Chapter 2

Bimatrix Games, Polytopes, and the Lemke–Howson algorithm

2.1 Nash equilibria and the best response condition

We use the following notation. Given a bimatrix game (A,B) with $m \times n$ payoff matrices A and B, a mixed strategy for player 1 is a vector x in $\mathbb{R}^m$ with nonnegative components that sum to one. A mixed strategy for player 2 is such a vector y in $\mathbb{R}^n$. All vectors are column vectors; the row vector corresponding to x is written as the transpose $x^\top$. The support of a mixed strategy is the set of pure strategies that have positive probability. A best response to y is a mixed strategy x of player 1 that maximizes his expected payoff $x^\top Ay$, and a best response to x is a mixed strategy y of player 2 that maximizes her expected payoff $x^\top By$. A Nash equilibrium is a pair of mutual best responses. Best responses are characterized by the following combinatorial condition, which we state only for a mixed strategy x of player 1.

Lemma 2.1 (Nash (1951)) *Let x and y be mixed strategies of player 1 and 2, respectively. Then x is a best response to y if and only if all strategies in the support of x are pure best responses to y.*

Proof: Let $(Ay)_i$ be the ith component of Ay, which is the expected payoff to player 1 when playing row i. Let $u = \max_i (Ay)_i$. Then

$$x^\top Ay = \sum_i x_i (Ay)_i = u - \sum_i x_i (u - (Ay)_i).$$

Since the sum $\sum_i x_i (u - (Ay)_i)$ is nonnegative, $x^\top Ay \leq u$. The expected payoff $x^\top Ay$ achieves the maximum u if and only if that sum is zero, that is, if $x_i > 0$ implies $(Ay)_i = u$, as claimed. □

A game (A, B) is *symmetric* if $A = B^\top$, so it does not change when the players change roles. The game of "chicken" with $A = B^\top = \begin{pmatrix} 2 & 2 \\ 4 & 1 \end{pmatrix}$ is an example. Its equilibria, in terms of the probability vectors, are the bottom left pure strategy pair $\big((0,1)^\top, (1,0)^\top\big)$ with expected payoffs $4, 2$ to players 1 and 2, respectively, the top right pure strategy pair $\big((1,0)^\top, (0,1)^\top\big)$ with expected payoffs $2, 4$, and the mixed strategy pair $\big((1/3, 2/3)^\top, (1/3, 2/3)^\top\big)$ with expected payoffs $2, 2$. The mixed strategy equilibrium is the only symmetric equilibrium. Its probabilities are uniquely determined by the condition that the pure strategies in the support of the opponent's strategy must both be best responses (by Lemma 2.1) and hence have equal expected payoff.

In a mixed equilibrium, the probabilities are uniquely given by the pair of supports if the corresponding sub-matrices have full rank; the support sizes are then equal. This holds if the game is *nondegenerate*, defined by the property that the number of pure best reponses to any mixed strategy never exceeds the size of its support (see von Stengel (2002) for a detailed discussion). The LH algorithm can be extended to degenerate games by standard lexicographic perturbation techniques, which we discuss in Section 4.3. All of the games we construct are nondegenerate.

By Lemma 2.1, an equilibrium is given if any pure strategy of a player is either a best response (to the opponent's mixed strategy) or is played with probability zero (by the player himself). This can be captured by polytopes (see Ziegler (1995), Grünbaum (2003)) whose facets represent pure strategies, either as best responses or having probability zero. We explain first the simpler case of symmetric equilibria of a symmetric game with $d \times d$ payoff matrix C to player 1, say. We then extend this easily to nonsymmetric games. Let

$$S = \{ z \in \mathbb{R}^d \mid z \geq \mathbf{0},\ \ Cz \leq \mathbf{1} \} \tag{2.1}$$

where $\mathbf{0}$ and $\mathbf{1}$ denote vectors with all entries 0 and 1, respectively, and inequalities hold for all components. The right hand sides of the inequality $Cz \leq \mathbf{1}$, which are 1, represent *normalized* payoffs, as explained in Lemma 2.2 and subsequently. We can assume that *C is nonnegative and has no zero column* by adding a constant to all payoffs, which does not change the best response structure, so that the polyhedron S is bounded and thus a polytope. We assume there are no redundant inequalities in $Cz \leq \mathbf{1}$, which would correspond to dominated strategies. Then the game is nondegenerate if and only if the polytope S is *simple*, that is, every vertex lies on exactly d facets of the polytope. A facet is obtained by making one of the inequalities defining the polytope *binding*, that is, by converting it into an equality. The following lemma characterizes Nash equilibria in terms of polytope vertices, as already shown by Vorob'ev (1958).

Lemma 2.2 *A mixed strategy pair (x,y) is a symmetric Nash equilibrium of the game $(C,C^\top)$ if and only if $x = y = u \cdot z$ and $z \in S$ in (2.1), $z \neq \mathbf{0}$, $u = 1/\sum_i z_i$, and $z^\top(\mathbf{1} - Cz) = 0$, where z must be a vertex of S by nondegeneracy.*

Proof: Let $z \in S$, $z \neq \mathbf{0}$ and $u = 1/\sum_i z_i$. Then $u > 0$, and zu is a mixed strategy x. The condition $Cz \leq \mathbf{1}$ is equivalent to $Cx \leq \mathbf{1}u$. The orthogonality condition $z^\top(\mathbf{1} - Cz) = 0$ is equivalent to $x^\top(\mathbf{1}u - Cx) = 0$, so that for each positive component x_i of x (of which there is at least one), $(Cx)_i = u = \max_k (Cx)_k$. Thus, by Lemma 2.1, x is a best response to itself, that is, (x,x) is a symmetric equilibrium. Conversely, any such equilibrium (x,x), with $u = \max_k (Cx)_k > 0$, and $z = x \cdot 1/u$, gives a vector z with the stated properties.

The vector z is on d facets of S since for each i, either $z_i = 0$ or $(Cz)_i = 1$. If z was not a vertex but on a higher-dimensional face of S, any vertex of that face would be on additional facets, contradicting nondegeneracy of the game as S would then not be a simple polytope. □

In the game of chicken above, $z = (1/6, 1/3)^\top$ gives the symmetric equilibrium. The vector z has to be re-scaled to become a mixed strategy x. The equilibrium payoff u, normalized to 1 in $Cz \leq \mathbf{1}$, is the scaling factor. The converse mapping from x to z defines a projective transformation of a polyhedron representing the "upper envelope" of expected payoffs to the polytope S (see von Stengel (2002)).

2.2 The Lemke–Howson algorithm

The conditions in Lemma 2.2 define a *linear complementarity problem* (LCP) (see Cottle, Pang, and Stone (1992)), usually stated as follows: find z so that

$$z \geq \mathbf{0}, \qquad q + Mz \geq \mathbf{0}, \qquad z^\top (q + Mz) = 0, \tag{2.2}$$

here with data $M = -C$, $q = \mathbf{1}$. This LCP has a trivial solution $z = \mathbf{0}$, which is not a Nash equilibrium. However, $\mathbf{0}$ is an *artificial equilibrium*, which is the starting point of what we call the *symmetric LH algorithm*.

Given a nondegenerate symmetric game $(C, C^\top)$, the symmetric LH algorithm finds a nonzero vertex z of the polytope S in (2.1) so that $z^\top(\mathbf{1} - Cz) = 0$, giving a symmetric Nash equilibrium by Lemma 2.2.

It is useful to *label* the facets of S, as done by Shapley (1974). For each pure strategy i, the facets defined by $z_i = 0$ and by $(Cz)_i = 1$ both get label i. Every vertex has the label of the facets it lies on. The complementarity condition $z^\top(\mathbf{1} - Cz) = 0$ then means that z is *completely labelled* (has all labels i), since then either $z_i = 0$ or $(Cz)_i = 1$ (or both, but this cannot occur since S is simple, so a completely labelled vertex has each label exactly once).

The LH algorithm is started from the completely labelled vertex $z = \mathbf{0}$ by choosing one label k that is initially *dropped*, meaning that label k is no longer required. This is the only free choice of the algorithm, which from then on proceeds in a unique manner. By leaving the facet with label k, a unique edge is traversed whose endpoint is another vertex, which lies on a new facet. The label, say j, of that facet, is said to be *picked up*. If this is the *missing* label k, the algorithm terminates at a completely labelled vertex. Otherwise, j is clearly *duplicate*, and the next edge is (uniquely) chosen by leaving the facet that so far had label j, and the process is repeated. The LH method generates a sequence of *k-almost complementary* edges and vertices (having all labels except possibly k, where k occurs only at the starting point and endpoint). The resulting path cannot repeat a vertex as this would offer a second way to proceed when that vertex is first encountered, which is not the case (since S is simple). Hence, it terminates at a Nash equilibrium. This is illustrated in Figure 2.1 for dimension 3.

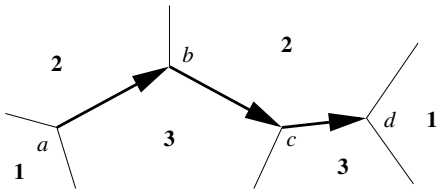

Figure 2.1 The symmetric LH algorithm in dimension 3. Point a is completely labelled, being adjacent to facets with labels $1,2,3$. Dropping label 1, it proceeds to point b picking up label 2, now duplicate. The next point is c with duplicate label 3, and finally d where the missing label 1 is picked up, which terminates the path.

As in the simplex algorithm (Dantzig (1963)), edge traversal is implemented algebraically by *pivoting* with variables entering and leaving a basis, the nonbasic variables representing the current facets. The only difference is the rule for choosing the next entering variable, which in the simplex algorithm depends on the objective function. Here, the *complementary pivoting* rule chooses the nonbasic variable with duplicate label to enter the basis.

In Figure 2.1, the starting point a has an *orientation*, with the labels $1,2,3$ in clockwise order. When label 1 is dropped, the remaining labels keep their orientation (in one dimension less) relative to the edges of the path. In Figure 2.1, label 2 is always to the left and label 3 always to the right of the edge. At the endpoint of the path, the missing label is picked up at the other end of the edge, so that the orientation of that vertex is opposite to that of the starting vertex of the path; in Figure 2.1, point d has labels $1,2,3$ in anticlockwise order. This generalizes to higher dimensions, where orientation is defined as the sign of a certain determinant. The endpoints of any LH path have opposite orientation, which leads to an *index theory* of equilibria (see Shapley (1974) and Garcia and Zangwill (1981)). Knowing the orientation of the artificial equilibrium, the orientation of any k-almost complementary edge can be determined directly.

For nonsymmetric bimatrix games (A,B), or even for finding nonsymmetric equilibria of symmetric games as in the game of "chicken" above, the LH algorithm is applied as follows, which is its standard form. Let

$$z = \begin{pmatrix} x \\ y \end{pmatrix}, \qquad C = \begin{pmatrix} 0 & A \\ B^\top & 0 \end{pmatrix}. \tag{2.3}$$

The polytope S of dimension $d = m+n$ in (2.1) is then the *product* $P \times Q$ of the polytopes

$$P = \{\, x \in \mathbb{R}^m \mid x \geq \mathbf{0},\ \ B^\top x \leq \mathbf{1} \,\}, \qquad Q = \{\, y \in \mathbb{R}^n \mid Ay \leq \mathbf{1},\ \ y \geq \mathbf{0} \,\}. \tag{2.4}$$

Any Nash equilibrium (x,y) of (A,B) is again given by $z^\top(\mathbf{1}-Cz) = 0$, which is equivalent to $x^\top(\mathbf{1}-Ay) = 0$ and $y^\top(\mathbf{1}-B^\top x) = 0$. These conditions state that x is a best response to y and vice versa, where x and y have to be normalized to represent mixed strategies. The only difference to Lemma 2.2 is that this normalization has to be done separately for x and y, rather than for the entire vector z. It is easy to see that in equilibrium $x = \mathbf{0}$ if and only if $y = \mathbf{0}$, and then $(\mathbf{0},\mathbf{0})$ is the artificial equilibrium.

The LH algorithm is applied as before, where a label corresponds either to a strategy i of player 1 or a strategy j of player 2. These have to be distinct, so it is convenient to number the n strategies of player 2 as $m+1,\ldots,m+n$, as suggested by Shapley (1974). A label then represents a pure strategy that has probability zero or is a best response. A label i with $1 \leq i \leq m$ is a strategy of player 1 and determines the facet $x_i = 0$ of P or $(Ay)_i = 1$ of Q, corresponding to the respective ith inequality in (2.4) that becomes binding. A label j with $m+1 \leq j \leq m+n$ is a strategy of player 2 and determines the facet $(B^\top x)_j = 1$ of P or $y_j = 0$ of Q, which is the respective jth binding inequality in (2.4).

The LH path using the edges of $S = P \times Q$ is a subgraph of the *product graph* of the edge graphs of P and Q. This means that edges are alternately traversed in P and Q, keeping the vertex in the other polytope fixed. A duplicate label picked up in P is dropped in Q and vice versa. This is the standard view of the LH algorithm; for further details see von Stengel (2002).

We end this section with two similar, simple, but useful lemmas concerning the LH algorithm, which, as far as we know, have not appeared explicitly in the literature thus far, although they are implicit in the work of McLennan and Tourky (2005). Assume that P and Q are simple, which is true if the game is nondegenerate. Then, any vertex of P is defined by exactly m facets, and so has m labels. Likewise, any vertex of Q has n labels.

Lemma 2.3 *Every LH path on $P \times Q$ induces a simple path in each polytope P and Q, that is, no vertex of P or Q is ever left and visited again on an LH path.*

Proof: Suppose to the contrary that a vertex x of P is left and visited again on an LH path. This means that there are three vertex pairs (x,y), (x,y'), and (x,y'') with $y \neq y' \neq y''$ on the

LH path on $P \times Q$, and all three are k-almost-complementary, for some k in $\{1, \ldots, m+n\}$. The m labels of x define $n-1$ labels that y, y', and y'' must share. However, this is impossible, since the $n-1$ labels correspond to $n-1$ equations in n-space that define a line, which can only contain two vertices. The situation for Q is analogous. □

Lemma 2.4 *Let (x,y) be a completely labelled vertex pair of $P \times Q$. If (x,y') is part of an LH path on $P \times Q$, either $y' = y$ or (x,y) and (x,y') are connected by an edge on that LH path (i.e. y' and y connected by an edge in Q).*

Lemma 2.4 states that if a vertex x of P is part of some complementary vertex pair, then, for any LH path, it only appears as the endpoint of the induced path in P. The analogous statement is true for a vertex y of Q that is part of a complementary vertex pair.

Proof: Since (x,y') is part of an LH path, it is a k-almost-complementary vertex pair, for some k in $\{1, \ldots, m+n\}$. This means that x and y' have at least $m+n-1$ distinct labels. Either they have all $m+n$ labels and $y' = y$, or y' shares exactly one label with x, and then this is a duplicate label. Dropping this label in Q leads to the vertex y, along the edge connecting y' and y. □