Computing Nash Equilibria in Two Player Strategic Form Games

A Thesis presented

by

Sam Ganzfried

to

the Department of Mathematics

in partial fulfillment of the honors requirements

for the degree of

Bachelor of Arts

Harvard University

Advisor: Avi Pfeffer

April 4, 2005

ABSTRACT

Game theory has been studied by economists and applied mathematicians for almost a century, and recently computer scientists have increasingly been interested in the field as well. This is a very important addition, as algorithmic and complexity issues must be addressed in order to be able to apply classical theoretical results. In particular, the Nash Existence Theorem, while of great theoretical interest, would be of limited practical use without efficient methods for actually finding equilibria. This paper begins by defining basic game theory concepts and presenting proofs of some classical results, including Nash's theorem and the Minimax Theorem. Section 4 begins with a discussion of the linear programming. I prove some preliminary results leading up to the Duality Theorem, and then show that every zero-sum game can be reduced to a pair of dual linear programs. Next I show that any two-player game can be reduced to another type of optimization problem called a linear complementarity problem. Section 5 presents the Lemke-Howson algorithm – which computes a Nash equilibrium in any two-player game – from a geometric perspective. In section 6, I show that the algorithm can also be interpreted algebraically to give a solution to the linear complementarity problem.

1. INTRODUCTION

A game is any social situation involving the interaction of two or more individuals. While tic-tac-toe, chess, and rock-paper-scissors are more traditional examples of games in the everyday usage of the term, games do not need to have such a recreational nature; military strategy, biological competition, and voting are examples of games with more "real-world" applications. Myerson (2004) defines game theory as "the study of mathematical models of conflict and cooperation between intelligent rational decision-makers." Game theory has been successfully applied to every social science discipline, the life sciences, and to many common, everyday situations. In addition, it has been a major area of study for mathematicians, economics, and more recently computer scientists.

Because of the field's widespread applicability and the variety of mathematical and computational issues it encompasses, it is hard to place game theory within any single discipline (although it has traditionally been viewed as a branch of economics). While the field is clearly benefitting from being analyzed from many different perspectives, it is also important to make sure that it doesn't become disorganized as a result. When I started doing research for my thesis, I was surprised at how difficult it was to find a basic introduction to the fundamental mathematical and computational results. I had to turn to game theory textbooks for proofs of classical results, operations research and optimization books for results in linear programming and linear complementarity, and more recent computer science and economics papers for algorithms and complexity results. Thus, the major contribution of this paper is to present the basic mathematical and computational results related to computing Nash equilibria in a coherent form that can benefit people from all fields.

In terms of background, no knowledge of game theory or optimization theory is necessary, and very little specific mathematical knowledge is assumed. However, a general mathematical maturity is necessary and familiarity with linear algebra and real analysis would be helpful.

In addition to my advisor Avi Pfeffer, I would also like to thank Daniel Goroff and David Parkes for taking the time to meet with me and suggesting useful references.

2. GAME THEORY BACKGROUND

In the introduction I mentioned that game theorists generally assume all players are *rational* and *intelligent*. We will make these two assumptions throughout this paper. A decision-maker is *rational* if he makes decisions in pursuit of maximizing his own well-being. Without getting too involved in the technical foundations of decision theory, we assume that each player's well-being can be formally measured by a *utility* scale and that each player's goal is to maximize his expected utility. A player is *intelligent* if he knows everything that we – as outside observers – know about the game and he can make any inferences about the situation that we can make. For example, he is aware of the possible strategies available to all players and the payoffs associated with each outcome.

Definition 2.1. *A **strategic-form** (or **normal-form**) game is any* Γ *of the form*

$$\Gamma = (N, (S_i)_{i\in N}, (u_i)_{i\in N}),$$

where N *is a nonempty set, and for each* i *in* N*,* S_i *is a nonempty set and* u_i *is a function from* $\times_{j\in N} S_j$ *into the set of real numbers* $\mathbb{R}$*.*

N denotes the set of players in the game Γ. For each player $i \in N$, S_i is the set of *strategies* (or *pure-strategies*) available to player i. A strategic-form game is *finite* if the set of players N and all the strategy sets S_i are finite. Let $N = \{1,2,...,I\}$, so that players are denoted by positive integers, and for each $i \in N$ let $S_i = \{s_{i1}, s_{i2}, \ldots, s_{ik_i}\}$. Often, I will slightly abuse this notation by letting s_i be an arbitrary element of S_i. This paper is only concerned with finite strategic-form games. Two player strategic-form games are often depicted using two $m \times n$ matrices A and B, where $m = k_1$, $n = k_2$, $a_{ij} = u_1(s_{1i}, s_{2j})$, and $b_{ij} = u_2(s_{1i}, s_{2j})$. For this reason, two-player strategic-form games are also known as *bimatrix* games.

When dealing with bimatrix games, we will sometimes assume that all entries of the payoff matrices are positive. Informally, we can always make this assumption without loss of generality because we can always add a sufficiently large number to all payoffs such that all payoffs become positive and the new game is fundamentally the same as the old one. In particular, both games have the same Nash equilibria (defined later in this section). The proof of this result is trivial, and I will omit it. Unless otherwise specified, assume all payoff matrices contain positive entries.

A *strategy profile* is a possible combination of strategies that the players in N might choose, where each player i chooses one pure-strategy in S_i. $S = \times_{j \in N} S_j$ denotes the set of all possible strategy profiles. Let $s = (s_1, s_2, \ldots, s_I)$ denote an arbitrary element of S. For any strategy profile $s \in S$, let $u_i(s) = u_i(s_1, \ldots, s_I)$.

I will refer to all players other than a given player i as "player i's opponents" and denote them by "$-i$." This does not mean that the other players are all trying to "beat" player i; they are trying to maximize their individual utility functions, which may or may not coincide with decreasing player i's utility. Let $u_i(s_i, s_{-i})$ denote the utility payoff to player i when he plays strategy $s_i \in S_i$ and his opponents together play $s_{-i} \in \times_{k \in N, k \neq i} S_k$.

In a two-player *zero-sum* game, maximizing one's utility is equivalent to minimizing the other player's utility: that is, $u_2(s) = -u_1(s)$ for all $s \in S$. So in matrix form $B = -A$, and the game is fully specified just by the matrix A. We will assume without loss of generality that all entries of A are positive in a zero-sum game (and therefore all entries of B are negative).

To demonstrate the definition of a strategic-form game, consider the following example of Rock-Paper-Scissors. Since there are two players, $N = \{1,2\}$. Each player has three available strategies: Rock (R), Paper (P), and Scissors (S). So $S_1 = S_2 = \{R, P, S\}$. Assume that the winner gains 1 unit of utility, the loser loses 1 unit of utility, and that utility does not change for either player in a tie. We can imagine that the loser pays the winner one dollar each round, in which case each unit of utility corresponds to a dollar. Then

$$\begin{aligned}
u_1(R,S) = u_1(S,P) = u_1(P,R) = u_2(S,R) = u_2(P,S) = u_2(R,P) = 1 \\
u_1(S,R) = u_1(P,S) = u_1(R,P) = u_2(R,S) = u_2(S,P) = u_2(P,R) = -1 \\
u_1(R,R) = u_1(P,P) = u_1(S,S) = u_2(R,R) = u_2(P,P) = u_2(S,S) = 0.
\end{aligned}$$

Notice that u_i is specified for each strategy profile and for each player. It is clear that Rock-Paper-Scissors is a zero-sum game and is fully specified by the matrix

$$A = \begin{bmatrix} 0 & -1 & 1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \end{bmatrix}.$$

Definition 2.2. *Given a strategic-form game $\Gamma = (N, (S_i)_{i \in N}, (u_i)_{i \in N})$, a* ***mixed strategy*** σ_i *for player i is a probability distribution over pure strategies. That is, it is a map $\sigma_i : S_i \to \mathbb{R}$ such that $\sigma_i(s_i) \geq 0$ for each $s_i \in S_i$ and $\sum_{s_i \in S_i} \sigma_i(s_i) = 1$.*

Pure strategies are a trivial subset of mixed strategies. I will denote the set of player i's possible mixed strategies by Σ_i, and the space of mixed-strategy profiles by $\Sigma = \times_i \Sigma_i$, with an individual element of Σ denoted by $\sigma = (\sigma_1, \ldots, \sigma_I)$. In a two-player game, we can equivalently represent a mixed strategy for player 1 by a vector $x \in \mathbb{R}^m$, where $x_i = \sigma_1(s_{1i})$, and a mixed strategy for player 2 by a vector $y \in \mathbb{R}^n$, where $y_j = \sigma_2(s_{2j})$. In this case, let $X = \Sigma_1$ and $Y = \Sigma_2$ for notational convenience.

The payoffs to a profile of mixed strategies are the expected values of the corresponding pure-strategy payoffs. So player i's payoff to profile σ is

$$\sum_{s \in S} \left(\prod_{j=1}^{I} \sigma_j(s_j) \right) u_i(s),$$

which I will denote by $u_i(\sigma)$. Notice that u_i has been defined in several different ways; however, it will be clear by the number and type of arguments which definition applies in a given context. For any τ_i in Σ_i, let (τ_i, σ_{-i}) denote the mixed-strategy profile in which player i plays τ_i and all other players play the same strategy as in σ.

A very natural question to ask about a game is how it should be played, or what strategies are in some sense "optimal." While there are several reasonable criteria for judging whether a strategy is "optimal," the concept of a Nash equilibrium solution has come to dominate much of game theory literature for several reasons.

Definition 2.3. *A mixed-strategy profile σ^* of Γ is a* ***Nash equilibrium*** *iff $u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(\sigma_i, \sigma_{-i}^*)$ for all $\sigma_i \in \Sigma_i$, for all $i \in N$.*

A Nash equilibrium is a profile of strategies such that each player's strategy is an optimal response to the other players' strategies. That is, no player can profit by deviating unilaterally from his strategy assuming his opponents' strategies remain fixed. If we suppose that an external observer – such as a social planner – publicly specificies a mixed strategy for each player before the game is actually played, then one would expect this profile to be followed if and only if no player could increase his utility by playing a different strategy. Equivalently, we could imagine that the players are allowed to communicate with each other and assign a mixed strategy to each player before playing. Then we expect the strategy profile would actually be played if and only if it is a Nash equilibrium. As with any possible solution concept, there are some obvious drawbacks of the Nash equilibrium. First, it might seem unrealistic to assume that each player "knows" what strategy each of his opponents will play in advance and that he can change his strategy without having any effect on his opponents' strategies. In this sense the the concept of Nash equilibria might seem to require too much.

Additionally, many games have multiple Nash equilibria, and it can be difficult to predict which one will (or should) be played. Most critics have supported the latter objection, and several refinements of the Nash equilibrium have been proposed, such as stable, perfect, and proper equilibria. In certain situations, different equilibrium concepts might have more natural interpretations than others. For those interested, Fudenberg and Myerson provide an analysis of several solution concepts, and McKelvey and von Stengel discuss methods for extending the algorithms discussed in this paper to find various equilibrium refinements. This paper will focus exclusively on techniques for computing Nash equilibria – which has been widely accepted as the "standard" solution concept in game theory. In section 3 we will prove Nash's famous result stating that every finite strategic-form game contains at least one Nash equilibrium. One problem with the various equilibrium refinements is that some of them are not guaranteed to exist in a given game, making them less desirable general solution concepts. In this paper, I will sometimes refer to Nash equilibria just as "equilibria."

The following lemma shows that the highest expected utility that any player can obtain against any combination of other players' mixed strategies does not depend on whether he uses mixed or only pure strategies.

Lemma 2.4. *For any σ in Σ and any player i in N,*

$$\max_{s_i \in S_i} u_i(s_i, \sigma_{-i}) = \max_{\tau_i \in \Sigma_i} u_i(\tau_i, \sigma_{-i}).$$

Proof. Let $A = \max_{s_i \in S_i} u_i(s_i, \sigma_{-i})$ and $B = \max_{\tau_i \in \Sigma_i} u_i(\tau_i, \sigma_{-i})$. It is clear that $B \geq A$, since pure strategies are a subset of mixed strategies. If $B > A$, then there exists $\tau_i \in \Sigma_i$ such that $u_i(\tau_i, \sigma_{-i}) > A$. Suppose $u_i(s_i^*, \sigma_{-i}) = A$, where $s_i^* \in S_i$. Then $u_i(s_i, \sigma_{-i}) \leq u_i(s_i^*, \sigma_{-i})$ for all $s_i \in S_i$. So

$$u_i(\tau_i, \sigma_{-i}) = \sum_{s_i \in S_i} \tau_i(s_i) u_i(s_i, \sigma_{-i}) \leq \sum_{s_i \in S_i} \tau_i(s_i) u_i(s_i^*, \sigma_{-i}).$$

$$= u_i(s_i^*, \sigma_{-i}) \sum_{s_i \in S_i} \tau_i(s_i) = u_i(s_i^*, \sigma_{-i}).$$

This contradicts the fact that $u_i(\tau_i, \sigma_{-i}) > A$. So $A = B$. $\square$

The following theorem shows that the optimal mixed strategies for each player are the strategies that assign positive probability only to his optimal pure strategies. Thus, it provides an equivalent definition of Nash equilibria which turns out to be quite useful.

Theorem 2.5. *The mixed strategy σ is a Nash equilibrium of Γ if and only if $\sigma_i(s_i) > 0$ implies $s_i \in$ $argmax_{c_i \in S_i} u_i(c_i, \sigma_{-i})$, for all $i \in N$ and $s_i \in S_i$.*

Proof. Suppose σ is a Nash equilibrium of Γ, and let $A = argmax_{c_i \in S_i} u_i(c_i, \sigma_{-i})$. Suppose there exists a player i and a pure strategy $s_i \in S_i$ such that $\sigma_i(s_i) > 0$ and $s_i \notin A$. Suppose $d_i \in A$. Then $u_i(d_i, \sigma_{-i}) > u_i(s_i, \sigma_{-i})$. Define $\rho_i \in \Sigma_i$ as follows: $\rho_i(c_i) = \sigma_i(c_i)$ if $c_i \neq s_i$ or d_i, $\rho_i(s_i) = 0$, and $\rho_i(d_i) = \sigma_i(d_i) + \sigma_i(s_i)$. So ρ_i is the same strategy as σ_i except with all of the weight given to s_i shifted to d_i. Then

$$u_i(\rho_i, \sigma_{-i}) = u_i(\sigma_i, \sigma_{-i}) - \sigma_i(s_i) u_i(s_i, \sigma_{-i}) + \sigma_i(s_i) u_i(d_i, \sigma_{-i}) > u_i(\sigma_i, \sigma_{-i}).$$

So σ_i is not a best response of player i to σ_{-i}. This contradicts the fact that σ is a Nash equilibrium, and we have a contradiction.

Conversely, pick $\sigma \in \Sigma$ and suppose that for all $i \in N$ and $s_i \in S_i$, $\sigma_i(s_i) > 0$ implies $s_i \in A$. If σ is not a Nash equilibrium, then by the previous lemma for any $s_i \in A$ we have $\bar{u} = u_i(s_i, \sigma_{-i}) >$

$u_i(\sigma_i, \sigma_{-i})$. By assumption,

$$u_i(\sigma_i, \sigma_{-i}) = \sum_{s_i \in S_i} \sigma_i(s_i) u_i(s_i, \sigma_{-i}) = \sum_{s_i \in A} \sigma_i(s_i) u_i(s_i, \sigma_{-i}) = \bar{u}.$$

So $\bar{u} > \bar{u}$, which is a contradiction. So σ is a Nash equilibrium. $\square$

The *support* of a mixed strategy σ_i is the set of pure strategies $s_i \in S_i$ for which $\sigma_i(s_i) > 0$. The following example illustrates how Theorem 1.5 can be helpful in computing equilibria.

Example 2.6. *The unique equilibrium of rock-paper-scissors is* $((\frac{1}{3}, \frac{1}{3}, \frac{1}{3}), (\frac{1}{3}, \frac{1}{3}, \frac{1}{3}))$.

Proof. First we must show that this profile is a Nash equilibrium. Suppose player 1 could profit by deviating to the profile $\sigma_1 = \alpha_1 R + \beta_1 P + \gamma_1 S$. Player 1's new payoff would be

$$\frac{\beta_1 - \gamma_1}{3} + \frac{\gamma_1 - \alpha_1}{3} + \frac{\alpha_1 - \beta_1}{3} = 0.$$

So any deviation will produce the same payoff. Similarly, player 2 cannot profitably deviate from his strategy. So the given profile is an equilibrium.

Now suppose there is another equilibrium in which each player plays the strategy $\sigma_i = \alpha_i R + \beta_i P + \gamma_i S$. Suppose $\beta_1 = \gamma_1 = 0$. Then if $\beta_2 < 1$ it is clear that player 2 can do better by playing P. But if player 2 plays P, player 1 can do better by playing S than R. Therefore, there is no Nash equilibrium in which player 1 plays R with probability 1. Similar logic shows that in a Nash equilibrium both player's supports must contain more than one pure strategy.

Now suppose there is an equilibrium in which player 1's support contains two strategies: without loss of generality assume $\gamma_1 = 0$, and $\alpha_1, \beta_1 > 0$. By Theorem 1.5, it follows that $u_1(R, \sigma_2) = u_1(P, \sigma_2) >= u_1(S, \sigma_2)$. This implies $\gamma_2 - \beta_2 = \alpha_2 - \gamma_2 \geq \beta_2 - \alpha_2$. The first two equalities imply that $\alpha_2 + \beta_2 = 2\gamma_2$. Since $\alpha_2 + \beta_2 + \gamma_2 = 1$, it follows that $\gamma_2 = 1/3$. Similarly, it follows that $\alpha_2 \geq 1/3$ and $\beta_2 \geq 1/3$, and therefore that $\alpha_2 = \beta_2 = \gamma_2 = 1/3$. So player 2's expected payoff is $-\frac{\beta_1}{3} + \frac{\alpha_1}{3} - \frac{\alpha_1}{3} + \frac{\beta_1}{3} = 0$. However, if player 2 instead played the pure strategy P, then his expected payoff would be $\alpha_1 > 0$. Similar logic shows that both player's supports must contain all three strategies.

By Theorem 1.5, $u_1(R, \sigma_2) = u_1(P, \sigma_2) = u_1(S, \sigma_2)$, which implies $\alpha_2 = \beta_2 = \gamma_2 = 1/3$ from the above analysis. Similarly we have $u_2(R, \sigma_1) = u_2(P, \sigma_1) = u_2(S, \sigma_1)$ and therefore $\alpha_1 = \beta_1 = \gamma_1 = 1/3$. So the Nash equilibrium is unique. $\square$

3. Nash Existence Theorem

In this section we prove the general existence theorem of Nash (1950):

Theorem 3.1. *Every finite strategic-form game has a mixed-strategy equilibrium.*

Notice that the theorem refers to mixed strategies, and every finite game does not need to contain a pure-strategy Nash equilibrium (as the rock-paper-scissors example demonstrated). Most of the definitions and results in this section are based on section 3.12 of Myerson (2004).

For any finite set M, let $\mathbb{R}^M$ denote the set of all vectors of the form $(x_m)_{m \in M}$ such that $x_m \in \mathbb{R}$ for each m in M. We can equivalently define of $\mathbb{R}^M$ to be the set of all functions from M into the set of real numbers $\mathbb{R}$; in this case I'll write the m-component of $x \in \mathbb{R}^M$ as $x(m)$ instead of x_m. It is clear that $\mathbb{R}^M$ is a finite-dimensional vector space.

Let T be a subset of $\mathbb{R}^M$. T is *convex* iff $\lambda x + (1-\lambda)y \in T$ for all vectors $x, y \in T$ and every $\lambda \in \mathbb{R}$ such that $0 \le \lambda \le 1$. T is *closed* iff $\lim_{j\to\infty} x(j) \in T$ for every convergent sequence of vectors $(x(j))_{j=1}^{\infty}$ such that $x(j) \in T$ for every j. T is *bounded* iff there exists some positive real number K such that $\sum_{m\in M} |x_m| \le K$ for every vector $x \in T$.

A point-to-set *correspondence* $G : X \to 2^Y$ is a mapping that sends each point $x \in X$ to a set $G(x) \subset Y$. Suppose that X and Y are normed linear spaces, so that the concepts of convergence and limits are defined for sequences in X and Y. A correspondence $G : X \to Y$ is *upper-hemicontinuous* if and only if, for every sequence $(x(j), y(j))_{j=1}^{\infty}$, if $x(j) \in X$ and $y(j) \in G(x(j))$ for every j, and the sequence $(x(j))_{j=1}^{\infty}$ converges to some point $\bar{x}$, and the sequence $(y(j))_{j=1}^{\infty}$ converges to some point $\bar{y}$, then $\bar{y} \in G(\bar{x})$. Thus $G : X \to Y$ is upper-hemicontinuous iff the set $\{(x,y)|x \in X, y \in G(x)\}$ is a closed subset of $X \times Y$. In particular, if $g : X \to Y$ is a continuous function from X to Y and $G(x) = \{g(x)\}$ for every x in X, then $G : X \to Y$ is an upper-hemicontinuous point-to-set correspondence. So upper-hemicontinuous correspondences can be viewed as a generalization of continuous functions.

A *fixed point* of a point-to-set correspondence $F : T \to T$ is any $x \in T$ such that $x \in F(x)$. I will now state the Kakutani fixed point theorem, which is central to the proof of Nash's theorem. A proof can be found in Scarf (1973).

Theorem 3.2. ***Kakutani Fixed-Point Theorem****. Let T be a nonempty, convex, bounded, and closed subset of a finite-dimensional vector space. Let $F : T \to T$ be an upper-hemicontinuous point-to-set correspondence such that, for $F(x)$ is a nonempty convex subset of T for each $x \in T$. Then F has a fixed point.*

To see how the various assumptions in the theorem come into play, consider the following example. Let

$$T = [0,1] = \{x \in \mathbb{R} | 0 \le x \le 1\},$$

and let $F_1 : T \to T$ be defined as

$$\begin{aligned} F_1(x) &= \{1\} \quad \text{if} \quad 0 \le x \le 0.5, \\ &= \{0\} \quad \text{if} \quad 0.5 < x \le 1. \end{aligned}$$

Then F_1 has no fixed points, and it satisfies all of the assumptions of the Kakutani fixed-point theorem except for upper-hemicontinuity. Specifically, the set $T' = \{(x,y)|x \in S, y \in F_1(x)\}$ is not closed at $(0.5, 0)$ since

$$\lim_{x\to 0.5^+} (x, F_1(x)) = (0.5, 0),$$

but $(0.5, 0) \notin T'$. To satisfy upper-hemicontinuity, we must extend this correspondence to $F_2 : T \to T$, where

$$\begin{aligned} F_2(x) &= \{1\} \quad \text{if} \quad 0 \le x < 0.5, \\ &= \{0,1\} \quad \text{if} \quad x = 0.5, \\ &= \{0\} \quad \text{if} \quad 0.5 < x \le 1. \end{aligned}$$

F_2 now satisfies all of the assumptions of the Kakutani fixed-point theorem except convex-valuedness, because $F_2(0.5)$ is not a convex set. In particular, $0,1 \in F_2(0.5)$, but if we set $\lambda = 0.5$ then $\lambda * 0 + \lambda * 1 = 0.5 \notin F_2(0.5)$. To satisfy convex-valuedness, we must extend the correspondence to $F_3 : T \to T$, where

$$\begin{aligned} F_3(x) &= \{1\} && \text{if} \quad 0 \leq x < 0.5, \\ &= [0,1] && \text{if} \quad x = 0.5, \\ &= \{0\} && \text{if} \quad 0.5 < x \leq 1. \end{aligned}$$

F_3 now satisfies all the assumptions of the Kakutani fixed-point theorem and has 0.5 as a fixed point, since $0.5 \in F_3(0.5)$.

We can now prove the Nash existence theorem.

Proof. Let $\Gamma = (N, (S_i)_{i \in N}, (u_i)_{i \in N})$. Consider the set $\Sigma = \times_{i \in N} \Sigma_i$ of all mixed-strategy profiles. Σ is clearly nonempty since S_i is nonempty for each player i. Suppose σ and τ are elements of Σ and let λ be a real number in $[0,1]$. Then $\pi = \lambda\sigma + (1-\lambda)\tau$ corresponds to the profile in which each player i plays the mixed strategy $\pi_i = \lambda\sigma_i + (1-\lambda)\tau_i$. For each player i, we know that $\sum_{s_i \in S_i} \sigma_i(s_i) = 1$ and $\sum_{s_i \in S_i} \tau_i(s_i) = 1$. It follows that

$$\sum_{s_i \in S_i} (\lambda\sigma_i(s_i) + (1-\lambda)\tau_i(s_i)) = \lambda + (1-\lambda) = 1.$$

Since $\sigma_i(s_i)$ and $\tau_i(s_i)$ are nonnegative for every $i \in N$ and $s_i \in S_i$, it follows that $\pi_i(s_i) \geq 0$ for each $i \in N$ and $s_i \in S_i$. Therefore, $\lambda\sigma + (1-\lambda)\tau \in \Sigma$, and Σ is convex. Σ is a subset of the finite-dimensional vector space $\mathbb{R}^M$, where $M = S$. Σ is bounded since

$$\sum_{m \in M} |\sigma(m)| = \sum_{i \in N, s_i \in S_i} \sigma_i(s_i) = I.$$

To see that Σ is a closed subset of $\mathbb{R}^M$, let $(\sigma^j)_{j=1}^{\infty}$ be a convergent sequence of vectors such that $\sigma^j \in \Sigma$ for every j. This sequence converges to some limit $\rho \in \mathbb{R}^M$. Suppose $\rho \notin \Sigma$. Then either $\rho_i(s_i) < 0$ for some $i \in N$ and $s_i \in S_i$, or $\sum\limits_{s_i \in S_i} \rho_i(s_i) \neq 1$ for some $i \in N$. Suppose we are using the 1-norm in $\mathbb{R}^M$, where $||\sigma|| = \sum\limits_{i \in N, s_i \in S_i} |\sigma_i(s_i)|$. In the first case, let $k = \rho_i(s_i) < 0$ for some $i \in N$. Then

$$||\rho - \sigma^j|| \geq |\sigma_i^j(s_i) - \rho_i(s_i)| \geq -k$$

for each σ^j in our sequence, and the sequence cannot converge to ρ. In the second case, let $c = \sum\limits_{s_i \in S_i} \rho_i(s_i) - 1 \neq 0$. If $c > 0$ then

$$||\rho - \sigma^j|| \geq \sum_{s_i \in S_i} |\rho_i(s_i) - \sigma_i^j(s_i)| \geq \sum_{s_i \in S_i} |\rho_i(s_i)| - \sum_{s_i \in S_i} |\sigma_i^j(s_i)| = c.$$

Similar logic shows that $||\rho - \sigma^j|| \geq -c$ if $c < 0$, and in either case we have a contradiction. So ρ must lie in Σ, and it follows that Σ is closed.

For any $\sigma \in \Sigma$ and any player j in N, let

$$R_j(\sigma_{-j}) = argmax_{\tau_j \in \Sigma_i} u_j(\tau_j, \sigma_{-j}).$$

Then $R_j(\sigma_{-j})$ is the set of player j's best responses in Σ_j to the combination σ_{-j} of mixed-strategies of his opponents. By Theorem 2.5, $R_j(\sigma_{-j})$ is the set of all $\rho_j \in \Sigma_j$ such that

$$\rho_j(s_j) = 0 \text{ for every } s_j \text{ such that } s_j \notin \mathit{argmax}_{d_j \in S_j} u_j(d_j, \sigma_{-j}).$$

$R_j(\sigma_{-j})$ is nonempty since it includes every pure strategy s_j in $\mathit{argmax}_{d_j \in S_j} u_j(d_j, \sigma_{-j})$, which is nonempty. Suppose ρ_j and τ_j are elements of $R_j(\sigma_{-j})$, and let λ be a real number in [0,1]. Let

$$\pi_j = \lambda\rho_j + (1-\lambda)\tau_j.$$

By previous analysis $\pi_j \in \Sigma_j$. For every $s_j \in S_j$ such that $s_j \notin \mathit{argmax}_{d_j \in S_j} u_j(d_j, \sigma_{-j})$, we know that $\rho_j(s_j) = \tau_j(s_j) = 0$. It follows that $\pi_j(s_j) = 0$ also, and therefore that $\pi_j \in R_j(\sigma_{-j})$. So $R_j(\sigma_{-j})$ is convex.

Let $R : \Sigma \to \Sigma$ be the point-to-set correspondence such that

$$R(\sigma) = \times_{j \in N} R_j(\sigma_{-j}) \text{ for each } \sigma \in \Sigma.$$

Then $R(\sigma) = \{\tau \in \Sigma | \tau_j \in R_j(\sigma_{-j}) \text{ for every } j \in N\}$. For each $\sigma \in \Sigma$, $R(\sigma)$ is nonempty and convex, because it is the Cartesian product of nonempty convex sets.

To show that R is upper-hemicontinuous, suppose that $(\sigma^k)_{k=1}^{\infty}$ and $(\tau^k)_{k=1}^{\infty}$ are convergent sequences where $\sigma^k \in \Sigma$ and $\tau^k \in R(\sigma^k)$ for all k, $\overline{\sigma} = \lim\limits_{k\to\infty} \sigma^k$, and $\overline{\tau} = \lim\limits_{k\to\infty} \tau^k$. These conditions imply that, for every player $j \in N$ and every $\rho_j \in \Sigma_j$,

$$u_j(\tau_j^k, \sigma_{-j}^k) \geq u_j(\rho_j, \sigma_{-j}^k) \text{ for all } k.$$

Taking the limit on the left yields

$$\begin{aligned} \lim_{k\to\infty} u_j(\tau_j^k, \sigma_{-j}^k) &= \lim_{k\to\infty} \sum_{s_j \in S_j} \tau_j^k(s_j) u_j(s_j, \sigma_{-j}^k) \\ &= \lim_{k\to\infty} \sum_{s_j \in S_j, s_{-j} \in S_{-j}} \tau_j^k(s_j)\sigma_{-j}^k(s_{-j}) u_j(s_j, s_{-j}) \\ &= u_j(\overline{\tau}_j, \overline{\sigma}_{-j}). \end{aligned}$$

Similarly, $\lim_{k\to\infty} u_j(\rho_j, \sigma_{-j}^k) = u_j(\rho_j, \overline{\sigma}_{-j})$. It follows that for every j in N and ρ_j in Σ_j,

$$u_j(\overline{\tau}_j, \overline{\sigma}_{-j}) \geq u_j(\rho_j, \overline{\sigma}_{-j}).$$

So $\overline{\tau}_{-j} \in R_j(\overline{\sigma}_{-j})$ for every j in N, and therefore $\overline{\tau} \in R(\overline{\sigma})$. Thus, $R : \Sigma \to \Sigma$ is an upper-hemicontinuous correspondence.

By the Kakutani fixed-point theorem, there exists $\sigma \in \Sigma$ such that $\sigma \in R(\sigma)$. So $\sigma_j \in R_j(\sigma_{-j})$ for every j in N, and hence σ is a Nash equilibrium of Γ. □

This theorem has an important consequence in two-player zero-sum games, known as the *Minimax Theorem*. Consider the zero-sum game defined by the following payoff matrix for player 1:

	H	T
H	-2	-4
T	3	5

It can easily be calculated that the unique Nash equilibrium is (T,H). The same conclusion could also be obtained by the following intuition. First note that the minimum payoff in row 1 is -4, while the minimum payoff in row 2 is 3. These payoffs represent the smallest payoff that player 1 can guarantee himself from each of his two available pure strategies. Knowing that player 2's goal is to minimize his utility, it seems logical for player 1 to prefer the strategy that achieves the maximum of these two minimum payoffs; choosing pure strategy T will guarantee that he will obtain at least 3, while choosing H will only guarantee him -4. Similarly, player 2 observes that the maximum payoff (to player 1) in column 1 is 3, while the maximum payoff in column 2 is 5. These represent the largest payoffs player 2 can guarantee player 1 will receive from each of his pure strategies. By similar reasoning, he prefers the minimum of these two payoffs; choosing pure strategy H will prevent player 1 from getting any more than this minimum amount. Thus, we will expect that (T,H) will be played, since the payoff is simultaneously a maximum of row minimums as well as a minimum of column maximums. We will call a strategy for player 1 that maximizes his expected minimum payoff a *maximin strategy* for player 1 (here T), and a strategy for player 2 that minimizes player 1's expected maximal payoff a *minimax strategy* for player 2 (here H).

Theorem 3.3. $\sigma = (\sigma_1, \sigma_2)$ *is a Nash equilibrium of* Γ *if and only if both of the following conditions hold:*

$$\sigma_1 \in argmax_{\tau_1 \in \Sigma_1} \min_{\tau_2 \in \Sigma_2} u_1(\tau_1, \tau_2), \tag{3.1}$$

$$\sigma_2 \in argmin_{\tau_2 \in \Sigma_2} \max_{\tau_1 \in \Sigma_1} u_1(\tau_1, \tau_2). \tag{3.2}$$

Furthermore, if σ *is an equilibrium then*

$$u_1(\sigma) = \max_{\tau_1 \in \Sigma_1} \min_{\tau_2 \in \Sigma_2} u_1(\tau_1, \tau_2) = \min_{\tau_2 \in \Sigma_2} \max_{\tau_1 \in \Sigma_1} u_1(\tau_1, \tau_2). \tag{3.3}$$

Proof. Suppose σ is a Nash equilibrium. Then

$$u_1(\sigma) = \max_{\tau_1 \in \Sigma_1} u_1(\tau_1, \sigma_2) \geq \max_{\tau_1 \in \Sigma_1} \min_{\tau_2 \in \Sigma_2} u_1(\tau_1, \tau_2) \geq \min_{\tau_2 \in \Sigma_2} u_1(\sigma_1, \tau_2) = u_1(\sigma),$$

where the final equality follows from the fact that player 2's utility is maximized when player 1's utility is minimized. Similarly, we have

$$u_1(\sigma) = \min_{\tau_2 \in \Sigma_2} u_1(\sigma_1, \tau_2) \leq \min_{\tau_2 \in \Sigma_2} \max_{\tau_1 \in \Sigma_1} u_1(\tau_1, \tau_2) \leq \max_{\tau_1 \in \Sigma_1} u_1(\tau_1, \sigma_2) = u_1(\sigma).$$

So all of the expressions are equal, which shows (3.3). (3.1) and (3.2) follow from the fact that

$$\min_{\tau_2 \in \Sigma_2} u_1(\sigma_1, \tau_2) = \max_{\tau_1 \in \Sigma_1} \min_{\tau_2 \in \Sigma_2} u_1(\tau_1, \tau_2) \tag{3.4}$$

and

$$\max_{\tau_1 \in \Sigma_1} u_1(\tau_1, \sigma_2) = \min_{\tau_2 \in \Sigma_2} \max_{\tau_1 \in \Sigma_1} u_1(\tau_1, \tau_2). \tag{3.5}$$

Conversely, suppose that σ_1 and σ_2 satisfy (3.1) and (3.2). (3.1) implies that (3.4) holds and (3.2) implies that (3.5) holds. Since Γ has an equilibrium by Theorem 3.1, the final equality in (3.3) still holds. So all four expressions in (3.4) and (3.5) are equal. So

$$u_1(\sigma) \geq \min_{\tau_2 \in \Sigma_2} u_1(\sigma_1, \tau_2) = \max_{\tau_1 \in \Sigma_1} u_1(\tau_1, \sigma_2) \geq u_1(\sigma).$$

It follows that all of these expressions are equal, and σ is a Nash equilibrium of Γ. $\square$

Notice that the assumption that an equilibrium exists was necessary to prove the converse.

4. Linear Programming and Complementarity

4.1. **Linear Programming.** From now on our goal will be to find a single Nash equilibrium of a given two-player game: we will refer to this as "solving" the game. In this section we will present two types of problems from mathematical programming called the linear programming problem (LPP) and the linear complementarity problem (LCP). We will then show that the problem of solving a zero-sum game is equivalent to solving a LPP and solving a general two-player game is equivalent to solving a LCP; here *equivalence* means that there are mappings between the solution sets of the two problems (solving one problem allows us to easliy solve the other, and vica versa).

Linear programming deals with finding a vector in a real vector space that maximizes (or minimizes) a given linear function subject to a set of linear constraints. We will call a LPP *standard* if all of the constraints are inequalities, and *general* if they include both inequalities and equalities. It turns out that both of these forms are equivalent, where equivalence here means that the programs have equal values (defined below). One direction is trivial, since standard problems are a subset of general problems; the first major result of this section will be to show that a general LPP can be transformed into an equivalent standard LPP. This transformation is important because solving zero-sum games can naturally be formulated as solving a general LPP, but it is easier to prove results about standard LPP's than general LPP's.

Assume that all vectors in this section are column vectors. Equalities and inequalities between two vectors apply to all components of the given vectors. That is, if x and y belong to $\mathbb{R}^m$ then $x = y$ means $x_i = y_i$ for $1 \le i \le m$. The vector $\mathbf{0}$ denotes the vector of zeroes of appropriate dimension, and $\mathbf{1}_m$ denotes the m-vector consisting of all ones.

Definition 4.1. *Suppose $A \in \mathbb{R}^{m\times n}$, $b \in \mathbb{R}^n$, and $c \in \mathbb{R}^m$. A* ***standard linear programming problem*** *is that of finding a nonnegative vector $x = (x_i) \in \mathbb{R}^m$ that either maximizes or minimizes the given linear function*

$$\sum_{i=1}^{m} x_i c_i \tag{4.1}$$

subject to the inequalities

$$\sum_{i=1}^{m} x_i a_{ij} \le b_j, \ \ j = 1, \ldots, n. \tag{4.2}$$

The vector $x \in \mathbb{R}^m$ is *feasible* if it satisfies (4.2), and a linear program is *feasible* if it has a feasible solution. A feasible solution which maximizes (or minimizes) (4.1) is called an *optimal* solution. The optimal value of the function (4.1) is called the *value* of the linear program.

Assuming that the preceding LPP is a maximization problem, define the following minimization problem to be its *dual*: to find a nonnegative vector $y = (y_j) \in \mathbb{R}^n$ that minimizes

$$\sum_{j=1}^{n} y_j b_j$$

subject to the inequalities

$$\sum_{j=1}^{n} y_j a_{ij} \geq c_i,\ i = 1, \ldots, m. \tag{4.3}$$

Now let $M = \{1, \ldots, m\}$ and $N = \{1, \ldots, n\}$. Let I be a subset of M and let $I' = M - I$. Similarly, let J be a subset of N and let $J' = N - J$. Then the *general maximum LPP* is to find a vector $x \in \mathbb{R}^m$ such that

$$\sum_{i=1}^{m} x_i c_i \text{ is maximized}$$

subject to

$$x_i \geq 0, i \in I,$$
$$\sum_{i=1}^{m} a_{ij} x_i \leq b_j, j \in J,$$
$$\sum_{i=1}^{m} a_{ij} x_i = b_j, j \in J'.$$

This problem has the following dual: to find a vector $y \in \mathbb{R}^n$ such that

$$\sum_{j=1}^{n} y_j b_j \text{ is minimized}$$

subject to

$$y_j \geq 0, j \in J$$
$$\sum_{j=1}^{n} a_{ij} y_j \leq c_i, i \in I,$$
$$\sum_{j=1}^{n} a_{ij} y_j = c_i, i \in I'.$$

Lemma 4.2. *Any general maximum problem can be transformed to a standard maximum problem that has the same solutions.*

Proof. Consider the general problem of finding an m-vector x such that

$$x^T c \text{ is a maximum} \tag{4.4}$$

subject to

$$x^T a^j \leq b_j \text{ for } \ j \in J, \tag{4.5}$$
$$x^T a^j = b_j \text{ for } \ j \in J', \tag{4.6}$$

where a^j denotes the j'th column of matrix A.

We first obtain constraints involving only inequalities by replacing the equations in (4.6) by the inequalities

$$x^T a^j \leq b_j, \tag{4.7}$$
$$-x^T a^j \leq -b_j. \tag{4.8}$$

It is clear that $x^T a^j = b_j$ if and only if both of the above inequalities hold.

Next, introduce new unknown nonnegative m-vectors $x' = (x'_i)$ and $x'' = (x''_i)$ and replace inequalities (4.5), (4.7), and (4.8) with

$$
\begin{aligned}
(x' - x'')^T a^j &\le b_j \quad \text{for } j \in J, \\
(x' - x'')^T a^j &\le b_j \quad \text{for } j \in J', \\
-(x' - x'')^T a^j &\le -b_j \quad \text{for } j \in J',
\end{aligned}
\tag{4.9}
$$

requiring that

$$(x' - x'')^T c \text{ be a maximum.} \tag{4.10}$$

Then it is clear that the vector $z = (x', x'')$ maximizes (4.10) subject to (4.9) if and only if the vector $x = x' - x''$ solves the original problem, and that these solutions have the same values. So we have transformed the general problem to an equivalent standard problem. □

I will now prove some results about standard LPP's, culminating in the Duality Theorem.

Lemma 4.3. *Let x be a feasible solution of a standard maximization problem and let y be a feasible solution of the dual problem. Then*

$$\sum_{i=1}^{m} x_i c_i \le \sum_{j=1}^{n} y_j b_j.$$

Proof. Multiplying the j-th inequality of (4.2) by y_j and summing over j gives

$$\sum_{j=1}^{n} y_j b_j \ge \sum_{j=1}^{n} (y_j \sum_{i=1}^{m} x_i a_{ij}) = \sum_{i=1}^{m} \sum_{j=1}^{n} x_i y_j a_{ij}.$$

Similarly, multiplying the i-th inequality of (4.3) by x_i and summing over i gives

$$\sum_{i=1}^{m} x_i c_i \le \sum_{i=1}^{m} (x_i \sum_{j=1}^{n} y_j a_{ij}) = \sum_{i=1}^{m} \sum_{j=1}^{n} x_i y_j a_{ij}.$$

Combining these two equations gives the desired result. □

Theorem 4.4. *If there exist feasible solutions x and y for the standard maximum problem and its dual such that*

$$\sum_{i=1}^{m} x_i c_i = \sum_{j=1}^{n} y_j b_j, \tag{4.11}$$

then the solutions are optimal.

Proof. Suppose $x' = (x'_i)$ is another feasible solution of the standard maximization problem. Then from the previous lemma we have

$$\sum_{i=1}^{m} x'_i c_i \le \sum_{j=1}^{n} y_j b_j.$$

Combining this with equation (4.11) gives

$$\sum_{i=1}^{m} x_i' c_i \leq \sum_{i=1}^{m} x_i c_i.$$

So x is an optimal solution. An identical argument proves the optimality of y. □

The following theorem follows from some basic linear algebra results. A proof can be found in Gale (1989), and I will omit it:

Theorem 4.5. *Let $A \in \mathbb{R}^{m\times n}, b \in \mathbb{R}^n$, and let $x \in \mathbb{R}^m, y \in \mathbb{R}^n$ be vectors of unknowns. Then exactly one of the following alternatives holds. Either the equation*

$$x^T A \leq b^T \tag{4.12}$$

has a nonnegative solution, or the inequalities

$$Ay \geq \boldsymbol{0} \qquad b^T y < \boldsymbol{0} \tag{4.13}$$

have a nonnegative solution.

4.2. **The Duality Theorem.**

Theorem 4.6. ***Duality Theorem.*** *If both a standard LPP and its dual are feasible, then both have optimal vectors and the values of the two programs are the same.*

Proof. Suppose that a standard maximum problem and its dual are feasible. This means that we have nonnegative solutions x and y to the inequalities

$$x^T A \leq b^T \tag{4.14}$$

$$Ay \geq c. \tag{4.15}$$

By Lemma 4.3, we know that if x and y satisfy these two inequalities then they also satisfy $x^T c \leq y^T b$. So if we can find a solution (x,y) of (4.14) and (4.15) that also satisfies

$$x^T c - y^T b \geq \mathbf{0}, \tag{4.16}$$

then it follows that $x^T c = y^T b$, and the solutions are optimal by Theorem 4.4.

So in order to derive a contradiction, suppose that the system (4.14), (4.15), and (4.16) has no nonnegative solution. We will now manipulate this system by defining a matrix A' and new vectors x', b' so that we can apply Theorem 4.5. Writing out the three inequalities componentwise yields:

$$\sum_{i=1}^{m} x_i a_{ij} \leq b_j \qquad \text{for } j = 1, \ldots, n, \tag{4.17}$$

$$\sum_{j=1}^{n} y_j(-a_{ij}) \leq -c_i \qquad \text{for } i = 1, \ldots, m, \tag{4.18}$$

$$\sum_{j=1}^{n} y_j b_j - \sum_{i=1}^{m} x_i c_i \leq 0. \tag{4.19}$$

Now construct the $(m+n+1)\times(m+n+1)$ matrix A' as follows:

$$\begin{aligned}
&a'_{ij} = a_{ij} \text{ for } 1 \le i \le m \text{ and } 1 \le j \le n,\\
&a'_{ij} = -a_{ji} \text{ for } m+1 \le i \le m+n \text{ and } n+1 \le j \le n+m,\\
&a'_{m+n+1,m+n+1} = 1,\\
&a'_{ij} = 0 \text{ otherwise.}
\end{aligned}$$

Let x' be the $(m+n+1)$-vector where

$$\begin{aligned}
&x'_i = x_i \text{ for } 1 \le i \le m,\\
&x'_i = y_{i-m} \text{ for } m+1 \le i \le m+n,\\
&x'_{m+n+1,m+n+1} = \sum_{j=1}^{n} y_j b_j - \sum_{i=1}^{m} x_i c_i.
\end{aligned}$$

Finally, let b' be the $(m+n+1)$-vector defined by

$$\begin{aligned}
&b'_i = b_i \text{ for } 1 \le i \le n,\\
&b'_i = -c_{i-n} \text{ for } n+1 \le i \le m+n,\\
&b'_{m+n+1,m+n+1} = 0.
\end{aligned}$$

It is clear that the system of inequalities (4.17), (4.18), and (4.19) is equivalent to the statement $x'^T A' \le b'^T$. If this inequality has no nonnegative solutions x', then Theorem 4.5 says that $A'y' \ge \mathbf{0}$ and $b'^T y' < \mathbf{0}$ are both satisfied by a nonnegative vector $y' \in \mathbb{R}^{(m+n+1)}$.

Now let z be the n-vector where $z_i = y'_i$, and let w be the m-vector such that $w_i = y'_{i+n}$. Then the following inequalities hold:

$$\sum_{j=1}^{n} z_j a_{ij} \ge 0 \qquad \text{for } i = 1,\ldots,m,$$

$$-\sum_{i=1}^{m} w_i a_{ij} \ge 0 \qquad \text{for } j = 1,\ldots,n,$$

$$\sum_{j=1}^{n} z_j b_j - \sum_{i=1}^{m} w_i c_i < 0.$$

Now let

$$m_1 = \sum_{j=1}^{n} \frac{z_j a_{ij}}{\overline{c}} \text{ where } \overline{c} = \max\{c_i : c_i > 0\}$$

if at least one of the c_i's is positive, and let $m_1 = 1$ otherwise. Also, let

$$m_2 = \sum_{i=1}^{m} \frac{w_i a_{ij}}{\overline{b}} \text{ where } \overline{b} = \min\{b_i : b_i < 0\}$$

if at least one of the b_j's is negative, and let $m_2 = 1$ otherwise.

Now define $\theta = \min\{m_1, m_2\}$. Then θ is clearly nonnegative, and the following inequalities hold:

$$Az \ge \theta c, \tag{4.20}$$

$$w^T A \leq \theta b^T, \tag{4.21}$$

$$z^T b - w^T c < 0. \tag{4.22}$$

If $\theta = 0$, then combining equations (4.14), (4.15), (4.20), and (4.21) yields $0 \leq x^T Az \leq z^T b$, and $0 \geq w^T Ay \geq w^T c$. This implies that $z^T b \geq w^T c$, which contradicts (4.22).

So θ must be positive and w/θ and z/θ are a pair of feasible solutions to the LPP of (4.14) and (4.15). So Lemma 4.3 says that

$$(\frac{w}{\theta})^T c \leq (\frac{z}{\theta})^T b,$$

or $w^T c \leq z^T b$, which again contradicts (4.22). It follows that the original set of inequalities (4.14), (4.15), and (4.16) has a solution. By our above reasoning, this solution provides the desired pair of optimal vectors and the values of the two LPP's are the same. □

We have already shown that any general LPP can be transformed into an equivalent standard LPP. So the Duality Theorem will also apply to general LPP's if we can show that the dual of the general LPP is equivalent to the dual of the corresponding standard LPP.

Theorem 4.7. *The duality theorem holds for the general problem.*

Proof. Suppose we are looking for a vector $x \in \mathbb{R}^m$ such that

$$\sum_{i=1}^{m} x_i c_i \text{ is a maximum} \tag{4.23}$$

subject to

$$\begin{aligned} x_i &\geq 0 \quad \text{for } i \in I, \\ x^T a^j &\leq b_j \quad \text{for } j \in J \\ x^T a^j &= b_j \quad \text{for } j \in J'. \end{aligned} \tag{4.24}$$

Now for each $i \in I'$, introduce two new nonnegative variables x_i' and x_i'' and consider the standard problem of

$$\text{maximizing } \sum_{i \in I} x_i c_i + \sum_{i \in I'} x_i' c_i - \sum_{i \in I'} x_i'' c_i$$

subject to

$$\begin{aligned} \sum_{i \in I} x_i a_{ij} + \sum_{i \in I'} x_i' a_{ij} - \sum_{i \in I'} x_i'' a_{ij} &\leq b_j \qquad j \in J, \\ \sum_{i \in I} x_i a_{ij} + \sum_{i \in I'} x_i' a_{ij} - \sum_{i \in I'} x_i'' a_{ij} &\leq b_j \qquad j \in J', \\ -\sum_{i \in I} x_i a_{ij} - \sum_{i \in I'} x_i' a_{ij} + \sum_{i \in I'} x_i'' a_{ij} &\leq -b_j \qquad j \in J'. \end{aligned}$$

The dual of this standard problem is to find nonnegative numbers y_j for $j \in J$, and y_j', y_j'' for $j \in J'$ such that

$$\sum_{j \in J} y_j b_j + \sum_{j \in J'} y_j' b_j - \sum_{j \in J'} y_j'' b_j \tag{4.25}$$

is a minimum subject to

$$
\begin{aligned}
\sum_{j\in J} y_j a_{ij} + \sum_{j\in J'} y'_j a_{ij} - \sum_{j\in J'} y''_j a_{ij} &\ge c_i \qquad i \in I,\\
\sum_{j\in J} y_j a_{ij} + \sum_{j\in J'} y'_j a_{ij} - \sum_{j\in J'} y''_j a_{ij} &\ge c_i \qquad i \in I',\\
-\sum_{j\in J} y_j a_{ij} - \sum_{j\in J'} y'_j a_{ij} + \sum_{j\in J'} y''_j a_{ij} &\ge -c_i \qquad i \in I'.
\end{aligned}
\tag{4.26}
$$

Now the dual of problem (4.23), (4.24) is by definition that of finding a vector $z \in \mathbb{R}^n$ which

$$\text{minimizes } \sum_{j=1}^{n} z_j b_j$$

subject to

$$
\begin{aligned}
z_j &\ge 0, \ \text{ for } j \in J,\\
\sum_{j=1}^{n} a_{ij} z_j &\ge c_i, \ \text{ for } i \in I,\\
\sum_{j=1}^{n} a_{ij} z_j &= c_i, \ \text{ for } i \in I'.
\end{aligned}
$$

By the same reasoning used earlier in this section, (y, y', y'') is a solution of the problem (4.25), (4.26) if and only if z is a solution of this problem, where $z_j = y_j$ for $j \in J$ and $z_j = y'_j - y''_j$ for $j \in J'$; so the problems are equivalent. It follows that if the original problem and its dual are feasible, the new equivalent standard problems are also feasible. But then from the standard duality theorem just proven the dual standard problems have equal values and hence so do the original problems, completing the proof. □

4.3. **Equivalence Between Solving Games and Programming Problems.** Let Γ be a two-player game with payoff matrices B and C (I won't use A since it already has an interpretation in linear programming), and let X and Y be the sets of mixed-strategy vectors. Given a fixed mixed-strategy $y \in Y$, a best response of player 1 is a vector $x \in X$ that maximizes the expression $x^T(By)$. That is, x is a solution to the LPP

$$\text{maximize } x^T(By) \tag{4.27}$$

subject to

$$
\begin{aligned}
x &\ge 0,\\
\mathbf{1}_m^T x &= 1.
\end{aligned}
\tag{4.28}
$$

It is easy to see this satisfies the definition of a general LPP with $c = By$, $b = 1$, and $A = \mathbf{1}_m^T$. The dual of this LPP (again assuming y is fixed) is to find a variable $u \in \mathbb{R}$ that

$$\text{minimizes } u \tag{4.29}$$

subject to

$$u \geq 0,$$
$$\mathbf{1}_m u \geq By. \tag{4.30}$$

Both LPP's are clearly feasible (pick any $x \in X$ and any positive u exceeding the largest entry of By), so by the duality theorem they have the same optimal value.

Now suppose Γ is a zero-sum game. Player 2, when choosing y, must assume that his opponent plays rationally and maximizes $x^T By$. This maximum payoff to player 1 is the optimal value of the LPP (4.27), (4.28), which is equal to the optimal value u of the dual LPP (4.29), (4.30). Then player 2's goal is to choose y that minimizes the value of u satisfying (4.30). Now define vectors y', b', c', and matrix A' as follows. Let y' be the $(n+1)$-vector where

$$y'_j = y_j \text{ for } 1 \leq j \leq n,$$
$$y'_{n+1} = u.$$

Choose $b' \in \mathbb{R}^{n+1}$ such that $b'_i = 0$ if $1 \leq i \leq n$ and $b'_{n+1} = 1$. Let $c' \in \mathbb{R}^{m+2}$ where $c'_i = 0$ for $1 \leq i \leq m+1$ and $c'_{m+2} = 1$. And define $A' \in \mathbb{R}^{(m+2)\times(n+1)}$ such that

$$\begin{aligned}
a'_{ij} &= b_{ij} && \text{for } 1 \leq i \leq m, 1 \leq j \leq n,\\
a'_{ij} &= -1 && \text{for } 1 \leq i \leq m+1, j = n+1,\\
a'_{ij} &= 0 && \text{for } i = m+1, 1 \leq j \leq n,\\
a'_{ij} &= 1 && \text{for } i = m+2, 1 \leq j \leq n,\\
a'_{ij} &= 0 && \text{for } i = m+2, j = n+1.
\end{aligned}$$

Then the LPP (4.29), (4.30) is equivalent to finding y' that minimizes $y'^T b'$ subject to

$$a'_i y' \leq c'_i \text{ for } 1 \leq i \leq m+1,$$
$$a'_{m+2} y' = c_{m+2},$$

where a'_i denotes the i-th row of A'. Notice that we even incorporated the fact that $y' \in Y$ into A' and c'. Here we take $I = m+1$, so that only the final constraint ($y \in Y$) is an equality. So in the case of a zero-sum game, the constraints of (4.30) are linear in u and y even if y is treated as a variable (whereas before we assumed y was fixed).

So by the above reasoning, a minimax strategy y of player 2 is a solution to the LPP

$$\min_{u,y} u \text{ subject to } \mathbf{1}_n y = 1, \mathbf{1}_m u - By \geq \mathbf{0}, y \geq \mathbf{0}. \tag{4.31}$$

The dual of the LPP (4.31) then has the form

$$\max_{v,y} \mathbf{1}_n v \text{ subject to } \mathbf{1}_m x = 1, \mathbf{1}_n v - x^T B \leq \mathbf{0}, x \geq \mathbf{0},$$

where $v \in \mathbb{R}$. This LPP similarly describes the problem of finding a maximin strategy x for player 1. By the duality theorem, both of these problems have solutions; so the minimax theorem states that (x,y) is a solution to the pair of dual LPP's if and only if (x,y) is a Nash equilibrium. This allows us to reformulate the problem of finding a Nash equilibrium of a zero-sum game as a pair of dual LPP's.

Even if Γ is not zero-sum, a best response x of player 1 to the mixed strategy y of player 2 is still a solution to the LPP (4.27), (4.28). By the duality theorem, a feasible solution x is optimal if and only if there is a dual solution u satisfying $\mathbf{1}_m u \geq By$ and $x^T(By) = u$. Since $x^T \mathbf{1}_m = 1$, this is the same as $x^T(By) = (x^T \mathbf{1}_m)u$, or equivalently

$$x^T(\mathbf{1}_m u - By) = 0. \tag{4.32}$$

Because the vectors x and $\mathbf{1}_m u - By$ are nonnegative, (4.32) states that both vectors cannot have positive components in the same position (at least one of them must have a zero entry at each position). So pure strategy i in S_1 is a best response to z if and only if the i-th component of $\mathbf{1}_m u - By$ is zero.

For player 2, strategy y is a best response to x if and only if it maximizes $(x^T C)y$ subject to $y \in Y$. The dual of this LPP is to minimize v subject to $\mathbf{1}_n v \geq x^T C$. Here, a pair y, v of feasible solutions is optimal if and only if

$$y^T(\mathbf{1}_n^T v - x^T C) = 0. \tag{4.33}$$

Combining these results shows the following:

Theorem 4.8. *The game $\Gamma = (A, B)$ has the Nash equilibrium (x,y) if and only if for suitable u, v,*

$$\begin{aligned} \mathbf{1}_m^T x &= 1 \\ \mathbf{1}_n^T y &= 1 \\ \mathbf{1}_m u - Ay &\geq \mathbf{0} \\ \mathbf{1}_n v - x^T B &\geq \mathbf{0} \\ x, y &\geq \mathbf{0} \end{aligned} \tag{4.34}$$

and (4.32), (4.33) *hold.*

Cottle (1992) defines a *linear complementarity problem* (LCP) as follows: to find a vector $z \in \mathbb{R}^m$ given a vector $q \in \mathbb{R}^n$ and matrix $M \in \mathbb{R}^{n \times n}$ such that

$$z \geq \mathbf{0} \tag{4.35}$$

$$q + Mz \geq \mathbf{0} \tag{4.36}$$

$$z^T(q + Mz) = \mathbf{0}. \tag{4.37}$$

A vector z satisfying (4.35) and (4.36) is called *feasible*, and a vector satisfying (4.37) is called *complementary*. The LCP is therefore to find a vector that is both feasible and complementary; such a vector is called a *solution* of the LCP.

Now consider the LCP with M, q, and z defined as follows. Let $q = \mathbf{1}_{m+n}$. Let z be the $(m+n)$-vector where $z_i = x_i$ for $1 \leq i \leq m$ and $z_i = y_{i-m}$ for $m+1 \leq i \leq m+n$. And let $M \in \mathbb{R}^{(m+n)\times(m+n)}$ be defined as follows:

$$\begin{aligned} &m_{ij} = -a_{ij} \text{ for } 1 \leq i \leq m, m+1 \leq j \leq m+n, \\ &m_{ij} = -b_{ji} \text{ for } m+1 \leq i \leq m+n, 1 \leq j \leq m, \\ &m_{ij} = 0 \text{ otherwise.} \end{aligned}$$

In the remainder of this paper, we will denote the LCP just described by (q, M).

Theorem 4.9. *There exist mappings between the set of solutions* (x,y,u,v) *of the problem in Theorem 4.8 and the set of nonzero solutions* (x',y') *of the LCP* (q,M) *defined by*

$$(x,y,u,v) \to (x/v, y/u) \text{ and } \tag{4.38}$$

$$(x',y') \to \left(\frac{x'}{\mathbf{1}_m^T x'}, \frac{y'}{\mathbf{1}_n^T y'}, \frac{1}{\mathbf{1}_n^T y'}, \frac{1}{\mathbf{1}_m^T x'} \right). \tag{4.39}$$

Proof. Suppose (x,y,u,v) is a solution to the problem in Theorem 4.8, and let $(x',y') = (x/v, y/u)$. If we suppose without loss of generality that all elements of A and B are positive, then both u and v must be positive and (4.35) is satisfied. It is clear that $\mathbf{1}_m - Ay' \geq \mathbf{0}$ iff $\mathbf{1}_m u - Ay \geq \mathbf{0}$ and $\mathbf{1}_n - Bx' \geq \mathbf{0}$ iff $\mathbf{1}_n v - Bx \geq \mathbf{0}$. So (4.36) is satisfied. Similarly, $x'^T(\mathbf{1}_m - Ay') = \mathbf{0}$ iff (4.32) holds and $y'^T(\mathbf{1}_n - Bx') = \mathbf{0}$ iff (4.33) holds. The other direction is similar. □

So if we can solve the LCP (q,M) to obtain a solution (x',y'), all we need to do is normalize both vectors to obtain a solution to the original problem.

In this section I have shown that finding a Nash equilibrium of a two-player zero-sum game can be reduced (in polynomial-time) to solving a pair of dual linear programs. There are known polynomial-time algorithms for solving linear programs (see the ellipsoid method in Chvatal, 1983). In practice, the simplex algorithm (which can be found in any optimization or operations research textbook) also solves linear programs efficiently, although it runs in worst-case exponential time. In section 6 we will show how the Lemke-Howson algorithm gives a solution to the LCP. Unlike linear programming problems, there are currently no known polynomial time algorithms for solving LCP's, and the Lemke-Howson algorithm runs in worst-case exponential time.

5. The Lemke-Howson Algorithm

In this section I will present the Lemke-Howson algorithm for computing a Nash equilibrium in any two-player strategic-form game. This section is based on Shapely (1974) and von Stengel's (2002) expositions of Lemke's original work. Suppose Γ is a two-player strategic-form game with payoff matrices A and B. For notational convenience, let $S_1 = \{1,\dots,m\}$, $S_2 = \{m+1,\dots,m+n\}$, and $S^* = S_1 \cup S_2$. Mixed strategies are represented by vectors $x = (x_1,\dots,x_m) \in \Sigma_1$ and $y = (y_{m+1},\dots,y_{m+n}) \in \Sigma_2$. Geometrically, the sets Σ_1 and Σ_2 are simplexes of dimension $m-1$ and $n-1$ respectively.

Now define regions Σ_1^k in Σ_1 for $k \in S^*$ as follows:

$$\Sigma_1^i = \{x \in \Sigma_1 : x_i = 0\} \text{ for } i \in S_1,$$

$$\Sigma_1^j = \{x \in \Sigma_1 : \sum_{i \in S_1} b_{ij} x_i = \max_{l \in S_2} \sum_{i \in S_1} b_{il} x_i\} \text{ for } j \in S_2.$$

The sets Σ_1^j where $j \in S_2$ (some of which may be empty) consist of all of player 1's strategies to which pure strategy j is a best response for player 2 (although it might not be the unique best response). Since there is always at least one pure strategy best response for player 2 to any strategy of player 1, they cover all of Σ_1.

Define the *label* of $x \in \Sigma_1$ to be the nonempty set

$$L'(x) = \{k : x \in \Sigma_1^k\}.$$

Define regions Σ_2^i and Σ_2^j and the label $L''(y)$ similarly. Also define the label of the pair $(x,y) \in \Sigma^*$ to be $L(x,y) = L'(x) \cup L''(y)$. (x,y) is *completely labeled* if $L(x,y) = S^*$, and *k-almost completely labeled* if $L(x,y) = S^* - \{k\}$ for some $k \in S^*$.

Lemma 5.1. *The profile $(x,y) \in \Sigma$ is a Nash equilibrium iff (x,y) is completely labeled.*

Proof. Suppose (x,y) is a Nash equilibrium and let $i \in S_1$ be arbitrary. If i is not in the support of x, then $x_i = 0$. So $x \in \Sigma_1^i \to i \in L'(x) \to i \in L(x,y)$. If i is in the support of x, then i must be a best response to y and hence $y \in \Sigma_2^i$; so again we have $i \in L(x,y)$. Similar logic shows that $j \in L(x,y)$ for all $j \in S_2$. So (x,y) is completely labeled. Now suppose (x,y) is completely labeled. Then for each $i \in S_1$ either $i \in \Sigma_1^i$ or $i \in \Sigma_2^i$. The first case implies that i is not in the support of x, and the second implies that i is a best response to y. So every strategy in the support of x is a best response to y. Similarly every strategy in the support of y is a best response to x. So (x,y) is a Nash equilibrium by Theorem 1.5. □

Consider the game defined by the following two matrices:

$$A = \begin{bmatrix} 0 & 6 \\ 2 & 5 \\ 3 & 3 \end{bmatrix} \quad B = \begin{bmatrix} 1 & 0 \\ 0 & 2 \\ 4 & 3 \end{bmatrix}. \tag{5.1}$$

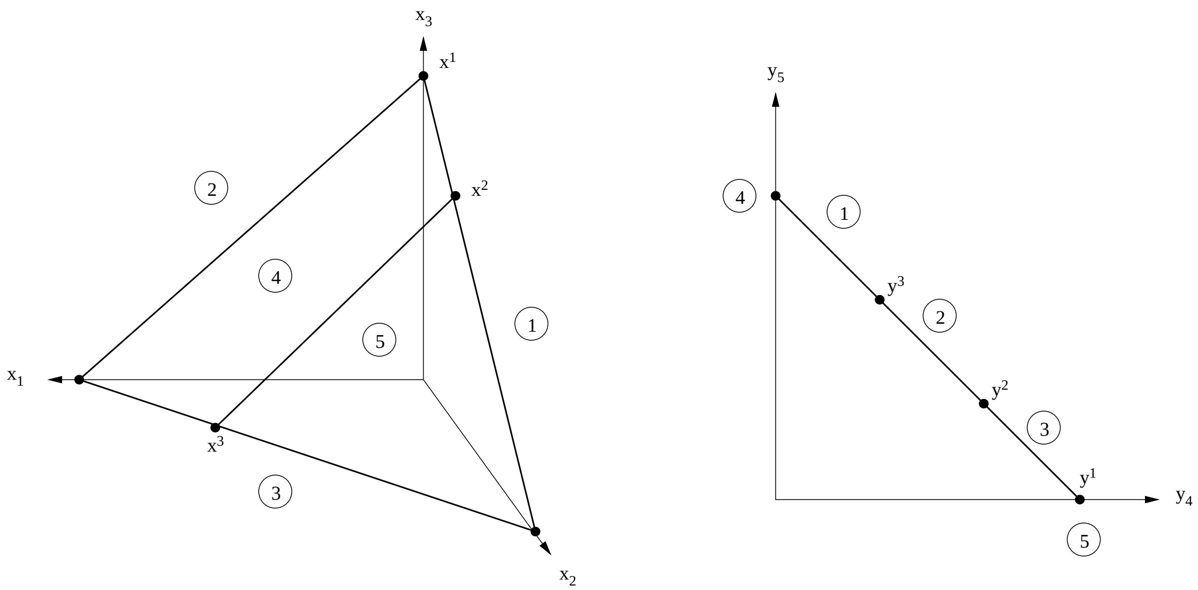

FIGURE 1. Mixed strategy sets X (left) and Y (right) for the game (5.1).

Figure 1 shows the mixed strategy sets X and Y of this game. In the left diagram, the large equilateral triangle with blackened edges represents the set X of possible mixed strategies, while the non-darkened lines denote the axes x_i of pure-strategies. The labels 1-5 are drawn as circled numbers. Labels 1, 2, and 3 represent pure strategies of player 1 and are marked in the left diagram when the corresponding strategy has probability zero and in the right diagram when they are best responses to the strategies of player 2. The pure strategies of player 2 are similarly labeled by 4 and 5. One can determine the labels of any mixed strategy by checking which labels are adjacent to it in the diagram. The nodes denote strategies $x \in X$ and $y \in Y$ that have the maximum number of labels in their respective strategy sets.

It is not difficult to apply Lemma 5.1 to find the equilibria of this game using Figure 1. The equilibria are $(x^1, y^1) = ((0,0,1),(1,0))$, where x^1 has the labels 1, 2, 4 (and y^1 has the remaining labels 3 and 5), $(x^2, y^2) = ((0, \frac{1}{3}, \frac{2}{3}), (\frac{2}{3}, \frac{1}{3}))$, with labels 1, 4, 5 for x^2, and $(x^3, y^3) = ((\frac{2}{3}, \frac{1}{3}, 0), (\frac{1}{3}, \frac{2}{3}))$ with labels 3, 4, 5 for x^3.

This geometric-qualitative inspection is effective at finding equilibria of games of size up to 3×3. It works by inspecting any point $x \in \Sigma_1$ with m labels and checking if there is a point $y \in \Sigma_2$ having the remaining n labels. A game is called *nondegenerate* if any $x \in \Sigma_1$ has at most m labels and any $y \in \Sigma_2$ has at most n labels. In the remainder of this section we will assume that Γ satisfies this nondegeneracy assumption. The following informal discussion explains why "most" games are nondegenerate, so that this assumption does not pose a significant limitation. Notice that every additional label imposes an additional equation that usually reduces the dimension of the set of points having these labels by one. Since the complete set X has dimension $m-1$, we expect no points to have more than m labels. This reasoning will only fail in exceptional circumstances if there is a special relationship between the elements of A or B. This suggests the following equivalent definition of nondegeneracy, which is similar to that used by Shapely (1996):

Lemma 5.2. *For any $x \in \Sigma_1$ with set of labels K and $y \in \Sigma_2$ with set of labels L, the set $\{x' \in \Sigma_1 : K \subseteq L'(x')\}$ is convex and has dimension $m - |K|$, and the set $\{y' \in \Sigma_2 : L \subseteq L''(y')\}$ is convex and has dimension $n - |L|$.*

A convex set X has *dimension* $d \geq 0$ iff there exist $d+1$ linearly independent points $p_1, \ldots, p_{d+1} \in X$ such that

$$X = \{\alpha_1 p_1 + \ldots + \alpha_{d+1} p_{d+1} : \alpha_i \geq 0, \sum_{i=1}^{d+1} \alpha_i = 1\},$$

and X is empty if its dimension is negative. A proof of this result as well as several other equivalent definitions of nondegeneracy can be found in von Stengel (1996).

As an obvious corollary, the set of elements of Σ_1 with a given set K of m labels is either empty or contains a single element (and similarly for elements in Σ_2 with a given set of n labels). Also, let $x \in \Sigma_1$ contain m labels, and let K be a subset of $m-1$ of these labels. Since x clearly contains all of these labels, Lemma 5.2 asserts that the set of elements of Σ_1 containing all the labels in K is one-dimensional. So there exists $x' \in \Sigma_1$ such that the set of elements containing all labels in K equals

$$\{x'' = \alpha x + (1-\alpha)x' : 0 \leq \alpha \leq 1\}.$$

It is clear that x' must be an endpoint of this line segment because either setting $\alpha > 1$ causes a component of x'' to be negative or because strategy j of player 2 is no longer a best response to x'', where $j \in K$. The first case implies that an additional element of x'' has become zero at $\alpha = 1$, in which case x' has m labels. The second case implies that there is an element $j' \in S_2 - K$ such that j' is a best response to x'' when $\alpha > 1$, but not when $\alpha < 1$. So j and j' are both best responses when $\alpha = 1$, and x' has m labels. So in either case x' must contain an additional m-th label, just as x does. It follows that there are exactly two elements of Σ_1 that contain m labels, which include the $m-1$ labels in K. Similar logic holds for all $y \in \Sigma_2$ containing n labels.

Lemma 5.3. *In a nondegenerate game only finitely many points in Σ_1 have m labels and only finitely many points in Σ_2 have n labels.*

Proof. Let K and L be subsets of S^* with $|K| = m$ and $|L| = n$. There are only finitely many such pairs (K, L). Consider the set of points in Σ_1 having the labels in K, and the set of points in Σ_2 having the labels in L. By the above lemma, these sets are empty or singletons. So the desired conclusion follows. □

We can now define two graphs $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ as follows. Let G_1 be the graph whose vertices are the points $x \in \Sigma_1$ that have exactly m labels, with an additional vertex $\mathbf{0}_1$ that has all labels i in S_1. Any two such vertices x and x' are joined by an edge if they differ in exactly one label (and have the other $m-1$ labels in common). Similarly, let G_2 be the graph with vertices $y \in \Sigma_2$ that have n labels, with the extra vertex $\mathbf{0}_2$ having all labels $j \in S_2$, and edges joining those vertices that have $n-1$ labels in common. The *product graph* $G = (V, E) = G_1 \times G_2$ has vertices $(x, y) \in V_1 \times V_2$. There is an edge between two vertices (x, y) and (x', y') iff $y = y'$ and $\{x, x'\} \in E_1$, or $x = x'$ and $\{y, y'\} \in E_2$.

The Lemke-Howson algorithm can be defined in terms of these graphs. Define

$$P = \{(x, y) \in V : L(x, y) = S^*\},$$

and for each $k \in S^*$ let

$$P^k = \{(x, y) \in V : L(x, y) \supseteq S^* - \{k\}\}.$$

These are the node pairs that are completely or k-almost completely labeled; note that $P^k \supseteq P$, and that $k \neq l$ implies $P^k \cap P^l = P$. It is clear from the nondegeneracy assumption and Lemma 5.1 that the members of P are the Nash equilibria of Γ and the node pair $(\mathbf{0}_1, \mathbf{0}_2)$. The following two results follow pretty much directly from the discussion following Lemma 5.2:

Lemma 5.4. *For each $k \in S^*$, any completely labeled $(x, y) \in V$ is adjacent to exactly one vertex pair $(x', y') \in V$ that is belongs to P^k.*

Proof. Every $k \in S^*$ must be a label of either x or y (and not both). If k is a label of x, we know that there is a unique element $x' \in \Sigma_1$ that shares the other $m-1$ labels of x and also has an m-th label: so x' is a vertex of G adjacent to x. So (x', y) is the only vertex of G that is k-almost completely labeled and adjacent to (x, y). Similarly if y has label k, then the unique adjacent vertex in G will be (x, y'). □

Lemma 5.5. *For each $k \in S^*$, any vertex (x, y) in G that belongs to $P^k - P$ is adjacent to exactly two vertices that belong to P^k.*

Proof. Let h denote the unique duplicate label that x and y have in common. We know there is exactly one vertex x' in G_1 adjacent to x that shares the other $m-1$ labels (excluding h), and exactly one vertex y' in G_2 adjacent to y that shares the other $n-1$ labels. So (x, y'), and (x', y) are both adjacent to (x, y) and are also k-almost completely labeled. If instead we omit a label from x or y that is not the shared label, then it will be impossible to obtain a vertex that is k-almost completely labeled: so there are only two such vertices. □

These results suggest the following strategy for finding a Nash equilibrium if we are given some completely labeled vertex v of G. First, pick a $k \in S^*$ and travel to the unique k-almost completely labeled vertex v' that is adjacent to v. If v' is completely labeled, then it is either an equilibrium or it is $(\mathbf{0}_1,\mathbf{0}_2)$. Otherwise, v' must be adjacent to one other k-almost completely labeled vertex besides v : so we can continue this process indefinitely until we arrive at a completely labeled vertex w, which has only one k-almost completely labeled neighbor. It can easily be shown by induction that this path has no cycles; so this strategy determines a unique path to a completely labeled vertex, which is either a Nash equilibrium or $(\mathbf{0}_1,\mathbf{0}_2)$. If we start at the known completely-labeled vertex $(\mathbf{0}_1,\mathbf{0}_2)$, then the path must terminate at a Nash equilibrium. This algorithm was discovered by Lemke and Howson (1964).

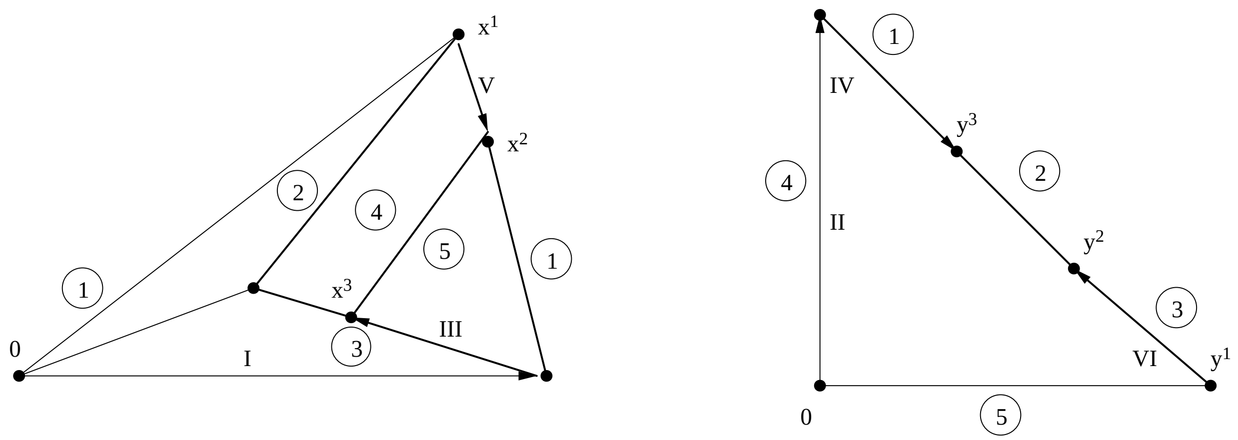

FIGURE 2. The graphs G_1 and G_2 for the game (5.1).

Figure 2 demonstrates the algorithm on the game (5.1) defined above with $k = 2$. The algorithm starts with $x = (0,0,0)$ and $y = (0,0)$. Step I: since x contains label 2, y will remain the same and we must switch x in G_1. It is clear that we must change x to $(0,1,0)$, which causes label 5 to be duplicated. Step II: dropping label 5 in G_2 changes y to $(0,1)$, which picks up label 1. Step III: dropping label 1 in G_1 changes x to $(\frac{2}{3},\frac{1}{3},0)$, which duplicates label 4. Step IV: dropping label 4 in G_2 changes y to $(\frac{1}{3},\frac{2}{3})$, which has the missing label 2. So the algorithm terminates at the Nash equilibrium $((\frac{2}{3},\frac{1}{3},0),(\frac{1}{3},\frac{2}{3}))$. Similarly, steps V and VI in the figure join the equilibria (x^1,y^1) and (x^2,y^2) on a 2-almost completely labeled path.

In addition to its computational power, the Lemke-Howson algorithm also provides an alternative constructive proof that every nondegenerate game contains an equilibrium, independent of Nash's result. In fact, it shows that the number of equilibria in any game must be odd. For, consider the set of all k-almost completely labeled vertices of G and the edges that connect them, for some fixed k. It is clear that this set of vertices and edges consists of disjoint paths and cycles. The cycles consist solely of elements of $P^k - P$, and the paths have completely labeled vertices as endpoints. It follows that there are an even number of completely labeled vertices, and since $(\mathbf{0}_1,\mathbf{0}_2)$ is the only one that is not an equilibrium, the number of Nash equilibria must be odd. Unfortunately, the Lemke-Howson algorithm cannot, in general, find all Nash equilibria of a given game. That is, it is possible that some equilibria are not endpoints of the k-almost completely labeled path from $(\mathbf{0}_1,\mathbf{0}_2)$ for all k. McKelvey and McLennan (1996) discuss some techniques for finding all equilibria of a game, which turns out to be much more difficult than finding a single equilibrium.

6. SOLVING THE LCP

In the previous section we presented the Lemke-Howson geometrically. In this section we show that it can also be interpreted algebraically as a procedure for finding a solution to the LCP (q,M) presented in section 4. Let

$$P_1 = \{x' \in \mathbb{R}^m : x' \geq \mathbf{0}, x'^T B \leq \mathbf{1}_n\}, \tag{6.1}$$

$$P_2 = \{y' \in \mathbb{R}^n : y' \geq \mathbf{0}, Ay' \leq \mathbf{1}_m\}. \tag{6.2}$$

Then by the analysis at the end of section 4, the elements of $(P_1 - \{\mathbf{0}\}) \times (P_2 - \{\mathbf{0}\})$ are feasible solutions to the LCP (q,M).

$P_1 \times P_2$ is the polyhedron defined by

$$Ay' + r = \mathbf{1}_M \tag{6.3}$$
$$B^T x' + s = \mathbf{1}_N \tag{6.4}$$

with $x', y', r, s \geq \mathbf{0}$, where $r \in \mathbb{R}^M$ and $s \in \mathbb{R}^N$ are vectors of *slack* variables. The system (6.3), (6.4) is of the form

$$Cz = q, \tag{6.5}$$

where C, q, and z are defined as follows: z is the $2(m+n)$-vector of nonnegative variables defined by

$$\begin{aligned}
&z_i = y'_i \text{ for } 1 \leq i \leq n,\\
&z_i = x'_{i-n} \text{ for } m+1 \leq i \leq m+n,\\
&z_i = r_{i-(m+n)} \text{ for } m+n+1 \leq i \leq 2m+n,\\
&z_i = s_{i-(2m+n)} \text{ for } 2m+n+1 \leq i \leq 2m+2n.
\end{aligned}$$

C is the $(m+n) \times 2(m+n)$ matrix defined by

$$C = \begin{bmatrix} A & 0 & I_M & 0 \\ 0 & B^T & 0 & I_N \end{bmatrix},$$

where I_M denotes the $m \times m$ identity matrix, and I_n denotes the $n \times n$ identity matrix. And $q = \mathbf{1}_{m+n}$.

Now let us assume that all rows of A are distinct and all rows of B^T are distinct. If this were not the case, then one player would have several "identical" strategies in the sense that he would receive the same payoff playing either one no matter what his opponent played. Thus we lose nothing my eliminating these redundant strategies. Under this assumption, it is clear that matrix C has full rank equal to $m+n$ (the number of rows). So q belongs to the space spanned by the columns C_j of C. A *basis* β is given by a basis $\{C_j | j \in \beta\}$ of this column space, so that the square matrix C_β formed by these columns is invertible. The corresponding *basic solution* is the unique vector $z_\beta = (z_j)_{j \in \beta}$ with $C_\beta z_\beta = q$, where the variables z_j for $j \in \beta$ are called *basic variables*, and $z_j = 0$ for all *nonbasic variables* $z_j, j \notin \beta$, so that (6.5) holds. The solution is unique since C_β is invertible. If this solution also satisfies $z \geq \mathbf{0}$, then the basis β is called *feasible*. If β is a basis for (6.5), then the corresponding basic solution can be read directly from the equivalent system $C_\beta^{-1}Cz = C_\beta^{-1}q$, called a *tableau*, since the columns of $C_\beta^{-1}C$ for the basic variables form the identity matrix. The tableau is equivalent to the system

$$z_\beta = C_\beta^{-1}q - \sum_{j\notin\beta} C_\beta^{-1}C_j z_j, \tag{6.6}$$

which shows how the basic variables depend on the nonbasic variables.

Pivoting is a change of the basis where a nonbasic variable z_j for some j not in β *enters* and a basic variable z_i for some i in β *leaves* the set of basic variables. The pivot step is possible if and only if the coefficient of z_j in the i-th row of the current tableau is nonzero, and is performed by solving the ith equation for z_j and then replacing z_j by the resulting expression in each of the remaining equations.

For a given entering variable z_j, the leaving variable is chosen to preserve feasibility of the basis. Let the components of $C_\beta^{-1}q$ be $\bar{q}_i$ and of $C_\beta^{-1}C_j$ be $\bar{c}_{ij}$, for $i \in \beta$. Then the largest value of z_j such that $z_\beta = C_\beta^{-1}q - C_\beta^{-1}C_j z_j \geq \mathbf{0}$ in (6.6) is

$$\min\{\bar{q}_i/\bar{c}_{ij} \mid i \in \beta, \bar{c}_{ij} > 0\}. \tag{6.7}$$

This is called a *minimum ratio test*. The following lemma says that the minimum in (6.7) will be unique and determines the leaving variable z_i uniquely under the nondegeneracy assumption. After pivoting, the new basis is $\beta \cup \{j\} - \{i\}$.

Lemma 6.1. *The game Γ determined by A and B is nondegenerate if and only if all basic variables have positive values in any basic feasible solution to* (6.3) *and* (6.4)

I won't give a proof of this equivalence, and refer the reader to von Stengel (1996).

The choice of the entering variable depends on the solution being sought. The Simplex method for linear programming is defined by pivoting with an entering variable that improves the value of the objective function. In the system (6.3), (6.4), we are looking for a *complementary* solution where

$$x'^T r = 0, \; y'^T s = 0; \tag{6.8}$$

this implies that (x', y') is a solution of the LCP (q, M), and therefore that its normalization is a Nash equilibrium by Theorem 4.8. In a basic solution to (6.3), (6.4), every nonbasic variable has value zero. Hence, each basis defines a vertex of G which is labeled with the indices of the nonbasic variables. The variables of the system come in *complementary* pairs (x_i, r_i) for the indices $i \in M$ and (y_j, s_j) for $j \in N$. Recall that the Lemke-Howson algorithm follows a path of solutions that have all labels in $S_1 \cup S_2$ except for a missing label k. Thus a k-almost completely labeled vertex is a basis that has exactly one basic variable from each complementary pair, except for a pair of variables (x_k, r_k) (if $k \in M$) or (y_k, s_k) (if $k \in N$) that are both basic. Correspondingly, there is another pair of complementary variables that are both nonbasic, representing the duplicate label. One of them is chosen as the entering variable, depending on the direction of the path being computed. The two possibilities represent the two k-almost completely labeled edges incident to that vertex. The algorithm is started with all components of r and s as basic variables and nonbasic variables $(x', y') = (\mathbf{0}_1, \mathbf{0}_2)$. This initial solution satisfies (6.8) and represents the artificial equilibrium.

Algorithm 6.2. *(Complementary pivoting.) For a bimatrix game with positive payoff matrices A,B compute a sequence of basic feasible solutions to the system* (6.3), (6.4) *as follows.*
(a) Initialize with basic variables $r = 1_M, s = 1_N$. Choose $k \in S_1 \cup S_2$, and let the first entering variable be x'_k if $k \in S_1$ and y'_k if $k \in S_2$.
(b) Pivot such as to maintain feasibility using the minimum ratio test.
(c) If the variable z_i that has just left the basis has index k, halt. Then (6.8) *holds and (x,y) defined by the mapping in Theorem 4.9 is a Nash equilibrium. Otherwise, choose the complement of z_i as the next entering variable and go to (b).*

Now we will demonstrate Algorithm 6.1 for the example of the previous section. The initial basic solution in the form (6.6) is given by

$$\begin{aligned} r_1 &= 1 - 6y'_5 \\ r_2 &= 1 - 2y'_4 - 5y'_5 \\ r_3 &= 1 - 3y'_4 - 3y'_5. \end{aligned} \tag{6.9}$$

and

$$\begin{aligned} s_4 &= 1 - x'_1 - 4x'_3 \\ s_5 &= 1 - 2x'_2 - 3x'_3. \end{aligned} \tag{6.10}$$

Pivoting can be performed separately for these two systems since they have no variables in common. With the missing label 2 as in Figure 2, the first entering variable is x'_2. Then the second equation of (6.10) is rewritten as $x'_2 = \frac{1}{2} - \frac{3}{2}x'_3 - \frac{1}{2}s_5$, and s_5 leaves the basis. Next, the complement y'_5 of s_5 enters the basis. The minimum ratio (6.7) in (6.9) is 1/6, so that r_1 leaves the basis and (6.9) is replaced by the system

$$\begin{aligned} y'_5 &= \frac{1}{6} - \frac{1}{6}r_1 \\ r_2 &= \frac{1}{6} - 2y'_4 + \frac{5}{6}r_1 \\ r_3 &= \frac{1}{2} - 3y'_4 + \frac{1}{2}r_1. \end{aligned} \tag{6.11}$$

Then the complement x'_1 of r_1 enters the basis and s_4 leaves, so that the system replacing (6.10) is now

$$\begin{aligned} x'_1 &= 1 - 4x'_3 - s_4 \\ x'_2 &= \frac{1}{2} - \frac{3}{2}x'_3 - \frac{1}{2}s_5. \end{aligned} \tag{6.12}$$

With y'_4 entering, the minimum ratio (6.7) in (6.11) is 1/12, where r_2 leaves the basis and (6.11) is replaced by

$$\begin{aligned} y'_5 &= \frac{1}{6} - \frac{1}{6}r_1 \\ y'_4 &= \frac{1}{12} + \frac{5}{12}r_1 - \frac{1}{2}r_2 \\ r_3 &= \frac{1}{4} - \frac{3}{4}r_1 - \frac{3}{2}r_2. \end{aligned} \tag{6.13}$$

Then the algorithm terminates since the variable r_2, with the missing label 2 as index, has become nonbasic. The solution defined by the final systems (6.12) and (6.13), with the nonbasic variables on the right hand side equal to zero, satisfies (6.8). Renormalizing x' and y' by the mapping in Theorem 4.9 gives the equilibrium $(x,y) = (x^3, y^3)$ mentioned after example (5.1), with payoffs 4 to player 1 and 2/3 to player 2.

7. Conclusion

We have seen that a two-player zero-sum game in strategic-form can be solved in polynomial time by converting it to a pair of dual linear programs, and a general two-player game in strategic form can be solved by the Lemke-Howson Algorithm, which runs in worst-case exponential time. It turns out that these are among the most efficient known algorithms for solving these problems, and it still remains an important open question whether there exist polynomial-time algorithms for finding a Nash equilibrium in any finite strategic-form game (it has not been proven to be NP-hard). It is surprising that so little is known about the complexity of this problem, despite the fact that a solution is guaranteed to exist by the Nash Existence Theorem. In fact, Berkeley Professor Christos Papadimitriou (2001) stated, “the complexity of finding a Nash equilibrium is in my opinion the most important concrete open question on the boundary of P today.”

While the results presented in this paper are very powerful, they only represent the simplest computational problems in game theory. We only looked at two player games, rather than n-player games; we assumed all games were nondegenerate; we only computed Nash equilibria, rather than other equilibrium refinements; we only computed a single Nash equilibrium, rather than multiple or all equilibria; we only considered games in strategic-form, rather than more powerful models such as extensive form and sequence form; and we only considered two particular linear techniques, rather than other algorithms. These are all interesting and very important questions, which unfortunately I do not have time to address in this paper and refer the reader to McKelvey (1996) and von Stengel (2002). However, many of these results involve relatively simple extensions of the Lemke-Howson algorithm and other techniques developed in this paper, and it should not be too difficult to understand them after this introduction.

References

[1] Chvatal, Vasek. 1983. *Linear Programming*. W. H. Freeman and Company: New York, NY.

[2] Cottle, Richard W., Jong-Shi Pang, and Richard E. Stone. 1992. *The Linear Complementarity Problem*. Academic Press: San Diego, CA.

[3] Fudenberg, Drew and Jean Tirole. 1991. *Game Theory*. The MIT Press: Cambridge, MA.

[4] Gale, David. 1989. *The Theory of Linear Economic Models*. The University of Chicago Press: Chicago, IL.

[5] Glicksman, Abraham M. 2001. *An Introduction to Linear Programming and the Theory of Games*. Dover Publications, Inc: Mineola, NY.

[6] Lemke, C.E. and J. T. Howson Jr. 1964. "Equilibrium Points of Bimatrix Games," in *Journal of the Society for Industrial and Applied Mathematics, Vol. 12, No. 2,* 413-423.

[7] McKelvey, Richard D. and Andrew McLennan. 1996. "Computation of Equilibria in Finite Games," in *Handbook of Computational Economics, Vol. I,* eds. H.M. Amman, D.A. Kendrick, and J. Rust. Elsevier: Amsterdam, pp. 87-142.

[8] Myerson, Roger B. 2004. *Game Theory: Analysis of Conflict*. Harvard University Press: Cambridge, MA.

[9] Nash, John F. 1950. "Equilibrium Points in n-Person Games." *Proceedings of the National Academy of Sciences of the U.S.A., Vol. 36, No. 1,* 48-49.

[10] Papadimitriou, Christos H. 2001. "Algorithms, Games, and the Internet." *STOC.*

[11] Scarf, Herbert E. 1973. *The Computation of Economic Equilibria*. Yale University Press: New Haven, CT.

[12] Shapely, Lloyd S. 1974. "A Note on the Lemke-Howson Algorithm." *Mathematical Programming Study 1,* 175-189.

[13] von Stengel, Bernhard. 1996. "Computing Equilibria for Two-Person Games." Technical Report 253, Dept. of Computer Science, ETH Zurich.

[14] von Stengel, Bernhard. 2002. "Computing Equilibria for Two-Person Games," in *Handbook of Game Theory, Vol. 3,* ed. by R.J. Aumann and S. Hart. Elsevier: Amsterdam, Chapter 45.