

Characterization and Computation of Equilibria in Infinite Games

by

Noah D. Stein

B.S., Electrical and Computer Engineering (2005)
Cornell University

Submitted to the Department of Electrical Engineering and Computer
Science in partial fulfillment of the requirements for the degree of

Master of Science in Electrical Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2007

© Massachusetts Institute of Technology 2007. All rights reserved.

Author
Department of Electrical Engineering and Computer Science
May 11, 2007

Certified by
Asuman Ozdaglar
Assistant Professor
Thesis Supervisor

Certified by
Pablo A. Parrilo
Associate Professor
Thesis Supervisor

Accepted by
Arthur C. Smith
Chairman, Department Committee on Graduate Theses

Characterization and Computation of Equilibria in Infinite Games

by
Noah D. Stein

Submitted to the Department of Electrical Engineering and Computer Science
on May 11, 2007, in partial fulfillment of the requirements for the degree of
Master of Science in Electrical Engineering

Abstract

Broadly, we study continuous games (those with continuous strategy spaces and utility functions) with a view towards computation of equilibria. We cover all of the game-theoretic background needed to understand these results in detail. Then we present new work, which can be divided into three parts.

First, it is known that arbitrary continuous games may have arbitrarily complicated equilibria, so we investigate some properties of games with polynomial utility functions and a class of games with polynomial-like utility functions called separable games. We prove new bounds on the complexity of equilibria of separable games in terms of the complexity of the utility functions. In order to measure this complexity we propose a new definition for the rank of a continuous game; when applied to the case of finite games this improves on the results known in that setting. Furthermore, we prove a characterization theorem showing that several conditions which are necessary for a game to possess a finite-dimensional representation each define the class of separable games precisely, providing evidence that separable games are the natural class of continuous games in which to study computation. The characterization theorem also provides a natural connection between separability and the notion of the rank of a game.

Second, we apply this theory to give an algorithm for computing ϵ -Nash equilibria of two-player separable games with continuous strategy spaces. While a direct comparison to corresponding algorithms for finite games is not possible, the asymptotic running time in the complexity of the game grows slower for our algorithm than for any known algorithm for finite games. Nonetheless, as in finite games, computing ϵ -Nash equilibria still appears to be difficult for infinite games.

Third, we consider computing approximate correlated equilibria in polynomial games. To do so, we first prove several new characterizations of correlated equilibria in continuous games which may be of independent interest. Then we introduce three algorithms for approximating correlated equilibria of polynomial games arbitrarily accurately. These include two discretization algorithms for computing a sample correlated equilibrium: a naive linear programming approach called static discretization which operates without regard to the structure of the game, and a semidefinite programming approach called adaptive discretization which exploits the structure of the game to achieve far better performance in practice. The third algorithm consists of a nested sequence of semidefinite programs converging to a description of the entire set of correlated equilibria.

Thesis Supervisor: Asuman Ozdaglar
Title: Assistant Professor

Thesis Supervisor: Pablo A. Parrilo
Title: Associate Professor

Acknowledgments

Most of all I would like to thank my advisors Asu and Pablo; without all their guidance and suggestions I could not have accomplished the work below. I am especially grateful for their shared focus on helping me develop as a student and mathematician – not just a member of their research groups – by taking me to conferences even when I had nothing to present and by encouraging me to follow my interests, such as taking many time-consuming math classes which were at best tangentially related to our research. I also would like to acknowledge their dedication in reading draft after draft of my papers. Individually, I want to thank Asu for her boundless energy and enthusiasm and Pablo for making me use MATLAB occasionally, but not too often. Last but not least, I appreciate that both Asu and Pablo have generously tolerated my quirks, which they were unaware of when initially accepting me as a student.

I also wish to thank Professor Muhamet Yildiz for a fruitful and far-reaching discussion about game theory. This conversation inspired the new characterizations of correlated equilibria in infinite games as well as the semidefinite programming relaxation methods for computing correlated equilibria of polynomial games. Both are integral parts of the work below.

Finally, I thank Laura for her willingness to follow me to Cambridge, her emotional support, and her continuing patience with my decision to be a student even though we are both “adults” now. Most of all, I thank her for always making life interesting, whether are out on an adventure together or sitting at home, and for being herself.

Contents

1	Introduction	9
1.1	Context	9
1.2	Outline and Contributions	10
2	Background	13
2.1	Classes of Games and Equilibria	13
2.1.1	Finite Games	13
2.1.2	Continuous Games	20
2.1.3	Separable Games	26
2.2	Computation of Nash Equilibria in Finite Games	28
2.2.1	Optimization Formulation	28
2.2.2	Support Enumeration	29
2.2.3	The Lemke-Howson Algorithm	31
2.3	Previous Work	35
3	Rank and Equilibria in Infinite Games	39
3.1	Rank of Finite Games	39
3.2	Rank of Continuous and Separable Games	41
3.3	Characterizations and Rank Bounds for Correlated Equilibria	46
3.4	Characterizations of Separable Games	50
3.5	Computing Ranks	54
4	Computation of Equilibria in Separable Games	61
4.1	Computing Nash Equilibria	61
4.2	Computing ϵ -Nash Equilibria	62
4.3	Approximating Correlated Equilibria	64
4.3.1	Static Discretization Methods	65
4.3.2	Adaptive Discretization Methods	67
4.3.3	Semidefinite Relaxation Methods	70
5	Conclusions	73
5.1	Summary	73
5.2	Future Work	74
A	Sum of Squares Techniques	77
	References	79

Chapter 1

Introduction

1.1 Context

In this section we briefly overview the goals of game theory in order to introduce the problems we have studied. In Chapter 2 we will formulate the background material precisely and describe related work in detail.

Broadly speaking a **game** is a mathematical formalization of a strategic interaction. The agents involved are called the **players**. Examples include Rock, Paper, Scissors; rival firms choosing production levels in an economic market; and users competing for bandwidth on the internet. The players are assumed to act rationally in some sense. The objective of game theory is to predict the outcome of games when possible, or to at least find conditions that the outcome should be expected to satisfy. Of course, the statement of this goal presupposes the existence – and perhaps uniqueness – of a well-defined “solution” for a game. As we will see several definitions of solution are possible and some assumptions are required even to obtain existence. When these assumptions are satisfied, the question arises whether such a solution can be computed efficiently, or at all. The answer has substantial philosophical implications because if a solution exists but the players cannot find it, then the “solution” lacks predictive power about the outcome of the game.

Different strategic situations lend themselves to different game-theoretic models. The players may act simultaneously or sequentially, they may act independently or cooperate, they may have complete or partial knowledge of the others’ preferences, and so forth. Generally we assume no communication between the players is allowed outside the game, because any such interaction could simply be included in the game model. The obvious examples of games tend to have a finite number of players, but for the purpose of analyzing limiting cases games with infinitely many players can be considered as well. There has been extensive research on many possible combinations of these parameters, but we will restrict attention throughout to the case of **strategic form games with finitely many players**, i.e. games in which the players choose their actions simultaneously and the entire structure of the game, including all players’ preferences over outcomes, is common knowledge among the players. While there exist arguments against this choice of model, it is used extensively in Game Theory and

we will use it without further comment.

A variety of **solution concepts** have been proposed to characterize “optimal” outcomes of games in strategic form, but we will consider only two. The most widely known and used is the **Nash equilibrium**, which is defined to be a (possibly random) choice of strategy for each player which has the stability property that a player who knows the other players will play according to the equilibrium strategies has no incentive not to do the same. Such an equilibrium can be viewed as a product probability distribution over the product of the player’s strategy spaces and corresponds to the case when the players do not have a jointly observable trusted source of randomness with which to correlate their actions. In real life such sources are abundant in the form of the weather, stock prices, traffic signals, etc. Therefore the **correlated equilibrium** has been defined as a generalization of the Nash equilibrium which allows for arbitrary probability distributions over the product of the players’ strategy sets rather than just independent distributions.

With the game theoretic setting in place, we now motivate and introduce the new problems we study below. There has been much work on computing Nash and correlated equilibria in finite games (i.e. games with finite strategy sets), including algorithms, complexity bounds, and hardness results. To briefly summarize the computational results for finite games: correlated equilibria of general finite games and Nash equilibria of two-player zero-sum (i.e. what’s good for one player is bad for the other and vice versa) can generally be computed efficiently, but there is evidence to suggest that Nash equilibria of general finite games probably cannot (for more detail see Section 2.3). On the other hand, most of the work on infinite games has focused on existence and characterization results, with little known about computation. Recently, Parrilo has shown that in a certain class of two-player zero-sum infinite games, Nash equilibria can again be computed efficiently [33]. Together, these past results have led us to consider the questions:

- What (if any) characteristics are necessary for an infinite game to be amenable to computation?
- Can we construct explicit algorithms to compute equilibria in classes of games with these characteristics?

The previous work on finite games provided guidance on the second question, suggesting that we were unlikely to find efficient algorithms to compute Nash equilibria of nonzero-sum games, and that correlated equilibria might be more computationally tractable.

1.2 Outline and Contributions

We conclude this introduction with an outline of the rest of the thesis, including summaries of our answers to the above questions. Chapter 2 defines and presents some fundamental known results about Nash and correlated equilibria. In Section 2.1 we define several classes of games and prove existence and characterization theorems

for Nash and correlated equilibria in these games. In particular we study finite games as well as continuous games, which may have infinite strategy sets but have additional topological structure making them amenable to analysis. We also consider a class of continuous games called separable games which have some additional algebraic structure that allows us to establish stronger results, and in particular relate the complexity of equilibria to the algebraic complexity of the game. Section 2.2 surveys some algorithms for computing Nash equilibria in finite games. Finally, in Section 2.3 we briefly discuss a variety of other related literature which led to the formulation of the problems we consider and the results we present in later chapters.

After the background chapter, all the material presented is original, except for that which is repeated to show how it can be generalized by or compared to our new results. In Chapter 3 we cover those contributions which are of a theoretical nature.

- We introduce the concept of the ranks of a continuous game, a list of integers which measure the complexity of the game for purposes of computation of (mixed strategy) Nash and correlated equilibria. This definition generalizes one for two-player finite games studied by Lipton et al. [26] to cover an arbitrary finite number of players (a problem those authors explicitly left open) as well as infinite strategy sets. Our definition also has the advantage of being independent of the representation chosen for the game.
- Using the ranks, we derive representation-independent bounds on the complexity of Nash and correlated equilibria, measured by the number of strategies played with positive probability. Doing so connects and generalizes the results of Lipton et al. on Nash equilibria [26] with the results of Germano and Lugosi on correlated equilibria [16].
- In the process, we prove several characterizations of correlated equilibria in infinite games. The most fundamental of these was conjectured to be false (under slightly weaker hypotheses) in the classic paper by Hart and Schmeidler [18]. This leads to another characterization, which to our knowledge is the first characterization of correlated equilibria in infinite games to allow computation without resorting to discretizing the players' strategy spaces.
- We prove several characterizations of separable games. For example, separable games are exactly those continuous games for which the ranks we have defined are finite, and hence which may be amenable to computation. We also show that a certain necessary condition for a continuous game to have a finite-dimensional representation is equivalent to the condition that the game be separable, hence the condition is also sufficient.
- We construct explicit formulas for the ranks of finite, polynomial, and arbitrary separable games as the ranks of associated matrices. These formulas are surprising because at first glance the problem of computing the ranks appears to be one of finding a minimal (in some sense) decomposition of certain tensors. Such problems tend to be awkward to formulate precisely and hard to solve

algorithmically. Nonetheless we are able to reduce the problem of computing ranks to a matrix decomposition, which is easily computed.

In Chapter 4 we apply these theoretical results to compute equilibria of separable games with infinite strategy sets.

- We give a discretization algorithm for approximating Nash equilibria of a large class of two-player separable games with continuous strategy sets, including games with polynomial utility functions, whose running time is bounded in terms of the rank of the game. While this algorithm does not apply to finite games (one cannot control the behavior of the payoffs of a finite game under discretization), it is interesting to note that its asymptotic behavior in terms of the complexity of the game is better than any known algorithms for finite games.
- We also introduce three types of methods for computing correlated equilibria.
 - Static discretization: Included primarily for benchmarking purposes, in these methods we discretize each player’s strategy set without regard to the game structure. Then we compute correlated equilibria of the induced finite game exactly by existing methods for finite games. The simplicity of this algorithm allows us to prove its convergence (as the number of points in the discretization increases) and also to exactly compute the worst-case convergence rate.
 - Adaptive discretization: In these methods we iteratively add points to the discretized strategy sets in a fashion which exploits the structure of the game. At each iteration we compute an approximate correlated equilibrium which is optimal with respect to the current discretization, then we use this to infer additional strategies which some player prefers over those available in the current discretization. We add these to the discretization at the next stage and iterate. To perform these computations algorithmically we require that the games have a polynomial structure and we use sum of squares techniques / semidefinite programming.
 - Semidefinite relaxation: These methods are the first correlated equilibrium algorithms for infinite games which do not rely on discretization. Instead they are based on a new characterization of correlated equilibria in infinite games. For polynomial games this characterization can be shown via sum of squares techniques to be equivalent to a nested sequence of semidefinite programs, describing a nested sequence of sets converging to a description of the entire set of correlated equilibria.

While explicit performance bounds are difficult to construct, in practice the adaptive discretization and semidefinite relaxation methods seem to converge rapidly.

We close with Chapter 5 which summarizes our results and gives directions for future research.

Chapter 2

Background

2.1 Classes of Games and Equilibria

2.1.1 Finite Games

The simplest class of strategic form games (see Chapter 1) are the **finite games**, defined as those in which each player only has a finite set of possible actions. We will use this setting to present the basic definitions and theorems which we will later generalize to infinite games. Except where noted, all the theorems and proofs in this chapter are known but are included for completeness.

We begin with some notation and definitions. We denote the number of players by n , and the set of actions available to player i , the so-called **pure strategies**, by $C_i = \{1, \dots, m_i\}$. A single pure strategy for player i will usually be denoted s_i or t_i . An n -tuple of pure strategies, one for each player, is called a **strategy profile** and will be written as s , while the corresponding $(n - 1)$ -tuple of pure strategies, one for each player except i , will be written s_{-i} . The sets of all such n - and $(n - 1)$ -tuples will be denoted by C and C_{-i} , respectively.

We make the standard utility-theoretic assumption that player i 's preferences can be captured by a **utility** or **payoff function** $u_i : C \rightarrow \mathbb{R}$ such that player i prefers the outcome under strategy profile s to that under t if and only if $u_i(s) > u_i(t)$ and is indifferent between the two if and only if $u_i(s) = u_i(t)$. For simplicity of notation, we will write $u_i(t_i, s_{-i})$ for the utility of player i when he plays t_i and the other players choose their strategies according to s_{-i} .

For an example, consider the game Odds and Evens (also known as Matching Pennies), pictured in Table 2.1.1. The numbers on the left are the strategies of the row player and the numbers along the top are the strategies of the column player. The ordered pairs in the table represent the payoffs to the row and column player respectively. In this case, the goal of the row player is for the sum of both players' strategies to be odd, while the column player wants it to be even. The loser pays 1 unit of utility to the winner.

Having defined a finite game, we turn to the question of what a "solution" of the game ought to be. There are a variety of plausible **solution concepts**. The most well known, and the first one we consider, is the **Nash equilibrium**. The idea is that

	1	2
1	(-1,1)	(1,-1)
2	(1,-1)	(-1,1)

Table 2.1.1: Odds and Evens, a two-player finite game

no communication, either direct or indirect, is allowed between the players, so they will choose their strategies independently. A Nash equilibrium is a strategy profile s which is self-enforcing in the sense that no single player can unilaterally deviate from s and improve his payoff. Specifically, s must satisfy

$$u_i(t_i, s_{-i}) \leq u_i(s)$$

for all i and all $t_i \in C_i$.

None of the pure strategy profiles for Odds and Evens form a Nash equilibrium, because if $s_1 + s_2$ is odd then player 2 could gain by deviating and if it is even then player 1 could gain by deviating. However, suppose player 1 were to choose randomly between his two strategies with equal probability; this type of strategy is called a **mixed strategy**. If we assume that the players' utilities for such a random choice equal their expected utilities, then player 2 will receive a utility of zero regardless of which strategy he chooses or how he randomizes. So, if he also chooses randomly between his strategies, then he cannot gain by deviating to a different strategy. By symmetry player 1 will not be able to gain by unilaterally deviating either. Therefore the mixed strategy profile in which each player mixes among his two strategies with equal probability forms a Nash equilibrium of Odds and Evens. It is straightforward to show that this is in fact the only Nash equilibrium of Odds and Evens, but for general finite games Nash equilibria need not be unique.

This example illustrates that a Nash equilibrium in pure strategies may not exist, so in many cases we are forced to consider mixed strategies. Define $\Delta(S)$ to be the set of probability distributions over the set S , so $\Delta_i = \Delta(C_i)$ is the set of probability distributions over C_i , i.e. the set of vectors of length m_i whose components $\sigma_i^1, \dots, \sigma_i^{m_i}$ are nonnegative and sum to unity. We will use $\sigma_i \in \Delta_i$ to denote a mixed strategy for player i and define $\sigma, \sigma_{-i}, \Delta$, and Δ_{-i} analogously to the symbols s, s_{-i}, C , and C_{-i} defined above. We will identify the pure strategy $s_i \in C_i$ with the mixed strategy in which s_i is chosen with probability one, so we may view $C_i \subseteq \Delta_i$. We will also make the obvious identification of σ with the product distribution $\sigma_1 \times \dots \times \sigma_n$. Note that this means Δ represents the set of independent probability distributions over C , whereas $\Delta(C)$ represents the set of all probability distributions over C .

We will always make the standard **von Neumann-Morgenstern expected utility assumption** that each player's utility due to a choice of mixed strategies is the expected utility when all players choose according to their respective distributions in

an independent fashion [47]. We can therefore write

$$u_i(\sigma) = \int u_i(\cdot) d\sigma = \sum_{j_1=1}^{m_1} \cdots \sum_{j_n=1}^{m_n} \sigma_1^{j_1} \cdots \sigma_n^{j_n} u_i(j_1, \dots, j_n)$$

and it is natural to extend our definition of an equilibrium similarly.

Definition 2.1.1. A mixed strategy profile σ is a **Nash equilibrium** if

$$u_i(\tau_i, \sigma_{-i}) \leq u_i(\sigma)$$

for all i and all $\tau_i \in \Delta_i$.

It is easy to show that it is necessary and sufficient that this condition be satisfied for all pure strategies $\tau_i \in C_i$.

While we have shown by example that a Nash equilibrium in pure strategies need not exist, it is true that a Nash equilibrium in mixed strategies exists for all finite games. This is one of the foundational results of game theory and is known as Nash's theorem. Considering the simplicity of the definitions involved, the theorem itself is surprisingly deep. To prove it we will need the following fixed point theorem.

Theorem 2.1.2 (Brouwer's Fixed Point Theorem). *Let K be a compact convex subset of a finite dimensional real vector space and let $f : K \rightarrow K$ be continuous. Then f has a fixed point, i.e. there exists $x \in K$ such that $f(x) = x$.*

Proof. Below we give the standard proof via algebraic topology; see [19] for an introduction. There also exist (longer) purely combinatorial proofs, including one by Gale with a game theoretic flavor [14].

The proof is by contradiction and consists of two parts. The first part is elementary and amounts to showing that the nonexistence of a fixed point would imply the existence of a continuous function retracting a disk in \mathbb{R}^k onto its boundary sphere. The second part is to show that such a map cannot exist. Intuitively this is because a retraction from a disk onto its boundary must involve a "tear" in the interior of the disk. For completeness we will show this rigorously using homology, but algebraic topology will not be invoked again after this proof.

Since K is a compact convex subset of a finite dimensional vector space, it is homeomorphic to the unit k -disk D^k for some $k \geq 0$. It suffices to consider the case $K = D^k$. If $k = 0$ the statement is trivial so we assume $k \geq 1$. Suppose f has no fixed points and let S^{k-1} be the unit $(k-1)$ -sphere, the boundary of D^k . For each $x \in D^k$, draw a ray from $f(x)$ through x and let $r(x)$ denote the point where this ray intersects S^{k-1} . This defines a retraction $r : D^k \rightarrow S^{k-1}$, a continuous map which restricts to the identity $r|_{S^{k-1}} = \mathbf{1}_{S^{k-1}}$ on S^{k-1} .

Let $i : S^{k-1} \rightarrow D^k$ denote the inclusion, so $ri = \mathbf{1}_{S^{k-1}}$. The identity map on a space induces the identity map on its reduced homology groups, so $r_*i_* = (\mathbf{1}_{S^{k-1}})_*$ is the identity on $\tilde{H}_{k-1}(S^{k-1})$. Therefore $r_* : \tilde{H}_{k-1}(D^k) \rightarrow \tilde{H}_{k-1}(S^{k-1})$ is surjective, which is a contradiction since $\tilde{H}_{k-1}(D^k) \cong 0$ and $\tilde{H}_{k-1}(S^{k-1}) \cong \mathbb{Z}$. Hence r cannot exist and f must have a fixed point. \square

Theorem 2.1.3 (Nash [28]). *Every finite game has a Nash equilibrium in mixed strategies.*

Proof. Let u_1, \dots, u_n be the utilities of an n -player finite game, with $C_i = \{1, \dots, m_i\}$ for all i . Let $[x]_+ = \max(x, 0)$, which is a continuous function for $x \in \mathbb{R}$. Define a continuous function $f : \Delta \rightarrow \Delta$ which maps n -tuples of mixed strategies to n -tuples of mixed strategies as follows:

$$[f(\sigma)]_i^k = \frac{\sigma_i^k + [u_i(k, \sigma_{-i}) - u_i(\sigma)]_+}{1 + \sum_{j=1}^{m_i} [u_i(j, \sigma_{-i}) - u_i(\sigma)]_+}$$

The function f has a fixed point by Theorem 2.1.2. Let σ be any such fixed point. Suppose for a contradiction that σ is not a Nash equilibrium, so there exist i and k such that $u_i(k, \sigma_{-i}) > u_i(\sigma)$. Then $\sum_{j=1}^{m_i} [u_i(j, \sigma_{-i}) - u_i(\sigma)]_+ > 0$ and since $\sigma_i^l = [f(\sigma)]_i^l$ for all l by definition of a fixed point, we have

$$\sigma_i^l = \frac{[u_i(l, \sigma_{-i}) - u_i(\sigma)]_+}{\sum_{j=1}^{m_i} [u_i(j, \sigma_{-i}) - u_i(\sigma)]_+}$$

for all l . In particular $\sigma_i^l > 0$ implies that $u_i(l, \sigma_{-i}) > u_i(\sigma)$, so

$$u_i(\sigma) = \sum_{l=1}^{m_i} \sigma_i^l u_i(l, \sigma_{-i}) = \sum_{l: \sigma_i^l > 0} \sigma_i^l u_i(l, \sigma_{-i}) > \sum_{l: \sigma_i^l > 0} \sigma_i^l u_i(\sigma) = u_i(\sigma),$$

a contradiction. Therefore σ is a Nash equilibrium. \square

Since this proof is nonconstructive it doesn't immediately provide any obvious algorithm for computing Nash equilibria of finite games. In Section 2.2 we will review several such algorithms, including the Lemke-Howson algorithm which provides a constructive proof of Nash's Theorem in the two-player case.

Now we turn to another well-known solution concept, that of a correlated equilibrium. The idea is that even if direct communication between the players is forbidden, they will still observe some common information, such as stories in the news, stock quotes, traffic signals, and the weather. If we assume that the players' strategy choices are random variables which depend on such environmental data, the players' actions will be correlated random variables distributed according to some joint distribution $\pi \in \Delta(C)$.

Let R be a random variable distributed according to π . A realization of R is a pure strategy profile and the i^{th} component of the realization will be called the recommendation to player i . Given such a recommendation, player i can use conditional probability to form a posteriori beliefs about the recommendations given to the other players. A distribution π is defined to be a correlated equilibrium if no player can ever expect to unilaterally gain by deviating from his recommendation, assuming the other players play according to their recommendations.

Definition 2.1.4. A **correlated equilibrium** of a finite game is a joint probability distribution $\pi \in \Delta(C)$ such that if R is a random variable distributed according to π

then

$$\sum_{s_{-i} \in C_{-i}} \text{Prob}(R = s | R_i = s_i) [u_i(t_i, s_{-i}) - u_i(s)] \leq 0 \quad (2.1)$$

for all players i , all $s_i \in C_i$ such that $\text{Prob}(R_i = s_i) > 0$, and all $t_i \in C_i$.

While this definition captures the idea we want, the following characterization is easier to apply and visualize.

Proposition 2.1.5. *A joint distribution $\pi \in \Delta(C)$ is a correlated equilibrium of a finite game if and only if*

$$\sum_{s_{-i} \in C_{-i}} \pi(s) [u_i(t_i, s_{-i}) - u_i(s)] \leq 0 \quad (2.2)$$

for all players i and all $s_i, t_i \in C_i$ such that $s_i \neq t_i$.

Proof. Using the definition of conditional probability we can rewrite the definition of a correlated equilibrium as the condition that

$$\sum_{s_{-i} \in C_{-i}} \frac{\pi(s)}{\sum_{t_{-i} \in C_{-i}} \pi(s_i, t_{-i})} [u_i(t_i, s_{-i}) - u_i(s)] \leq 0$$

for all i , all $s_i \in C_i$ such that $\sum_{t_{-i} \in C_{-i}} \pi(s_i, t_{-i}) > 0$, and all $t_i \in C_i$. The denominator does not depend on the variable of summation so it can be factored out of the sum and cancelled, yielding the simpler condition that (2.2) holds for all i , all $s_i \in C_i$ such that $\sum_{t_{-i} \in C_{-i}} \pi(s_i, t_{-i}) > 0$, and all $t_i \in C_i$. But if $\sum_{t_{-i} \in C_{-i}} \pi(s_i, t_{-i}) = 0$ then the left hand side of (2.2) is zero regardless of i and t_i , so the equation always holds trivially in this case. Equation (2.2) also holds trivially when $s_i = t_i$, so we only need to check it in the case $s_i \neq t_i$. \square

We can also think of correlated equilibria as joint distributions corresponding to recommendations which will be given to the players as part of an extended game. The players are then free to play any function of their recommendation (this is called a **departure function**) as their strategy in the game. If it is a (pure strategy) Nash equilibrium of this extended game for each player to play his recommended strategy, then the distribution is a correlated equilibrium. This interpretation is justified by the following alternative characterization of correlated equilibria.

Proposition 2.1.6. *A joint distribution $\pi \in \Delta(C)$ is a correlated equilibrium of a finite game if and only if*

$$\sum_{s \in C} \pi(s) [u_i(\zeta_i(s_i), s_{-i}) - u_i(s)] \leq 0 \quad (2.3)$$

for all players i and all functions $\zeta_i : C_i \rightarrow C_i$ (called *departure functions* in this setting).

Proof. By substituting $t_i = \zeta_i(s_i)$ into (2.2) and summing over all $s_i \in C_i$ we obtain (2.3) for any i and any $\zeta_i : C_i \rightarrow C_i$. For the converse, define ζ_i for any $s_i, t_i \in C_i$ by

$$\zeta_i(r_i) = \begin{cases} t_i & r_i = s_i \\ r_i & \text{else.} \end{cases}$$

Then all the terms in (2.3) except the s_i terms cancel yielding (2.2). □

These two propositions both show that the set of correlated equilibria is defined by a finite number of linear equations and inequalities (those in (2.2) or (2.3) along with $\pi(s) \geq 0$ for all $s \in C$ and $\sum_{s \in C} \pi(s) = 1$) and is therefore polyhedral. Next we show that the set of correlated equilibria is also nonempty.

It can be shown that Nash equilibria are the same as correlated equilibria which are product distributions, so Theorem 2.1.3 immediately implies the existence of a correlated equilibrium. However, it is also possible to prove this theorem in an elementary fashion without resorting to a fixed point argument as used in Theorem 2.1.3, and for completeness we will do so here.

Theorem 2.1.7. *Every finite game has a correlated equilibrium.*

Proof. The proof follows [18] and [30] and consists of two applications of linear programming duality. Consider the linear program

$$\begin{aligned} \max \quad & \sum_{s \in C} \pi(s) \\ \text{s.t.} \quad & \sum_{s_{-i} \in C_{-i}} \pi(s) [u_i(t_i, s_{-i}) - u_i(s)] \leq 0 \quad \text{for all } 1 \leq i \leq n \text{ and } s_i, t_i \in C_i \\ & \pi(s) \geq 0 \quad \text{for all } s \in C. \end{aligned} \tag{2.4}$$

Regardless of the choice of utility functions $\pi \equiv 0$ is feasible, so (2.4) is always feasible. If the LP is unbounded then there exists a feasible solution with positive objective value. Any such solution can be scaled to produce a feasible solution with objective value 1, which is a correlated equilibrium by Proposition 2.1.5. Therefore if (2.4) is unbounded then the game has a correlated equilibrium.

Since (2.4) is feasible, strong linear programming duality implies that it is unbounded if the dual linear program

$$\begin{aligned} \min \quad & 0 \\ \text{s.t.} \quad & \sum_i \sum_{t_i \in C_i} y_i^{s_i t_i} [u_i(t_i, s_{-i}) - u_i(s)] \geq 1 \quad \text{for all } s \in C \\ & y_i^{s_i t_i} \geq 0 \quad \text{for all } 1 \leq i \leq n \text{ and } s_i, t_i \in C_i \end{aligned} \tag{2.5}$$

is infeasible [3].

Suppose for a contradiction that (2.5) is feasible and let y be some fixed feasible solution. We will show that there exists a product probability distribution $\pi^y(s) =$

$\pi_1^y(s_1) \cdots \pi_n^y(s_n)$ satisfying

$$\sum_{s \in C} \sum_i \sum_{t_i \in C_i} \pi^y(s) y_i^{s_i t_i} [u_i(t_i, s_{-i}) - u_i(s)] = 0 \quad (2.6)$$

which shows that the chosen y is in fact not a feasible solution of (2.5) to obtain the desired contradiction. For any $1 \leq i \leq n$ and $s \in C$, the coefficient of $u_i(s)$ in the left hand side of (2.6) is

$$\sum_{t_i \in C_i} [\pi^y(t_i, s_{-i}) y_i^{t_i s_i} - \pi^y(s) y_i^{s_i t_i}] \quad (2.7)$$

If for each i we can construct a probability vector π_i^y satisfying

$$\sum_{t_i \in C_i} \pi_i^y(t_i) y_i^{t_i s_i} - \pi_i^y(s_i) \sum_{t_i \in C_i} y_i^{s_i t_i} = 0 \quad (2.8)$$

for all $s_i \in C_i$, then the product distribution $\pi^y(s) = \pi_1^y(s_1) \cdots \pi_n^y(s_n)$ would make all the terms of the form (2.7) equal to zero and hence it would satisfy (2.6) because the coefficient of $u_i(s)$ would be zero for all i and all $s \in C$.

Now for each i consider the linear program

$$\begin{aligned} \max \quad & \sum_{s_i \in C_i} \pi_i^y(s_i) \\ \text{s.t.} \quad & \sum_{t_i \in C_i} \pi_i^y(t_i) y_i^{t_i s_i} - \pi_i^y(s_i) \sum_{t_i \in C_i} y_i^{s_i t_i} = 0 \quad \text{for all } s_i \in C_i \\ & \pi_i^y(s_i) \geq 0 \quad \text{for all } s_i \in C_i. \end{aligned} \quad (2.9)$$

The constraints are always feasible as the assignment $\pi_i^y \equiv 0$ shows. If (2.9) is unbounded then any feasible solution with positive objective value can be scaled to yield one with unit objective value, which is exactly a probability vector satisfying (2.8). Therefore to prove the existence of such a probability vector and obtain our desired contradiction, it suffices to show that (2.9) is unbounded. Since it is feasible, we can again do this using linear programming duality by proving that the dual linear program

$$\begin{aligned} \min \quad & 0 \\ \text{s.t.} \quad & \sum_{s_i \in C_i} (z_i^{t_i} - z_i^{s_i}) y_i^{s_i t_i} \geq 1 \quad \text{for all } t_i \in C_i \end{aligned} \quad (2.10)$$

is infeasible. Suppose for a contradiction that (2.10) were feasible and let z_i be a feasible vector. Let $t_i \in C_i$ be an index such that $z_i^{t_i}$ is smallest. Then $z_i^{t_i} - z_i^{s_i} \leq 0$ for all s_i and $y_i^{s_i t_i} \geq 0$ for all s_i by definition of y , so $\sum_{s_i \in C_i} (z_i^{t_i} - z_i^{s_i}) y_i^{s_i t_i} \leq 0$, contradicting the feasibility of z_i .

This proves that (2.10) is infeasible, so (2.9) is unbounded and hence there is a probability vector satisfying (2.8) for each i . Therefore the product probability distribution $\pi^y(s) = \pi_1^y(s_1) \cdots \pi_n^y(s_n)$ satisfies (2.6), contradicting the feasibility of y and proving that (2.5) is infeasible. Thus (2.4) is unbounded, so in particular must

have a feasible solution with unit objective value, which is a correlated equilibrium by Proposition 2.1.5. \square

It is worth noting that in the proof if we had constructed a product distribution π^y satisfying (2.6) *which did not depend on y* then it would have to be a feasible solution of (2.4), hence it would be a correlated equilibrium. Being a product distribution it would also be a Nash equilibrium, and we would obtain a simple proof of the existence of a Nash equilibrium. However, the construction shows that π^y does depend on y and therefore need not be a feasible solution of (2.4). For this reason the proof given is not strong enough to imply the existence of a Nash equilibrium.

2.1.2 Continuous Games

Having discussed the basic definitions and existence theorems in the context of finite games, we now move to the larger class of continuous games. We will see that similar theorems are true in this context, though several complications arise.

Definition 2.1.8. A **continuous game** is an n -player game in which the i^{th} player's strategy space C_i is a compact metric space and the utility functions $u_i : C \rightarrow \mathbb{R}$ are continuous.

Note that any finite set is a compact metric space under the discrete metric and any function whose domain is endowed with the discrete metric is automatically continuous. Therefore all finite games are continuous games. Of course there are many other examples, such as:

Definition 2.1.9. A **polynomial game** is an n -player game in which all the strategy spaces are compact intervals in \mathbb{R} and the utility functions are multivariate polynomials in the strategies.

We have seen in the finite case that pure strategies do not always suffice to ensure the existence of Nash equilibria, so we must consider probability distributions. For this reason we define Δ_i to be the set of Borel probability measures over the compact metric space C_i . The condition that the measures be Borel is added here so that the measures will be compatible with the topology on the strategy space. We will identify each pure strategy $s_i \in C_i$ with the corresponding atomic probability distribution which assigns unit probability to s_i . We can then view C_i as a subset of Δ_i . Having made this identification, the convex hull $\text{conv } C_i \subseteq \Delta_i$ represents the set of finitely supported probability measures, i.e. those which assign probability unity to some finite set. We will define the topology on Δ_i as follows.

Definition 2.1.10. The **weak* topology** on Δ_i is defined to be the weakest topology which makes $\int f(\cdot) d\sigma_i$ a continuous function of σ_i whenever $f : C_i \rightarrow \mathbb{R}$ is continuous.

The weak* topology is defined in general for the dual of a Banach space; in this case we are viewing Δ_i as a subset of the dual of the Banach space of continuous real-valued functions on C_i . This dual is the space V_i of finite signed measures on C_i , a corollary of the Riesz representation theorem, and hence contains Δ_i [38]. We will need the following properties of this topology, which are standard results [34].

Proposition 2.1.11. *The weak* topology makes Δ_i a compact metric space with $\text{conv } C_i$ a dense subset.*

Again we will make the standard von Neumann-Morgenstern expected utility assumption so we can write

$$u_i(\sigma) = \int u_i(\cdot) d\sigma$$

for all i and all $\sigma \in \Delta$, where we have identified the mixed strategy profile $\sigma = (\sigma_1, \dots, \sigma_n)$ with the corresponding product distribution $\sigma_1 \times \dots \times \sigma_n$ as above. We can now define a Nash equilibrium of a continuous game in a way which is exactly analogous to the finite game case.

Definition 2.1.12. A strategy profile σ is an ϵ -Nash equilibrium if $u_i(\tau_i, \sigma_{-i}) \leq u_i(\sigma) + \epsilon$ for all i and all $\tau_i \in \Delta_i$. If $\epsilon = 0$ we call this a **Nash equilibrium**.

As in the finite case, it is straightforward to show that it suffices to check this condition for all pure strategies $\tau_i \in C_i$; the other $\tau_i \in \Delta_i$ will then satisfy it automatically. Also like the finite case, it is true that every continuous game has a Nash equilibrium. This can be proven directly by a fixed point argument using a generalization of Brouwer's fixed point theorem, but here we will prove it as a consequence of Nash's theorem.

Lemma 2.1.13. *Every continuous game has an ϵ -Nash equilibrium for all $\epsilon > 0$.*

Proof. Fix $\epsilon > 0$. Since the utilities u_i are continuous functions defined on a compact metric space, they are uniformly continuous. Thus we can find a $\delta > 0$ such that $|u_i(s_i, s_{-i}) - u_i(s'_i, s_{-i})| < \epsilon$ whenever the distance $d(s_i, s'_i) < \delta$ for $s_i, s'_i \in C_i$, for all players i and all $s_{-i} \in C_{-i}$. This relation also holds with s_{-i} replaced by a mixed strategy σ_{-i} because

$$\begin{aligned} |u_i(s_i, \sigma_{-i}) - u_i(s'_i, \sigma_{-i})| &= \left| \int u_i(s_i, \cdot) d\sigma_{-i}(\cdot) - \int u_i(s'_i, \cdot) d\sigma_{-i}(\cdot) \right| \leq \\ &\int |u_i(s_i, \cdot) - u_i(s'_i, \cdot)| d\sigma_{-i}(\cdot) < \epsilon \int d\sigma_{-i} = \epsilon \end{aligned} \quad (2.11)$$

by the triangle inequality for integrals.

Since each C_i is a compact metric space, we can choose a finite subset $D_i \subseteq C_i$ such that for all $s_i \in C_i$ there is an $s'_i \in D_i$ with $d(s_i, s'_i) < \delta$. Consider the n -player finite game with strategy spaces D_i and utilities $u_i|_D$. Theorem 2.1.3 states that this game has a Nash equilibrium, so there exist mixed strategies $\sigma_i \in \Delta(D_i) \subseteq \Delta_i$ such that $u_i(t_i, \sigma_{-i}) \leq u_i(\sigma)$ for all i and all $t_i \in D_i$. Let $s_i \in C_i$ be arbitrary and let $s'_i \in D_i$ be chosen so that $d(s_i, s'_i) < \delta$. Then

$$\begin{aligned} u_i(s_i, \sigma_{-i}) - u_i(\sigma) &= u_i(s_i, \sigma_{-i}) - u_i(s'_i, \sigma_{-i}) + u_i(s'_i, \sigma_{-i}) - u_i(\sigma) \\ &\leq u_i(s_i, \sigma_{-i}) - u_i(s'_i, \sigma_{-i}) \leq \epsilon, \end{aligned}$$

where the first inequality follows because σ is a Nash equilibrium of the game when the strategies are restricted to the sets D_i and the second inequality follows from (2.11). Thus σ is an ϵ -Nash equilibrium. \square

Theorem 2.1.14 (Glicksberg [17]). *Every continuous game has a Nash equilibrium in mixed strategies.*

Proof. Choose any sequence of positive reals $\epsilon^k \rightarrow 0$. By Lemma 2.1.13 there exists an ϵ^k -Nash equilibrium σ^k for each k . Proposition 2.1.11 implies that by passing to a subsequence if necessary we can assume σ_i^k weak* converges to some σ_i for each i . It is straightforward to see that σ^k weak* converges to σ and σ_{-i}^k weak* converges to σ_{-i} for all i . Therefore

$$u_i(s_i, \sigma_{-i}) - u_i(\sigma) = \lim_{k \rightarrow \infty} [u_i(s_i, \sigma_{-i}^k) - u_i(\sigma^k)] \leq \lim_{k \rightarrow \infty} \epsilon^k = 0$$

for all i and $s_i \in C_i$, so σ is a Nash equilibrium. \square

Theorem 2.1.14 is the main reason we have chosen to consider the class of continuous games. If either the continuity or compactness conditions are removed from the definition of a continuous game, then there are games with no mixed strategy Nash equilibria.

Example 2.1.15. To construct such a game it suffices to consider the one-player case, i.e. optimization problems. Consider the game with compact strategy space $C_1 = [0, 1]$ and discontinuous utility function

$$u_1(s_1) = \begin{cases} s_1, & s_1 < 1; \\ 0, & s_1 = 1. \end{cases}$$

Then for any $\sigma_1 \in \Delta_1$ we have $u_1(\sigma_1) < 1$ since $u_1(\sigma_1) \geq 1$ would imply that $u_1(s_1) \geq 1$ for some $s_1 \in C_1$, a contradiction. Therefore for any $\sigma_1 \in \Delta_1$ we can find s_1 such that $u_1(\sigma_1) < s_1 < 1$. Then $u_1(\sigma_1) < u_1(s_1)$, so σ_1 is not a Nash equilibrium and the game has no Nash equilibria. The same argument shows that the game with noncompact strategy space $C_1 = [0, 1)$ and continuous utility function $u_1(s_1) = s_1$ has no Nash equilibrium. Furthermore this argument goes through for any one-player game in which u_1 is bounded above but does not achieve its supremum on C_1 . For multi-player games with this property Nash equilibria can exist, however.

Several authors have taken up the problem of finding conditions to guarantee the existence of equilibria in games with discontinuous utility functions. Motivated by several specific economic examples, Dasgupta and Maskin study weaker semicontinuity-type properties which suffice to ensure the existence of a Nash equilibrium [8], but are not satisfied by this example. There is also a novel solution by Simon and Zame in which a game with discontinuous utility functions is considered to be underspecified; nonexistence of an equilibrium is then taken to be a flaw in the model of the game rather than the equilibrium concept [43]. In particular, the utilities are thought of as being indeterminate at the points of discontinuity, and it is shown that under mild topological conditions, the utilities can be (re)defined at these points in such a way that an equilibrium exists.

Defining correlated equilibria in continuous games requires somewhat more care than in finite games. The standard definition as used in e.g. [18] is a straightforward generalization of the characterization of correlated equilibria for finite games

in Proposition 2.1.6. In this case we must add the additional assumption that the departure functions be Borel measurable to ensure that the integrals are defined.

Definition 2.1.16. A **correlated equilibrium** of a continuous game is a joint probability measure $\pi \in \Delta(C)$ such that

$$\int [u_i(\zeta_i(s_i), s_{-i}) - u_i(s)] d\pi(s) \leq 0$$

for all i and all Borel measurable functions $\zeta_i : C_i \rightarrow C_i$.

The obvious approach to proving the existence of correlated equilibria in continuous games would be to follow a discretization and limiting argument as used to prove the existence of Nash equilibria in Theorem 2.1.14. Indeed, if we define an ϵ -**correlated equilibrium** by changing the zero in Definition 2.1.16 to an ϵ , then the analog of Lemma 2.1.13 can be proven using the same argument, invoking the existence theorem for correlated equilibria of finite games instead of the existence theorem for Nash equilibria. However, the same limiting argument does not go through because the integrands in the the definition of a correlated equilibrium are not continuous functions. Even though $\int f d\sigma^k \rightarrow \int f d\sigma$ holds for all continuous functions f when σ^k weak* converges to σ , this may fail for discontinuous functions f .

It is possible to make this proof technique work by using a more complicated limiting argument as in [18], but we will not do so here. Instead we will prove the existence of correlated equilibria by proving that Nash equilibria are also correlated equilibria.

Theorem 2.1.17. *Every continuous game has a correlated equilibrium.*

Proof. Let σ be a Nash equilibrium of the given continuous game, which exists by Theorem 2.1.14. Define $\pi = \sigma_1 \times \dots \times \sigma_n$. Choose an arbitrary player i and any measurable departure function $\zeta_i : C_i \rightarrow C_i$. By definition of a Nash equilibrium we have $u_i(t_i, \sigma_{-i}) \leq u_i(\sigma)$ for all $t_i \in C_i$. Letting $t_i = \zeta_i(s_i)$ and integrating with respect to the probability measure $\sigma_i(s_i)$ we obtain $\int u_i(\zeta_i(s_i), \sigma_{-i}) d\sigma_i(s_i) \leq u_i(\sigma)$. Applying Fubini's theorem [38] and the definition of π yields the defining condition for a correlated equilibrium. \square

Despite the fact that Nash and correlated equilibria are guaranteed to exist in continuous games, the equilibrium measures may be arbitrarily complicated and so are not in general amenable to computation. To demonstrate this, we will explicitly construct continuous games whose only equilibria are complex, following the construction of games with prescribed Nash equilibria in [21].

The observation and proof below that the same construction yields a unique correlated equilibrium appear to be new. Since we have proven that every continuous game has a Nash equilibrium and that Nash equilibria are correlated equilibria which are product measures, it is only possible to exhibit a continuous game with a unique correlated equilibrium π if π is a product measure.

Theorem 2.1.18. *Let σ_1 and σ_2 be any probability measures on $[0, 1]$ which are not finitely supported, i.e. for which $\sigma_1(F) < 1$ and $\sigma_2(F) < 1$ for all finite sets $F \subset [0, 1]$. Then there exists a two-player continuous game with strategy spaces $C_1 = C_2 = [0, 1]$ whose unique Nash equilibrium is (σ_1, σ_2) and whose unique correlated equilibrium is $\pi = \sigma_1 \times \sigma_2$.*

Proof. Define the moments of σ_i by $\sigma_i^k = \int s_i^k d\sigma_i(s_i)$, with the convention that the k in σ_i^k is an index and the k in s_i^k is an exponent. Clearly $0 \leq \sigma_i^k \leq 1$ for all i and k . Define utilities by

$$u_1(s_1, s_2) = -u_2(s_1, s_2) = \sum_{k=1}^{\infty} \frac{1}{2^k} (s_1^k - \sigma_1^k)(s_2^k - \sigma_2^k).$$

The terms in parentheses always have absolute value less than or equal to unity for $s_1, s_2 \in [0, 1]$, so the series converges uniformly as a function of s_1 and s_2 and the convergence is absolute. Thus the utilities are continuous and we may apply Fubini's theorem to compute

$$\begin{aligned} u_1(s_1, \sigma_2) &= \int \sum_{k=1}^{\infty} \frac{1}{2^k} (s_1^k - \sigma_1^k)(s_2^k - \sigma_2^k) d\sigma_2(s_2) = \sum_{k=1}^{\infty} \frac{1}{2^k} (s_1^k - \sigma_1^k) \int (s_2^k - \sigma_2^k) d\sigma_2(s_2) \\ &= \sum_{k=1}^{\infty} \frac{1}{2^k} (s_1^k - \sigma_1^k)(0) = 0 \end{aligned}$$

for all $s_1 \in [0, 1]$. Therefore $u_1(s_1, \sigma_2) - u_1(\sigma_1, \sigma_2) = 0 - 0 \leq 0$. The same argument works with the players interchanged, so (σ_1, σ_2) is a Nash equilibrium of this game.

Next we prove that (σ_1, σ_2) is the unique Nash equilibrium. Let (τ_1, τ_2) be any Nash equilibrium of the game. Consider the function $f(s_1) = u_1(s_1, \tau_2) = \sum_{k=1}^{\infty} \frac{\tau_2^k - \sigma_2^k}{2^k} (s_1^k - \sigma_1^k)$. By definition of a Nash equilibrium $\int f d\tau_1 \geq \int f d\sigma_1 = u_1(\sigma_1, \tau_2) = 0$. Similarly $-\int f d\tau_1 = u_2(\tau_1, \tau_2) \geq u_2(\tau_1, \sigma_2) = 0$, so $\int f d\tau_1 = \int f d\sigma_1 = 0$. By definition of a Nash equilibrium, $f(s_1) = u_1(s_1, \tau_2) \leq u_1(\tau_1, \tau_2) = \int f d\tau_1 = 0$ for all $s_1 \in [0, 1]$.

The function f is defined by a power series which converges for $s_1 \in [-2, 2]$ and hence is analytic in a neighborhood of $[0, 1]$. Therefore the set $F \subseteq [0, 1]$ of zeros of the function f on the interval $[0, 1]$ is either finite or equal to $[0, 1]$. Suppose for a contradiction that F is finite. Since F is finite we have $\sigma_1([0, 1] \setminus F) > 0$ by assumption. But $f(s_1) < 0$ for $s_1 \in [0, 1] \setminus F$, so

$$\int f d\sigma_1 = \int_F f d\sigma_1 + \int_{[0,1] \setminus F} f d\sigma_1 = \int_{[0,1] \setminus F} f d\sigma_1 < 0,$$

a contradiction. Therefore $f \equiv 0$ and since f is given by a power series, all the coefficients in the series must be zero. That means that $\tau_2^k = \sigma_2^k$ for all $k \geq 1$. In fact this equation holds for all $k \geq 0$ because both σ_2 and τ_2 are probability measures, hence satisfy $\tau_2^0 = \sigma_2^0 = 1$. It is a straightforward consequence of the Weierstrass approximation theorem [37] and the Riesz representation theorem [38] that equality

of all the moments for $k \geq 0$ implies equality of the measures $\sigma_2 = \tau_2$. Interchanging players and applying the same argument shows that $\sigma_1 = \tau_1$.

Finally we prove the uniqueness of the correlated equilibrium. Let π be any correlated equilibrium of the game. Then taking $\zeta_1(s_1) \equiv t_1$ in the definition of a correlated equilibrium and defining the marginal $\pi_2(B) = \pi(C_1 \times B)$ for all Borel sets $B \subseteq [0, 1]$ we have

$$u_1(t_1, \pi_2) = \int u_1(t_1, s_2) d\pi(s) \leq \int u_1 d\pi$$

Integrating both sides with respect to $\sigma_1(t_1)$ yields $0 = u_1(\sigma_1, \pi_2) \leq \int u_1 d\pi$. Repeating this analysis with the players interchanged shows that $0 \leq \int u_2 d\pi = -\int u_1 d\pi$, so $\int u_1 d\pi = 0$. Therefore $u_1(t_1, \pi_2) \leq 0$ for all $t_1 \in [0, 1]$, so (σ_1, π_2) is a Nash equilibrium of the game, and by the uniqueness of the Nash equilibrium we have $\pi_2 = \sigma_2$. Similarly if we define the marginal π_1 by $\pi_1(B) = \pi(B \times C_2)$ for all Borel sets B then $\pi_1 = \sigma_1$.

The probability measure $\sigma_1 \times \sigma_2$ is uniquely defined by the condition that $(\sigma_1 \times \sigma_2)(B_1 \times B_2) = \sigma_1(B_1)\sigma_2(B_2)$ for all Borel sets B_1 and B_2 [13]. Therefore to prove $\pi = \sigma_1 \times \sigma_2$ it suffices to prove that $\pi(B_1 \times B_2) = \sigma_1(B_1)\sigma_2(B_2)$. If $\sigma_1(B_1) = 0$ then we have

$$0 \leq \pi(B_1 \times B_2) \leq \pi(B_1 \times C_2) = \pi_1(B_1) = \sigma_1(B_1) = 0$$

so $0 = \pi(B_1 \times B_2) = \sigma_1(B_1)\sigma_2(B_2)$ holds in this case. On the other hand, suppose $\sigma_1(B_1) > 0$. For $t_1, t'_1 \in [0, 1]$, define

$$\zeta_1(s_1) = \begin{cases} t_1 & s_1 \in B_1 \\ t'_1 & s_1 \notin B_1 \end{cases}$$

Then by the definition of a correlated equilibrium we have

$$\int u_1(\zeta_1(s_1), s_2) d\pi(s) = \int_{B_1 \times C_2} u_1(t_1, s_2) d\pi(s) + \int_{B_1^c \times C_2} u_1(t'_1, s_2) d\pi(s) \leq \int u_1 d\pi = 0$$

Integrating with respect to the probability measure $\sigma_1(t'_1)$ and applying Fubini's theorem yields $\int_{B_1 \times C_2} u_1(t_1, s_2) d\pi(s) \leq 0$ since $u_1(\sigma_1, s_2) = 0$ for all s_2 . Now define the measure $\pi_2^{B_1}(B_2) = \pi(B_1 \times B_2)$ for all Borel sets B_2 . Then $\pi_2^{B_1}(C_2) = \pi_1(B_1) = \sigma_1(B_1) > 0$, so we can define the probability measure $\tilde{\pi}_2^{B_1} = \frac{1}{\sigma_1(B_1)} \pi_2^{B_1}$. Under this definition

$$u_1(t_1, \tilde{\pi}_2^{B_1}) = \frac{1}{\sigma_1(B_1)} \int u_1(t_1, s_2) d\pi_2^{B_1}(s_2) = \int_{B_1 \times C_2} u_1(t_1, s_2) d\pi(s) \leq 0$$

for all $t_1 \in [0, 1]$, so $(\sigma_1, \tilde{\pi}_2^{B_1})$ is a Nash equilibrium. By the uniqueness of the Nash equilibrium we have $\tilde{\pi}_2^{B_1} = \sigma_2$. Substituting in the definition of $\tilde{\pi}_2^{B_1}$ shows that $\pi(B_1 \times B_2) = \sigma_1(B_1)\sigma_2(B_2)$. This equation holds for all Borel sets B_1, B_2 , hence $\pi = \sigma_1 \times \sigma_2$. \square

An analogous construction exists if σ_1 and σ_2 are finitely supported, but we do not include it here [15].

2.1.3 Separable Games

In general there is no way to represent an arbitrary probability measure over $[0, 1]$ on a computer; there is simply too much information. Even approximating an arbitrary measure to some specified degree of accuracy could require many bits, depending on the measure. In light of Theorem 2.1.18 we should not expect computing exact equilibria of arbitrary continuous games to be possible on a computer and even approximating equilibria may be difficult. Therefore we will focus on a class of continuous games in which the players may restrict their choice of mixed strategies to a set of strategies which admit “simple” descriptions. In particular these will be the strategies in which the players randomize over some finite set of bounded cardinality.

Definition 2.1.19. A **separable game** is a continuous game with utility functions $u_i : C \rightarrow \mathbb{R}$ taking the form

$$u_i(s) = \sum_{j_1=1}^{m_1} \cdots \sum_{j_n=1}^{m_n} a_i^{j_1 \cdots j_n} f_1^{j_1}(s_1) \cdots f_n^{j_n}(s_n), \quad (2.12)$$

where $a_i^{j_1 \cdots j_n} \in \mathbb{R}$ and the $f_i^j : C_i \rightarrow \mathbb{R}$ are continuous.

Every finite set is a compact metric space under the discrete metric and any function from a finite set to \mathbb{R} is continuous and can be written in the form (2.12) by taking the f_i^j to be Kronecker δ functions, so finite games are automatically separable games. Another important example is the class of polynomial games defined above. When it is convenient to do so, and always for polynomial games, we will begin the summations in (2.12) at $j_i = 0$. For polynomial games we can then use the convention that $f_i^j(s_i) = s_i^j$, where the superscript on the right hand side denotes an exponent rather than an index.

Example 2.1.20. Consider a two player game with $C_1 = C_2 = [-1, 1] \subset \mathbb{R}$. Letting x and y denote the pure strategies of players 1 and 2, respectively, we define the utility functions

$$\begin{aligned} u_1(x, y) &= 2xy + 3y^3 - 2x^3 - x - 3x^2y^2, \\ u_2(x, y) &= 2x^2y^2 - 4y^3 - x^2 + 4y + x^2y. \end{aligned} \quad (2.13)$$

This is a polynomial game, and we will return to it periodically to apply the results presented.

Given a separable game and a mixed strategy $\sigma_i \in \Delta_i$ we will define the **generalized moments** $\sigma_i^k = \int f_i^k d\sigma_i$. Sometimes we will instead write $f_i^k(\sigma_i)$ for σ_i^k . In the case of polynomial games, these are just the classical moments. Having defined the moments this way, the equation (2.12) holds when s is replaced by a mixed strategy profile σ . Analyzing the sets $f_i(\Delta_i)$ of all possible moments due to measures in Δ_i

will allow us to prove general theorems about separable games. See Figure 2.1.1 for an example of such a set.

If two mixed strategies $\sigma_i, \tau_i \in \Delta_i$ satisfy $\sigma_i^k = \tau_i^k$ for $1 \leq k \leq m_i$ we will call the strategies **moment equivalent**. Then $u_j(\sigma_i, \sigma_{-i}) = u_j(\tau_i, \sigma_{-i})$ for all players j and all $\sigma_{-i} \in \Delta_{-i}$, so σ_i and τ_i are also **payoff equivalent** in the sense that switching between σ_i and τ_i doesn't affect the outcome of the game in any way. In separable games, every mixed strategy is moment equivalent to a "simple" mixed strategy, as we prove below. This theorem has been proven using a separating hyperplane argument [21], but here we give a new topological proof.

Theorem 2.1.21. *In a separable game every mixed strategy σ_i is moment equivalent to a finitely-supported mixed strategy τ_i with $|\text{supp}(\tau_i)| \leq m_i + 1$. Moreover, if σ_i is countably-supported τ_i can be chosen with $\text{supp}(\tau_i) \subset \text{supp}(\sigma_i)$.*

Proof. Note that the map

$$f_i : \sigma_i \mapsto (f_i^1(\sigma_i), \dots, f_i^{m_i}(\sigma_i))$$

is linear and weak* continuous. Therefore

$$f_i(\Delta_i) = f_i(\overline{\text{conv } C_i}) \subseteq \overline{f_i(\text{conv } C_i)} = \overline{\text{conv } f_i(C_i)} = \text{conv } f_i(C_i) = f_i(\text{conv } C_i) \subseteq f_i(\Delta_i).$$

The first three steps follow from Proposition 2.1.11, continuity of f_i , and linearity of f_i , respectively. The next equality holds because $\text{conv } f_i(C_i)$ is compact, being the convex hull of a compact subset of a finite-dimensional space. The final two steps follow from linearity of f_i and the containment $\text{conv } C_i \subseteq \Delta_i$, so we have $f_i(\Delta_i) = \text{conv } f_i(C_i) = f_i(\text{conv } C_i)$. Thus any mixed strategy is moment equivalent to a finitely-supported mixed strategy, and applying Carathéodory's theorem [2] to the set $\text{conv } f_i(C_i) \subset \mathbb{R}^{m_i}$ yields the uniform bound. Since a countable convex combination of points in a bounded subset of \mathbb{R}^{m_i} can always be written as a finite convex combination of at most $m_i + 1$ of those points, the final claim follows. \square

This theorem allows the players in a separable game to restrict their attention to "simple" strategies, regardless of what the other players do. If we replace the strategies in a Nash equilibrium by payoff equivalent strategies, we obtain a Nash equilibrium. Therefore combining Theorem 2.1.21 with Theorem 2.1.14 yields the following.

Corollary 2.1.22. *Every separable game has a Nash equilibrium in which player i mixes among at most $m_i + 1$ pure strategies.*

This is the most fundamental result about separable games. We will construct tighter bounds on the size of the support of Nash and correlated equilibria in Chapter 3, but this corollary is what makes computing or approximating equilibria theoretically possible.

Example 2.1.20 (cont'd). Apply the standard definition of the f_i^j to the polynomial game with payoffs given in (2.13). The set of moments $f_i(\Delta_i)$ as described in Theorem 2.1.21 is shown in Figure 2.1.1. In this case the set is the same for both players.




Figure 2.1.1: The space $f_i(\Delta_i)$ of possible moments for either player's mixed strategy under the payoffs given in (2.13) due to a measure σ_i on $[-1, 1]$. The zeroth moment, which is identically unity, has been omitted.

For each player the range of the indices in (2.12) is $0 \leq j_i \leq 3$, so by Corollary 2.1.22, this game has an equilibrium in which each player mixes among at most $4+1 = 5$ pure strategies. To produce this bound, we have not used any information about the payoffs except for the degree of the polynomials. However, there is extra structure here to be exploited. For example, u_2 depends on the expected value $\int x^2 d\sigma_1(x)$, but not on $\int x d\sigma_1(x)$ or $\int x^3 d\sigma_1(x)$. In particular, player 2 is indifferent between the two strategies $\pm x$ of player 1 for all x , insofar as this choice does not affect his payoff (though it does affect what strategy profiles are equilibria). The bounds in Chapter 3 make these simplifications in a systematic manner.

2.2 Computation of Nash Equilibria in Finite Games

2.2.1 Optimization Formulation

There are several ways to formulate the Nash equilibria of a finite game as solutions to an optimization problem. Here we will focus on one presented by Başar and Olsder [1].

Proposition 2.2.1. *Consider the optimization problem*

$$\begin{aligned}
 \max \quad & \sum_{i=1}^n [u_i(\sigma) - p_i] \\
 \text{s.t.} \quad & \sigma_i \in \Delta_i \text{ for all } i \\
 & u_i(s_i, \sigma_{-i}) \leq p_i \text{ for all } i, \text{ all } s_i \in C_i
 \end{aligned} \tag{2.14}$$

where p_i is an auxiliary variable representing the equilibrium payoff to player i . The

optimum objective value of this problem is zero and is attained exactly when σ is a Nash equilibrium with payoff p_i to player i .

Proof. The constraints imply that $u_i(\sigma) - p_i \leq 0$ for all i , so the objective function is bounded above by zero. Given any Nash equilibrium σ , let $p_i = u_i(\sigma)$ for all i . Then the objective function evaluates to zero and all the constraints are satisfied by definition of a Nash equilibrium. Therefore the optimal objective function value is zero and it is attained at all Nash equilibria.

Conversely suppose some feasible σ and p achieve objective function value zero. Then the inequality $u_i(\sigma) - p_i \leq 0$ implies that $u_i(\sigma) = p_i$ for all i . Also, the final constraint implies that player i cannot achieve a payoff of more than p_i by unilaterally changing his strategy. Therefore σ is a Nash equilibrium. \square

Several observations about (2.14) are in order. The objective function is multi-affine of degree n in the σ_i . The first constraint can be expressed via finitely many linear inequalities, and the second constraint can be expressed via finitely many inequalities bounding multiaffine functions of degree $n - 1$. Therefore the problem is in general nonconvex, but for small n it can be solved by nonlinear programming methods or general methods for polynomial inequalities.

The two-player case has even more structure. All the constraints are linear (since $n = 2$) and the objective function is biaffine, so while the objective function is nonconvex, the feasible set is polyhedral. If the game also satisfies the zero sum condition $u_1 + u_2 \equiv 0$, meaning that the game is strictly competitive because what is good for one player is bad for the other, then the objective function is linear as well. Thus Nash equilibria of two-player zero-sum finite games can be computed efficiently because (2.14) reduces to a linear program [3].

2.2.2 Support Enumeration

For two-player nonzero-sum finite games, the conceptually simplest way to compute Nash equilibria is by selecting a set of strategies which each player will be allowed to use with positive probability, and then checking whether there exists an equilibrium with these supports. It turns out that for fixed supports, an equilibrium can be computed, or proven not to exist over those supports, using a linear program.

We will refer to the two players as the row player and the column player. Their strategy sets will be $C_r = \{1, \dots, m_r\}$ and $C_c = \{1, \dots, m_c\}$, respectively, and their payoffs will be given by the matrices $U_r \in \mathbb{R}^{m_r \times m_c}$ and $U_c \in \mathbb{R}^{m_r \times m_c}$, respectively. That is to say, if the row player chooses strategy $s_r \in C_r$ and the column player chooses strategy $s_c \in C_c$, then they will receive payoffs $U_r^{s_r, s_c}$ and $U_c^{s_r, s_c}$, respectively. For simplicity of exposition we will write $*$ in place of r or c when a statement applies to either player.

Definition 2.2.2. A **support** for player $*$ is any nonempty subset $S_* \subseteq C_*$.

In this section we will view mixed strategies as probability vectors $\sigma_* \in \mathbb{R}^{m_*}$. For any such mixed strategy, we define its support $S(\sigma_*) = \{i \in C_* \mid \sigma_*^i > 0\}$. The following alternative characterization of Nash equilibria is standard.

Lemma 2.2.3. *A mixed strategy profile (σ_r, σ_c) is a Nash equilibrium if and only if*

$$\begin{aligned} S(\sigma_r) &\subseteq \arg \max_k [U_r \sigma_c]^k \\ S(\sigma_c) &\subseteq \arg \max_k [\sigma_r U_c]^k. \end{aligned}$$

Proof. Follows directly from the definition of a Nash equilibrium. □

We can restate this lemma as:

Corollary 2.2.4. *Given supports S_r and S_c , a mixed strategy profile satisfying the linear constraints*

$$\begin{aligned} [U_r \sigma_c]^k &\leq [U_r \sigma_c]^l \text{ for all } k \text{ if } l \in S_r \\ [\sigma_r U_c]^k &\leq [\sigma_r U_c]^l \text{ for all } k \text{ if } l \in S_c \\ \sigma_*^k &> 0 \text{ if } k \in S_* \\ \sigma_*^k &= 0 \text{ if } k \notin S_* \\ \sum_{k \in S_*} \sigma_*^k &= 1 \end{aligned} \tag{2.15}$$

is a Nash equilibrium with supports $S(\sigma_) = S_*$ (if the strict inequality is made non-strict we instead get $S(\sigma_*) \subseteq S_*$). If these constraints are infeasible then no equilibrium with the given supports exists.*

While the above form is a compact way to write these constraints on paper, it involves quadratically many constraints and is quite redundant. It is easy to see that the number of constraints can be reduced so that the number of constraints and variables in this linear program is linear in the total number of strategies $m_r + m_c$ available to the two players. For fixed supports, this linear program can be solved efficiently using standard linear programming algorithms [3].

By Theorem 2.1.3 every finite game has a Nash equilibrium, so there exists some support pair (S_r, S_c) for which the linear constraints in Corollary 2.2.4 are feasible. In general the number of possible support pairs to be checked is exponential in $m_r + m_c$, so in practice checking support pairs at random may not find an equilibrium quickly. However, if it can be proven that the game in question has an equilibrium in which the cardinalities of the supports are bounded by some (small) constants, then only supports up to this size need be considered. Hence, the support of an equilibrium can be found in time polynomial in $m_r + m_c$ in this case. This technique will be used in Section 4.2 to compute ϵ -Nash equilibria of two-player separable games with continuous strategy sets.

Finally, it is worth noting that there exist analogs of Lemma 2.2.3 and Corollary 2.2.4 which apply to games with more than two players. However, with three or more players the constraints in Corollary 2.2.4 become nonlinear and nonconvex, so a feasible point can no longer be found efficiently.

2.2.3 The Lemke-Howson Algorithm

Next we sketch another algorithm for computing Nash equilibria of two-player finite games which amounts to doing support enumeration in a structured way. This method, called the Lemke-Howson algorithm, is similar in many respects to the simplex method for solving linear programs, so we will begin by quickly reviewing a simplified version of that algorithm. A linear program is a problem of the form $\arg \max_{Ax \leq b} c'x$ where $A \in \mathbb{R}^{m \times n}$, $b \in \mathbb{R}^m$, and $c \in \mathbb{R}^n$ are the data of the problem and $x \in \mathbb{R}^n$ is a vector of variables.

The constraints $Ax \leq b$ define a convex polyhedron P of values for x and the function $c'x$ amounts to a projection of this polyhedron onto a line. We will assume for simplicity of exposition that this polyhedron P is in fact a bounded polytope, though this is not necessary. The simplex method works by beginning at a vertex x_1 of P and checking to see along which, if any, of the edges leaving x_1 the objective function $c'x$ is increasing. Any such edge may be chosen to be the one followed; to make the algorithm deterministic we can use the pivoting rule which selects the edge along which the fastest improvement in the objective function $c'x$ will be made. By our boundedness assumption, this edge ends at another vertex x_2 of P , and we have $c'x_2 > c'x_1$.

The process then repeats, terminating when no outgoing edge can be found along which the objective function increases. The objective function values $c'x_i$ strictly increase, so no vertex can be visited more than once. Since P is a polytope it has finitely many vertices, hence the algorithm terminates. Under some nondegeneracy assumptions on A , b , and c , the polytope P will have no edges along which the objective function $c'x$ remains exactly constant. When the algorithm terminates at the final vertex x_* , the objective function must strictly decrease along all edges leaving x_* . Therefore x_* is a local maximizer of the concave function $c'x$ over the convex set P , hence also a global maximizer. This completes the construction of the simplex algorithm and proves its correctness.

Of course, to actually implement this algorithm, many more details are needed. Most importantly, a significant amount of linear algebra is required to find a starting vertex, then at each step to find an edge along which the objective function increases, to follow it, and to find the vertex at the other end. By adding a few extra steps one can easily handle the case when P is unbounded, and with significantly more work one can also modify the algorithm to remove the nondegeneracy assumptions. For a complete discussion of the simplex algorithm which includes these technical details, see [3].

For our purposes, the idea of the algorithm is the important part. It begins with a vertex of the polytope, then follows a path from vertex to vertex along edges of the polytope, using a deterministic pivoting rule to select the edge followed at each iteration. The graph formed by the vertices and edges of the polytope is finite and the algorithm cannot cycle, so termination is guaranteed. The Lemke-Howson algorithm uses the same piecewise linear path following idea but a different pivoting rule to solve two-player finite games instead of linear programs.

We follow Shapley's exposition [42] to present a geometric sketch of the Lemke-

Howson algorithm in the spirit of the overview of the simplex algorithm given above. We use notation similar to the previous section, with U_r and U_c the matrices of payoffs to the row and column player, respectively. The one minor change is that we number the row player's strategies $C_r = \{1, \dots, m_r\}$ and the column player's strategies $C_c = \{m_r + 1, \dots, m_r + m_c\}$. For reasons which will become clear below, we need these sets to be disjoint. As above, we use $*$ in place of r or c in statements which apply to both players.

Let $\Delta_* = \{\sigma_* \geq 0 \mid \sum_k \sigma_*^k = 1\}$ be the simplex of mixed strategy vectors for player $*$. Define

$$\bar{\Delta}_* = \Delta_* \cup \{\sigma_* \geq 0 \mid \sum_k \sigma_*^k \leq 1 \text{ and } \sigma_*^l = 0 \text{ for some } l\},$$

so each $\bar{\Delta}_*$ is the boundary of a simplex with one dimension more than Δ_* . This extension of the mixed strategy spaces is not motivated by any game-theoretic interpretation, but simplifies the description and visualization of the algorithm. We also define the subsets

$$\begin{aligned} \bar{\Delta}_r^k &= \{\sigma_r \in \bar{\Delta}_r \mid \sigma_r^k = 0\} \text{ for } k \in C_r \\ \bar{\Delta}_r^k &= \{\sigma_r \in \Delta_r \mid [\sigma_r U_c]_k \geq [\sigma_r U_c]_l \text{ for all } l \in C_c\} \text{ for } k \in C_c \\ \bar{\Delta}_c^k &= \{\sigma_c \in \Delta_c \mid [U_r \sigma_c]_k \geq [U_r \sigma_c]_l \text{ for all } l \in C_r\} \text{ for } k \in C_r \\ \bar{\Delta}_c^k &= \{\sigma_c \in \bar{\Delta}_c \mid \sigma_c^k = 0\} \text{ for } k \in C_c \end{aligned}$$

Note all of these $\bar{\Delta}_*$ are bounded sets defined by linear equations and inequalities, i.e. they are polytopes. Furthermore, every point in $\bar{\Delta}_r \setminus \Delta_r$ is contained in at least one of the $\bar{\Delta}_r^k$ for $k \in C_r$ and each mixed strategy $\sigma_r \in \Delta_r$ has some best response $l \in C_c$ for the column player so that $\sigma_r \in \bar{\Delta}_r^l$. Therefore the polytopes $\bar{\Delta}_r^k$ cover $\bar{\Delta}_r$ and for similar reasons the same is true with r replaced by c .

For each point $\sigma_* \in \bar{\Delta}_*$, define its **label** $L_*(\sigma_*) = \{k \in C_r \cup C_c \mid \sigma_* \in \bar{\Delta}_*^k\}$. To each pair (σ_r, σ_c) we also assign a label $L(\sigma_r, \sigma_c) = L_r(\sigma_r) \cup L_c(\sigma_c)$. We call a pair (σ_r, σ_c) **completely labeled** if $L(\sigma_r, \sigma_c) = C_r \cup C_c$. The game-theoretic motivation for this definition is:

Lemma 2.2.5. *A pair $(\sigma_r, \sigma_c) \in \Delta_r \times \Delta_c$ is completely labeled if and only if it is a Nash equilibrium.*

Proof. This is just a matter of rewriting definitions. A pair is completely labeled if and only if $\sigma_r \in \bar{\Delta}_r^k$ or $\sigma_c \in \bar{\Delta}_c^k$ for each $k \in C_r \cup C_c$. This happens if and only if for each player $*$ and each k in C_* either $\sigma_*^k = 0$ or k is a best response to σ_{-*} . That is to say, if and only if each player only plays best responses with positive probability. \square

For simplicity we will make a “nondegeneracy” assumption on the payoff matrices. This amounts to linear independence conditions on certain submatrices. It is possible to extend the algorithm to cover these “degenerate” cases, but doing so tends to obscure the main ideas of the algorithm.

Assumption 2.2.6. The utility matrices U_r and U_c are such that an intersection of l distinct sets $\overline{\Delta}_*^k$ is either empty or has dimension $m_* - l$.

This assumption will implicitly remain in force throughout the rest of this subsection. It endows the sets $\overline{\Delta}_*^k$ with a rich geometric structure, which we now study.

Lemma 2.2.7. *Every $\sigma_* \in \overline{\Delta}_*$ satisfies $|L_*(\sigma_*)| \leq m_*$ for all σ_* , so completely labeled pairs must satisfy $|L_*(\sigma_*)| = m_*$ and $L_r(\sigma_r) \cap L_c(\sigma_c) = \emptyset$.*

Proof. The intersection $\bigcap_{k \in L_*(\sigma_*)} \overline{\Delta}_*^k$ is nonempty since it contains σ_* , hence it has dimension $m_* - |L_*(\sigma_*)| \geq 0$. The conclusions follow immediately from this fact. \square

Lemma 2.2.8. *The set $V_* = \{\sigma_* \in \overline{\Delta}_* : |L_*(\sigma_*)| = m_*\}$ is finite and the set $E_* = \{\sigma_* \in \overline{\Delta}_* : |L_*(\sigma_*)| = m_* - 1\}$ is the union of finitely many nonintersecting open line segments whose endpoints lie in V_* .*

Proof. For any element $\sigma_* \in V_*$, Assumption 2.2.6 implies that the set of points with the same label as σ_* is a polytope of dimension 0, i.e. it is the singleton $\{\sigma_*\}$. There are finitely many labels of size m_* , so finitely many elements of V_* . Similarly we see that E_* is the union of finitely many open line segments. The points in each segment have the same label, whereas the points in different segments have different labels, so distinct segments cannot intersect. Proving the last claim involves verifying that $L : \overline{\Delta}_* \rightarrow 2^{\{1, \dots, m_r + m_c\}}$ is upper semi-continuous with respect to the set inclusion order and that any neighborhood of an endpoint of a segment in E_i contains a point with a label not shared by the points on the segment. This involves several cases, so we omit the detailed proof. \square

By Lemma 2.2.8, we can view the points V_* as the vertices of an undirected graph G_* with edges defined by the set E_* . That is to say, two points in V_* are adjacent if they are the two endpoints of one of the segments in E_* . We will say that two nodes in V_* **differ by one label** if they share exactly $m_* - 1$ labels. It is easy to see from Assumption 2.2.6 and the definitions of the labels and the $\overline{\Delta}_*^k$ in terms of linear equations and inequalities that two nodes are adjacent in G_* if and only if they differ by one label.

Proposition 2.2.9. *The graph G_* is m_* -regular, i.e. each node is adjacent to exactly m_* other nodes.*

Proof. Given any node $\sigma_* \in V_*$, we have $|L_*(\sigma_*)| = m_*$ by Lemma 2.2.7. This set of labels has m_* distinct subsets of cardinality $m_* - 1$. By Assumption 2.2.6 and Lemma 2.2.8, Each such subset corresponds to a unique edge in G_* connecting σ_* to some other node. Furthermore, all edges incident on σ_* must be of this form. All the edges are straight line segments in $\overline{\Delta}_*$, so the endpoints of all m_* edges must be distinct. \square

Using the nondegeneracy assumption we can improve Lemma 2.2.5 as follows.

Lemma 2.2.10. *A pair $(\sigma_r, \sigma_c) \in \overline{\Delta}_r \times \overline{\Delta}_c$ is completely labeled if and only if it is a Nash equilibrium in $\Delta_r \times \Delta_c$ or it is $(0, 0)$.*

Proof. By Lemma 2.2.5, it suffices to prove that a completely labeled point is either $(0, 0)$ or in $\Delta_r \times \Delta_c$. Suppose (σ_r, σ_c) is completely labeled and $\sigma_r \in \overline{\Delta_r} \setminus (\Delta_r \cup \{0\})$. Then $L_r(\sigma_r) \cap C_c = \emptyset$ since $\sigma_r \notin \Delta_r$, but $\sigma_r \neq 0$ so $L_r(\sigma_r) \neq C_r$. Therefore $L_r(\sigma_r) \subsetneq C_r$ and $L(\sigma_r, \sigma_c) = C_r \cup C_c$ implies $|L_c(\sigma_c)| > m_c$, contradicting Lemma 2.2.7. Similar logic applies to σ_c and shows that $\sigma_* \in \Delta_* \cup \{0\}$.

Now suppose $\sigma_r = 0$. Then $L_r(\sigma_r) = C_r$, so we must have $L_c(\sigma_c) \supseteq C_c$, i.e. $\sigma_c = 0$. Thus (σ_r, σ_c) must be $(0, 0)$ or in $\Delta_r \times \Delta_c$. \square

Now we form the Cartesian product graph $G = G_r \square G_c$, which has vertex set $V_r \times V_c$. Two vertices $(\sigma_r, \sigma_c), (\tau_r, \tau_c) \in V_r \times V_c$ are defined to be adjacent in G if $\sigma_r = \tau_r$ and σ_c is adjacent to τ_c in V_c , or vice versa (with r and c interchanged). Let $N \subseteq V_r \times V_c$ be the set of completely labeled vertices; $N \setminus (0, 0)$ is the set of Nash equilibria of the game by Lemma 2.2.10. For each $k \in \{1, \dots, m_r + m_c\}$ define the set of k -almost completely labeled vertices $N^k = \{(\sigma_r, \sigma_c) \in V_r \times V_c \mid L(\sigma) \cup \{k\} = \{1, \dots, m_r + m_c\}\}$, so $N \subseteq N^k$ for all k .

Proposition 2.2.11. *For $k = 1, \dots, m_r + m_c$, each node in N is adjacent to a unique node in N^k and each node in $N^k \setminus N$ is adjacent to exactly two nodes in N^k .*

Proof. Let (σ_r, σ_c) be a node in N . Then by Lemma 2.2.7, $|L_*(\sigma_*)| = m_*$ and $L_r(\sigma_r) \cap L_c(\sigma_c) = \emptyset$. A node (τ_r, τ_c) adjacent to σ would have to be equal in one coordinate and differ by one label in the other. For τ to also be in N^k , we are only allowed to have $L_r(\tau_r) \cup L_c(\tau_c) \cup \{k\} = \{1, \dots, m_r + m_c\}$. Therefore if $L_r(\sigma_r)$ contains k we must change σ_r to the unique τ_r adjacent to σ_r in G_r which is missing the label k and let $\tau_c = \sigma_c$, or if $L_c(\sigma_c)$ contains k we must do the opposite. This proves the existence and uniqueness of a node $\tau \in N^k$ adjacent to σ .

Now let $(\sigma_r, \sigma_c) \in N^k \setminus N$. By definition $|L_*(\sigma_*)| = m_*$ but $L(\sigma) = \{1, \dots, m_c + m_r\} \setminus \{k\}$, so $L_r(\sigma_r) \cap L_c(\sigma_c) = \{l\}$ for some $l \neq k$. Therefore if we remove the label l from either σ_r or σ_c , the resulting pair would still be k -almost completely labeled. On the other hand if we were to remove any other label, this would not be true. Therefore σ is adjacent to $\tau \in N^k$ if and only if $\sigma_r = \tau_r$ and τ_c is the node joined to σ_c by the unique edge incident on σ_c which is missing label l , or vice versa with r and c interchanged. Therefore there are exactly two nodes in N^k adjacent to σ . \square

Define the graph G^k to be the subgraph of G induced by the set of nodes N^k . Now we can put everything together.

Theorem 2.2.12 (Lemke and Howson [25]). *If a two player finite game satisfies Assumption 2.2.6 then*

- *Each graph G^k is a disjoint union of cycles and chains, the endpoints of which form the set N ;*
- *The set $N \setminus \{(0, 0)\}$ of Nash equilibria has an odd number of elements and so in particular is nonempty;*
- *A Nash equilibrium can be computed by fixing any k and following the chain in N^k starting at $(0, 0)$ to its other endpoint, which will be an element of $N \setminus \{(0, 0)\}$ and hence a Nash equilibrium.*

Proof. By Proposition 2.2.11, each node in the graph G^k has degree 1 or 2, and the nodes with degree 1 are the elements of the set N . It can be proven by a simple induction on the cardinality $|G^k|$ that a graph with this property is a disjoint union of cycles and chains, and the endpoints of the chains are the elements of degree 1. Each chain has exactly two endpoints, so the cardinality $|N|$ must be even. Since $(0,0) \in N$, $|N \setminus \{(0,0)\}|$ is odd. The point $(0,0) \in N$ is the endpoint of a chain and must be distinct from the other endpoint of that chain, so the final statement is clear. \square

We close this subsection with several notes about the algorithm described. We did not use Nash’s theorem at any point, so we have obtained a constructive proof of Nash’s theorem in the case of two-player finite games satisfying Assumption 2.2.6. The algorithm can be extended to arbitrary two-player finite games by lexicographic pivoting rules similar to those used for resolving degeneracies in the simplex algorithm (see [27]). Doing so yields a constructive proof of the existence of a single Nash equilibrium for arbitrary two-player finite games, but necessarily complicates the geometry and relaxes the other conclusions of Theorem 2.2.12. In particular it need not be true for degenerate games that the number of equilibria is odd or even finite.

One can imagine that using different values of k may allow additional equilibria to be found, and that changing the value of k from one step of the algorithm to the next may allow even more equilibria to be found. In fact, both are true. Taking this idea to its logical conclusion, we can form the subgraph of G induced by the set of nodes $\bigcup_{k=1}^{m_r+m_c} N^k$. The equilibria lying in the same connected component of this graph as the node $(0,0)$ are called **accessible equilibria**, and it can be shown that there exist games in which not all equilibria are accessible [42] in this sense. This notion of accessibility is not known to correspond to a more intrinsic game-theoretic property.

Savani and von Stengel have constructed an explicit example of a game for which the Lemke-Howson algorithm requires an exponential number of steps regardless of which value of k is used, so the Lemke-Howson algorithm can be quite slow in some cases [40]. However, no asymptotically faster exact algorithm is known for finding a single Nash equilibrium of a two-player finite game.

2.3 Previous Work

The theorems and algorithms presented in Chapters 3 and 4 are related to previous work in several areas, which we categorize roughly as existence results for equilibria, theory of separable and low-rank games, and computation of equilibria in finite games. The first category has been discussed in detail above, so we will focus on the latter two, briefly mentioning some of the key papers and their contributions. Some are direct predecessors of our work; the rest are listed to provide context for these.

Separable games were first studied around the 1950’s by Dresher, Karlin, and Shapley in papers such as [12], [11], and [22], which were later combined in Karlin’s book [21]. The work focuses on the zero-sum case, which contains some of the key

ideas of the nonzero-sum case, such as the fact that the infinite-dimensional mixed strategy spaces can be replaced with finite-dimensional moment spaces and the connection between the complexity of the payoffs and the complexity of equilibria. The goal of developing this theory was as a means of attack for general continuous games [21], but it was realized that finding the solutions of polynomial games could be quite complicated as well. Supporting this hypothesis, Gale and Gross showed that arbitrary finitely supported measures can be realized as the unique equilibrium of some zero-sum polynomial game [15].

There are formal similarities between separable games and finite games satisfying low-rank conditions. Lipton et al. [26] consider two-player finite games and provide bounds on the cardinality of the support of Nash equilibrium strategies using the ranks of the payoff matrices of the players. Using a similar proof, Germano and Lugosi [16] bound the cardinality of the support of a correlated equilibrium in an arbitrary finite game in terms of the number of strategies available to each player.

There has also been considerable work on finding algorithms for computing mixed-strategy Nash equilibria in finite games. Lemke and Howson give a path-following algorithm for two-player finite games which can be viewed as the simplex method for linear programming operating with a different pivoting rule, as described in Subsection 2.2.3 [25]. To find equilibria of games with more players, Scarf constructs a simplicial subdivision algorithm which also works for more general fixed point problems [41]. Lipton et al. investigate the problem of computing approximate Nash equilibria in finite games and present the first algorithm for computing approximate equilibria which is quasi-polynomial in the number of strategies [26]. These methods all rely on the polyhedral structure of the mixed strategy spaces of finite games. For surveys of known algorithms which compute equilibria of finite games, see [27] and [48].

Another growing literature has been investigating the computational complexity of finding mixed strategy Nash equilibria of finite games. Daskalakis, Goldberg, and Papadimitriou settle this question for finite normal form games with four or more players, showing that the problem of computing a single Nash equilibrium is PPAD-complete [9]. In essence this means that it is computationally equivalent to a number of other fixed point problems which are believed to be computationally difficult. These problems share the feature that a solution can be proven to exist, but the known proofs of existence are inefficient; for more about the complexity class PPAD, see [29]. Daskalakis and Papadimitriou later improve this result by proving PPAD-completeness in the case of three players [10]. Chen and Deng also prove this independently in [4] and finally complete this line of work proving PPAD-completeness of the problem for two players in [5]. In this literature, there has been no work on the analogous problems for continuous games.

The problem of algorithmically finding correlated equilibria of finite games seems to be much easier. As shown in Proposition 2.1.5, we can cast this as a linear program, and linear programs can be solved efficiently [3]. However, note that the amount of data needed to explicitly write the payoffs of the game is exponential in n , the number of players, so the corresponding linear program is exponential in size if n is allowed to vary. Nonetheless, for many interesting classes of games there exist succinct repre-

sentations of the utilities with polynomial length. This raises the natural question of whether correlated equilibria of such games can also be represented succinctly and if so whether they can be computed in polynomial time. These questions were first considered by Papadimitriou and Roughgarden who addressed a few special cases [31]. Shortly thereafter, Papadimitriou showed that correlated equilibria can be computed in polynomial time for any class of finite games with the property that the expected utilities of the players can be computed in polynomial time if each player chooses a mixed strategy independently from the other players [30]. This includes a variety of important classes such as graphical games [23] and congestion games [36].

Our work is also related to that of Kannan and Theobald, who study a different notion of rank in two-player finite games [20]. They take an algorithmic perspective and view zero-sum games as the simplest type of games. To generalize these, they propose a hierarchy of classes of two-player finite games in which the rank of the sum $U_r + U_c$ of the players' payoff matrices is bounded by a constant k ; the case $k = 0$ corresponds to zero-sum games. For fixed k , they show how to compute approximate Nash equilibria of two-player finite games in time polynomial in the description length of the game. This algorithm relies on an approximation result for quadratic programming due to Vavasis [46] which depends on the polyhedral structure of the problem.

Very little work has been done at the intersection of these areas, namely on algorithms for computing equilibria of games with infinite strategy sets. Some simple special cases were considered by Karlin [21] and others in the 1950's, but no general purpose algorithms were known at that point. The problem of computing mixed strategy Nash equilibria of zero-sum games with polynomial utilities was considered to be an important open question in the theory of infinite games by Kuhn and Tucker [24] and Karlin [21], despite the fact that the analogous problem for finite games is easily solved by linear programming. In fact the problem of computing Nash equilibria of zero-sum polynomial games was not settled until Parrilo's 2006 paper showing that equilibria of such games can be computed efficiently using semidefinite programming and algorithmic sum of squares techniques [33], which were not available until the mid 1990's and early 2000's, respectively [45, 32].

Chapter 3

Rank and Equilibria in Infinite Games

It was shown in Chapter 2 that the complexity of strategies played in an equilibrium can be bounded for separable games. This feature, combined with the existence of finite-dimensional representations of the mixed strategy spaces in terms of moments, is what makes these equilibrium problems potentially amenable to computation. The goal of this chapter is to prove tighter bounds on the cardinality of the support of equilibrium strategies, which can yield better performance for algorithms. Such bounds are also of interest on a purely conceptual level, because they show how the complexity of the payoffs in a game is related to the complexity of the strategies played in equilibrium. Theorem 2.1.18 can be viewed as a first result in this direction, because it shows that for games in which the utilities have unbounded complexity, the equilibrium strategies can also have unbounded complexity.

We also cover several new formulations of correlated equilibria in continuous games.

3.1 Rank of Finite Games

In this section we review known bounds on the number of strategies played with positive probability in equilibria of finite games. We begin with a bound on the cardinality of the support of Nash equilibria for two-player finite games.

Theorem 3.1.1 (Lipton et al. [26]). *Suppose we are given a two-player finite game defined by matrices U_r and U_c of payoffs to the row and column player, respectively, and any Nash equilibrium σ of the game. Then there exists a Nash equilibrium τ which yields the same payoffs to both players as σ , but in which the column player mixes among at most $\text{rank } U_r + 1$ pure strategies and the row player mixes among at most $\text{rank } U_c + 1$ pure strategies.*

In the following section we will extend this theorem to arbitrary separable games, thereby weakening the restriction that the strategy spaces be finite and treating the multiplayer case which was left open in [26]. The extended theorem also yields a

slightly tighter bound of $\text{rank } U_r$ instead of $\text{rank } U_r + 1$ (respectively for U_c) in some cases, depending on the structure of U_r and U_c . Before laying the framework for the extended version of the theorem, we will prove Theorem 3.1.1. The original proof was essentially an algorithmic version of Carathéodory's theorem. Here we give a shorter nonalgorithmic proof which will illustrate some of the ideas to be used later in this section.

Proof. Let r and c be probability column vectors corresponding to the mixed strategies of the row and column players in the given equilibrium σ . Then the payoffs to the row and column players are $r'U_r c$ and $r'U_c c$. Since c is a probability vector, we can view $U_r c$ as a convex combination of the columns of U_r . These columns all lie in the column span of U_r , which is a vector space of dimension $\text{rank } U_r$. By Carathéodory's theorem, we can therefore write any convex combination of these vectors using only $\text{rank } U_r + 1$ terms. That is to say, there is a probability vector d such that $U_r d = U_r c$, d has at most $\text{rank } U_r + 1$ nonzero entries, and a component of d is nonzero only if the corresponding component of c was nonzero.

Since r was a best response to c and $U_r c = U_r d$, r is a best response to d . On the other hand, since (r, c) was a Nash equilibrium c must have been a mixture of best responses to r . But d only assigns positive probability to strategies to which c assigned positive probability. Thus d is a best response to r , so (r, d) is a Nash equilibrium which yields the same payoffs to both players as (r, c) , and d only assigns positive probability to $\text{rank } U_r + 1$ pure strategies. Applying the same procedure to r we can find an s which only assigns positive probability to $\text{rank } U_c + 1$ pure strategies and such that (s, d) is a Nash equilibrium with the same payoffs to both players as (r, c) . \square

Observe that $U_r, U_c \in \mathbb{R}^{m_r \times m_c}$ where m_r and m_c are the number of strategies of the row and column players, respectively. Therefore $\text{rank } U_r \leq m_c$ and $\text{rank } U_c \leq m_r$, so Theorem 3.1.1 implies that any two-player finite game has an equilibrium in which the row player mixes among at most $\min(m_r, m_c + 1)$ pure strategies and the column player mixes among at most $\min(m_r + 1, m_c)$ pure strategies. This means that if one of the players has a small number of strategies, then there will be a Nash equilibrium in which both players mix among a small number of pure strategies.

We noted above that in some cases we can improve the bound from $\text{rank } U_r + 1$ to $\text{rank } U_r$ and similarly for U_c . This can be seen directly from the proof given here. We considered convex combinations of the columns of U_r and noted that these all lie in the column span of U_r . In fact, they all lie in the affine hull of the set of columns of U_r . If this affine hull does not include the origin, then it will have dimension $\text{rank } U_r - 1$. The rest of the proof goes through using this affine hull instead of the column span, and so in this case we get a bound of $\text{rank } U_r$ on the number of strategies played with positive probability by the row player. Alternatively, the dimension of the affine hull can be computed directly as the rank of the matrix produced by subtracting a fixed column of U_r from all the other columns of U_r .

We can use a similar argument to give a bound on the support of correlated equilibria in n -player finite games. It is possible to tighten this bound by writing

in terms of the ranks of certain matrices associated with the game. The following sections will prove these tighter bounds, but first we will prove the original result due to Germano and Lugosi, which is stated in terms of the number of strategies available to each player [16]. The original proof of this result can also be seen as an algorithmic version of Carathéodory's theorem, but here we give a shorter linear algebraic proof which is similar to the proof of Theorem 3.1.1.

Theorem 3.1.2 (Germano and Lugosi [16]). *Every n -player finite game in which player i has m_i pure strategies possesses a correlated equilibrium in which at most $1 + \sum_{i=1}^n m_i(m_i - 1)$ pure strategy profiles are played with positive probability.*

Proof. We will use the characterization of correlated equilibria in terms of linear equations and inequalities presented in Proposition 2.1.5. Let $\pi \in \Delta(C)$ be a joint probability distribution over C written as a probability vector. Then the linear inequalities

$$\sum_{s_{-i} \in C_{-i}} \pi(s) [u_i(t_i, s_{-i}) - u_i(s)] \leq 0$$

for all i and all $s_i, t_i \in C_i$ with $s_i \neq t_i$ can be written together as an inequality of vectors $D\pi \leq 0$ for some matrix D with $\sum_{i=1}^n m_i(m_i - 1)$ rows and $\prod_{i=1}^n m_i$ columns.

A correlated equilibrium is then a probability vector π which satisfies $D\pi \leq 0$. By Theorem 2.1.7 there exists such a probability vector π . But $D\pi$ is just a convex combination of the columns of D , which all lie in the column span of D . This column span is a vector space whose dimension is the rank of D , which is bounded by the number of rows of D . Therefore $D\pi$ is a convex combination of points in a vector space of dimension at most $\sum_{i=1}^n m_i(m_i - 1)$. By Carathéodory's theorem any such convex combination can be written as a convex combination using at most $1 + \sum_{i=1}^n m_i(m_i - 1)$ of these points, so there exists a probability vector $\tilde{\pi}$ with at most $1 + \sum_{i=1}^n m_i(m_i - 1)$ nonzero components such that $D\tilde{\pi} = D\pi \leq 0$. Hence $\tilde{\pi}$ is a correlated equilibrium with the desired property. \square

In addition to simplifying the original proof of this result, the argument using Carathéodory's theorem shows how the conclusion of the theorem could be strengthened in terms of the rank of D . In particular the same proof shows that there exists a correlated equilibrium with at most $1 + \text{rank } D$ nonzero components, this result can be improved slightly more by replacing $\text{rank } D$ with the dimension of the affine hull of the columns of D , a number which is either the same as or one less than $\text{rank } D$.

3.2 Rank of Continuous and Separable Games

The goal of this section is to generalize the theorems in the previous section to arbitrary separable games. The utilities of separable games can frequently be written in several different forms, which may yield different bounds on the cardinality of equilibrium strategies when applying Corollary 2.1.22. For this reason we will focus on representation-independent bounds. The bounds we produce will always be at least

as tight as those from Corollary 2.1.22. To define the bounds in a representation-independent fashion we will define them for all continuous games, but we will show in Section 3.4 that they are only finite for separable games.

Before presenting the bounds we will cover some preliminary definitions. We have already introduced the set $\Delta_i = \Delta(C_i)$ of all Borel probability measures over C_i . Now let V_i denote the set of all finite signed Borel measures (henceforth simply called measures) over C_i . This is a vector space which includes Δ_i as a convex subset whose affine hull is a hyperplane, i.e. has codimension 1. The same definition of the weak* topology on Δ_i defines a topology on V_i , which is also called the **weak* topology** and makes V_i into a Hausdorff topological vector space [39].

As in the case of C and Δ we define $V = \prod_{i=1}^n V_i$. We can define the utility $u_i(\sigma)$ for all $\sigma \in V_i$ to be the integral $\int u_i d\sigma$ as we did for $\sigma \in \Delta$. Though this definition lacks a concrete game-theoretic interpretation, it is analytically convenient because it makes the utilities into multilinear functionals defined on the spaces of signed measures, which will allow us to apply linear algebraic techniques. Naturally, for separable games we extend the moment functionals f_i^j to all of V_i by the definition $f_i^j(\sigma_i) = \int f_i^j d\sigma_i$, so the formula (2.12) remains valid with $s \in C$ replaced by any $\sigma \in V$.

In Section 2.1.3 we defined the notion of moment equivalence for separable games and mentioned that moment equivalent mixed strategies are payoff equivalent in the sense that switching between them does not affect the outcome of the game under any circumstances. It is clear that moment equivalence depends on the way the utility functions have been written down, so to obtain representation-independent bounds we will work directly with the concept of payoff equivalence.

Definition 3.2.1. Two measures σ_i and τ_i are called **payoff equivalent** if $u_j(\sigma_i, \sigma_{-i}) = u_j(\tau_i, \sigma_{-i})$ for all players j and all mixed strategy profiles $\sigma_{-i} \in V_{-i}$ or, equivalently, for all pure strategy profiles $\sigma_{-i} \in C_i$.

In the proof of Theorem 3.1.1 the column player switches between mixed strategy vectors c and d which are equivalent with respect to the row player's payoff in the sense that $Rc = Rd$, but which may not be equivalent with respect to the column player's own payoff. To generalize Theorem 3.1.1 we must generalize this notion of equivalence, to what we call almost payoff equivalence. In the generalization of Theorem 3.1.2 to separable games we will also require the opposite notion of equivalence, namely two strategies for a player will be equivalent if the choice between the two can affect his own payoff but never the payoffs of other players. To formalize these concepts we make the following definitions.

Definition 3.2.2. Two measures σ_i and τ_i are called **almost payoff equivalent** (respectively **own payoff equivalent**) if $u_j(\sigma_i, \sigma_{-i}) = u_j(\tau_i, \sigma_{-i})$ for all players $j \neq i$ (respectively $j = i$) and all mixed strategy profiles $\sigma_{-i} \in V_{-i}$ or, equivalently, for all pure strategy profiles $\sigma_{-i} \in C_i$.

Clearly two measures in V_i are payoff equivalent if and only if they are both almost payoff equivalent and own payoff equivalent. We will consider V_i mod these three equivalence relations, so we define the following subspaces of V_i .

Definition 3.2.3. Let 0 denote the zero measure in V_i and view $u_j(\sigma_i, s_{-i})$ as a linear functional of $\sigma_i \in V_i$.

- $U_{i,*} = \{\text{measures payoff equivalent to } 0\} = \bigcap_{\substack{s_{-i} \in C_{-i} \\ 1 \leq j \leq n}} \ker u_j(\cdot, s_{-i})$
- $U_{i,-i} = \{\text{measures almost payoff equivalent to } 0\} = \bigcap_{\substack{s_{-i} \in C_{-i} \\ i \neq j}} \ker u_j(\cdot, s_{-i})$
- $U_{i,i} = \{\text{measures own payoff equivalent to } 0\} = \bigcap_{s_{-i} \in C_{-i}} \ker u_i(\cdot, s_{-i})$

These definitions allow us to express the condition that two measures σ_i and τ_i are, for example, almost payoff equivalent by writing $\sigma_i - \tau_i \in U_{i,-i}$. The fact that payoff equivalence is the same as almost payoff equivalence and own payoff equivalence together is represented by the fact that $U_{i,*} = U_{i,-i} \cap U_{i,i}$. Since the players are only allowed to choose probability measures, we will consider the sets Δ_i modulo these equivalence relations, i.e. the image of Δ_i in $V_i/U_{i,-i}$, etc. To avoid defining excessively many symbols we will denote these images by $\Delta_i/U_{i,-i}$ and so forth, even though Δ_i itself is not a vector space. In what follows the dimension of a set will refer to the dimension of its affine hull.

Definition 3.2.4. The **ranks** of a game are defined by $\rho_{i,*} = \dim \Delta_i/U_{i,*}$, $\rho_{i,-i} = \dim \Delta_i/U_{i,-i}$, and $\rho_{i,i} = \dim \Delta_i/U_{i,i}$. We say that a game has **finite rank** if $\rho_{i,*} < \infty$ for all i .

It is immediate from the definitions that $\rho_{i,-i}, \rho_{i,i} \leq \rho_{i,*} \leq \rho_{i,-i} + \rho_{i,i}$. If we define W_i to be the set of measures moment equivalent to the zero measure in some separable game, then clearly $W_i \subseteq U_{i,*}$. The moment map $f_i : \sigma_i \mapsto (f_i^1(\sigma_i), \dots, f_i^{m_i}(\sigma_i))$ has kernel W_i and a range of dimension m_i , so $\rho_{i,*} \leq \dim V_i/W_i \leq m_i$ for a separable game. We will show in Section 3.4 that the condition of finite rank is equivalent to the condition that $\rho_{i,-i} < \infty$ for all i and in fact also equivalent to the condition that the game be separable.

In Theorem 3.1.1 we bounded the number of strategies played in a Nash equilibrium of a finite game in terms of the ranks of the payoff matrices. In particular we bounded the number of strategies played by the row player in terms of the rank of the column player's payoff matrix and vice versa. The natural generalization of this argument to a multiplayer separable game is to begin with a Nash equilibrium and have each player switch to an almost payoff equivalent strategy, so we will obtain a bound on the support of Nash equilibria in terms of the $\rho_{i,-i}$.

Theorem 3.2.5. *Given a Nash equilibrium σ of a separable game, there exists a Nash equilibrium τ in which each player i mixes among at most $\rho_{i,-i} + 1$ pure strategies and $u_i(\sigma) = u_i(\tau)$. If $\rho_{i,*} = 1$ and the metric space C_i is connected, then this bound can be improved so that τ_i is a pure strategy. Furthermore, all strategy profiles in the box $\text{conv}\{\sigma_1, \tau_1\} \times \dots \times \text{conv}\{\sigma_n, \tau_n\}$ formed by σ and τ are Nash equilibria with payoff $u_i(\sigma)$ to player i .*

Proof. While the setting of this theorem is more abstract, the proof uses essentially the same argument as the proof of Theorem 3.1.1. By Theorem 2.1.21, we can assume without loss of generality that each player's mixed strategy σ_i is finitely supported. Fix i , let $\psi_i : V_i \rightarrow V_i/U_{i,-i}$ denote the canonical projection transformation, and let $\sigma_i = \sum_j \lambda^j s_i^j$ be a finite convex combination of pure strategies. By linearity of ψ_i we have

$$\psi_i(\sigma_i) = \sum_j \lambda^j \psi_i(s_i^j).$$

Carathéodory's theorem states that (renumbering the s_i^j and adding some zero terms if necessary) we can write

$$\psi_i(\sigma_i) = \sum_{j=0}^{\rho_{i,-i}} \mu^j \psi_i(s_i^j),$$

a convex combination but perhaps with fewer terms. Let $\tau_i = \sum_{j=0}^{\rho_{i,-i}} \mu^j s_i^j$. Then $\psi_i(\sigma_i) = \psi_i(\tau_i)$. Since σ was a Nash equilibrium, and σ_i is almost payoff equivalent to τ_i , σ_j is a best response to $(\tau_i, \sigma_{-i,j})$ for all $j \neq i$. On the other hand σ_i was a mixture among best responses to the mixed strategy profile σ_{-i} , so the same is true of τ_i , making it a best response to σ_{-i} . Thus (τ_i, σ_{-i}) is a Nash equilibrium.

If $\rho_{i,*} = 1$ and C_i is connected, then $C_i/U_{i,*}$ is connected, compact, and one-dimensional, i.e. it is an interval. Therefore it is convex, so $\Delta_i/U_{i,*} = \text{conv}(C_i/U_{i,*}) = C_i/U_{i,*}$. This implies that there exists a pure strategy s_i which is payoff equivalent to σ_i , so we may take $\tau_i = s_i$ and (τ_i, σ_{-i}) is a Nash equilibrium.

Beginning with this equilibrium and repeating the above steps for each player in turn completes the construction of τ and the final statement of the theorem is clear. \square

While the preceding theorem was the original reason for our choice of the definition of the ranks, the definition turns out to have other interesting properties which we study below. The following alternative characterization of the ranks of a continuous game is more concrete than the definition given above. This theorem simplifies the proofs of many rank-related results and will be applied to the problem of computing the ranks of separable games in Section 3.5.

Theorem 3.2.6. *The rank $\rho_{i,-i}$ of a continuous game is given by the smallest $r_{i,-i}$ such that there exist continuous functions $g_i^k : C_i \rightarrow \mathbb{R}$ and $h_{i,j}^k : C_{-i} \rightarrow \mathbb{R}$ which satisfy*

$$u_j(s) = h_{i,j}^0(s_{-i}) + \sum_{k=1}^{r_{i,-i}} g_i^k(s_i) h_{i,j}^k(s_{-i}) \quad (3.1)$$

for all $s \in C$ and $j \neq i$ ($\rho_i = \infty$ if and only if no such representation exists). Furthermore, the minimum value of $r_{i,-i} = \rho_{i,-i}$ is achieved by functions $g_i^k(s_i)$ of the form $u_j(s_i, s_{-i})$ for some $s_{-i} \in C_{-i}$ and $j \neq i$ and functions $h_{i,j}^k(s_{-i})$ of the form $\int u_j(\cdot, s_{-i}) d\sigma_i$ for some $\sigma_i \in V_i$ and $j \neq i$.

If instead we let j range from 1 to n or restrict it to $j = i$ we get similar characterizations of $\rho_{i,*}$ and $\rho_{i,i}$, respectively.

Proof. Throughout the proof we will automatically extend any functions $g_i^k : C_i \rightarrow \mathbb{R}$ to all of V_i in the canonical way. Suppose we are given a representation of the form (3.1). Let $g_i : C_i \rightarrow \mathbb{R}^{r_{i,-i}}$ be defined by $g_i(s_i) = (g_i^1(s_i), \dots, g_i^{r_{i,-i}}(s_i))$. By definition, $\rho_{i,-i}$ is the dimension of $\Delta_i/U_{i,-i}$. Let Z_i denote the subspace of V_i parallel to Δ_i , i.e. the space of all signed measures σ_i such that $\int \sigma_i = 0$. Then $\rho_{i,-i} = \dim Z_i / (Z_i \cap U_{i,-i})$. By (3.1) any signed measure which is in Z_i and in $\ker g_i$ is almost payoff equivalent to the zero measure, so $Z_i \cap \ker g_i \subseteq Z_i \cap U_{i,-i}$ and therefore

$$\begin{aligned} \rho_{i,-i} &= \dim Z_i / (Z_i \cap U_{i,-i}) \leq \dim Z_i / (Z_i \cap \ker g_i) \\ &= \dim g_i(Z_i) \leq r_{i,-i}. \end{aligned}$$

It remains to show that if $\rho_{i,-i} < \infty$ then there exists a representation of the form (3.1) with $r_{i,-i} = \rho_{i,-i}$. Recall that $U_{i,-i}$ is defined to be

$$U_{i,-i} = \bigcap_{\substack{s_{-i} \in C_{-i} \\ j \neq i}} \ker u_j(\cdot, s_{-i})$$

where $u_j(\cdot, s_{-i})$ is interpreted as a linear functional on V_i . Since $\rho_{i,-i} = \dim Z_i / (Z_i \cap U_{i,-i})$ we can choose $\rho_{i,-i}$ linear functionals, call them $g_i^1, \dots, g_i^{\rho_{i,-i}}$, from the collection of functionals whose intersection forms $U_{i,-i}$ in such a way that $Z_i \cap U_{i,-i} = Z_i \cap \ker g_i$, where $g_i = (g_i^1, \dots, g_i^{\rho_{i,-i}})$ as above. We cannot choose a smaller collection of linear functionals and achieve $Z_i \cap U_{i,-i} = Z_i \cap \ker g_i$, because $\rho_{i,-i} = \dim Z_i / (Z_i \cap U_{i,-i})$. Note that $Z_i \cap \ker g_i = \ker (1, g_i^1, \dots, g_i^{\rho_{i,-i}})$ where 1 is the linear functional $1(\sigma_i) = \int d\sigma_i$. Therefore no functional can be removed from the list $(1, g_i) = (1, g_i^1, \dots, g_i^{\rho_{i,-i}})$ without affecting the kernel of the transformation $(1, g_i)$, so the functionals $1, g_i^1, \dots, g_i^{\rho_{i,-i}}$ are linearly independent.

This means that any of the linear functionals $u_j(\cdot, s_{-i})$ (the intersection of whose kernels yields $U_{i,-i}$) can be written uniquely as a linear combination of the functionals $1, g_i^1, \dots, g_i^{\rho_{i,-i}}$. That is to say, there are unique functions $h_{i,j}^k$ such that (3.1) holds with the functions g_i^k constructed here and $r_{i,-i} = \rho_{i,-i}$. The g_i^k are continuous by construction, so to complete the proof we must show that the functions $h_{i,j}^k$ are continuous as well. Since the functionals $1, g_i^1, \dots, g_i^{\rho_{i,-i}}$ are linearly independent, we can choose a measure $\sigma_i^k \in V_i$ which makes the k^{th} of these functionals evaluate to unity and all the others zero. Substituting these values into (3.1) shows that $h_{i,j}^k(s_{-i}) = \int u_j(\cdot, s_{-i}) d\sigma_i^k$. Since u_j is continuous, $h_{i,j}^k$ is therefore also continuous.

The same argument works for characterizing $\rho_{i,*}$ and $\rho_{i,i}$. \square

Note that in the statement of Theorem 3.2.6 we have distinguished the component $h_{i,j}^0(s_{-i})$ in u_j . We have shown that this distinction formally follows from the definition of $\rho_{i,-i}$, but there is also an intuitive game theoretic reason why this separation is natural. As mentioned above, $\rho_{i,-i}$ is intended to capture the number of essential degrees of freedom that player i has in his choice of strategy when playing a Nash equilibrium. Theorems 3.2.5 and 3.2.6 taken together show that player i only needs to take the other players' utilities into account to compute this number, and not his own. But player i is only concerned with the other players' utilities insofar as his own

strategic choice affects them. The function $h_{i,j}^0(s_{-i})$ captures the part of player j 's utility which does not depend on player i 's strategy, so whether this function is zero or not it has no effect on the rank $\rho_{i,-i}$.

We close this section with an application. If a submatrix is formed from a matrix by “sampling,” i.e. selecting a subset of the rows and columns, the rank of the submatrix is bounded by the rank of the original matrix. Theorem 3.2.6 shows that the same is true of continuous games, because a factorization of the form (3.1) for a game immediately provides a factorization for any smaller game produced by restricting the players' choices of strategies.

Corollary 3.2.7. *Let $(\{C_i\}, \{u_i\})$ be a continuous game with rank ρ and \tilde{C}_i be a nonempty compact subset of C_i for each i , with $\tilde{u}_i = u_i|_{\tilde{C}_i}$. Then the game $(\{\tilde{C}_i\}, \{\tilde{u}_i\})$ satisfies $\tilde{\rho}_{i,-i} \leq \rho_{i,-i}$, $\tilde{\rho}_{i,*} \leq \rho_{i,*}$, and $\tilde{\rho}_{i,i} \leq \rho_{i,i}$ for all i .*

Definition 3.2.8. The game $(\{\tilde{C}_i\}, \{\tilde{u}_i\})$ in Corollary 3.2.7 is called a **sampled game** or a **sampled version** of $(\{C_i\}, \{u_i\})$.

Note that if we take \tilde{C}_i to be finite for each i , then the sampled game is a finite game. If the original game is separable and hence has finite rank, then Corollary 3.2.7 yields a uniform bound on the complexity of finite games which can arise from this game by sampling. This fact is applied to the problem of computing approximate Nash equilibria in Section 4.2 below. Note that while the proof of Corollary 3.2.7 is trivial, there exist other kinds of bounds on the cardinality of the support of equilibria (e.g. for special classes of polynomial games as studied by Karlin [21]) which do not share this sampling property.

3.3 Characterizations and Rank Bounds for Correlated Equilibria

In this section we prove several characterizations of correlated equilibria in continuous games. These characterizations will be used at the end of the section to prove the analog of Theorem 3.2.5 for correlated equilibria. They are also fundamental pieces of the correlated equilibrium algorithms presented in Chapter 4.

We can prove statements about Nash equilibria of separable games using the finite dimensional formulation in terms of moments as in Theorem 3.2.5. On the other hand no exact finite-dimensional characterization of correlated equilibria in separable games is known. Given the characterization of Nash equilibria in terms of moments, a natural attempt would be to try to characterize correlated equilibria in terms of the joint moments, i.e. the values $\int f_1^{k_1} \cdots f_n^{k_n} d\pi$ for integers k_i and joint measures π . The reason this attempt fails to yield a finite dimensional formulation is that the definition of a correlated equilibrium imposes conditions on the conditional distributions of the equilibrium measure. A finite set of moments does not seem to contain enough information about these conditional distributions to check the required conditions. Therefore we are forced to consider approximate correlated equilibria.

Definition 3.3.1. An ϵ -correlated equilibrium of a continuous game is a joint probability measure $\pi \in \Delta(C)$ such that

$$\int [u_i(\zeta_i(s_i), s_{-i}) - u_i(s)] d\pi(s) \leq \epsilon$$

for all i and all Borel measurable functions $\zeta_i : C_i \rightarrow C_i$.

As in the case of Nash and ϵ -Nash equilibria, this definition reduces to that of a correlated equilibrium when $\epsilon = 0$.

Before considering any algorithms for computing approximate correlated equilibria, we will prove several alternative characterizations of exact and approximate correlated equilibria which are more amenable to analysis than the definition. These characterizations may also be of independent interest.

We begin with a technical lemma which we will use to prove the more interesting characterization theorems. A more general version of this lemma has appeared as Lemma 20 in [44].

Lemma 3.3.2. *Simple departure functions (those with finite range) suffice to define ϵ -correlated equilibria. That is to say, a joint measure π is a correlated equilibrium if and only if*

$$\int [u_i(\zeta_i(s_i), s_{-i}) - u_i(s)] d\pi(s) \leq \epsilon$$

for all i and all Borel measurable simple functions $\zeta_i : C_i \rightarrow C_i$.

Proof. The forward direction is immediate from the definitions. To prove the reverse, first fix i . Then choose any measurable function $\zeta_i : C_i \rightarrow C_i$ and let $\xi_i^k : C_i \rightarrow C_i$ be a sequence of measurable simple functions converging to ζ_i pointwise; such a sequence exists because C_i is a compact metric space. Then $u_i(\xi_i^k(s_i), s_{-i})$ converges to $u_i(\zeta_i(s_i), s_{-i})$ pointwise since u_i is continuous. Thus

$$\begin{aligned} \int [u_i(\zeta_i(s_i), s_{-i}) - u_i(s)] d\pi(s) &= \int \lim_{k \rightarrow \infty} [u_i(\xi_i^k(s_i), s_{-i}) - u_i(s)] d\pi(s) \\ &= \lim_{k \rightarrow \infty} \int [u_i(\xi_i^k(s_i), s_{-i}) - u_i(s)] d\pi(s) \leq \lim_{k \rightarrow \infty} \epsilon = \epsilon, \end{aligned}$$

where the second equality follows from Lebesgue's dominated convergence theorem [38] and the inequality is by assumption. \square

The following characterization is a generalization of the standard formulation of correlated equilibria in finite games in terms of linear constraints presented in Proposition 2.1.5.

Theorem 3.3.3. *A joint measure π is a correlated equilibrium if and only if*

$$\int_{B_i \times C_{-i}} [u_i(t_i, s_{-i}) - u_i(s)] d\pi(s) \leq 0 \tag{3.2}$$

for all i , $t_i \in C_i$, and measurable subsets $B_i \subseteq C_i$.

Proof. (\Rightarrow) Fix i , $t_i \in C_i$ and a measurable set $B_i \subseteq C_i$. Define the measurable function $\xi_i : C_i \rightarrow C_i$ by $\xi_i(s_i) = t_i$ if $s_i \in B_i$ and $\xi_i(s_i) = s_i$ otherwise. By definition of a correlated equilibrium we have

$$\int_{B_i \times C_{-i}} [u_i(t_i, s_{-i}) - u_i(s)] d\pi(s) = \int [u_i(\xi_i(s_i), s_{-i}) - u_i(s)] d\pi(s) \leq 0.$$

(\Leftarrow) By Lemma 3.3.2 it suffices to show that the condition defining correlated equilibria holds for all simple measurable departure functions ξ_i . Fix such a function and let $\{t_i^1, \dots, t_i^k\}$ be its range. Let $B_i^k = \xi_i^{-1}(\{t_i^k\})$, which is measurable by assumption. Note that the sets B_i^1, \dots, B_i^k partition C_i , so we have

$$\begin{aligned} \int [u_i(\xi_i(s_i), s_{-i}) - u_i(s)] d\pi(s) &= \sum_{j=1}^k \int_{B_i^j \times C_{-i}} [u_i(\xi_i(s_i), s_{-i}) - u_i(s)] d\pi(s) \\ &= \sum_{j=1}^k \int_{B_i^j \times C_{-i}} [u_i(t_i^j, s_{-i}) - u_i(s)] d\pi(s) \leq \sum_{j=1}^k 0 = 0, \end{aligned}$$

where the inequality follows from (3.2). \square

The preceding lemma and theorem can be extended to an arbitrary set of players. This result is interesting on its own since in [18] it was conjectured that for an arbitrary set of players and compact Hausdorff strategy spaces, the analog of Theorem 3.3.3 does not hold.

The next theorem is an alternative characterization of correlated equilibria in continuous games, which we will use in Subsection 4.3.3 to develop a class of algorithms for computing (approximate) correlated equilibria.

Theorem 3.3.4. *A joint measure π is a correlated equilibrium if and only if*

$$\int f_i(s_i) [u_i(t_i, s_{-i}) - u_i(s)] d\pi(s) \leq 0 \tag{3.3}$$

for all i and $t_i \in C_i$ as f_i ranges over any of the following sets of functions from C_i to $[0, \infty)$:

1. Borel measurable functions,
2. Continuous functions,
3. Squares of polynomials (if $C_i \subset \mathbb{R}^{k_i}$ for some k_i).

Proof. Define the signed measure $\mu_{i,t_i}(B_i) = \int_{B_i \times C_{-i}} [u_i(s) - u_i(t_i, s_{-i})] d\pi(s)$ for each player i , strategy $t_i \in C_i$, and Borel measurable set $B_i \subseteq C_i$. Theorem 3.3.3 can be restated as follows: π is a correlated equilibrium if and only if μ_{i,t_i} is a positive measure for all i and t_i . If this condition holds, then $\int f_i d\mu_{i,t_i} \geq 0$ for all measurable functions $f_i : C_i \rightarrow [0, \infty)$.

Now fix any i and $t_i \in C_i$ and suppose $\int f_i d\mu_{i,t_i} \geq 0$ for all continuous functions $f_i : C_i \rightarrow [0, \infty)$. Then by the Riesz representation theorem there is a unique Borel measure ν_{i,t_i} such that $\int f_i d\mu_{i,t_i} = \int f_i d\nu_{i,t_i}$ for all continuous f_i and ν_{i,t_i} is a positive measure [38]. Therefore $\mu_{i,t_i} = \nu_{i,t_i}$ is a positive measure for all i and $t_i \in C_i$, so π is a correlated equilibrium.

Finally, assume that $C_i \subseteq \mathbb{R}^{k_i}$ for some k_i ; we will show that 3 \Rightarrow 2. Let $f_i : C_i \rightarrow [0, \infty)$ be a continuous function. By the Stone-Weierstrass theorem $\sqrt{f_i}$ can be approximated arbitrary well by a polynomial p , with respect to the sup norm. Thus f_i can be approximated arbitrarily well by a square of a polynomial with respect to the sup norm. Lebesgue's dominated convergence theorem completes the proof. \square

Finally, we will consider ϵ -correlated equilibria which are supported on some finite subset. In this case, we obtain another generalization of Proposition 2.1.5.

Theorem 3.3.5. *A probability measure $\pi \in \Delta(\tilde{C})$, where $\tilde{C} = \prod_{j=1}^n \tilde{C}_j$ is a finite subset of C , is an ϵ -correlated equilibrium of a continuous game if and only if there exist ϵ_{i,s_i} such that*

$$\sum_{s_{-i} \in \tilde{C}_{-i}} \pi(s) [u_i(t_i, s_{-i}) - u_i(s)] \leq \epsilon_{i,s_i}$$

for all players i , all $s_i \in \tilde{C}_i$, and all $t_i \in C_i$, and also

$$\sum_{s_i \in \tilde{C}_i} \epsilon_{i,s_i} \leq \epsilon$$

for all players i .

Proof. If we replace t_i with $\zeta_i(s_i)$ in the first inequality then sum over all $s_i \in \tilde{C}_i$ and combine with the second inequality, we get the equivalent condition that

$$\sum_{s \in \tilde{C}} \pi(s) [u_i(\zeta_i(s_i), s_{-i}) - u_i(s)] \leq \epsilon$$

holds for all i and any function $\zeta_i : \tilde{C}_i \rightarrow C_i$. This is exactly the definition of an ϵ -correlated equilibrium in the case when π is supported on \tilde{C} . \square

We can now prove a support bound for correlated equilibria in terms of the ranks, analogous to Theorem 3.2.5 for Nash equilibria.

Theorem 3.3.6. *Every separable game has a correlated equilibrium with support of size at most $1 + \sum_{i=1}^n \rho_{i,i}(\rho_{i,-i} + 1)$.*

Proof. First, apply Theorem 3.2.5 to obtain a Nash equilibrium in which each player mixes among at most $\rho_{i,-i} + 1$ pure strategies. Since Nash equilibria are correlated equilibria, this is automatically a correlated equilibrium supported on a finite cartesian product set \tilde{C} which satisfies $|\tilde{C}_i| \leq \rho_{i,-i} + 1$. By Theorem 3.2.6 we can write

the utilities of the game as

$$u_i(s) = h_i^0(s_{-i}) + \sum_{k=1}^{\rho_{i,i}} g_i^k(s_i) h_i^k(s_{-i})$$

for some continuous functions g_i^k and h_i^k . Applying Theorem 3.3.5 with $\epsilon = 0$ and $\epsilon_{i,s_i} = 0$ for all i , $s_i \in \tilde{C}_i$, we can write the conditions defining a correlated equilibrium on \tilde{C} as

$$\sum_{s_{-i} \in \tilde{C}_{-i}} \pi(s) [u_i(t_i, s_{-i}) - u_i(s)] = \sum_{k=1}^{\rho_{i,i}} [g_i^k(t_i) - g_i^k(s_i)] \sum_{s_{-i} \in \tilde{C}_{-i}} \pi(s) h_i^k(s_{-i}) \leq 0$$

for all players i , all $s_i \in \tilde{C}_i$, and all $t_i \in C_i$. In each term of the sum the values s_i and s_{-i} are fixed as t_i varies, so this inequality can be rewritten as

$$\sum_{k=1}^{\rho_{i,i}} L_{i,s_i}^k(\pi) [g_i^k(t_i) - g_i^k(s_i)] \leq 0 \quad (3.4)$$

for some linear functionals L_{i,s_i}^k .

By our assumption when defining \tilde{C} there exists some $\tilde{\pi} \in \Delta(\tilde{C})$ which is a correlated equilibrium of the game and so satisfies (3.4) for all i , all $s_i \in \tilde{C}_i$, and all $t_i \in C_i$. Therefore, any $\pi \in \Delta(\tilde{C})$ for which $L_{i,s_i}^k(\pi) = L_{i,s_i}^k(\tilde{\pi})$ for all players i , all $s_i \in \tilde{C}_i$, and all $1 \leq k \leq \rho_{i,i}$ satisfies (3.4) and hence is automatically a correlated equilibrium. The total number of linear functionals L_{i,s_i}^k is $\sum_{i=1}^n \rho_{i,i} |\tilde{C}_i|$, so by Caratheódory's theorem there exists a probability distribution $\pi \in \Delta(\tilde{C})$ which assigns positive probability to at most $1 + \sum_{i=1}^n \rho_{i,i} |\tilde{C}_i| \leq 1 + \sum_{i=1}^n \rho_{i,i} (1 + \rho_{i,-i})$ strategy profiles and satisfies $L_{i,s_i}^k(\pi) = L_{i,s_i}^k(\tilde{\pi})$ for all i , $s_i \in \tilde{C}_i$, and $t_i \in C_i$. This π is the desired correlated equilibrium. \square

3.4 Characterizations of Separable Games

In this section we prove a characterization theorem for separable games. Then we give a counterexample to show that the assumptions cannot be weakened.

Theorem 3.4.1. *For a continuous game, the following are equivalent:*

1. *The game is separable.*
2. *The game has finite rank (i.e. $\rho_{i,*} < \infty$ for all i).*
3. *The rank $\rho_{i,-i} < \infty$ for all i .*
4. *For each player i , every countably supported mixed strategy σ_i is almost payoff equivalent to a finitely supported mixed strategy τ_i with $\text{supp}(\tau_i) \subset \text{supp}(\sigma_i)$.*

To prove that finite rank implies separability we repeatedly apply Theorem 3.2.6. The proof that the technical condition (4) implies (3) uses a linear algebraic argument to show that $\text{span } C_i/U_{i,-i}$ is finite dimensional and then a topological argument along the lines of the proof of Theorem 2.1.21 to show that $V_i/U_{i,-i}$ is also finite dimensional.

After the proof of Theorem 3.4.1 we will give an explicit example of a game in which all mixed strategies are payoff equivalent to pure strategies, but for which the containment $\text{supp}(\tau_i) \subset \text{supp}(\sigma_i)$ in condition (4) fails. In light of Theorem 3.4.1 this will show that the constructed game is nonseparable and that the containment $\text{supp}(\tau_i) \subset \text{supp}(\sigma_i)$ cannot be dropped from condition (4).

Proof. (1 \Rightarrow 4) This was proven in Theorem 2.1.21.

(1 \Rightarrow 2) This follows from the proof of Theorem 2.1.21.

(2 \Rightarrow 3) This is immediate from the definitions.

(3 \Rightarrow 1) We will prove this by induction on the number of players n . When $n = 1$ the statement is trivial and the case $n = 2$ follows immediately from Theorem 3.2.6. Suppose we have an n -player continuous game with $\rho_{i,-i} < \infty$ for all i and that we have proven that $\rho_{i,-i} < \infty$ for all i implies separability for $(n - 1)$ -player games. By fixing any signed measure $\sigma_n \in V_n$ we can form an $(n - 1)$ -player continuous game from the given game by removing the n^{th} player and integrating all payoffs of players $i < n$ with respect to σ_n , yielding a new game with payoffs $\tilde{u}_i(s_{-n}) = \int u_i(s_n, s_{-n}) d\sigma_n(s_n)$.

From the definition of $U_{i,-i}$ it is clear that $U_{i,-i} \subseteq \tilde{U}_{i,-i}$ for all $1 \leq i < n$. Therefore $\tilde{\rho}_{i,-i} = \dim \Delta_i/\tilde{U}_{i,-i} \leq \dim \Delta_i/U_{i,-i} = \rho_{i,-i} < \infty$ for $1 \leq i < n$. By the induction hypothesis, that means that the function \tilde{u}_1 is a separable function of the strategies s_1, \dots, s_{n-1} . Theorem 3.2.6 states that there exist continuous functions $g_{n,1}^k$ and $h_{n,1}^k$ such that

$$u_1(s) = h_{n,1}^0(s_{-n}) + \sum_{k=1}^{\rho_{n,-n}} g_{n,1}^k(s_n) h_{n,1}^k(s_{-n}) \quad (3.5)$$

where $h_{n,1}^k = \int u_1(s) d\sigma_n^k$ for some $\sigma_n^k \in V_n$. Therefore by choosing σ_n appropriately we can make $\tilde{u}_1 = h_{n,1}^k$ for any k , so $h_{n,1}^k(s_{-n})$ is a separable function of s_1, \dots, s_{n-1} for all k . By (3.5) u_1 is a separable function of s_1, \dots, s_n . The same argument works for all the u_i so the given game is separable and the general case is true by induction.

(4 \Rightarrow 3) Let $\psi_i : V_i \rightarrow V_i/U_{i,-i}$ be the canonical projection transformation. First we will prove that $\text{span } \psi_i(C_i)$ is finite dimensional. It suffices to prove that for every countable subset $\tilde{C}_i = \{s_i^1, s_i^2, \dots\} \subseteq C_i$, the set $\psi_i(\tilde{C}_i)$ is linearly dependent. Let $\{p^k\}$ be a sequence of positive reals summing to unity. Define the mixed strategy

$$\sigma_i = \sum_{k=1}^{\infty} p^k s_i^k.$$

By assumption there exists an M and q^1, \dots, q^M summing to unity such that

$$\psi_i(\sigma_i) = \psi_i \left(\sum_{k=1}^M q^k s_i^k \right) = \sum_{k=1}^M q^k \psi_i(s_i^k).$$

Let $\alpha = \sum_{k=M+1}^{\infty} p^k > 0$ and define the mixed strategy

$$\tau_i = \sum_{k=M+1}^{\infty} \frac{p^k}{\alpha} s_i^k.$$

Applying the assumption again shows that there exist N and r^{M+1}, \dots, r^N such that

$$\psi_i(\tau_i) = \psi_i \left(\sum_{k=M+1}^N r^k s_i^k \right) = \sum_{k=M+1}^N r^k \psi_i(s_i^k).$$

Therefore

$$\begin{aligned} \sum_{k=1}^M p^k \psi_i(s_i^k) &= \psi_i \left(\sum_{k=1}^M p^k s_i^k \right) = \psi_i(\sigma_i - \alpha \tau_i) = \psi_i(\sigma_i) - \alpha \psi_i(\tau_i) \\ &= \sum_{k=1}^M q^k \psi_i(s_i^k) - \sum_{k=M+1}^N \alpha r^k \psi_i(s_i^k), \end{aligned}$$

and rearranging terms shows that $\sum_{k=1}^M (p^k - q^k) \psi_i(s_i^k) + \sum_{k=M+1}^N \alpha r^k \psi_i(s_i^k) = 0$. Also $\sum_{k=1}^M (p^k - q^k) = -\alpha < 0$, so $p^k - q^k \neq 0$ for some k . Therefore $\psi_i(\tilde{C}_i)$ is linearly dependent, so $\text{span } \psi_i(C_i)$ is finite dimensional.

Since $U_{i,-i}$ is defined as the intersection of the kernels of a family of continuous linear functionals on V_i , it is a closed subspace. Therefore $V_i/U_{i,-i}$ is a Hausdorff topological vector space under the quotient topology and ψ_i is continuous with respect to this topology [39]. Being finite dimensional, the subspace $\text{span } \psi_i(C_i) \subseteq V_i/U_{i,-i}$ is also closed [39]. Thus we have

$$V_i/U_{i,-i} = \psi_i(V_i) = \psi_i(\overline{\text{span } C_i}) \subseteq \overline{\psi_i(\text{span } C_i)} = \overline{\text{span } \psi_i(C_i)} = \text{span } \psi_i(C_i) \subseteq V_i/U_{i,-i}$$

where the first step is by definition, the second follows from 2.1.11, the next two are by continuity and linearity of ψ_i , and the final two are because $\text{span } \psi_i(C_i)$ is a closed subspace of $V_i/U_{i,-i}$. Therefore $\rho_{i,-i} = \dim \Delta_i/U_{i,-i} \leq \dim V_i/U_{i,-i} = \dim \text{span } \psi_i(C_i) < \infty$. \square

The following counterexample shows that the containment $\text{supp } \tau_i \subset \text{supp } \sigma_i$ is a necessary part of condition 4 in Theorem 3.4.1 by showing that there exists a nonseparable continuous game in which every mixed strategy is payoff equivalent to a pure strategy.

Example 3.4.2. Consider a two-player game with $C_1 = C_2 = [0, 1]^\omega$, the set of all infinite sequences of reals in $[0, 1]$, which forms a compact metric space under the metric

$$d(x, x') = \sup_i \frac{|x_i - x'_i|}{i}.$$

Define the utilities

$$u_1(x, y) = u_2(x, y) = \sum_{i=1}^{\infty} 2^{-i} x_i y_i.$$

To show that this is a continuous game we must show that u_1 is continuous. Assume $d(x, x'), d(y, y') \leq \delta$. Then $|x_i - x'_i| \leq \delta i$ and $|y_i - y'_i| \leq \delta i$, so

$$\begin{aligned} |u_1(x, y) - u_1(x', y')| &= \left| \sum_{i=1}^{\infty} 2^{-i} (x_i y_i - x'_i y'_i) \right| \\ &= \left| \sum_{i=1}^{\infty} 2^{-i} (x_i y_i - x'_i y_i + x'_i y_i - x'_i y'_i) \right| \\ &\leq \sum_{i=1}^{\infty} 2^{-i} (y_i |x_i - x'_i| + x'_i |y_i - y'_i|) \\ &\leq \sum_{i=1}^{\infty} 2^{-i} (2\delta i) = \left(2 \sum_{i=1}^{\infty} 2^{-i} i \right) \delta. \end{aligned}$$

Thus $u_1 = u_2$ is continuous (in fact Lipschitz), making this a continuous game.

Let σ and τ be mixed strategies for the two players. By the Tonelli-Fubini theorem,

$$u_1(\sigma, \tau) = \int u_1 d(\sigma \times \tau) = \sum_{i=1}^{\infty} 2^{-i} \left(\int x_i d\sigma \right) \left(\int y_i d\tau \right).$$

Thus σ is payoff equivalent to the pure strategy $(\int x_1 d\sigma, \int x_2 d\sigma, \dots) \in C_1$ and similarly for τ , so this game has the property that every mixed strategy is payoff equivalent to a pure strategy.

Finally we will show that this game is nonseparable. Let $e^i \in C_1$ be the element having component i equal to unity and all other components zero. Let $\{p_i\}$ be a sequence of positive reals summing to unity and define the probability distribution $\sigma = \sum_{i=1}^{\infty} p_i e^i \in \Delta_1$. Suppose σ were almost payoff equivalent to some mixture among finitely many of the e^i , call it $\tau = \sum_{i=1}^{\infty} q_i e^i$ where $q_i = 0$ for i greater than some fixed N . Let e_{N+1} be the strategy for player 2. Then the payoff if player 1 plays σ is

$$u_2(\sigma, e_{N+1}) = \int 2^{-(N+1)} x_{N+1} d\sigma = 2^{-(N+1)} p_{N+1}.$$

Similarly, if he chooses τ the payoff is $2^{-(N+1)} q_{N+1}$. Since $p_{N+1} > 0$ and $q_{N+1} = 0$, this contradicts the hypothesis that σ and τ are almost payoff equivalent. Thus condition 4 of Theorem 3.4.1 does not hold, so this game is not separable.

Therefore the condition that all mixed strategies be payoff equivalent to finitely supported strategies does not imply separability, even if a uniform bound on the size of the support is assumed. Hence the containment $\text{supp } \tau_i \subset \text{supp } \sigma_i$ cannot be removed from condition (4) of Theorem 3.4.1.

3.5 Computing Ranks

In this section we construct formulas for the ranks of an arbitrary separable game and specialize them to get formulas for the ranks of polynomial and finite games. In the proofs we will focus on the problem of computing the rank $\rho_{i,-i}$, but the same arguments work for the other ranks and we will state how they can be computed similarly. For clarity of presentation we first prove a bound on the ranks of a separable game which uses an argument which is similar to but simpler than the argument for the exact formula. While it is possible to prove all the results in this section directly from the definition of the ranks, we will give proofs based on the alternative characterization in Theorem 3.2.6 because they are easier to understand and provide more insight into the structure of the problem.

Given a separable game in the standard form (2.12), construct a matrix $S_{i,j}$ for players i and j which has m_i columns and $\prod_{k \neq i} m_k$ rows and whose elements are defined as follows. Label each row with an $(n-1)$ -tuple $(l_1, \dots, l_{i-1}, l_{i+1}, \dots, l_n)$ such that $1 \leq l_k \leq m_k$; the order of the rows is irrelevant. Label the columns $l_i = 1, \dots, m_i$. Each entry of the matrix then corresponds to an n -tuple (l_1, \dots, l_n) . The entry itself is given by the coefficient $a_j^{l_1 \dots l_n}$ in the utility function u_j .

Let $f_i(s_i)$ denote the column vector whose components are $f_i^1(s_i), \dots, f_i^{m_i}(s_i)$ and $f_{-i}(s_{-i})$ denote the row vector whose components are the products

$$f_1^{l_1}(s_1) \cdots f_{i-1}^{l_{i-1}}(s_{i-1}) f_{i+1}^{l_{i+1}}(s_{i+1}) \cdots f_n^{l_n}(s_n)$$

ordered in the same way as the $(n-1)$ -tuples $(l_1, \dots, l_{i-1}, l_{i+1}, \dots, l_n)$ were ordered above. Then $u_j(s) = f_{-i}(s_{-i}) S_{i,j} f_i(s_i)$.

Example 3.5.1. Consider a three player polynomial game with strategy spaces $C_1 = C_2 = C_3 = [-1, 1]$ and payoffs

$$\begin{aligned} u_1(x, y, z) &= 1 + 2x + 3x^2 + 2yz + 4xyz + 6x^2yz \\ &\quad + 3y^2z^2 + 6xy^2z^2 + 9x^2y^2z^2 \\ u_2(x, y, z) &= 7 + 2x + 3x^2 + 2y + 4xy + 6x^2y \\ &\quad + 3z^2 + 6xz^2 + 9x^2z^2 \\ u_3(x, y, z) &= -z - 2xz - 3x^2z - 2yz - 4xyz - 6x^2yz \\ &\quad - 3yz^2 - 6xyz^2 - 9x^2yz^2 \end{aligned} \tag{3.6}$$

where x , y , and z are the strategies of player 1, 2, and 3, respectively. Order the functions f_k^l so that $f_1(x) = [1 \ x \ x^2]'$ and similarly for f_2 and f_3 with x replaced by y and z , respectively. If we wish to write down the matrices $S_{1,2}$ and $S_{1,3}$ we must choose an order for the pairwise products of the functions f_2^l and f_3^l . Here we will choose the order $f_{-1}(y, z) = [1 \ y \ y^2 \ z \ yz \ y^2z \ z^2 \ yz^2 \ y^2z^2]$. We can write

down the desired matrices immediately from the given utilities.

$$S_{1,1} = \begin{bmatrix} 1 & 2 & 3 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 2 & 4 & 6 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 3 & 6 & 9 \end{bmatrix}, S_{1,2} = \begin{bmatrix} 7 & 2 & 3 \\ 2 & 4 & 6 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ 3 & 6 & 9 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \text{ and } S_{1,3} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ -1 & -2 & -3 \\ -2 & -4 & -6 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \\ -3 & -6 & -9 \\ 0 & 0 & 0 \end{bmatrix}$$

This yields $u_j(x, y, z) = f_{-1}(y, z)S_{1,j}f_1(x)$ for all j as claimed.

Define $S_{i,-i}$ to be the matrix with m_i columns and $(n-1)\prod_{j \neq i} m_j$ rows which consists of all the matrices $S_{i,j}$ for $j \neq i$ stacked vertically (in any order). Construct $S_{i,*}$ similarly by stacking the matrices $S_{i,j}$ for all j . In the example above, $S_{1,-1}$ would be the 18×3 matrix obtained by placing $S_{1,2}$ above $S_{1,3}$ on the page.

Theorem 3.5.2. *The ranks of a separable game are bounded by $\rho_{i,-i} \leq \text{rank } S_{i,-i}$, $\rho_{i,i} \leq \text{rank } S_{i,i}$, and $\rho_{i,*} \leq \text{rank } S_{i,*}$.*

Proof. Using any of a variety of matrix factorization techniques (e.g. the singular value decomposition), we can write $S_{i,-i}$ as

$$S_{i,-i} = \sum_{k=1}^{\text{rank } S_{i,-i}} v^k w^k$$

for some column vectors v^k and row vectors w^k . The vectors v^k will have length $(n-1)\prod_{j \neq i} m_j$ since that is the number of rows of $S_{i,-i}$. Because of the definition of $S_{i,-i}$, we can break each v^k into $n-1$ vectors of length $\prod_{j \neq i} m_j$, one for each player except i , and let v_j^k be the vector corresponding to player j . Then we have

$$S_{i,j} = \sum_{k=1}^{\text{rank } S_{i,-i}} v_j^k w^k$$

for all $j \neq i$. Define the linear combinations $g_i^k(s_i) = w^k f_i(s_i)$ and $h_{i,j}^k = f_{-i}(s_{-i})v_j^k$, which are obviously continuous functions. Then

$$u_j(s) = f_{-i}(s_{-i})S_{i,j}f_i(s_i) = \sum_{k=1}^{\text{rank } S_{i,-i}} g_i^k(s_i)h_{i,j}^k(s_{-i})$$

for all $s \in C$ and $j \neq i$, so $\rho_{i,-i} \leq \text{rank } S_{i,-i}$ by Theorem 3.2.6. \square

Example 3.5.3. To demonstrate the power of the bound in Theorem 3.5.2 we will use it to give an immediate proof of Theorem 3.1.1. Consider any two-player finite game,

where one player (labeled r) chooses rows in the payoff matrices and the other player (labeled c) chooses columns. Let $C_* = \{1, \dots, m_*\}$ be the set of strategies and U_* be the matrix of payoffs for players $* = r, c$. We can then define $f_*^l(s_*)$ to be unity if $s_* = l$ and zero otherwise. This gives

$$\begin{aligned} u_r(s_1, s_2) &= f_r(s_r)' U_r f_c(s_c) \\ u_c(s_1, s_2) &= f_c(s_c)' U_c' f_r(s_r) \end{aligned}$$

so $S_{r,-r} = U_c'$ and $S_{c,-c} = U_r$. Therefore by Theorem 3.5.2, $\rho_{r,-r} \leq \text{rank } S_{r,-r} = \text{rank } U_c$ and $\rho_{c,-c} \leq \text{rank } S_{c,-c} = \text{rank } U_r$. Substituting these bounds into Theorem 3.2.5 yields Theorem 3.1.1, so we have in fact generalized the results of Lipton et al. [26].

It is easy to see that there are cases in which the bound in Theorem 3.5.2 is not tight. For example, this will be the case (for generic coefficients $a_i^{j_1 \dots j_n}$) if $m_i \geq 2$ for each i and f_i^k is the same function for all k .

Fortunately we can use a similar technique to compute the ranks exactly instead of just computing bounds. To do so we need to write the utilities in a special form. First, we add the new function $f_i^1(s_i) \equiv 1$ to the list of functions for player i appearing in the separable representation of the game if this function does not already appear, relabeling the other f_i^k as necessary. Next, we consider the set of functions $\{f_j^k\}$ for each player j in turn and choose a maximal linearly independent subset. For players $j \neq i$ any such subset will do; for player i we must include the function which is identically unity in the chosen subset. Finally, we rewrite the utilities in terms of these linearly independent sets of functions. This is possible because all of the f_j^k are linear combinations of those which appear in the maximal linearly independent sets.

From now on we will assume the utilities are in this form and that $f_i^1(s_i) \equiv 1$. Let $\bar{S}_{i,j}$, $\bar{S}_{i,-i}$, and $\bar{S}_{i,*}$ be the matrices $S_{i,j}$, $S_{i,-i}$, and $S_{i,*}$ defined above, where the bar denotes the fact that we have put the utilities in this special form. Let $T_{i,-i}$ be the matrix $\bar{S}_{i,-i}$ with its first column removed, and similarly for $T_{i,i}$ and $T_{i,*}$. Note that the first column corresponds to the function $f_i^1(s_i) \equiv 1$ which we have distinguished above, and therefore represents the components of the utilities which do not depend on player i 's choice of strategy. As mentioned in the note following Theorem 3.2.6, these components don't affect the ranks. This is exactly the reason we must remove the first column from $\bar{S}_{i,-i}$ in order to compute $\rho_{i,-i}$. We will prove that $\rho_{i,-i} = \text{rank } T_{i,-i}$, but first we need a lemma.

Lemma 3.5.4. *If the functions $f_j^1(s_j), \dots, f_j^{m_j}(s_j)$ are linearly independent for all j , then the set of all $\prod_{j=1}^n m_j$ product functions of the form $f_1^{k_1}(s_1) \cdots f_n^{k_n}(s_n)$ is a linearly independent set.*

Proof. It suffices to prove this in the case $n = 2$, because the general case follows by induction. We prove the $n = 2$ case by contradiction. Suppose the set were linearly dependent. Then there would exist $\lambda_{k_1 k_2}$ not all zero such that

$$\sum_{k_1=1}^{m_1} \sum_{k_2=1}^{m_2} \lambda_{k_1 k_2} f_1^{k_1}(s_1) f_2^{k_2}(s_2) = 0 \quad (3.7)$$

for all $s \in C$. Choose l_1 and l_2 such that $\lambda_{l_1 l_2} \neq 0$. By the linear independence assumption there exists a signed measure σ_2 such that $\int f_2^k d\sigma_2$ is unity for $k = l_2$ and zero otherwise. Integrating (3.7) with respect to σ_2 yields

$$\sum_{k_1=1}^{m_1} \lambda_{k_1 l_2} f_1^{k_1}(s_1) = 0,$$

contradicting the linear independence assumption for $f_1^1, \dots, f_1^{m_1}$. \square

Theorem 3.5.5. *If the representation of a separable game satisfies $f_i^1 \equiv 1$ and the set $\{f_j^1, \dots, f_j^{m_j}\}$ is linearly independent for all j then the ranks of the game are $\rho_{i,-i} = \text{rank } T_{i,-i}$, $\rho_{i,i} = \text{rank } T_{i,i}$, and $\rho_{i,*} = \text{rank } T_{i,*}$.*

Proof. The proof that $\rho_{i,-i} \leq \text{rank } T_{i,-i}$ follows essentially the same argument as the proof of Theorem 3.5.2. We use the singular value decomposition to write $T_{i,-i}$ as

$$T_{i,-i} = \sum_{k=1}^{\text{rank } T_{i,-i}} v^k w^k$$

for some column vectors v^k and row vectors w^k . The vectors v^k will have length $(n-1)\prod_{j \neq i} m_j$ since that is the number of rows of $S_{i,-i}$. Let v^0 be the first column of $\bar{S}_{i,-i}$, which was removed from $S_{i,-i}$ to form $T_{i,-i}$. Because of the definition of $T_{i,-i}$ and $\bar{S}_{i,-i}$, we can break each v^k into $n-1$ vectors of length $\prod_{j \neq i} m_j$, one for each player except i , and let v_j^k be the vector corresponding to player j . Putting these definitions together we get

$$\bar{S}_{i,j} = v_j^0 [1 \ 0 \ \dots \ 0] + \sum_{k=1}^{\text{rank } T_{i,-i}} v_j^k [0 \ w^k].$$

Define the linear combinations $g_i^k(s_i) = [0 \ w^k] f_i(s_i)$ and $h_{i,j}^k(s_{-i}) = f_{-i}(s_{-i}) v_j^k$, which are obviously continuous functions. Then

$$u_j(s) = f_{-i}(s_{-i}) \bar{S}_{i,j} f_i(s_i) = h_{i,j}^0(s_{-i}) + \sum_{k=1}^{\text{rank } T_{i,-i}} g_i^k(s_i) h_{i,j}^k(s_{-i})$$

for all $s \in C$ and $j \neq i$, so $\rho_{i,-i} \leq \text{rank } T_{i,-i}$ by Theorem 3.2.6.

To prove the reverse inequality, choose continuous functions $g_i^k(s_i)$ and $h_{i,j}^k(s_{-i})$ such that

$$u_j(s) = h_{i,j}^0(s_{-i}) + \sum_{k=1}^{\rho_{i,-i}} g_i^k(s_i) h_{i,j}^k(s_{-i})$$

holds for all $s \in C$ and $j \neq i$. By Theorem 3.2.6 we can choose these so that $g_i^k(s_i)$ is of the form $u_j(s_i, s_{-i})$ for some $s_{-i} \in C_{-i}$, $j \neq i$ and $h_{i,j}^k(s_{-i})$ is of the form $\int u_j(\cdot, s_{-i}) d\sigma_i$ for some $\sigma_i \in V_i$. Substituting these conditions into equation (2.12) defining the form of a separable game shows that $g_i^k(s_i) = w^k f_i(s_i)$ for some row vectors w^k and

$h_{i,j}^k = f_{-i}(s_{-i})v_j^k$ for some column vectors v_j^k . Define $w^0 = [1 \ 0 \ \cdots \ 0]$. Then

$$u_j(s) = \sum_{k=0}^{\rho_{i,-i}} f_{-i}(s_{-i})' v_j^k w^k f_i(s_i)$$

for all $s \in C$ and $j \neq i$.

This expresses $u_j(s)$ as a linear combination of products of the form $f_1^{k_1}(s_1) \cdots f_n^{k_n}(s_n)$. By construction the sets $\{f_j^1, \dots, f_j^{m_j}\}$ are linearly independent for all j , and therefore the set of products of the form $f_1^{k_1}(s_1) \cdots f_n^{k_n}(s_n)$ is linearly independent by Lemma 3.5.4. Thus the expression of $u_j(s)$ as a linear combination of these products is unique.

But we also have $u_j(s) = f_{-i}(s_{-i})' \bar{S}_{i,j} f_i(s_i)$ by definition of $\bar{S}_{i,j}$, so uniqueness implies that $\bar{S}_{i,j} = \sum_{k=0}^{\rho_{i,-i}} v_j^k w^k$. Let v^k be the vector of length $(n-1)\prod_{j \neq i} m_j$ formed by concatenating the v_j^k in the obvious way. Then $\bar{S}_{i,-i} = \sum_{k=0}^{\rho_{i,-i}} v^k w^k$. Let \tilde{w}^k be w^k with its first entry removed. By definition of $T_{i,-i}$ we have $T_{i,-i} = \sum_{k=0}^{\rho_{i,-i}} v^k \tilde{w}^k$. But w^0 is the standard unit vector with a 1 in the first coordinate, so \tilde{w}^0 is the zero vector and we may therefore remove the $k=0$ term from the sum. Thus $T_{i,-i} = \sum_{k=1}^{\rho_{i,-i}} v^k \tilde{w}^k$, which proves that $\text{rank } T_{i,-i} \leq \rho_{i,-i}$. \square

As corollaries of Theorem 3.5.5 we obtain formulas for the rank of polynomial and finite games.

Corollary 3.5.6. *Consider a game with polynomial payoffs*

$$u_i(s) = \sum_{j_1=0}^{m_1-1} \cdots \sum_{j_n=0}^{m_n-1} a_i^{j_1 \cdots j_n} s_1^{j_1} \cdots s_n^{j_n} \quad (3.8)$$

and compact strategy sets $C_i \subset \mathbb{R}$ which satisfy the cardinality condition $|C_i| > m_i$ for all i . Then $T_{i,-i}$ is $S_{i,-i}$ with its first column removed and $\rho_{i,-i} = \text{rank } T_{i,-i}$. If $T_{i,i}$ and $T_{i,*}$ are formed in the same way then similarly $\rho_{i,i} = \text{rank } T_{i,i}$ and $\rho_{i,*} = \text{rank } T_{i,*}$.

Proof. Linear independence of the f_i^l follows from the cardinality condition and we have $f_i^0 \equiv 1$, so Theorem 3.5.5 applies. \square

Example 3.5.1 (cont'd). Applying this formula to the utilities in (3.6) shows that in this case $\rho_{1,-1} = 1$ and $\rho_{2,-2} = \rho_{3,-3} = 2$.

Corollary 3.5.7. *Consider an n -player finite game with strategy sets $C_i = \{1, \dots, m_i\}$ and payoff $a_i^{s_1 \cdots s_n}$ to player i if the players play strategy profile (s_1, \dots, s_n) . The utilities can be written as*

$$u_i(s) = \sum_{j_1=1}^{m_1} \cdots \sum_{j_n=1}^{m_n} a_i^{j_1 \cdots j_n} f_1^{j_1}(s_1) \cdots f_n^{j_n}(s_n)$$

where $f_i^l(s_i)$ is unity if $s_i = l$ and zero otherwise. Let $S_{i,-i}$ be the matrix for player i as defined above and let c_1, \dots, c_{m_i} be the columns of $S_{i,-i}$. Then we may take $T_{i,-i} = [c_2 - c_1 \ \cdots \ c_{m_i} - c_1]$ and $\rho_{i,-i} = \text{rank } T_{i,-i}$. If $T_{i,i}$ and $T_{i,*}$ are formed in the same way then similarly $\rho_{i,i} = \text{rank } T_{i,i}$ and $\rho_{i,*} = \text{rank } T_{i,*}$.

Proof. If we replace f_i^1 with the function which is identically unity then the linear independence assumption on the f_k^l will still be satisfied, so we can apply Theorem 3.5.5. After this replacement, the coefficients in the new separable representation for the game are

$$\bar{a}_k^{j_1 \cdots j_n} = \begin{cases} a_k^{j_1 \cdots j_n} & \text{if } j_i = 1, \\ a_k^{j_1 \cdots j_n} - a_k^{j_1 \cdots j_{i-1} 1 j_{i+1} \cdots j_n} & \text{if } j_i \neq 1. \end{cases}$$

Therefore if c_1, \dots, c_{m_i} are the columns of $S_{i,-i}$ from the original representation of the game we get $\bar{S}_{i,-i} = [c_1 \ c_2 - c_1 \ \cdots \ c_{m_i} - c_1]$, so $T_{i,-i}$ is as claimed and an application of Theorem 3.5.5 completes the proof. \square

Chapter 4

Computation of Equilibria in Separable Games

4.1 Computing Nash Equilibria

The moments of an equilibrium can in principle be computed by nonlinear programming techniques using the following generalization of the equilibrium formulation for finite games presented by Başar and Olsder [1] (see Subsection 2.2.1):

Proposition 4.1.1. *Consider the optimization problem*

$$\begin{aligned} \max \quad & \sum_{i=1}^n [v_i(x) - p_i] \\ \text{s.t.} \quad & x_i \in f_i(\Delta_i) \text{ for all } i \\ & v_i(f_i(s_i), x_{-i}) \leq p_i \text{ for all } i, \text{ all } s_i \in C_i \end{aligned} \tag{4.1}$$

where x_i are the moments, f_i is the moment function, and v_i is the payoff function on the moment spaces defined by $v_i(f_1(\sigma_1), \dots, f_n(\sigma_n)) = u_i(\sigma)$. Each player also has an auxiliary variable p_i . The optimum objective value of this problem is zero and is attained exactly when the x_i are the moments of a Nash equilibrium with payoff p_i to player i .

Proof. Same as Proposition 2.2.1. □

To compute equilibria by this method, we require an explicit description of the spaces of moments $f_i(\Delta_i)$. We also require a method for computing the payoff to player i if he plays a best response to an m_{-i} -tuple of moments for the other players.

While it seems doubtful that such descriptions could be found for arbitrary f_i^j , they do exist for two-player polynomial games in which the pure strategy sets are intervals. In this case the $x_i \in f_i(\Delta_i)$ constraints express that x_i are the moments of a measure on an interval in the real line. As shown in Appendix A, these constraints can be written in terms of linear matrix inequalities. The constraints of the form $v_i(f_i(s_i), x_{-i}) \leq p_i$ can be seen to express that the univariate polynomial $p_i - v_i(\cdot, x_{-i})$ is nonnegative on the interval $s_i \in C_i$. The coefficients of the polynomial are linear in the decision variables, so these constraints can also be written as linear matrix

inequalities as shown in Appendix A. Thus in this case the problem becomes one with biaffine objective function and linear matrix inequality constraints. This problem is still difficult to solve, due to the nonconvex objective function, but commercial solvers do exist for problems of this form.

If the game also satisfies the zero-sum condition $u_1(s) + u_2(s) \equiv 0$, then the objective function simplifies to the linear function $-(p_1 + p_2)$. Therefore the problem of computing Nash equilibria of two-player zero-sum games with polynomial payoffs can be cast exactly as a semidefinite program as shown in [33].

4.2 Computing ϵ -Nash Equilibria

The difficulties in computing Nash equilibria by general nonconvex optimization techniques suggest the need for more specialized systematic methods. As a step toward this, we present an algorithm for computing approximate Nash equilibria of two-player separable games. There are several possible definitions of approximate equilibrium, but here we will use the notion of ϵ -Nash equilibrium from Definition 2.1.12.

The algorithm we will present could be applied to many classes of separable games. For the sake of clarity we will make several assumptions, noting which could be relaxed and which could not.

Assumption 4.2.1.

- There are two players.
- The game is separable.
- The utilities can be evaluated efficiently.

Assumption 4.2.2. These conditions simplify the presentation but can be relaxed in a variety of ways.

- The strategy spaces are $C_1 = C_2 = [-1, 1]$.
- The utility functions are Lipschitz.

In the description of the algorithm we will emphasize why Assumption 4.2.1 is needed for our analysis. After presenting the algorithm we will discuss how Assumption 4.2.2 could be relaxed.

Theorem 4.2.3. *For $\epsilon > 0$, the following algorithm computes an ϵ -Nash equilibrium of a game of rank ρ satisfying Assumptions 4.2.1 and 4.2.2 in time polynomial in $\frac{1}{\epsilon}$ for fixed ρ and time polynomial in the components of ρ for fixed ϵ (for the purposes of asymptotic analysis of the algorithm with respect to ρ the Lipschitz condition is assumed to be satisfied uniformly by the entire class of games under consideration).*

Algorithm 4.2.4. By the Lipschitz assumption there are real numbers L_1 and L_2 such that

$$|u_i(s_i, s_{-i}) - u_i(s'_i, s_{-i})| \leq L_i |s_i - s'_i|$$

for all $s_{-i} \in C_{-i}$ and $i = 1, 2$. Clearly this is equivalent to requiring the same inequality for all $\sigma_{-i} \in \Delta_{-i}$. Divide the interval C_i into equal subintervals of length no more than $2\frac{\epsilon}{L_i}$; at most $\lceil \frac{L_i}{\epsilon} \rceil$ such intervals are required. Let \tilde{C}_i be the set of center points of these intervals and let \tilde{u}_i be the corresponding sampled payoffs. Call the resulting payoff matrices U_1 and U_2 . Compute a Nash equilibrium of the sampled game using support enumeration as described in Subsection 2.2.2. An equilibrium can be found by checking only supports of size up to $\rho_{i,-i}$ for player i , and this can be done in polynomial time.

Proof. For the purpose of analyzing the complexity of the algorithm we will view the Lipschitz constants as fixed, even as the ranks vary. Suppose σ is a Nash equilibrium of the sampled game. Choose any $s_i \in C_i$ and let s'_i be an element of \tilde{C}_i closest to s_i , so $|s_i - s'_i| \leq \frac{\epsilon}{L_i}$. Then

$$\begin{aligned} & u_i(s_i, \sigma_{-i}) - u_i(\sigma) \\ & \leq u_i(s_i, \sigma_{-i}) - u_i(s'_i, \sigma_{-i}) + u_i(s'_i, \sigma_{-i}) - u_i(\sigma) \\ & \leq |u_i(s_i, \sigma_{-i}) - u_i(s'_i, \sigma_{-i})| + \tilde{u}_i(s'_i, \sigma_{-i}) - \tilde{u}_i(\sigma) \\ & \leq L_i \frac{\epsilon}{L_i} + 0 = \epsilon \end{aligned}$$

so σ is automatically an ϵ -Nash equilibrium of the original separable game. Thus it will suffice to compute a Nash equilibrium of the finite sampled game.

To do so, first compute or bound the ranks $\rho_{i,-i}$ of the original separable game using Theorem 3.5.5 or 3.5.2. By Theorem 3.2.5 and Corollary 3.2.7, the sampled game has a Nash equilibrium in which player i mixes among at most $\rho_{i,-i} + 1$ pure strategies, independent of how large $|\tilde{C}_i|$ is. The separability assumption is fundamental because without it we would not obtain this uniform bound independent of $|\tilde{C}_i|$. The number of possible choices of at most $\rho_{i,-i} + 1$ pure strategies from \tilde{C}_i is

$$\sum_{k=1}^{\rho_{i,-i}+1} \binom{|\tilde{C}_i|}{k} \leq \binom{|\tilde{C}_i| + \rho_{i,-i}}{1 + \rho_{i,-i}} = \binom{|\tilde{C}_i| + \rho_{i,-i}}{|\tilde{C}_i| - 1},$$

which is a polynomial in $|\tilde{C}_i| \propto \frac{1}{\epsilon}$ for fixed $\rho_{i,-i}$ and a polynomial in the $\rho_{i,-i}$ for fixed ϵ . This leaves the step of checking whether there exists an equilibrium σ for a given choice of $S_i = \text{supp}(\sigma_i) \subseteq \tilde{C}_i$ with $|S_i| \leq \rho_{i,-i} + 1$ for each i , and if so, computing such an equilibrium. If the game has two players, the set of such equilibria for given supports is described by a number of linear equations and inequalities which is polynomial in $\frac{1}{\epsilon}$ for fixed $\rho_{i,-i}$ and polynomial in the $\rho_{i,-i}$ for fixed ϵ ; these equations and inequalities are given by (2.15). Using a polynomial time linear programming algorithm we can find a feasible solution to such inequalities or prove infeasibility in polynomial time. The two player assumption is key at this step, because with more

players the constraints would fail to be linear or convex and we could no longer use a polynomial time linear programming algorithm.

Thus we can check all supports and find an ϵ -Nash equilibrium of the sampled game in polynomial time as claimed. \square

We will now consider weakening Assumption 4.2.2. The Lipschitz condition could be weakened to a Hölder condition and the same proof would work, but it seems that we must require some quantitative bound on the speed of variation of the utilities in order to bound the running time of the algorithm. Also, the strategy space could be changed to any compact set which can be efficiently sampled, e.g. a box in \mathbb{R}^n . However, for the purpose of asymptotic analysis of the algorithm, the proof here only goes through when the Lipschitz constants and strategy space are fixed. A more complex analysis would be required if the strategy space were allowed to vary with the ranks, for example.

It should be noted that the requirement of a quantitative bound on the speed of variation of the utilities and the requirement that the strategy space be fixed for asymptotic analysis mean that Theorem 4.2.3 does not apply to finite games, at least not if the number of strategies is allowed to vary. For the sake of comparison and completeness we record the best known ϵ -Nash equilibrium algorithm for finite games here.

Theorem 4.2.5 (Lipton et al. [26]). *There exists an algorithm to compute an ϵ -Nash equilibrium of an m -player finite game with n strategies per player which is polynomial in $\frac{1}{\epsilon}$ for fixed m and n , polynomial in m for fixed n and ϵ , and quasipolynomial in n for fixed ϵ and m (assuming the payoffs of the games are uniformly bounded).*

In the case of two-player separable games which we have considered, the complexity of the payoffs is captured by the ranks, which are bounded by the cardinality of the strategy spaces in two-player finite games. Therefore in finite games the complexity of the payoffs and the complexity of the strategy spaces are intertwined, whereas in games with infinite strategy spaces they are decoupled. The best known algorithm for finite games stated in Theorem 4.2.5 has quasipolynomial dependence on the complexity of the game. Our algorithm is interesting because it has polynomial dependence on the complexity of the payoffs when the strategy spaces are held fixed. In finite games this type of asymptotic analysis is not possible due to the coupling between the two notions of complexity of a game, so a direct comparison between Theorem 4.2.3 and Theorem 4.2.5 cannot be made.

4.3 Approximating Correlated Equilibria

When faced with a game to analyze, we may be interested in any of the following problems:

- (P1) computing a single correlated equilibrium,
- (P2) computing a projection of the entire set of correlated equilibria, or

(P3) optimizing some objective function over the set of correlated equilibria.

Since computing exact correlated equilibria in continuous games is intractable (no finite-dimensional formulation is known; see Section 3.3), we focus in this section on developing algorithms that can compute approximate correlated equilibria with arbitrary accuracy. We consider three types of algorithms. In the first, the strategy sets are discretized without regard to the structure of the game. This algorithm applies to any continuous game, but may not use computational resources efficiently and may yield conservative performance estimates. The other two classes of algorithms apply only to polynomial games, but take advantage of the algebraic structure present and function better than the first in practice.

Example 4.3.1. We will use the following polynomial game to illustrate the algorithms presented below. The game has two players, x and y , who each choose their strategies from the interval $C_x = C_y = [-1, 1]$. Their utilities are given by

$$\begin{aligned} u_x(x, y) &= 0.596x^2 + 2.072xy - 0.394y^2 + 1.360x \\ &\quad - 1.200y + 0.554 \\ u_y(x, y) &= -0.108x^2 + 1.918xy - 1.044y^2 - 1.232x \\ &\quad + 0.842y - 1.886. \end{aligned}$$

where the coefficients have been selected at random. This example is convenient, because as Figure 4.3.3 shows, the game has a unique correlated equilibrium (the players choose $x = y = 1$ with probability one), which makes convergence easier to see. For the purposes of visualization and comparison, we will project the computed equilibria and approximations thereof into expected utility space, i.e. we will plot pairs $(\int u_x d\pi, \int u_y d\pi)$.

4.3.1 Static Discretization Methods

The techniques in this subsection are general enough to apply to arbitrary continuous games, so we will not restrict our attention to polynomial games here. The basic idea of **static discretization methods** is to select some finite subset $\tilde{C}_i \subset C_i$ of strategies for each player and limit his strategy choice to that set. Restricting the utility functions to the product set $\tilde{C} = \prod_{i=1}^n \tilde{C}_i$ produces a finite game, called a **sampled game** or **sampled version** of the original continuous game. The simplest computational approach is then to consider the set of correlated equilibria of this sampled game. This set is defined by the linear inequalities in Proposition 2.1.5 along with the conditions that π be a probability measure on \tilde{C} , so in principle it is possible to solve any of the problems (P1) - (P3) for the discretized game using standard linear programming techniques. The complexity of this approach in practice depends on the number of points in the discretization.

The question is then: what kind of approximation does this technique yield? In general the correlated equilibria of the sampled game may not have any relation to the set of correlated equilibria of the original game. The sampled game could, for example, be constructed by selecting a single point from each strategy set, in which

case the unique probability measure over \tilde{C} is automatically a correlated equilibrium of the sampled game but is a correlated equilibrium of the original game if and only if the points chosen form a pure strategy Nash equilibrium. Nonetheless, it seems intuitively plausible that if a large number of points were chosen such that any point of C_i were near a point of \tilde{C}_i then the set of correlated equilibria of the finite game would be “close to” the set of correlated equilibria of the original game in some sense, despite the fact that each set might contain points not contained in the other.

To make this precise, we will show how to choose a discretization so that the correlated equilibria of the finite game are ϵ -correlated equilibria of the original game.

Proposition 4.3.2. *Given a continuous game with strategy sets C_i and payoffs u_i along with any $\epsilon > 0$, there exist $\delta_i > 0$ such that if all points of C_i are within δ_i of a point of the finite set $\tilde{C}_i \subseteq C_i$ (such a \tilde{C}_i exists since C_i is a compact metric space, hence totally bounded) then all correlated equilibria of the sampled game with strategy spaces \tilde{C}_i and utilities $u_i|_{\tilde{C}_i}$ will be ϵ -correlated equilibria of the original game.*

Proof. Note that the utilities are continuous functions on a compact set, hence uniformly continuous. Therefore for any $\epsilon > 0$ we can choose $\delta_i > 0$ such that if we change any of the arguments of u_i by a distance of no more than δ_i , then u_i changes by no more than ϵ . Let \tilde{C} satisfy the stated assumption and let π be any correlated equilibrium of the corresponding finite game. Then by Proposition 2.1.5,

$$\sum_{s_{-i} \in \tilde{C}_{-i}} \pi(s) [u_i(t_i, s_{-i}) - u_i(s)] \leq 0$$

for all i and all $s_i, t_i \in \tilde{C}_i$. Any $t_i \in C_i$ is within a distance δ_i of some $\tilde{t}_i \in \tilde{C}_i$, so

$$\begin{aligned} & \sum_{s_{-i} \in \tilde{C}_{-i}} \pi(s) [u_i(t_i, s_{-i}) - u_i(s)] \\ & \leq \sum_{s_{-i} \in \tilde{C}_{-i}} \pi(s) [u_i(\tilde{t}_i, s_{-i}) - u_i(s) + \epsilon] \\ & \leq \epsilon \sum_{s_{-i} \in \tilde{C}_{-i}} \pi(s). \end{aligned}$$

Therefore the assumptions of Theorem 3.3.5 are satisfied with $\epsilon_{i,s_i} = \epsilon \sum_{s_{-i} \in \tilde{C}_{-i}} \pi(s)$. \square

The proof shows that if the utilities are Lipschitz functions, such as polynomials, then the δ_i can in fact be chosen proportional to ϵ , so the number of points needed in \tilde{C}_i is $O(\frac{1}{\epsilon})$. More concretely, if the strategy spaces are $C_i = [-1, 1]$ as in a polynomial game, then \tilde{C}_i can be chosen to be uniformly spaced within $[-1, 1]$, and if this is done ϵ will be $O(\frac{1}{d})$ where $d = \max_i |\tilde{C}_i|$.

Example 4.3.1 (continued). Figure 4.3.1 is a sequence of static discretizations for this game for increasing values of d , where d is the number of points in \tilde{C}_x and \tilde{C}_y . These




Figure 4.3.1: Convergence of a sequence of ϵ -correlated equilibria of the game in Example 4.3.1 computed by a sequence of static discretizations, each with some number d of equally spaced strategies chosen for each player. The axes represent the utilities received by players x and y . It can be shown that the convergence in this example happens at a rate $\epsilon(d) = \Theta\left(\frac{1}{d}\right)$.

points are selected by dividing $[-1, 1]$ into d subintervals of equal length and letting $\tilde{C}_x = \tilde{C}_y$ be the set of midpoints of these subintervals. For this game it is possible to show that the rate of convergence is in fact $\epsilon(d) = \Theta\left(\frac{1}{d}\right)$ so the worst case bound on convergence rate is achieved in this example.

4.3.2 Adaptive Discretization Methods

For the next two subsections, we restrict attention to the case of polynomial games. While static discretization methods are straightforward, they do not exploit the algebraic structure of polynomial games. Furthermore, the sampling of points in C_i to produce an ϵ -correlated equilibrium via Proposition 4.3.2 is conservative, in two senses. First, the ϵ -correlated equilibrium produced by that method may in fact be an ϵ^* -correlated equilibrium for some ϵ^* which is much less than ϵ ; below we will show how to compute the minimal value of ϵ^* for a given joint probability measure. Second, polynomial games are separable, so any polynomial game has a Nash equilibrium, and

hence a correlated equilibrium, which is supported on a finite set (see Chapters 2 and 3). Hence, at least in principle there is no need for the number of points in \tilde{C}_i to grow without bound as $\epsilon \rightarrow 0$. In this subsection we consider methods in which the points in the discretization are chosen more carefully.

An **adaptive discretization method** is an iterative procedure in which the finite set of strategies \tilde{C}_i available to player i changes in some way on each iteration; we let \tilde{C}_i^k denote the strategies available to player i on the k^{th} iteration. The goal of such a method is to produce a sequence of ϵ^k -correlated equilibria with $\epsilon^k \rightarrow 0$.

There are many possible update rules to generate \tilde{C}_i^{k+1} from \tilde{C}_i^k . The simplest are the **dense update rules** in which $\tilde{C}_i^1 \subseteq \tilde{C}_i^2 \subseteq \dots$ and $\bigcup_{k=1}^{\infty} \tilde{C}_i^k$ is dense in C_i for all i . However, if such a method adds points without regard to the problem structure many iterations may be wasted adding points which do not get the algorithm any closer to a correlated equilibrium. Furthermore, the size of the discretized strategy sets \tilde{C}_i^k may become prohibitively large before the algorithm begins to converge. Therefore it seems advantageous to choose the points to add to \tilde{C}_i^k in a structured way, and it may also be worthwhile to delete points which don't seem to be in use after a particular iteration.

To get a handle on the convergence properties of these algorithms, we will use the ϵ -correlated equilibrium characterization in Theorem 3.3.5 since we are dealing with sampled strategy spaces. By that theorem, we can begin with a product set $\tilde{C} \subseteq C$ and find the joint measures $\pi \in \Delta(\tilde{C})$ which correspond to ϵ -correlated equilibria with minimal ϵ values by solving the following optimization problem:

$$\begin{aligned}
& \text{minimize } \epsilon \\
& \text{s.t.} \\
& \sum_{s_{-i} \in \tilde{C}_{-i}} \pi(s) [u_i(t_i, s_{-i}) - u_i(s)] \leq \epsilon_{i,s_i} \quad \text{for all } i, s_i \in \tilde{C}_i, \text{ and } t_i \in C_i \\
& \sum_{s_i \in \tilde{C}_i} \epsilon_{i,s_i} \leq \epsilon \quad \text{for all } i \\
& \pi(s) \geq 0 \quad \text{for all } s \in \tilde{C} \\
& \sum_{s \in \tilde{C}} \pi(s) = 1
\end{aligned} \tag{4.2}$$

For fixed s_{-i} , the functions $u_i(t_i, s_{-i})$ are univariate polynomials in t_i , so this problem can be solved exactly as a semidefinite program (see Lemmas A.2 and A.3 in Appendix A).

If the sequence of optimal ϵ values tends to zero for all games under a given update rule, we say that rule converges. Dense update rules converge by Proposition 4.3.2. Given the problem (4.2), a natural category of update rules are those which select an optimal solution to the problem, remove any strategies which are assigned zero or nearly zero probability in this solution, then add some or all of the values t_i which make the inequalities tight in this optimal solution into \tilde{C}_i^k to obtain \tilde{C}_i^{k+1} . This corresponds to selecting constraints in Definition 2.1.16 which are maximally violated by the chosen optimal solution, so we call these **maximally violated constraint**

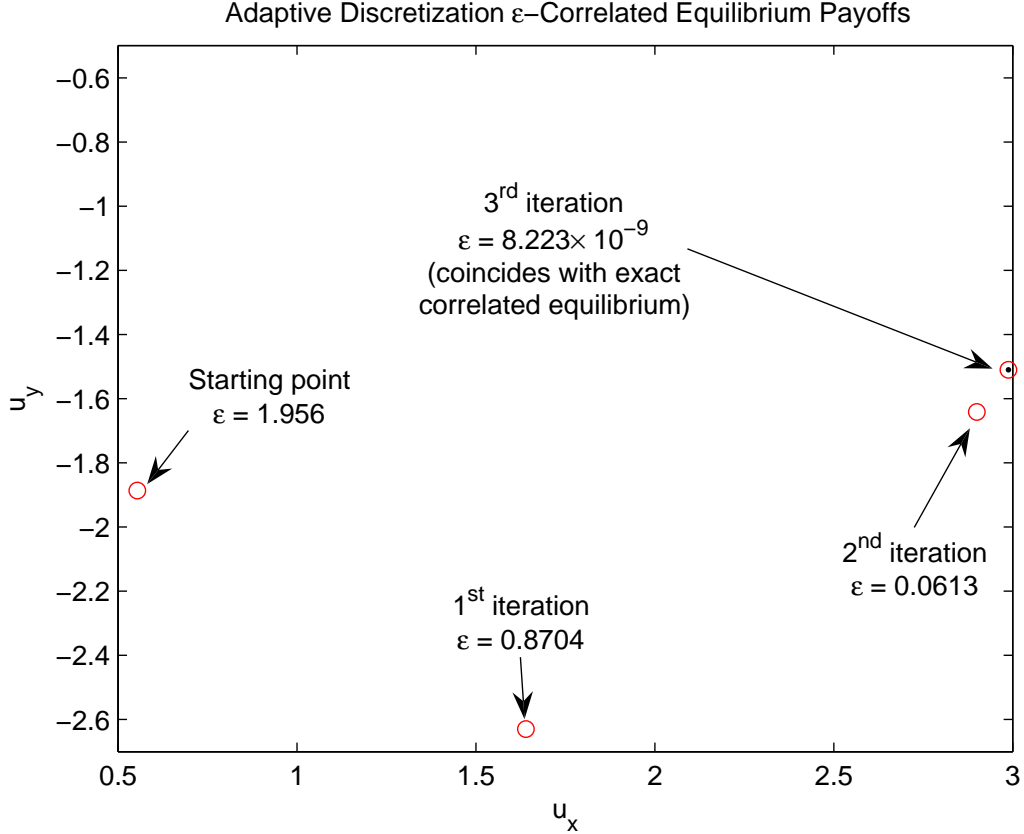


Figure 4.3.2: Convergence of an adaptive discretization method with a maximally violated constraint update rule (note the change in scale from Figure 4.3.1). At each iteration, the expected utility pair is plotted along with the computed value of ϵ for which that iterate is an ϵ -correlated equilibrium of the game. In this case convergence to $\epsilon = 0$ (to within numerical error) occurred in three iterations.

update rules. These rules seem to perform well in practice, but it is not known whether they converge in general.

Example 4.3.1 (continued). In Figure 4.3.2 we illustrate an adaptive discretization method using a maximally violated constraint update rule. The solver was initialized with $\tilde{C}_x^0 = \tilde{C}_y^0 = \{0\}$. At each iteration the ϵ -correlated equilibrium π of minimal ϵ -value was computed. Then ϵ was reported and one player's sampled strategy set was enlarged, the player for whom the constraint $\sum_{s_i \in \tilde{C}_i^k} \epsilon_{i,s_i} \leq \epsilon$ was tight. To choose which points to add to \tilde{C}_i^k , the algorithm identified the points $s_i \in \tilde{C}_i^k$ which were assigned positive probability under π . For each such s_i the values of $t_i \in C_i$ making the constraints in (4.2) tight were added to \tilde{C}_i^k to obtain \tilde{C}_i^{k+1} . The other player's strategy set was not changed.

In this case convergence happened in three iterations, significantly faster than the static discretization method. The resulting strategy sets were $\tilde{C}_x^3 = \{0, 1\}$ and $\tilde{C}_y^3 = \{0, 0.9131, 1\}$.

4.3.3 Semidefinite Relaxation Methods

In this subsection we again consider only polynomial games. The **semidefinite relaxation methods** for computing correlated equilibria have a different flavor from the discretization methods discussed above. Instead of using tractable finite approximations of the correlated equilibrium problem derived via discretizations, we begin with the alternative exact characterization given in Condition 3 of Theorem 3.3.4. In particular, a measure π on C is a correlated equilibrium if and only if

$$\int p^2(s_i) [u_i(t_i, s_{-i}) - u_i(s)] d\pi(s) \leq 0 \quad (4.3)$$

for all i , $t_i \in C_i$, and polynomials p . If we wish to check all these conditions for polynomials p of degree less than or equal to d , we can form the matrices

$$S_i^d = \begin{bmatrix} 1 & s_i & s_i^2 & \cdots & s_i^d \\ s_i & s_i^2 & s_i^3 & \cdots & s_i^{d+1} \\ s_i^2 & s_i^3 & s_i^4 & \cdots & s_i^{d+2} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ s_i^d & s_i^{d+1} & s_i^{d+2} & \cdots & s_i^{2d} \end{bmatrix}.$$

Let c be a column vector of length $d + 1$ whose entries are the coefficients of p , so $p^2(s_i) = c' S_i^d c$. If we define

$$M_i^d(t_i) = \int S_i^d [u_i(t_i, s_{-i}) - u_i(s)] d\pi(s),$$

then (4.3) is satisfied for all p of degree at most d if and only if $c' M_i^d(t_i) c \leq 0$ for all c , i.e. if and only if $M_i^d(t_i)$ is negative semidefinite.

The matrix $M_i^d(t_i)$ has entries which are polynomials in t_i with coefficients which are linear in the joint moments of π . To check the condition that $M_i^d(t_i)$ be negative semidefinite for all $t_i \in [-1, 1]$ for a given d we can use a semidefinite program (Lemma A.4 in Appendix A), so as d increases we obtain a sequence of semidefinite relaxations of the correlated equilibrium problem and these converge to the exact condition for a correlated equilibrium.

We can also let the measure π vary by replacing the moments of π with variables and constraining these variables to satisfy some necessary conditions for the moments of a joint measure on C (see Appendix A). These conditions can be expressed in terms of linear matrix inequalities and there is a sequence of these conditions which converges to a description of the exact set of moments of a joint measure π . Thus we obtain a nested sequence of semidefinite relaxations of the set of moments of measures which are correlated equilibria, and this sequence converges to the set of correlated equilibria. In this way we can use semidefinite relaxation methods to solve problems (P2) and (P3) given above.

Example 4.3.1 (continued). Figure 4.3.3 shows semidefinite relaxations of orders $d = 0, 1$, and 2. Since semidefinite relaxations are outer approximations of the set of

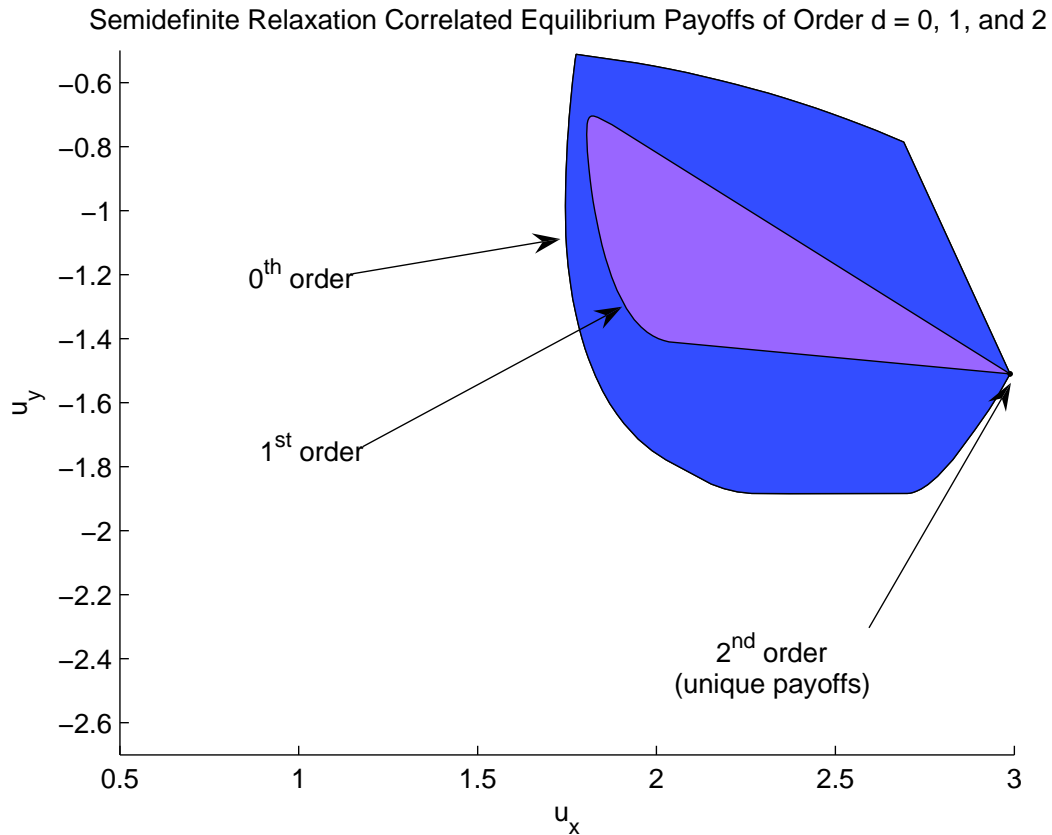


Figure 4.3.3: Semidefinite relaxations approximating the set of correlated equilibrium payoffs. The second order relaxation is a singleton, so this game has a unique correlated equilibrium.

correlated equilibria and the 2nd order semidefinite relaxation corresponds to a unique point in expected utility space, all correlated equilibria of the example game have exactly this expected utility. In fact, the set of points in this relaxation is a singleton (even before being projected into utility space), so this proves that the example game has a unique correlated equilibrium.

Chapter 5

Conclusions

5.1 Summary

We have shown that separable games, and in particular polynomial games, form a natural setting in which to study computational aspects of infinite games. To begin this argument, we considered properties which a game would need in order for the complexity of mixed equilibria to scale gracefully with the complexity of the game. This led to the concept of separability as a qualitative description of what it means for an infinite game to be “simple,” and to our new notion of the ranks of a game which can be viewed as quantitative measures of a game’s complexity.

We constructed simple bounds on the complexity of equilibria of separable games, measured as the number of strategies (or strategy profiles in the case of correlated equilibria) played with positive probability, which do in fact scale gracefully with the ranks. We also showed that there is an even closer connection between the qualitative idea of separability and the quantitative idea of ranks, in the sense that a game is separable if and only if its ranks are finite. This link confirms the fundamental theoretical importance of both concepts.

Viewing finite games as a subclass of separable games, we applied the rank bounds on equilibrium complexity in the setting of finite games as well. This yielded generalizations and improvements on the known rank-type bounds for finite games (see [26] and [16]) as immediate consequences of this new theory designed for infinite games. Furthermore, the abstraction needed to prove these results in the more general setting of infinite games made the structure implied by low-rank assumptions clearer and the arguments easier to visualize than in the original setting of finite games.

After these bounds and characterizations, we showed how separability and ranks apply to computation proper. In particular, the problem of computing an approximate Nash equilibrium of a two-player separable game reduces to the problem of computing exact Nash equilibria of finite games with low-rank payoff matrices. This problem can be solved efficiently by support enumeration. The ranks of the original separable game give an immediate bound on the running time of the resulting algorithm, for a fixed degree of approximation.

While we have given support bounds on correlated equilibria in terms of ranks,

the ranks do not appear to have a close connection to fast algorithms for computing correlated equilibria. The reason for this is that algorithmically, ranks seem to be closely tied to methods of support enumeration, as used to compute Nash equilibria. The ranks can bound the running time of static discretization, which is based on linear programming and hence can be viewed as purely combinatorial.

However, we have also constructed better techniques for computing correlated equilibria (at least in the case of polynomial games), and ranks do not appear to be key parameters of these algorithms. These faster algorithms are based on new characterizations of exact and approximate correlated equilibria in infinite games which mesh well with the algebraic structure of polynomial games. In particular, we have introduced adaptive discretization and semidefinite relaxation, both of which are based on semidefinite programming rather than linear programming. The time required to solve semidefinite programs is a function of the degree of the polynomials involved, and a polynomial game may have low ranks but payoffs of high degree. Thus it seems that novel techniques will likely be required to understand the behavior of these new algorithms in detail.

5.2 Future Work

Our primary motivation is to identify and solve game theoretic problems with infinite strategy sets. Viewing polynomial games as the class of infinite games with the most computationally tractable structure, we will restrict attention to this class when discussing possible avenues for future research. Any further generalizations to separable games or beyond would also be interesting.

The most pressing open questions regarding the above work concern the convergence of the semidefinite programming based algorithms for computing correlated equilibria of polynomial games. In particular, it is not known whether there exists an adaptive discretization method which can be proven to converge under a maximally violated constraint update rule. We conjecture that the answer is yes, since empirical evidence on random instances suggests that convergence occurs quickly. The semidefinite relaxation methods also seem to converge quickly in practice, and while we have shown above that these do converge, bounding the rate of convergence and obtaining error estimates remain important open problems.

Aside from the rank results of Lipton et al. [26] and Germano and Lugosi [16], other types of low-rank conditions have been studied for finite games as well. For example Kannan and Theobald have considered the condition on two-player games that the sum of the payoff matrices be low-rank [20]. It is likely that that the discretization techniques used in Chapter 4 can be applied in an analogous way to yield results about computing approximate equilibria of continuous games when the sum of the payoffs of the players satisfies the infinite game equivalent of a low-rank assumption, i.e. it is a separable function. It also may be possible to extend their result directly to the setting of polynomial games without discretization. Their procedure uses an algorithm for efficiently solving optimization problems with linear constraints and low-rank indefinite quadratic objective functions which is due to Vavasis [46]. If

this algorithm could be extended to handle linear matrix inequality constraints (i.e. those that appear in the optimization formulation of the Nash equilibria of a polynomial game), then Kannan and Theobald's procedure would immediately extend to an algorithm computing exact Nash equilibria of two-player polynomial games when the sum of the payoff polynomials has low rank.

There also exist many computational techniques for finite games which do not make low-rank assumptions. It may be possible to extend some of these techniques directly to polynomial games to yield algorithms for computing exact equilibria of polynomial games. While the lack of polyhedral structure in the moment spaces would most likely prohibit the use of a Lemke-Howson type algorithm, a variety of other finite game algorithms may be extendable to this setting; see McKelvey and McLennan for a survey of such algorithms [27].

Computing Nash equilibria of two-player zero-sum games and correlated equilibria of arbitrary games are two of the main equilibrium problems in game theory known to lead to convex optimization problems. In [33] and the present work, respectively, it has been shown that these can be solved using sum of squares methods in the case of polynomial games. We leave the task of extending these results to other convex equilibrium-type problems in polynomial games for future work.

Finally, there exist many solution concepts which we have not explored in this thesis. For example, it is trivial to search for pure strategy Nash equilibria of finite games, but it is not clear whether there is an efficient algorithm to solve the corresponding problem for polynomial games.

A seemingly trickier example is to perform iterated elimination of strictly dominated strategies on a polynomial game. This solution concept is a weakening of the Nash and correlated equilibria, which removes so-called strictly dominated strategies, those which are dominated by another strategy which is better in all situations. Doing so may allow more strategies to become dominated, so this procedure can be iterated. This is easily done in finite games, but in polynomial games it is much less obvious how to do even one iteration. Furthermore, there exist examples of polynomial games in which the procedure can be repeated infinitely often, strictly shrinking the set of available strategies at each stage. Computing the limit of this process thus appears to be even more difficult. An alternative characterization of this limit suggests that exact computation may nonetheless be possible [6].

Appendix A

Sum of Squares Techniques

Below we summarize the sum of squares results used in Chapter 4. Broadly, sum of squares methods allow nonnegativity conditions on polynomials to be expressed exactly or approximately as small semidefinite programs, and hence to be used in optimization problems which can be solved efficiently by interior point methods for semidefinite programming. The condition that a list of numbers correspond to the moments of a measure is dual to polynomial nonnegativity and can also be represented by similar semidefinite constraints.

The idea of sum of squares techniques is that the square of a real-valued function is nonnegative on its entire domain, and hence the same is true of a sum of squares of real-valued functions. In particular, any polynomial of the form $p(x) = \sum p_k^2(x)$, where p_k are polynomials, is guaranteed to be nonnegative for all x . This gives a sufficient condition for a polynomial to be nonnegative. It is a classical result that this condition is also necessary if p is univariate [35].

Lemma A.1. *A univariate polynomial p is nonnegative on \mathbb{R} if and only if it is a sum of squares.*

Frequently we are interested in polynomials which are nonnegative only on some interval such as $[-1, 1]$. These can be characterized almost as simply.

Lemma A.2. *A univariate polynomial p is nonnegative on $[-1, 1]$ if and only if $p(x) = s(x) + (1 - x^2)t(x)$ where s and t are sums of squares.*

These sum of squares conditions are easy to express using linear equations and semidefinite constraints. The proof of the following claim proceeds by factoring the positive semidefinite matrix P as a product $P = Q'Q$.

Lemma A.3. *A univariate polynomial $p(x) = \sum_{k=0}^d p_k x^k$ of degree d is a sum of squares if and only if there exists a $(d + 1) \times (d + 1)$ positive semidefinite matrix P which satisfies $p_k = \sum_{i+j=k} P_{i,j}$ when the rows and columns of P are numbered 0 through d .*

Similar semidefinite characterizations exist for multivariate polynomials to be sums of squares. While the condition of being a sum of squares does not characterize general nonnegative multivariate polynomials exactly, there exist sequences of

sum of squares relaxations which can approximate the set of nonnegative polynomials (on e.g. \mathbb{R}^k , $[-1, 1]^k$, or a more general semialgebraic set) arbitrarily tightly [35]. Furthermore, for some special classes of multivariate polynomials, the sum of squares condition is exact. For example, the condition that a square matrix $M(t)$ whose entries are polynomials be positive semidefinite for $-1 \leq t \leq 1$ amounts to checking whether the multivariate polynomial $x'M(t)x$ is nonnegative for $-1 \leq t \leq 1$ and all $x \in \mathbb{R}^k$ [7].

Lemma A.4 ([7]). *A matrix $M(t)$ whose entries are univariate polynomials in t is positive semidefinite on $[-1, 1]$ if and only if $x'M(t)x = S(x, t) + (1 - t^2)T(x, t)$ where S and T are polynomials which are sums of squares.*

Now suppose we wish to answer the question of whether a finite sequence (μ^0, \dots, μ^k) of reals correspond to the moments of a measure on $[-1, 1]$, i.e. whether there exists a positive measure μ on $[-1, 1]$ such that $\mu^i = \int x^i d\mu(x)$. Clearly if such a measure exists then we must have $\int p(x) d\mu(x) \geq 0$ for any polynomial p of degree at most k which is nonnegative on $[-1, 1]$. But any such integral is a linear combination of the moments (μ^0, \dots, μ^k) by definition and the polynomials p which are nonnegative on $[-1, 1]$ can be characterized with semidefinite constraints using Lemmas A.2 and A.3. Therefore this necessary condition for (μ_0, \dots, μ_k) to be the moments of a measure on $[-1, 1]$ can be written in terms of semidefinite constraints. It turns out that this condition is also sufficient [22], so we have:

Lemma A.5. *The condition that a finite sequence of numbers (μ^0, \dots, μ^k) be the moments of a positive measure on $[-1, 1]$ can be written in terms of linear equations and semidefinite matrix constraints.*

One can formulate similar questions about whether a finite sequence of numbers corresponds to the joint moments $\int x_1^{i_1} \dots x_k^{i_k} d\mu(x)$ of a positive measure μ on $[-1, 1]^k$ (or a more general semialgebraic set). Using a sequence of semidefinite relaxations of the set of nonnegative polynomials on $[-1, 1]^k$, a sequence of necessary conditions for joint moments can be obtained in the same way as the conditions for moments of univariate measures. While no single one of these conditions is sufficient for a list of numbers to be joint moments, these conditions approximate the set of joint moments arbitrarily closely.

References

- [1] T. Başar and G. J. Olsder. *Dynamic Noncooperative Game Theory*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 1999.
- [2] D. P. Bertsekas, A. Nedić, and A. E. Ozdaglar. *Convex Analysis and Optimization*. Athena Scientific, Belmont, MA, 2003.
- [3] D. Bertsimas and J. N. Tsitsiklis. *Introduction to Linear Optimization*. Athena Scientific, Belmont, MA, 1997.
- [4] X. Chen and X. Deng. 3-Nash is PPAD-complete. *Electronic Colloquium on Computational Complexity*, TR05-134, 2005.
- [5] X. Chen and X. Deng. Settling the complexity of two-player Nash equilibrium. In *Proceedings of the 47th annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2006.
- [6] Y. C. Chen, N. Y. Long, and X. Luo. Iterated strict dominance in general games. *Games and Economic Behavior*, to appear.
- [7] M. D. Choi, T. Y. Lam, and B. Reznick. Real zeroes of positive semidefinite forms i. *Mathematische Zeitschrift*, 171:1 – 26, 1980.
- [8] P. Dasgupta and E. Maskin. The existence of equilibrium in discontinuous economic games, I: Theory. *The Review of Economic Studies*, 53(1):1 – 26, January 1986.
- [9] C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou. The complexity of computing a Nash equilibrium. In *Proceedings of the 38th annual ACM Symposium on Theory of Computing (STOC)*, pages 71 – 78, New York, NY, 2006. ACM Press.
- [10] C. Daskalakis and C. H. Papadimitriou. Three-player games are hard. *Electronic Colloquium on Computational Complexity*, TR05-139, 2005.
- [11] M. Dresher and S. Karlin. Solutions of convex games as fixed points. In H. W. Kuhn and A. W. Tucker, editors, *Contributions to the Theory of Games II*, number 28 in Annals of Mathematics Studies, pages 75 – 86. Princeton University Press, Princeton, NJ, 1953.

- [12] M. Dresher, S. Karlin, and L. S. Shapley. Polynomial games. In H. W. Kuhn and A. W. Tucker, editors, *Contributions to the Theory of Games I*, number 24 in Annals of Mathematics Studies, pages 161 – 180. Princeton University Press, Princeton, NJ, 1950.
- [13] R. M. Dudley. *Real Analysis and Probability*. Cambridge University Press, New York, 2002.
- [14] D. Gale. The game of Hex and the Brouwer fixed-point theorem. *The American Mathematical Monthly*, 86(10), December 1979.
- [15] D. Gale and O. Gross. A note on polynomial and separable games. *Pacific Journal of Mathematics*, 8(4):735 – 741, 1958.
- [16] F. Germano and G. Lugosi. Existence of sparsely supported correlated equilibria. *Economic Theory*, To appear.
- [17] I. L. Glicksberg. A further generalization of the Kakutani fixed point theorem, with application to Nash equilibrium points. *Proceedings of the American Mathematical Society*, 3(1):170 – 174, February 1952.
- [18] S. Hart and D. Schmeidler. Existence of correlated equilibria. *Mathematics of Operations Research*, 14(1), February 1989.
- [19] A. Hatcher. *Algebraic Topology*. Cambridge University Press, New York, NY, 2002.
- [20] R. Kannan and T. Theobald. Games of fixed rank: A hierarchy of bimatrix games. In *Proceedings of the 18th annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1124–1132, 2007.
- [21] S. Karlin. *Mathematical Methods and Theory in Games, Programming, and Economics*, volume 2: Theory of Infinite Games. Addison-Wesley, Reading, MA, 1959.
- [22] S. Karlin and L. S. Shapley. *Geometry of Moment Spaces*. American Mathematical Society, Providence, RI, 1953.
- [23] M. Kearns, M. Littman, and S. Singh. Graphical models for game theory. In *Proceedings of the UAI*, 2001.
- [24] H. W. Kuhn and A. W. Tucker, editors. *Contributions to the Theory of Games*. Number 24 in Annals of Mathematics Studies. Princeton University Press, Princeton, NJ, 1950.
- [25] C. E. Lemke and J. T. Howson, Jr. Equilibrium points in bimatrix games. *SIAM Journal on Applied Math*, 12:413 – 423, 1964.

- [26] R. J. Lipton, E. Markakis, and A. Mehta. Playing large games using simple strategies. In *Proceedings of the 4th ACM Conference on Electronic Commerce*, pages 36 – 41, New York, NY, 2003. ACM Press.
- [27] R.D. McKelvey and A. McLennan. Computation of equilibria in finite games. In H. M. Amman, D. A. Kendrick, and J. Rust, editors, *Handbook of Computational Economics*, volume 1, pages 87 – 142. Elsevier, Amsterdam, 1996.
- [28] J. F. Nash. Non-cooperative games. *Annals of Mathematics*, 54(2):286 – 295, September 1951.
- [29] 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, June 1994.
- [30] C. H. Papadimitriou. Computing correlated equilibria in multi-player games. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC)*, New York, NY, 2005. ACM Press.
- [31] C. H. Papadimitriou and T. Roughgarden. Computing equilibria in multiplayer games. In *Proceedings of the 16th annual ACM-SIAM symposium on discrete algorithms (SODA)*, 2005.
- [32] P. A. Parrilo. *Structured Semidefinite Programs and Semialgebraic Geometry Methods in Robustness and Optimization*. PhD thesis, California Institute of Technology, May 2000.
- [33] P. A. Parrilo. Polynomial games and sum of squares optimization. In *Proceedings of the 45th IEEE Conference on Decision and Control (CDC)*, 2006.
- [34] K. R. Parthasarathy. *Probability Measures on Metric Spaces*. Academic Press, New York, NY, 1967.
- [35] B. Reznick. Some concrete aspects of Hilbert’s 17th problem. In C. N. Delzell and J. J. Madden, editors, *Real Algebraic Geometry and Ordered Structures*, pages 251 – 272. American Mathematical Society, 2000.
- [36] R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2:65 – 67, 1973.
- [37] W. Rudin. *Principles of Mathematical Analysis*. McGraw-Hill, New York, 1976.
- [38] W. Rudin. *Real & Complex Analysis*. WCB / McGraw-Hill, New York, 1987.
- [39] W. Rudin. *Functional Analysis*. McGraw-Hill, New York, 1991.
- [40] R. Savani and B. von Stengel. Exponentially many steps for finding a Nash equilibrium in a bimatrix game. In *Proceedings of the 45th annual IEEE Symposium on Foundations of Computer Science*, pages 258 – 267, 2004.

- [41] H. Scarf. The approximation of fixed points of a continuous mapping. *SIAM Journal of Applied Mathematics*, 15:1328 – 1343, 1967.
- [42] L. S. Shapley. A note on the Lemke-Howson algorithm. *Mathematical Programming Study I*, pages 175 – 189, 1974.
- [43] L. K. Simon and W. R. Zame. Discontinuous games and endogenous sharing rules. *Econometrica*, 58(4):861 – 872, July 1990.
- [44] G. Stoltz and G. Lugosi. Learning correlated equilibria in games with compact sets of strategies. *Games and Economic Behavior*, to appear.
- [45] L. Vandenberghe and S. Boyd. Semidefinite programming. *SIAM Review*, 38(1):49 – 95, 1996.
- [46] S. Vavasis. Approximation algorithms for indefinite quadratic programming. *Mathematical Programming*, 57:279 – 311, 1992.
- [47] J. Von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, NJ, 1944.
- [48] B. von Stengel. Computing equilibria for two-person games. In R. J. Aumann and S. Hart, editors, *Handbook of Game Theory with Economic Applications*, volume 3, chapter 45, pages 1723 – 1759. Elsevier, Amsterdam, 2002.