7. Infinite Games.

In this Chapter, we treat infinite two-person, zero-sum games. These are games (X, Y, A), in which at least one of the strategy sets, X and Y, is an infinite set. The famous example of Exercise 4.7.3, he-who-chooses-the-larger-integer-wins, shows that an infinite game may not have a value. Even worse, the example of Exercise 4.7.5 shows that the notion of a value may not even make sense in infinite games without further restrictions. This latter problem can be avoided if we assume that the function $A(x, y)$ is either bounded above or bounded below.

7.1 The Minimax Theorem for Semi-Finite Games. The minimax theorem for finite games states that every finite game has a value and both players have optimal mixed strategies. The first theorem below generalizes this result to the case where only one of the players has a finite number of pure strategies. The conclusion is that the value exists and the player with a finite number of pure strategies has an optimal mixed strategy. But first we must discuss mixed strategies and near optimal strategies for infinite games.

Mixed Strategies for Infinite Games: First note that for infinite games, the notion of a mixed strategy is somewhat open to choice. Suppose the strategy set, Y, of Player II is infinite. The simplest choice of a mixed strategy is a *finite distribution* over Y. This is a distribution that gives all its probability to a finite number of points. Such a distribution is described by a finite number of points of Y, say $y_1, y_2, \ldots, y_n$, and a set of probabilities, $q_1, q_2, \ldots, q_n$ summing to one with the understanding that point y_j is chosen with probability q_j. We will denote the set of finite distributions on Y by Y^*_F.

When Y is an interval of the real line, we may allow as a mixed strategy any distribution over Y given by its distribution function, $F(z)$. Here, $F(z)$ represents the probabiity that the randomly chosen pure strategy, y, is less than or equal to z. The advantage of enlarging the set of mixed strategies is that it then becomes more likely that an optimal mixed strategy will exist. The payoff for using such a strategy is denoted by $A(x, F)$ for $x \in X$.

Near Optimal Strategies for Infinite Games: When a game has a finite value and an optimal strategy for a player does not exist, that player must be content to choosing a strategy that comes within ϵ of achieving the value of the game for some small $\epsilon > 0$. Such a strategy is called an ϵ-optimal strategy and was discussed in Chapter 6.

In infinite games, we allow the value to be $+\infty$ or $-\infty$. For example, the value is $+\infty$ if for every number B, however large, there exists a mixed strategy $\boldsymbol{p}$ for Player I such that $A(\boldsymbol{p}, y) \geq B$ for all $y \in Y$. A simple example would be: $X = [0, \infty)$, Y arbitrary, and $A(x, y) = x$ independent of y. The value is $+\infty$, since for any B, Player I can guarantee winning at least B by choosing any $x \geq B$. Such a strategy might be called B-optimal. We will refer to both ϵ-optimal and B-optimal strategies as *near optimal strategies.*

The Semi-Finite Minimax Theorem. For finite $X = \{x_1, \ldots, x_m\}$, we denote the set of mixed strategies of Player I as usual by X^*. If Player I uses $\boldsymbol{p} \in X^*$ and Player II uses

$\boldsymbol{q} \in Y_F^*$, then the average payoff is denoted by

$$A(\boldsymbol{p}, \boldsymbol{q}) = \sum_i \sum_j p_i A(x_i, y_j) q_j. \tag{1}$$

We denote the set of mixed strategies of Player II by Y^*, but we shall always assume that $Y_F^* \subset Y^*$.

Consider the semi-finite two-person zero-sum game, (X, Y, A), in which X is a finite set, Y is an arbitrary set, and $A(x, y)$ is the payoff function — the winnings of Player I if he chooses $x \in X$ and Player II chooses $y \in Y$. To avoid the the possibility that the average payoff does not exist or that the value might be $-\infty$, we assume that the payoff function, $A(x, y)$, is bounded below. By bounded below, we mean that there is a number M such that $A(x, y) > M$ for all $x \in X$ and all $y \in Y$. This assumption is weak from the point of view of utility theory because, as mentioned in Appendix 1, it is customary to assume that utility is bounded.

It is remarkable that the minimax theorem still holds in this situation. Specifically the value exists and Player I has a minimax strategy. In addition, for every $\epsilon > 0$, Player II has an ϵ-minimax strategy within Y_F^*.

Theorem 7.1. *If X is finite and A is bounded below, then the game (X, Y, A) has a finite value and Player I has an optimal mixed strategy. In addition, if X has m elements, then Player II has near optimal strategies that give weight to at most m points of Y.*

This theorem is valid without the asumption that A is bounded below provided Player II is restricted to finite strategies, i.e. $Y^* = Y_F^*$. See Exercise 1. However, the value may be $-\infty$, and the notion of near optimal strategies must be extended to this case.

By symmetry, if Y is finite and A is bounded above, then the game (X, Y, A) has a value and Player II has an optimal mixed strategy.

Solving Semi-Finite Games. Here are two methods that may be used to solve semi-finite games. We take X to be the finite set, $X = \{x_1, \ldots, x_m\}$.

METHOD 1. The first method is similar to the method used to solve $2 \times n$ games presented in Section 2.2. For each fixed $y \in Y$, the payoff, $A(\boldsymbol{p}, y)$, is a linear function of $\boldsymbol{p}$ on the set $X^* = \{\boldsymbol{p} = (p_1, \ldots, p_m) : p_i \geq 0, \sum p_i = 1\}$. The optimal strategy for Player I is that value of $\boldsymbol{p}$ that maximizes the lower envelope, $f(\boldsymbol{p}) \equiv \inf_{y \in Y} A(\boldsymbol{p}, y)$. Note that $f(\boldsymbol{p})$, being the infimum of a collection of concave continuous (here linear) functions, is concave and continuous on X^*. Since X^* is compact, there exists a $\boldsymbol{p}$ at which the maximum of $f(\boldsymbol{p})$ is attained. General methods for solving concave maximization problems are available.

Example 1. Player I chooses $x \in \{x_1, x_2\}$, Player II chooses $y \in [0, 1]$, and the payoff is

$$A(x, y) = \begin{cases} y & \text{if } x = x_1 \\ (1 - y)^2 & \text{if } x = x_2 \end{cases}$$

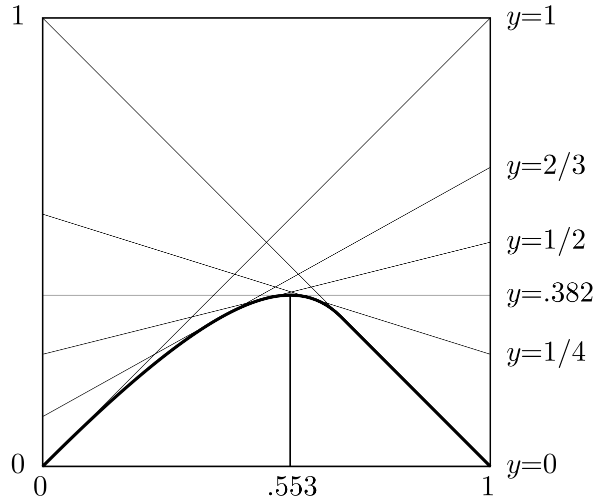

Figure 7.1

Let p denote the probability that Player I chooses x_1. For a given choice of $y \in Y$ by Player II, the expected payoff to Player I is $A(p,y) = py + (1-p)(1-y)^2$. The minimum of $A(p,y)$ over y occurs at $p - (1-p)2(1-y) = 0$, or $y = (2-3p)/(2-2p)$; except that for $p > 2/3$, the minimum occurs at $y = 0$. So,

$$f(p) = \min_y A(p,y) = \begin{cases} p\,\frac{4-5p}{4-4p} & \text{if } p \le 2/3 \\ 1-p & \text{if } p \ge 2/3 \end{cases}$$

The maximum of this function occurs for $p \le 2/3$, and is easily found to be $p = 1 - (1/\sqrt{5}) = .553\ldots$. The optimal strategy for Player II occurs at that value of y for which the slope of $A(p,y)$ (as a function of p) is zero. This occurs when $y = (1-y)^2$. We find $y = (3-\sqrt{5})/2 = .382\ldots$ is an optimal pure strategy for Player II. This is also the value of the game. See Figure 7.1.

METHOD 2: S-GAMES. (Blackwell and Girshick (1954).) Let $X = \{1, 2, \ldots, m\}$, and let S be a non-empty convex subset of m-dimensional Euclidean space, $\mathbb{R}^m$, and assume that S is bounded below. Player II chooses a point $\boldsymbol{s} = (s_1, \ldots, s_m)$ in S, and simultaneously Player I chooses a coordinate $i \in X$. Then Player II pays y_i to Player I. Such games are called S-games.

This game arises from the semi-finite game (X, Y, A) with $X = \{1, 2, \ldots, m\}$ by letting $S_0 = \{\boldsymbol{s} = (A(1,y), \ldots, A(m,y)) : y \in Y\}$. Choosing $y \in Y$ is equivalent to choosing $\boldsymbol{s} \in S_0$. Although S_0 is not necessarily convex, Player II can, by using a mixed strategy, choose a probability mixture of points in S_0. This is equivalent to choosing a point $\boldsymbol{s}$ in S, where S is the convex hull of S_0.

To solve the game, let W_c denote the "wedge" at the point $(c, \ldots, c)$ on the diagonal in $\mathbb{R}^m$,

$$W_c = \{\boldsymbol{s} : s_i \le c \quad \text{for all } i = 1, \ldots, m\}.$$

Start with some c such that the wedge W_c contains no points of s, i.e. $W_c \cap S = \emptyset$. Such a value of c exists from the assumption that S is bounded below. Now increase c and so push the wedge up to the right until it just osculates S. See Figure 7.2(a). This gives the value of the game:

$$v = \sup\{c : W_c \cap S = \emptyset\}.$$

Any point $\boldsymbol{s} \in W_v \cap S$ is an optimal pure strategy for Player II. It guarantees that II will lose no more than v. Such a strategy will exist if S is closed. If $W_v \cap S$ is empty, then any point $\boldsymbol{s} \in W_{v+\epsilon} \cap S$ is an ϵ-optimal strategy for Player II. The point $(v, \ldots, v)$ is not necessarily an optimal strategy for Player II. The optimal strategy could be on the side of the wedge as in Figure 7.2(b).

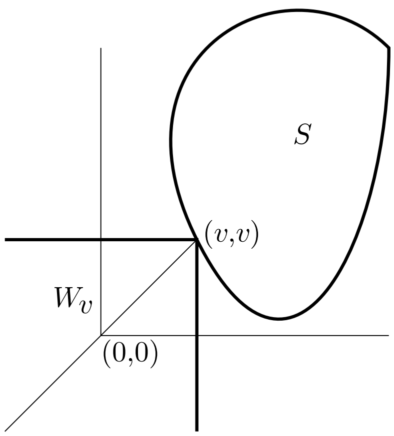

Figure 7.2(a)

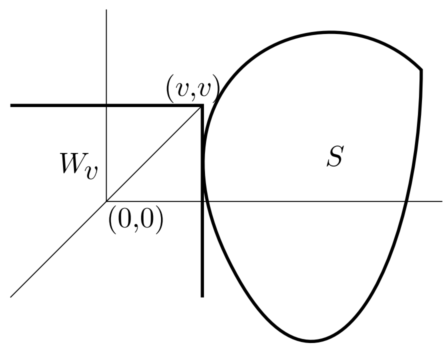

Figure 7.2(b)

To find an optimal strategy for Player I, first find a plane that separates W_v and S, i.e. that keeps W_v on one side and S on the other. Then find the vector perpendicular to this plane, i.e. the normal vector. The optimal strategy of Player I is the mixed strategy with components proportional to this normal vector.

Example 2. Let S be the set $S = \{(y_1, y_2) : y_1 \geq 0, y_2 \geq (1 - y_1)^2\}$. This is essentially the same as Example 1. The wedge first hits S at its vertex (v, v), that is, when $v = (1 - v)^2$. The solution to this equation gives tha vaslue, $v = (3 - \sqrt{5})/2$. The point (v, v) is optimal for Player II. To find Player I's optimal strategy, we find the slope of the curve $y_2 = (1 - y_1)^2$ at the point (v, v). The slope of the curve is $-2(1 - y_1)$, which at $y_1 = v$ is $-2(1 - v) = 1 - \sqrt{5}$. The slope of the normal is the negative of the reciprocal of this, namely $1/(\sqrt{5} - 1)$. So $p_2/p_1 = 1/(\sqrt{5} - 1)$, and since $p_1 + p_2 = 1$, we find $p_2(\sqrt{5} - 1) = 1 - p_2$, or $p_2 = 1/\sqrt{5}$ and $p_1 = 1 - (1/\sqrt{5})$. as found in Example 1.

Exercise 1. Let $X = \{-1, 1\}$, let $Y = \{\ldots, -2, -1, 0, 1, 2, \ldots\}$ be the set of all integers, and let $A(x, y) = xy$.

(a) Show that if we take $Y^* = Y_F^*$, the set of all finite distributions on Y, then the value exists, is equal to zero and both players have optimal strategies.

(b) Show that if Y^* is taken to be the set of all distributions on Y, then we can't speak of the value, because Player II has a strategy, $\boldsymbol{q}$, for which the expected payoff, $A(x, \boldsymbol{q})$ doesn't exist for any $x \in X$.

Exercise 2. Simultaneously, Player I chooses $x \in \{x_1, x_2\}$, and Player II chooses $y \in [0,1]$; then I receives

$$A(x,y) = \begin{cases} y & \text{if } x = x_1 \\ e^{-y} & \text{if } x = x_2 \end{cases}$$

from II. Find the value and optimal strategies for the players.

Exercise 3. Player II chooses a point (y_1, y_2) in the ellipse $(y_1-3)^2 + 4(y_2-2)^2 = 4$. Simultaneously, Player I chooses a coordinate $k \in \{1,2\}$ and receives y_k from Player II. Find the value and optimal strategies for the players.

Solutions.

1. (a) The value is 0. Player II has a pure optimal strategy, namely $y = 0$, since $A(x,0) = 0$ for all $x \in X$. Player I's optimal strategy is 1 and -1 with equal probability 1/2, since if $\boldsymbol{q} \in Y_F^*$,

$$A(1/2, \boldsymbol{q}) = (1/2)\sum_j jq_j - (1/2)\sum_j jq_j = 0$$

because the sums are finite.

(b) If $q_j = \begin{cases} (1/4)(1/|j|) & \text{for } |j| = 2^{-k} \text{ for some } k \ge 2 \\ 0 & \text{otherwise} \end{cases}$, then $\sum_j q_j = 1$ and

$$A(x, \boldsymbol{q}) = (\cdots - \frac{1}{4} - \frac{1}{4} - \frac{1}{4}) + (\frac{1}{4} + \frac{1}{4} + \frac{1}{4} + \cdots) = -\infty + \infty.$$

2. It is easier to use Method 2. The value occurs at the intersection of the curve $y_2 = e^{-y_1}$ and the line $y_1 = y_2$, namely, it is the solution of the equation $v = e^{-v}$. This is about $v = .5671$. The optimal strategy for II is $y = v$. The slope of the tangent line to the curve $y_2 = e^{-y_1}$ at the point $y_1 = v$ is $-e^{-v} = -v$. The normal to this is the negative of the reciprocal, namely $1/v$. The optimal strategy of Player I takes x_1 and x_2 in proportions $v : 1$. This is the mixed strategy $(v/(1+v), 1/(1+v)) = (.3619, .6381)$.

3. The value occurs where the curve $(y_1-3)^2 + 4(y_2-2)^2 = 4$ is first hit by the line $y_1 = y_2$. Thus it satisfies the equation $(v-3)^2 + 4(v-2)^2 = 4$ with $v < 2$. This leads to the equation, $v^5 - 22v + 21 = 0$, or $(5v-7)(v-3) = 0$, so $v = 7/5$ is the value. Player II's optimal pure strategy is (v,v). The slope of the curve $(y_1-3)^2 + 4(y_2-2)^2 = 4$ at the point (v,v) is $y_2' = \frac{-(y_1-3)}{4(y_2-2)} = \frac{-(v-3)}{4(v-2)}$. The normal is $\frac{4(2-v)}{(3-v)} = \frac{3}{2}$. So Player I's optimal strategy is $(2/5, 3/5)$.