

Game Theory

R. Chandrasekaran

1 Extensive Forms

A game tree (a rooted tree) consists of a finite set of nodes or vertices, (one of which is the root) and a set of pairs of nodes called edges. There are no simple closed loops in this structure and so it is called a tree. There is a direction (sort of) imposed by the root A . A vertex C follows (is a descendant of) another vertex B if the vertex B occurs on the unique path from A to C . If C follows B and there is no intermediate vertex between them on the path from A , then we think of (B, C) as an *edge* of the tree and we say C follows B immediately. A vertex is a terminal vertex (or a leaf node) if no vertex follows it.

Definition 1 *By an n -person game in extensive form is meant:*

1. *a (directed) tree Γ with a distinguished vertex A called the starting point of the game depicted by Γ (this is the root of the tree).*
2. *a function, called the payoff function, which assigns an n -vector to each terminal vertex of the tree Γ .*
3. *a partition of the nonterminal vertices of Γ into $n + 1$ sets S_0, S_1, \dots, S_n called player sets. S_0 refers to those sets when chance (nature) plays a role.*
4. *a probability distribution, defined at each vertex of the set S_0 , over its immediate followers.*
5. *for each $i = 1, 2, \dots, n$, a subpartition of S_i into subsets S_i^j , called information set such that the vertices in the same information sets have the same number of immediate followers and that no vertex can follow*

another in the same information set. We should not be able to distinguish between two vertices in the same information set of a player. This is because a player might not know the actions in that round of other players when (s)he makes a move.

6. for each information set S_i^j , an index set I_i^j , together with a one-to-one mapping of the set I_i^j onto the set of immediate followers of each vertex of S_i^j .

Definition 2 Player i in game Γ is said to have perfect information if $|S_i^j| = 1$ for all j . The game Γ is said to have perfect information if each player has perfect information in Γ .

1.1 Games in Normal/Strategic Form

The notion of a "Book of Game Plans" that specifies for each possible situation that a player finds himself in is called a "strategy". More formally:

Definition 3 By a strategy for player i is meant a function which assigns to each of player i 's information sets S_i^j , one of the edges that follows a representative vertex [please note that in any of these vertices it is the corresponding follower that will be chosen]. The set of all strategies for player i will be denoted by Σ_i .

We denote by the vector $\pi(\sigma_1, \sigma_2, \dots, \sigma_n) \in R^n$ the payoff (expected payoff if there are chance moves) if player i selects strategy $\sigma_i \in \Sigma_i$ for $i = 1, 2, \dots, n$. $\pi_i(\sigma_1, \sigma_2, \dots, \sigma_n)$ is the payoff to i^{th} player for this combination. This n -dimensional array of "numbers" is called the "normal/strategic form" of the game Γ . It is assumed in this form that this array is known to all in the game with complete information.

Definition 4 Γ is said to be finite if it has finite number of nodes.

Definition 5 A game Γ is said to decompose at a vertex v if there are no information sets which have vertices from both $\{\text{followers of } v\}$ and the rest of the game tree. The (sub)game represented by v and its followers is denoted by Γ_v and the game representing the remaining part tree of Γ with v replaced by Γ_v is denoted by Γ/v . The payoff in Γ/v at the vertex that replaces v (now a terminal vertex) is the subgame Γ_v .

If we decompose a game at v , then we can decompose the strategy σ into two parts: $\sigma_{\Gamma/v}$ by restricting σ to information sets in Γ_v , and σ_{Γ_v} by doing the same on σ for Γ_v . Conversely such parts can be combined to get an overall strategy for the whole game. We will come back to this when we show that every game with perfect information has an equilibrium (whatever this means!).

Domination and Iterated Domination

Definition 6 A strategy σ_i for player i is dominated by strategy σ'_i , if

$$\pi_i(\sigma_1, \sigma_2, \dots, \sigma_{i-1}, \sigma_i, \sigma_{i+1}, \dots, \sigma_n) < \pi_i(\sigma_1, \sigma_2, \dots, \sigma_{i-1}, \sigma'_i, \sigma_{i+1}, \dots, \sigma_n) \forall \sigma_j \in \Sigma_j; j \neq i$$

It is argued that "rational" players do not play dominated strategies, because there is no belief that a player could hold about the strategies of other players such that it would be optimal to play such a strategy.

Example 1 Consider a two player game with two strategies for each player commonly referred to as the Prisoner's Dilemma: [See the complementary issue discussed later].

1 ↓ 2 →	Mum	Fink
Mum	[-1, -1]	[-9, 0]
Fink	[0, -9]	[-6, -6]

Fink dominates Mum for each player. So it is argued by this theory that rational players will choose [Fink, Fink] combination.

Example 2 (TF Book)

1 ↓ 2 →	L	R
U	[1, 3]	[4, 1]
D	[0, 2]	[3, 4]

For player 1, U dominates D for player 2 (after D is eliminated by player 1), L dominates R. So the theory of this type (assuming that players eliminate dominated strategies and both know this of the other) predicts the outcome (U,L). Suppose however, player 1 changes his payoff for (U,L) to -1 and that of (U,R) to 2, so that the game looks like:

1 ↓ 2 →	L	R
U	[-1, 3]	[2, 1]
D	[0, 2]	[3, 4]

Now D dominates U for player 1 and player 2 knowing this, has R dominating L . So now the predicted outcome is (D,R) ! This yields more to player 1 than (U,L) in both games and so there is an incentive for players to "fudge" the payoff values if they can get away with this. It is not sufficient to change one's payoff values but you need to let the other player know this !

Example 3 Consider the following game:

$1 \downarrow 2 \rightarrow$	<i>Left</i>	<i>Middle</i>	<i>Right</i>
<i>Up</i>	[1, 0]	[1, 2]	[0, 1]
<i>Down</i>	[0, 3]	[0, 1]	[2, 0]

To begin with, neither strategy dominates other for player 1. However, for player 2, *Middle* dominates *Right*. So one argues that if player 2 is "rational", (s)he will not play *Right*. If player 1 "knows" that player 2 is a "rational" player who will not play the dominated strategy *Right*, player 1 now looks at the reduced game:

$1 \downarrow 2 \rightarrow$	<i>Left</i>	<i>Middle</i>
<i>Up</i>	[1, 0]	[1, 2]
<i>Down</i>	[0, 3]	[0, 1]

and finds that *Down* is dominated by *Up*. So if player 1 is a "rational" player, he will not play *Down*. So the game reduces further to:

$1 \downarrow 2 \rightarrow$	<i>Left</i>	<i>Middle</i>
<i>Up</i>	[1, 0]	[1, 2]

Now *Middle* dominates *left* for player 2 and so we arrive at the combination [*Up*,*Middle*] as the final outcome if both players are rational and each knows that the other is rational. This process is called "iterated elimination of strictly dominated strategies". If we want to do this arbitrary number of times, we assume "common knowledge" that players are rational and both know this of each other and both know that of each other and so on. But in a game given below this does not produce any outcome since no strategy is dominated:

$1 \downarrow 2 \rightarrow$	<i>L</i>	<i>C</i>	<i>R</i>
<i>T</i>	[0, 4]	[4, 0]	[5, 3]
<i>M</i>	[4, 0]	[0, 4]	[5, 3]
<i>B</i>	[3, 5]	[3, 5]	[6, 6]

A complementary question is also of interest. If there is no belief that player i could hold (about the strategies that the other players will choose) such that it would be optimal to play the strategy σ_i , can we conclude that there must be another strategy for player i that strictly dominates σ_i ? [See mixed strategies for an answer to this question].

Example 4 (TF Book): Second Price Auction: A seller has an object for sale. There are n buyers interested in buying this from the seller. The value of the object for i^{th} buyer is v_i and we suppose that the players are numbered so that $0 \leq v_1 \leq v_2 \leq \dots \leq v_n$ and all buyers know the value of the object to all buyers. This is called "common knowledge". Buyer i submits a sealed bid $s_i \in [0, \infty)$ which is the amount he is willing to pay the seller for the object. The buyer who "bids" the highest "wins" the "auction" but pays the "bid" by the second highest bidder. If bidder i "wins" (and hence $s_i > \max_{j \neq i} s_j$) the auction, he gains an amount $u_i = v_i - \max_{j \neq i} s_j$. The remaining bidders pay nothing and do not get the object and hence they gain 0. If several bidders have the highest bid, the winner is "randomly" chosen among them. We analyze this problem in what follows.

Let $\max_{j \neq i} s_j = r_i$. If player i bids $s_i > v_i$, then his gain is 0 if $r_i > s_i$. In this case his gain is 0 even if he bids v_i . So bidding v_i (weakly) dominates bidding s_i . If $r_i < s_i$, his gain is $v_i - r_i$ and he gains the same amount even if he bids v_i . Again bidding v_i dominates (weakly) bidding s_i . Thus, any bid higher than v_i is dominated (weakly) by the bid of v_i . If player i bids $s_i < v_i$: If $r_i > v_i$, then his gain is 0; but his gain would be the same if he bid v_i . If $r_i < s_i$ his gain is $v_i - r_i$ and it would be the same if he bid v_i . If $s_i < r_i < v_i$, his gain is 0 where as it is $v_i - r_i > 0$ if he bids v_i . Thus again the bid of v_i dominates the bid of s_i . Thus the only undominated bid is v_i for player i . Hence player n "wins" the auction and his gain is $v_n - v_{n-1}$. Bidding in auctions are very important applications of game theory.

Unfortunately, domination alone is not sufficient to arrive at a solution as in the game:

$1 \downarrow 2 \rightarrow$	L	R
U	$[-1, 1]$	$[0, 0]$
D	$[0, 0]$	$[-1, 1]$

where there is no domination. A stronger concept follows.

(Cournot) Nash Equilibrium

Definition 7 Given a game Γ in strategic form, a set of strategies $(\sigma_1^*, \sigma_2^*, \dots, \sigma_n^*)$ is said to be "in equilibrium" or an equilibrium n -tuple of strategies, if and only if

$$\pi_i(\sigma_1^*, \sigma_2^*, \dots, \sigma_{i-1}^*, \sigma_i, \sigma_{i+1}^*, \dots, \sigma_n^*) \leq \pi_i(\sigma_1^*, \sigma_2^*, \dots, \sigma_{i-1}^*, \sigma_i^*, \sigma_{i+1}^*, \dots, \sigma_n^*) \forall \sigma_i \in \Sigma_i; \forall i$$

In other words, no single player has an incentive to move away from the equilibrium strategy, assuming all others stick to their part of the equilibrium. Such an idea is often referred to as "self-enforcing" or "strategically stable". In two person zero-sum games these are also known as "saddle points".

Players' strategies that are part of some Nash equilibrium survive the iterated elimination process of removing dominated strategies. In the case of Prisoners' Dilemma there is only one Nash equilibrium – Fink-Fink. But there are games as the one below (known as the Battle of the Sexes") that have multiple Nash equilibria:

1 ↓ 2 →	Opera	Ballgame
Opera	[2, 1]	[0, 0]
Ballgame	[0, 0]	[1, 2]

So we may need to develop concepts that select among Nash equilibria. So far we have games that have zero, one, or many Nash equilibria among "pure" strategies. We extend the set of strategies to avoid the nonexistence problem.

The reason for calling these "equilibrium points" is that if one changes the strategies by a "small" perturbation, incentives on rationality will bring the players to equilibrium. The region from which the solution will come back to a particular equilibrium is known as the region of attraction. This is what Cournot was originally interested in systems.

Mixed Strategies

Definition 8 A "mixed strategy" for a player is a probability distribution over his strategies.

Definition 9 A game Γ in strategic form is called a "zero-sum" game if:

$$\sum_{i=1}^n \pi_i(\sigma_1, \sigma_2, \dots, \sigma_{i-1}, \sigma_i, \sigma_{i+1}, \dots, \sigma_n) = 0 \quad \forall (\sigma_1, \sigma_2, \dots, \sigma_{i-1}, \sigma_i, \sigma_{i+1}, \dots, \sigma_n) \in \Sigma_1 \times \Sigma_2 \times \dots \times \Sigma_n$$

Games in which

$$\sum_{i=1}^n \pi_i(\sigma_1, \sigma_2, \dots, \sigma_{i-1}, \sigma_i, \sigma_{i+1}, \dots, \sigma_n) = c \quad \forall (\sigma_1, \sigma_2, \dots, \sigma_{i-1}, \sigma_i, \sigma_{i+1}, \dots, \sigma_n) \in \Sigma_1 \times \Sigma_2 \times \dots \times \Sigma_n$$

where c is a constant are equivalent to zero-sum games and are often called 'strictly competitive' games.

Please note that every game can be made into a zero-sum game by adding one more player with just one strategy. But our interest is mostly in two person zero-sum games which have a rich connection with the theory of linear programming.

1.2 Two Person Zero-sum Games

Here we deal with games in which both players have finite number of pure strategies. These are often called matrix games since the payoff can be specified by one matrix A (since the payoff matrix B for the other player is equal to $-A$). If the number of pure strategies for the players are m and n respectively, then the matrix A is an $m \times n$ matrix. Thus a mixed strategy for player 1 is the set:

$$X = \{x \in R^m : x \geq 0; \sum_{i=1}^m x_i = 1\}$$

and that for player 2 is the set:

$$Y = \{y \in R^n : y \geq 0; \sum_{j=1}^n y_j = 1\}$$

The (expected) payoff to player 1 (from player 2) given that player 1 selects $x \in X$ and player 2 selects $y \in Y$ is given by

$$x^t A y = \sum_{i=1}^m \sum_{j=1}^n a_{i,j} x_i y_j$$

Guaranteed Payoffs and Minimax Strategies Given a strategy $x \in X$ for player 1, we call

$$v(x) = \min_{y \in Y} x^t A y = \min_{j=1}^n x^t A_{.j} = \min_{j=1}^n \sum_{i=1}^m a_{i,j} x_i$$

the guaranteed payoff to player 1.

$$v_I = \max_{x \in X} v(x) = \max_{x \in X} \min_{y \in Y} x^t A y$$

is the overall guaranteed payoff to player 1 that he can assure himself against all possibilities for player 2. The strategy that achieves this value is often referred to as the "optimal" strategy for player 1. In a similar manner we can define a strategy for player 2 that maximizes his payoff (or equally in a zero-sum situation one which minimizes payoff to player 1) by solving the problem:

$$v_{II} = \min_{y \in Y} v(y) = \min_{y \in Y} \max_{x \in X} x^t A y$$

where

$$v(y) = \max_{x \in X} x^t A y = \max_{i=1}^m A_{i.} y = \max_{i=1}^m \sum_{j=1}^n a_{i,j} y_j$$

For any function $f(x, y)$ defined over $(x, y) \in X \times Y$ we have:

$$\max_{x \in X} \min_{y \in Y} f(x, y) \leq \min_{y \in Y} \max_{x \in X} f(x, y)$$

and hence $v_I \leq v_{II}$.

In order to compute these "two" values, we resort to the following linear programming problems; this also gives us the set of "optimal" strategies (this is where the name comes from):

$$\begin{aligned} x_i &\geq 0; & i &= 1, 2, \dots, m \\ \sum_{i=1}^m x_i &= 1 \\ \sum_{i=1}^m a_{i,j} x_i &\geq v_I & j &= 1, 2, \dots, n \\ &\max v_I \end{aligned}$$

Here v_I is an unrestricted variable (it can take on positive or negative values). This is for player 1. And for player 2:

$$\begin{aligned}
 & y_j \geq 0; \quad j = 1, 2, \dots, n \\
 & \sum_{j=1}^n y_j = 1 \\
 & \sum_{j=1}^n a_{i,j} y_j \leq v_{II} \quad i = 1, 2, \dots, m \\
 & \min v_{II}
 \end{aligned}$$

It is straight forward to show that these two linear programs are dual to each other and that both are feasible. Hence by duality theorem, they both have optimal solutions (in the linear programming sense) and the values v_I and v_{II} are equal at optimality. That is :

$$\min_{y \in Y} \max_{x \in X} x^t A y = \max_{x \in X} \min_{y \in Y} x^t A y$$

and this result is known as the **minimax theorem** due to J. von Neumann. Thus, every two person zero-sum game reduces to a linear program and this can be used as a computational model to find optimal strategies for the players in the game. [Actually the converse is also true that every linear program can be converted to a two person zero-sum game and in this sense these two problems are equivalent.] Further, their connection with Nash equilibrium is quite strong as illustrated by:

Theorem 10 *Suppose x^* and y^* are optimal solutions to the above pair of linear programs. Then (i) (x^*, y^*) form a Nash equilibrium pair of strategies for this game; (ii) Moreover, every pair of Nash equilibrium strategies corresponds to a pair of optimal solutions to this pair of linear programs; (iii) If (x^*, y^*) and (x^{**}, y^{**}) are equilibrium pairs, then so are (x^*, y^{**}) and (x^{**}, y^*) also equilibrium pairs; (iv) all equilibria have the same payoffs.*

Proof. (i) Since x^* guarantees v_I^* , and y^* guarantees (a loss of no more than)

$v_{II}^* = v_I^*$, it follows that x^* is the best response for player I against y^* for player II. Similarly y^* is the best response for player II against x^* for player I. Hence (x^*, y^*) is a Nash equilibrium pair of strategies.

(ii) Suppose (x^*, y^*) is a Nash equilibrium. Then it follows that

$$\begin{aligned} x^t Ay^* &\leq (x^*)^t Ay^* & \forall x \in X \\ (x^*)^t Ay &\geq (x^*)^t Ay^* & \forall y \in Y \end{aligned}$$

Hence it follows that

$$\begin{aligned} v_{II}^* &= (x^*)^t Ay^* \geq \sum_{j=1}^n a_{i,j} y_j^* & \text{for } i = 1, 2, \dots, m \\ \sum_{j=1}^n y_j^* &= 1 \\ y_j^* &\geq 0 & \text{for } j = 1, 2, \dots, n \end{aligned}$$

Hence (y^*, v_{II}^*) is a feasible solution to the dual (Player II's LP). Similarly, (x^*, v_I^*) is a feasible solution to Primal (Player I's LP). Moreover, $v_I^* = v_{II}^*$ and hence these solutions are respectively optimal to the two LPs.

(iii) Suppose (x^*, y^*) and (x^{**}, y^{**}) are Nash equilibria. By (ii) both x^* and x^{**} are optimal solutions to player I's LP. Hence both guarantee v_I^* . Similarly both y^* and y^{**} guarantee to player II a loss of v_{II}^* . And $v_I^* = v_{II}^*$. Hence x^{**} is a best response to y^* and y^* is a best response to x^{**} . So (x^{**}, y^*) is also a Nash equilibrium. Similarly (x^*, y^{**}) is also an equilibrium.

(iv) By above, it follows that $(x^*)^t Ay^{**} = (x^{**})^t Ay^* = v_I^* = v_{II}^*$. ■

This works only for two person zero-sum games. So we consider such a game to be completely solved in some sense. All of these nice properties are not guaranteed to hold for more general games. Moreover, the methods used to prove existence for such games are quite different.

Constrained Two Person Zero-sum Games These are games in which the sets of strategies for the two players are modified as follows:

$$\begin{aligned} X &= \{x \in R^m : x \geq 0; Bx \leq d\} \\ Y &= \{y \in R^n : y \geq 0; E^t y \geq f\} \end{aligned}$$

Player 1 wants to solve the problem $\max_{x \in X} \min_{y \in Y} x^t Ay$ and player 2 wants to solve the problem $\min_{y \in Y} \max_{x \in X} x^t Ay$ to find a pair of strategies that maximize (resp. minimize) their payoffs. Again such a pair forms a pair of strategies in Nash equilibrium. To solve this problem player 1 solves the

following LP: if we let $v(x) = \min_{y \in Y} x^t A y$, then $v(x)$ is the solution to the linear program:

$$\begin{aligned} E^t y &\geq f \\ y &\geq 0 \\ \min(x^t A) y \end{aligned}$$

By duality theorem, this value also equals the value of the linear program:

$$\begin{aligned} E z &\leq A^t x \\ z &\geq 0 \\ \max f^t z \end{aligned}$$

Hence $\max_{x \in X} v(x)$ can be found by solving the linear program:

$$\begin{aligned} B x &\leq d \\ E z - A^t x &\leq 0 \\ x &\geq 0; z \geq 0 \\ \max f^t z \end{aligned}$$

Similarly player 2's problem is $\min_{y \in Y} \max_{x \in X} x^t A y$ and if we define $v(y) = \max_{x \in X} x^t A y$ then $v(y)$ is the solution to the linear program:

$$\begin{aligned} B x &\leq d \\ x &\geq 0 \\ \max(y^t A^t) x \end{aligned}$$

By duality theorem, this value can be found by solving the linear program

$$\begin{aligned} B^t s &\geq A y \\ s &\geq 0 \\ \min d^t s \end{aligned}$$

Hence $\min_{y \in Y} v(y)$ can be solved by

$$\begin{aligned} E^t y &\geq f \\ B^t s - A y &\geq 0 \\ y &\geq 0; s \geq 0 \\ \min d^t s \end{aligned}$$

It should be easy to verify that the final two problems for the two players are dual to each other in the linear programming sense.