978. [119, 123, 124]
- A. A. Schäffer and M. Yannakakis. Simple local search problems that are hard to solve. SIAM Journal on Computing, 20(1):56-87, 1991. [17, 72, 74]

- G. Schoenebeck and S. Vadhan. The computational complexity of Nash equilibria in concisely represented games. In *Proceedings of the 7th ACM Conference on Electronic Commerce (ACM-EC)*, pages 270–279. ACM Press, 2006. [37, 59, 63, 105, 111, 112]
- P. D. Seymour. On the two-colouring of hypergraphs. The Quarterly Journal of Mathematics, 25:303-312, 1974. [117, 118]
- L. Shapley. Order matrices. I. Technical Report RM-1142, The RAND Corporation, 1953a. [27, 137, 138]
- L. Shapley. Order matrices. II. Technical Report RM-1145, The RAND Corporation, 1953b. [137]
- L. Shapley. A condition for the existence of saddle-points. Technical Report RM-1598, The RAND Corporation, 1955. [137]
- L. Shapley. Some topics in two-person games. In M. Dresher, L. S. Shapley, and A. W. Tucker, editors, *Advances in Game Theory*, volume 52 of *Annals of Mathematics Studies*, pages 1–29. Princeton University Press, 1964. [10, 137, 140]
- P. Spirakis. Approximate equilibria for strategic two person games. In Proceedings of the 1st International Symposium on Algorithmic Game Theory (SAGT), pages 5-21, 2008. [24]
- D. Squires. Impossibility theorems for normal form games. *Theory and Decision*, 44: 67-81, 1998. [129]
- R. Szelepcsenyi. The method of forced enumeration for nondeterministic automata. Acta Informatica, 26(3):279-284, 1988. [15]
- M. Tennenholtz. Competitive safety analysis: Robust decision-making in multi-agent systems. Journal of Artificial Intelligence Research, 17:363–378, 2002. [32]
- C. Thomassen. Even cycles in directed graphs. *European Journal of Combinatorics*, 6: 85–89, 1985. [118]
- E. van Damme. Refinements of the Nash Equilibrium Concept. Springer-Verlag, 1983. [36, 128, 129, 132]
- E. van Damme. Stability and Perfection of Nash Equilibria. Springer-Verlag, 2nd edition, 1991. [129]
- E. van Damme. On the state of the art in game theory: An interview with Robert Aumann. Games and Economic Behavior, 24:181-210, 1998. [2]
- H. Vollmer. Introduction to Circuit Complexity. Springer-Verlag, 1999. [5]

- J. von Neumann. Zur Theorie der Gesellschaftspiele. Mathematische Annalen, 100: 295-320, 1928. [47, 58, 130]
- J. von Neumann. First draft of a report on the EDVAC. Technical report, Moore School of Electrical Engineering, University of Pennsylvania, 1945. [2]
- J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1944. [1, 5]
- J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 2nd edition, 1947. [6, 29, 58, 98]
- O. De Wolf. Optimal strategies in n-person unilaterally competitive games. Discussion Paper 9949, Center for Operations Research and Econometrics, Université catholique de Louvain, 1999. [32]

Lebenslauf

Felix Fischer wurde am 11. März 1978 in München geboren. Von 1984 bis 1997 besuchte er Grundschule und Gymnasium in Garching bei München. Nach Erfüllung der allgemeinen Wehrpflicht begann er 1998 ein Studium der Informatik mit Nebenfach Elektrotechnik an der Technischen Universität München, das er im Jahr 2003 mit dem Abschluß zum Diplom-Informatiker beendete. Nach einer gut einjährigen Tätigkeit als wissenschaftlicher Mitarbeiter am Lehrstuhl für Theoretische Informatik und Grundlagen der Künstlichen Intelligenz der Technischen Universität München begann er im Jahr 2005 schließlich mit der Promotion in der Forschungsgruppe Präferenzbündelung in Multiagentensystemen an der Lehr- und Forschungseinheit für Theoretische Informatik der Ludwig-Maximilians-Universität München. Im Jahr 2007 war er über einen Zeitraum von insgesamt fünf Monaten als Gastwissenschaftler an der Hebrew University of Jerusalem in Israel tätig.