

COS 340: Reasoning About Computation*

Game Theory and Linear Programming

Moses Charikar

March 24, 2014

In these notes, we introduce some basic concepts in game theory and linear programming (LP). We show a connection between equilibrium strategies in a certain kind of two player game and LP duality and connect this to the multiplicative weights update algorithm we saw earlier in the context of learning with experts.

1 Game Theory

1.1 Introduction

A *two player game* (or more correctly, a *two player normal-form game*) is specified by two $m \times n$ *payoff matrices* R and C corresponding to the *row* and *column* player respectively. Each of these matrices has m rows corresponding to the m strategies of the row player and n columns corresponding to the n strategies of the column payer. The row player picks a row $i \in [m]$, and the column player picks a column $j \in [n]$. We say that the outcome of the game is (i, j) . The payoff to the row player is $R_{i,j}$ and the payoff to the column player is $C_{i,j}$. The goal of each player is to maximize their own payoff.

The strategies for the row and column player described so far are referred to as *pure* strategies. Both the row and column player may also pick a distribution over rows and columns respectively, referred to as *mixed* strategies. A mixed strategy for the row player can be specified by a vector $\mathbf{x} = (x_1, \dots, x_m)$ such that $x_i \geq 0$ and $\sum_i x_i = 1$. Here the i th coordinate x_i specifies the probability that the (pure) strategy $i \in [m]$ is picked by the row player. The space of mixed strategies for the row player is denoted by Δ_m . This is just the set of vectors $\mathbf{x} = (x_1, \dots, x_m)$ such that $x_i \geq 0$ and $\sum_i x_i = 1$. Similarly, a mixed strategy for the column player is specified by a vector $\mathbf{y} = (y_1, \dots, y_n)$ such that $y_j \geq 0$ and $\sum_j y_j = 1$. The space of mixed strategies for the column player is denoted by Δ_n .

If the row and column player play mixed strategies \mathbf{x}, \mathbf{y} respectively, then the interpretation is that both players pick independently from the probability distributions specified by \mathbf{x} and \mathbf{y} , i.e. the probability that the outcome is (i, j) is $x_i y_j$. The expected payoff to

*Princeton University, Spring 2014

the row player is $\sum_{i,j} R_{i,j}x_iy_j = \mathbf{x}^T R \mathbf{y}$ and the expected payoff to the column player is $\sum_{i,j} C_{i,j}x_iy_j = \mathbf{x}^T C \mathbf{y}$. In this vector notation, a pure strategy for the row player can be represented by a vector e_i , $i \in [m]$, where e_i is a m -dimensional vector such that the i th coordinate is 1 and all other coordinates are zero. Similarly, a pure strategy for the column player can be represented by a vector e_j , $j \in [n]$, where e_j is a n -dimensional vector such that the j th coordinate is 1 and all other coordinates are zero. In these notes, we will reserve \mathbf{x} , e_i and i to refer to strategies of the row player and \mathbf{y} , e_j and j to refer to strategies of the column player.

Given a game, what strategies should we expect the players to play? This is a fundamental question in game theory. Given full information about the other player's strategy, each player attempts to pick a strategy that maximizes his/her payoff.

We say that a pair of (pure) strategies i and j are in *equilibrium* if, when the row player plays i and the column player plays j , neither can get better payoff by unilaterally switching to a different strategy. i.e. $R_{i',j} \leq R_{i,j}$ for all $i' \in [m]$ and $C_{i,j'} \leq C_{i,j}$ for all $j' \in [n]$.

If we have a two player game with a pair of strategies i and j in equilibrium, then it is reasonable to expect that one possible outcome of the game is (i, j) . Unfortunately, not all two player games have such an equilibrium. However, if we broaden our notion of equilibrium to include mixed strategies, then we always have a pair of mixed strategies in equilibrium. This is a consequence of a famous theorem of John Nash which shows that such equilibrium strategies exist in the more general multi-player setting – this is part of the work for which he was awarded the Nobel Prize in Economics in 1994.

We say that a pair of mixed strategies \mathbf{x} and \mathbf{y} are in Nash equilibrium if, when the row player plays \mathbf{x} and the column player plays \mathbf{y} , neither can get better payoff by unilaterally switching to a different strategy, i.e.

$$\mathbf{x}^T R \mathbf{y} \geq (\mathbf{x}')^T R \mathbf{y} \quad \text{for all } \mathbf{x}' \in \Delta_m, \text{ and} \quad (1)$$

$$\mathbf{x}^T C \mathbf{y} \geq \mathbf{x}^T C \mathbf{y}' \quad \text{for all } \mathbf{y}' \in \Delta_n \quad (2)$$

An equivalent definition is the following: A pair of *mixed* strategies \mathbf{x}, \mathbf{y} are in Nash equilibrium if neither player can unilaterally switch to a pure strategy that gives better payoff:

$$\mathbf{x}^T R \mathbf{y} \geq e_i^T R \mathbf{y} \quad \text{for all } i \in [m], \text{ and} \quad (3)$$

$$\mathbf{x}^T C \mathbf{y} \geq \mathbf{x}^T C e_j \quad \text{for all } j \in [n] \quad (4)$$

Note that $e_i^T R \mathbf{y}$ is the expected payoff when the row player plays pure strategy i and the column player plays mixed strategy \mathbf{y} . Thus $e_i^T R \mathbf{y} = \sum_j R_{i,j}y_j$. Similarly $\mathbf{x}^T C e_j$ is the expected payoff when the row player plays mixed strategy \mathbf{x} and the column player plays pure strategy j . Thus $\mathbf{x}^T C e_j = \sum_i R_{i,j}x_i$.

It is a good exercise to prove that the above two definitions are equivalent – we highly recommend this as a way to test your understanding of these definitions!

A game could have multiple such Nash equilibria (\mathbf{x}, \mathbf{y}) , but at least one is guaranteed to exist by Nash's theorem. Given a two player game, we do not know an efficient (polynomial

time) algorithm to find a Nash equilibrium. Designing such an algorithm (or proving that it does not exist) is a significant open problem in theoretical computer science. We have some evidence that suggests that finding equilibrium strategies is a hard problem, however it is not known whether the problem is NP-hard (a concept we will learn about in the last week of the course).

1.2 Examples

We start with a few classic examples of games and discuss their Nash equilibria.

TO BE COMPLETED

1.3 Zero sum games

Two player zero sum games are a special class of two player games where $R + C = 0$, i.e. $R_{i,j} + C_{i,j} = 0$ for all i, j . In this case, the payoffs can be viewed as a payment from one player to the other. The goal of the column player is to maximize the payoff $\mathbf{x}^T C \mathbf{y}$. On the other hand, the goal of the row player is to minimize the payoff $\mathbf{x}^T R \mathbf{y} = -\mathbf{x}^T C \mathbf{y}$. In other words, the goal of the row player is to minimize $\mathbf{x}^T C \mathbf{y}$.

Let's think about the optimization problem faced by the row player: If the row player picks (mixed) strategy \mathbf{x} , the column player will pick (mixed) strategy \mathbf{y} that maximizes $\mathbf{x}^T C \mathbf{y}$. In other words, for a fixed \mathbf{x} , the value to the row player is $\max_{\mathbf{y}} \mathbf{x}^T C \mathbf{y}$. Since the goal of the row player is to minimize $\mathbf{x}^T C \mathbf{y}$, the optimization problem that the row player should attempt to solve is

$$\min_{\mathbf{x}} \max_{\mathbf{y}} \mathbf{x}^T C \mathbf{y} \tag{5}$$

By similar reasoning, the optimization problem that the column player should attempt to solve is

$$\max_{\mathbf{y}} \min_{\mathbf{x}} \mathbf{x}^T C \mathbf{y} \tag{6}$$

In fact, it turns out that the values of the two optimization problems are exactly the same, and this was proved by Von Neumann:

Theorem 1.1 (Von Neumann minimax theorem).

$$\min_{\mathbf{x}} \max_{\mathbf{y}} \mathbf{x}^T C \mathbf{y} = \max_{\mathbf{y}} \min_{\mathbf{x}} \mathbf{x}^T C \mathbf{y} \tag{7}$$

We will prove this theorem later by appealing to linear programming duality, a concept we will introduce shortly. Before we do this, we first discuss the implication for Nash Equilibrium in 2 player zero sum games.

Corollary 1.2. *Every 2 player zero sum game has a Nash equilibrium.*

Proof. This proof looks heavy on notation, but is conceptually quite simple. The main point is that the strategies \mathbf{x}^* and \mathbf{y}^* that optimize the LHS and RHS of the equality (7) are in Nash equilibrium.

Consider a 2 player zero sum game where C is the payoff matrix for the column player. By Theorem 1.1,

$$\min_{\mathbf{x}} \max_{\mathbf{y}} \mathbf{x}^T C \mathbf{y} = \max_{\mathbf{y}} \min_{\mathbf{x}} \mathbf{x}^T C \mathbf{y} \quad (8)$$

Let λ^* be the value of the optimization problem on the LHS and RHS above. Further, let \mathbf{x}^* be the value of \mathbf{x} that achieves $\max_{\mathbf{y}} \mathbf{x}^T C \mathbf{y} = \lambda^*$ and let \mathbf{y}^* be the value of \mathbf{y} that achieves $\min_{\mathbf{x}} \mathbf{x}^T C \mathbf{y} = \lambda^*$. Hence

$$(\mathbf{x}^*)^T C \mathbf{y} \leq \lambda^* \quad \text{for all } \mathbf{y} \in \Delta_n \quad (9)$$

$$\mathbf{x}^T C \mathbf{y}^* \geq \lambda^* \quad \text{for all } \mathbf{x} \in \Delta_m \quad (10)$$

This implies that $(\mathbf{x}^*)^T C \mathbf{y}^* = \lambda^*$. We claim that the pair of strategies $(\mathbf{x}^*, \mathbf{y}^*)$ are in Nash equilibrium. Since $\mathbf{x}^T C \mathbf{y}^* \geq (\mathbf{x}^*)^T C \mathbf{y}^*$ for all strategies $\mathbf{x} \in \Delta_m$, the row player has no incentive to deviate from \mathbf{x}^* . Since $(\mathbf{x}^*)^T C \mathbf{y} \leq (\mathbf{x}^*)^T C \mathbf{y}^*$ for all strategies $\mathbf{y} \in \Delta_n$, the column player has no incentive to deviate from \mathbf{y}^* . Hence the pair of strategies $\mathbf{x}^*, \mathbf{y}^*$ are in Nash equilibrium. \square

In order to prepare for the proof of the minimax theorem, we restate the equality (7) in an equivalent form. This is similar to the alternate definition of Nash equilibrium we mentioned earlier.

First, we claim that for a fixed \mathbf{x} , $\mathbf{x}^T C \mathbf{y}$ is maximized at $\mathbf{y} = e_j$ for $j \in [n]$. In other words, $\max_{\mathbf{y}} \mathbf{x}^T C \mathbf{y} = \max_j \mathbf{x}^T C e_j$. Thus the optimization for the row player can be expressed as $\min_{\mathbf{x}} \max_j \mathbf{x}^T C e_j$. Similarly, for a fixed \mathbf{y} , $\mathbf{x}^T C \mathbf{y}$ is minimized at $\mathbf{x} = e_i$ for $i \in [m]$. In other words, $\min_{\mathbf{x}} \mathbf{x}^T C \mathbf{y} = \min_i e_i^T C \mathbf{y}$. Thus the optimization for the column player can be expressed as $\max_{\mathbf{y}} \min_i e_i^T C \mathbf{y}$. Hence the minimax theorem can be rephrased in the following equivalent form:

$$\min_{\mathbf{x}} \max_j \mathbf{x}^T C e_j = \max_{\mathbf{y}} \min_i e_i^T C \mathbf{y} \quad (11)$$

Alternately:

$$\min_{\mathbf{x}} \max_j \left\{ \sum_i C_{i,j} x_i \right\} = \max_{\mathbf{y}} \min_i \left\{ \sum_j C_{i,j} y_j \right\} \quad (12)$$

The minimax theorem finds many uses in theoretical computer science. In particular, it is the basis of a technique to show lower bounds on the competitive ratio for randomized online algorithms, and lower bounds for randomized algorithms in general. We next introduce linear programming (LP) and then use LP duality to prove the minimax theorem.

2 Linear Programming

A linear program (LP) is a particular kind of optimization problem. Linear programs are very useful because they arise in many different applications and we have efficient algorithms to solve them. A linear program consists of a set of variables with linear constraints on them (either equalities or inequalities where both sides are linear functions of the variables). The objective function is a linear function of the variables. The goal is to assign real values to the variables so as to optimize (either minimize or maximize) the objective function. Linear programming is a very useful tool in optimization because many problems can be phrased as linear programming problems. The optimum solution to a linear program can be computed in time polynomial in the size of the input. We will not actually discuss any algorithms to solve linear programs in these notes (and indeed such algorithms are outside the scope of the course) but we mention that the first polynomial time algorithm to solve linear programs optimally was developed by Khachiyan in 1979.

Here is an example of a linear program

$$\max \quad 3y_1 + 4y_2 + y_3 \quad (13)$$

subject to

$$y_1 + y_2 \leq 5 \quad (14)$$

$$y_2 + y_3 \leq 4 \quad (15)$$

$$y_i \geq 0 \quad (16)$$

A *feasible* solution for this linear program is a setting of values to the variables such that all the constraints are satisfied. An *optimum* solution to the linear program is a feasible solution that optimizes the value of the objective function. Note that the optimum solution need not be unique.

Given a solution to this particular linear program, how do we prove that it is optimum? We can do that by exhibiting an upper bound on the value of the objective function for any feasible solution to the LP. One way to obtain such an upper bound is to multiply the first inequality by 3 and the second by 1 and add them up.

$$3 \times (y_1 + y_2) + (y_2 + y_3) \leq 3 \times 5 + 4 \quad (17)$$

$$3y_1 + 4y_2 + y_3 \leq 19 \quad (18)$$

This gives an upper bound of 19 on the objective function of the LP without actually solving it. In fact this bound is tight. Consider the solution $y_1 = 5$, $y_2 = 0$, $y_3 = 4$. It is easy to check that this satisfies the constraints and the value of the objective function for this solution is 19. We have just proved that 19 is in fact the optimum value of this LP.

Consider now a slight modification of the LP we started out with where we change one

coefficient in the objective function:

$$\begin{aligned} \max \quad & 2y_1 + 4y_2 + y_3 & (19) \\ \text{subject to} \end{aligned}$$

$$y_1 + y_2 \leq 5 \quad (20)$$

$$y_2 + y_3 \leq 4 \quad (21)$$

$$y_i \geq 0 \quad (22)$$

How can we obtain an upper bound for the modified LP? The same bound of 19 derived earlier still holds:

$$2y_1 + 4y_2 + y_3 \leq 3 \times (y_1 + y_2) + (y_2 + y_3) \leq 3 \times 5 + 4 \leq 19 \quad (23)$$

Can we do better? Suppose we multiply the first inequality by $x_1 \geq 0$ and the second inequality by $x_2 \geq 0$ where x_1 and x_2 are unknowns to be determined. (We insist that x_1, x_2 be nonnegative to ensure that multiplying by these values does not reverse the direction of the inequality). What conditions are needed to ensure that we have a valid upper bound on the value of the LP? Here is what we need:

$$2y_1 + 4y_2 + y_3 \leq x_1 \times (y_1 + y_2) + x_2 \times (y_2 + y_3) \leq 5x_1 + 4x_2 \quad (24)$$

For this to be correct, note that the first inequality should hold for every $y_1, y_2, y_3 \geq 0$. One way to guarantee this is to make sure that the coefficient of y_1 on the LHS of the first inequality is \leq then coefficient of y_1 on the RHS and similarly for the coefficients of y_2 and y_3 . This gives three linear inequalities for x_1 and x_2 . The bound that we get from this is $5x_1 + 4x_2$ and the best bound from this method is obtained by minimizing this expression. This gives a (new) linear program on the variables x_1 and x_2 :

$$\min \quad 5x_1 + 4x_2 \quad (25)$$

subject to

$$x_1 \geq 2 \quad (26)$$

$$x_1 + x_2 \geq 4 \quad (27)$$

$$x_2 \geq 1 \quad (28)$$

$$x_j \geq 0 \quad (29)$$

Setting $x_1 = 2, x_2 = 2$ gives a solution that satisfies all the constraints of this new LP. The value of the solution is $5 \times 2 + 4 \times 2 = 18$, which is an upper bound on the value of the LP (19)-(22). In fact, this bound is tight. Consider the following solution to the LP (19)-(22): $y_1 = 1, y_2 = 4, y_3 = 0$. The value of this solution is $2y_1 + 4y_2 + y_3 = 18$. Thus we have proved that 18 is the optimal value of this LP.

In fact, what we have just seen are simple examples of *LP duality* that we now proceed to discuss. The two LPs (19)-(22) and (25)-(29) are duals of each other.

More generally, any linear program can be expressed in the following form:

$$\max \quad b^T \mathbf{y} \tag{30}$$

$$A\mathbf{y} \leq \mathbf{d} \tag{31}$$

$$\mathbf{y} \geq 0 \tag{32}$$

Here $\mathbf{y} = (y_1, \dots, y_n)$ is a column vector of variables and $b = (b_1, \dots, b_n)$ is a column vector of coefficients (i.e. real numbers). The objective function is $\max \sum_j b_j y_j$. A is an $m \times n$ matrix of coefficients and $\mathbf{d} = (d_1, \dots, d_m)$ is a column vector whose entries are real numbers. The linear program has m constraints – the i th constraint is given by $\sum_j A_{i,j} y_j \leq d_i$. All variables y_j are constrained to be non-negative.

We mentioned earlier that linear programs can be solved in polynomial time. In fact, there are three possibilities for a linear program:

- (1) the optimum value is bounded,
- (2) the optimum value is unbounded, or
- (3) there is no feasible solution.

In time polynomial in the size of a given LP, we can decide which of the three cases this LP falls into. In case the optimum is bounded, we can obtain an optimal solution in polynomial time.

We refer to the original linear program as the *primal* LP. The dual of this linear program is the following:

$$\min \quad d^T \mathbf{x} \tag{33}$$

$$A^T \mathbf{x} \geq b \tag{34}$$

$$\mathbf{x} \geq 0 \tag{35}$$

Here $\mathbf{x} = (x_1, \dots, x_m)$ is a vector of variables. A is the same matrix and b, d are the column vectors that appear in the description of the primal LP. Note that we have a variable x_i in the dual corresponding to the i th constraint in the primal. Think of the x_i variables as multipliers for the constraints in the primal LP. We would like to multiply the i th constraint of the primal by x_i and sum over all constraints to obtain a valid upper bound on the value of the primal. The bound that we obtain thus is $\sum_i d_i x_i$ which is the objective function of the dual LP. We have a constraint in the dual corresponding to every variable in the primal. Since we multiply the i th primal constraint by x_i and add up over all i , the coefficient of y_j in this linear combination is $\sum_i A_{i,j} x_i$. This ought to be \geq the coefficient of y_j in the primal objective, which is b_j . This gives the following constraint: $\sum_i A_{i,j} x_i \geq b_j$ which is exactly the j th constraint in the dual LP. The dual consists of all such constraints for $j \in [n]$.

We state two lemmas that relate the values of the primal and dual LPs:

Lemma 2.1 (Weak Duality). *For any feasible solution \mathbf{x} to the LP (33)-(35) and \mathbf{y} to the LP (30)-(32), we have $b^T \mathbf{y} \leq d^T \mathbf{x}$.*

In fact, we have already presented all the ideas needed to prove this lemma. Prove this as an exercise to test your understanding!

Lemma 2.2 (Strong Duality). *For any optimal solution \mathbf{x}^* to the LP (33)-(35) and optimal solution \mathbf{y}^* to the LP (30)-(32), we have $b^T \mathbf{y}^* = d^T \mathbf{x}^*$.*

Duality is a very useful property of linear programs. We have already seen how to exploit duality to prove that a solution to a linear program is optimal: simply exhibit an optimal solution to the dual linear program. We won't prove LP duality here. We will use it to prove the minimax theorem stated earlier.

In general, the form of LP duality we have stated here is general enough to write down the dual for any linear program. However, it is convenient to remember a few more rules that make the job of writing the dual easier in some cases – these can be derived from the form that we described earlier. Suppose the the primal LP has some equality constraints. In principle, each equality constraint $\sum_i A_{i,j} y_j = d_i$ can be written down as two inequalities: $\sum_j A_{i,j} y_j \geq d_i$ and $-\sum_j A_{i,j} y_j \geq -d_i$. Let $x_i^{(+)}$ and $x_i^{(-)}$ be the two corresponding dual variables. It is easier to write down the dual directly by using the variable $x_i = x_i^{(+)} - x_i^{(-)}$. In this case however, the dual variable x_i is unconstrained, i.e. can be either positive, negative or zero. Let's try to get some intuition for why this makes sense. Recall the motivation for the dual LP. The x_i variable is a multiplier for the i th constraint in the dual. The reason why we insisted on $x_i \geq 0$ earlier is to ensure that multiplying by x_i does not reverse the inequality. However, if the i th constraint is an equality, we can multiply it by any value – positive or negative – and we still have an equality. Hence the dual variable corresponding to an equality is unconstrained.

Similarly, if we had an unconstrained variable y_j in the primal LP, the corresponding j th constraint in the dual LP is an equality constraint. One way to see this is to convert to an LP where all variables are nonnegative by replacing y_j by $y_j^{(+)} - y_j^{(-)}$ where the two new variables $y_j^{(+)}, y_j^{(-)} \geq 0$. If we go through the calculations, we will see that that the dual constraints corresponding to $y_j^{(+)}$ and $y_j^{(-)}$ are of the form $\sum_i A_{i,j} x_i \geq b_j$ and $-\sum_i A_{i,j} x_i \geq -b_j$. Together they imply the following constraint: $\sum_i A_{i,j} x_i = b_j$, which is easier to write down directly. Again intuitively, it makes sense that the dual constraint corresponding to an unconstrained variable is an equality constraint. Recall the motivation for the dual LP and constraints. Suppose for $i \in [n]$, we multiply the i th constraint in the primal LP by x_i and add them up. The goal is to ensure that the LHS of this linear combination of constraints is an upper bound on the objective function of the primal LP. Then the j th constraint of the dual is obtained by examining the coefficient of primal variable y_j on the LHS of this linear combination and in the objective function of the primal LP. If the primal variable y_j is unconstrained, the only way we can guarantee that the value of the linear combination is an upper bound on the objective function is to ensure that the coefficients of y_j in the linear combination and the objective function match exactly. In other words, the constraint corresponding to the y_j should be an equality.

To summarize, we state the primal and dual LPs again allowing for unconstrained variables and equality constraints in addition to inequality constraints. It is convenient to remember the primal and dual LPs in this form. The primal LP is as follows: (Here

$S \subseteq [m], T \subseteq [n]$).

$$\max \sum_j b_j y_j \tag{36}$$

$$\forall i \in S \quad \sum_j A_{i,j} y_j \leq d_i \tag{37}$$

$$\forall i \in [m] \setminus S \quad \sum_j A_{i,j} y_j = d_i \tag{38}$$

$$\begin{aligned} \forall j \in T & \quad y_j \geq 0 \\ \forall j \in [n] \setminus T & \quad y_j \text{ unconstrained} \end{aligned}$$

The dual LP is as follows:

$$\min \sum_i d_i x_i \tag{39}$$

$$\forall j \in T \quad \sum_i A_{i,j} x_i \geq b_j \tag{40}$$

$$\forall j \in [n] \setminus T \quad \sum_i A_{i,j} x_i = b_j \tag{41}$$

$$\begin{aligned} \forall i \in S & \quad x_i \geq 0 \\ \forall i \in [m] \setminus S & \quad x_i \text{ unconstrained} \end{aligned}$$

Note that unconstrained variables in the primal correspond to equality constraints in the dual, and unconstrained variables in the dual correspond to equality constraints in the primal. We also remark that the dual of the dual is the primal LP.

2.1 Proof of minimax theorem

Recall the equivalent formulation of the minimax theorem we derived earlier:

$$\min_{\mathbf{x}} \max_j \left\{ \sum_i C_{i,j} x_i \right\} = \max_{\mathbf{y}} \min_i \left\{ \sum_j C_{i,j} y_j \right\} \tag{42}$$

In this section, we will prove the above equality. The LHS is the value of the optimization problem faced by the row player and the RHS is the value of the optimization problem faced by the column player. We will formulate both these problems as linear programs and show that they are duals of each other.

First we claim that the following linear program exactly captures the optimization prob-

lem for the row player:

$$\min \quad z \tag{43}$$

$$\forall j \in [n] \quad \sum_{i=1}^m C_{i,j} x_i \leq z \tag{44}$$

$$\sum_i x_i = 1 \tag{45}$$

$$\forall i \in [m] \quad x_i \geq 0$$

$$z \quad \text{unconstrained}$$

Note that the x_i variables are the coordinates of the mixed strategy vector \mathbf{x} used by the row player. For any fixing of the x_i variables, the minimum value of the objective function is achieved for $z = \max_j \{ \sum_i C_{i,j} x_i \}$, so the optimum value of the linear program is indeed the minimum of $\max_j \{ \sum_i C_{i,j} x_i \}$ over all choices of mixed strategies \mathbf{x} .

By similar reasoning, we claim that the following linear program exactly captures the optimization problem for the column player:

$$\max \quad w \tag{46}$$

$$\forall i \in [m] \quad \sum_{j=1}^n C_{i,j} y_j \geq w \tag{47}$$

$$\sum_j y_j = 1 \tag{48}$$

$$\forall j \in [n] \quad y_j \geq 0$$

$$w \quad \text{unconstrained}$$

Here the y_j variables are the coordinates of the mixed strategy vector \mathbf{y} used by the row player. For any fixing of the y_j variables, the maximum value of the objective function is achieved for $w = \min_i \{ \sum_j C_{i,j} y_j \}$, so the optimum value of the linear program is indeed the maximum of $\min_i \{ \sum_j C_{i,j} y_j \}$ over all choices of mixed strategies \mathbf{y} .

Let's reformulate the linear program (46)-(48) in the form of the LP (36)-(38):

$$\max \quad w \tag{49}$$

$$\forall i \in [m] \quad w - \sum_{j=1}^n C_{i,j} y_j \leq 0 \tag{50}$$

$$\sum_j y_j = 1 \tag{51}$$

$$\forall j \in [n] \quad y_j \geq 0$$

$$w \quad \text{unconstrained}$$

Now, let's write down the dual of this LP, using variables x_i , $i \in [m]$, for the constraints (50), and variable z for the constraint (51). Corresponding to each y_j , $j \in [n]$, we have a dual constraint which is an inequality. Corresponding to w , we have a dual constraint which is an equality. The dual LP is

$$\min \quad z \tag{52}$$

$$\forall j \in [n] \quad z - \sum_{i=1}^m C_{i,j} x_i \geq 0 \tag{53}$$

$$\sum_i x_i = 1 \tag{54}$$

$$\forall i \in [m] \quad x_i \geq 0$$

$$z \quad \text{unconstrained}$$

It is easy to see that the dual LP we derived above is exactly the same as the LP (43)-(45) we wrote down earlier to capture the optimization problem for the row player. Thus the linear programs we wrote down for the row and column player are duals of each other. By strong duality, the values of their optimum solutions are equal. Now let's go back to the reformulation (42) of the minimax theorem we stated at the beginning of this section. The optimum value of the row player's LP is the LHS of (42) and the optimum value of the column player's LP is the RHS of (42). Hence we have proved that both sides are equal, establishing the minimax theorem.