Duality

Primal program P:

Maximize $\sum_{j=1}^{n} c_j x_j$ **subject to**

$$\begin{array}{rcll} \sum_{j=1}^{n} a_{ij} x_j & \leq & b_i, & i = 1, 2, \ldots, m \\ x_j & \geq & 0, & j = 1, 2, \ldots, n \end{array}$$

Dual program D:

Minimize $\sum_{i=1}^{m} b_i y_i$ **subject to**

$$\begin{array}{rcll} \sum_{i=1}^{m} a_{ij} y_i & \geq & c_j, & j = 1, 2, \ldots, n \\ y_i & \geq & 0, & i = 1, 2, \ldots, m \end{array}$$

Weak Duality Theorem

If x is a feasible solution to P and y is a feasible solution to D then the value $c^T x$ is smaller than the value $b^T y$.

The Duality Theorem

Weak Duality Theorem

If x is a feasible solution to P and y is a feasible solution to D then the value $c^T x$ is smaller than the value $b^T y$.

(Strong) Duality Theorem

If P has an optimal solution x^* then D has an optimal solution y^* and $c^T x^* = b^T y^*$.

Complementary Slackness

P: **Maximize** $c^T x$
subject to $Ax \leq b, x \geq 0$

D: **Minimize** $b^T y$
subject to $A^T y \geq c,\ y \geq 0$

Suppose x and y are a pair of optimal solutions to P and D.

By Strong Duality: $c^T x = b^T y$.

Complementary Slackness

P: **Maximize** $c^T x$
subject to $Ax \leq b, x \geq 0$

D: **Minimize** $b^T y$
subject to $A^T y \geq c,\ y \geq 0$

Suppose x and y are a pair of optimal solutions to P and D.

By Strong Duality: $c^T x = b^T y$.

Consider the inequalities from the Weak Duality Theorem:

$$c^T x \leq (y^T A)x = y^T (Ax) \leq y^T b$$

Complementary Slackness

P: **Maximize** $c^T x$
subject to $Ax \leq b, x \geq 0$

D: **Minimize** $b^T y$
subject to $A^T y \geq c,\ y \geq 0$

Suppose x and y are a pair of optimal solutions to P and D.

By Strong Duality: $c^T x = b^T y$.

Consider the inequalities from the Weak Duality Theorem:

$$c^T x \leq (y^T A)x = y^T(Ax) \leq y^T b$$

We must have all inequalities are in fact *equalities*.

Complementary Slackness (II)

$$\sum_{j=1}^{n} c_j x_j = \sum_{j=1}^{n} \left(\sum_{i=1}^{m} a_{ij} y_i \right) x_j = \sum_{i=1}^{m} \left(\sum_{j=1}^{n} a_{ij} x_j \right) y_i = \sum_{i=1}^{m} b_i y_i$$

Complementary Slackness (II)

$$\sum_{j=1}^{n} c_j x_j = \sum_{j=1}^{n} \left(\sum_{i=1}^{m} a_{ij} y_i \right) x_j = \sum_{i=1}^{m} \left(\sum_{j=1}^{n} a_{ij} x_j \right) y_i = \sum_{i=1}^{m} b_i y_i$$

Consider last part

$$\sum_{i=1}^{m} \left(\sum_{j=1}^{n} a_{ij} x_j \right) y_i = \sum_{i=1}^{m} b_i y_i$$

Complementary Slackness (II)

$$\sum_{j=1}^{n} c_j x_j = \sum_{j=1}^{n} \left(\sum_{i=1}^{m} a_{ij} y_i \right) x_j = \sum_{i=1}^{m} \left(\sum_{j=1}^{n} a_{ij} x_j \right) y_i = \sum_{i=1}^{m} b_i y_i$$

Consider last part

$$\sum_{i=1}^{m} \left(\sum_{j=1}^{n} a_{ij} x_j \right) y_i = \sum_{i=1}^{m} b_i y_i$$

Rewrite as

$$\sum_{i=1}^{m} \left(b_i - \sum_{j=1}^{n} a_{ij} x_j \right) y_i = 0$$

Complementary Slackness (II)

$$\sum_{j=1}^{n} c_j x_j = \sum_{j=1}^{n} \left(\sum_{i=1}^{m} a_{ij} y_i \right) x_j = \sum_{i=1}^{m} \left(\sum_{j=1}^{n} a_{ij} x_j \right) y_i = \sum_{i=1}^{m} b_i y_i$$

Consider last part

$$\sum_{i=1}^{m} \left(\sum_{j=1}^{n} a_{ij} x_j \right) y_i = \sum_{i=1}^{m} b_i y_i$$

Rewrite as

$$\sum_{i=1}^{m} \left(b_i - \sum_{j=1}^{n} a_{ij} x_j \right) y_i = 0$$

This is a sum of non-negative numbers (Why?) summing to 0.

Complementary Slackness (II)

$$\sum_{j=1}^{n} c_j x_j = \sum_{j=1}^{n} \left(\sum_{i=1}^{m} a_{ij} y_i \right) x_j = \sum_{i=1}^{m} \left(\sum_{j=1}^{n} a_{ij} x_j \right) y_i = \sum_{i=1}^{m} b_i y_i$$

Consider last part

$$\sum_{i=1}^{m} \left(\sum_{j=1}^{n} a_{ij} x_j \right) y_i = \sum_{i=1}^{m} b_i y_i$$

Rewrite as

$$\sum_{i=1}^{m} \left(b_i - \sum_{j=1}^{n} a_{ij} x_j \right) y_i = 0$$

This is a sum of non-negative numbers (Why?) summing to 0.

Conclusion: All terms must be 0.

Complementary Slackness (III)

$$\sum_{i=1}^{m} \left(b_i - \sum_{j=1}^{n} a_{ij} x_j \right) y_i = 0$$

All terms must be 0. Thus for all $i = 1, \ldots, m$ we have

$$\left(b_i - \sum_{j=1}^{n} a_{ij} x_j \right) y_i = 0$$

Complementary Slackness (III)

$$\sum_{i=1}^{m} \left(b_i - \sum_{j=1}^{n} a_{ij} x_j \right) y_i = 0$$

All terms must be 0. Thus for all $i = 1, \ldots, m$ we have

$$\left(b_i - \sum_{j=1}^{n} a_{ij} x_j \right) y_i = 0$$

Hence (at least) one of the following must hold:

- $b_i - \sum_{j=1}^{n} a_{ij} x_j = 0$ (Note this is the slack variable x_{n+i})
- $y_i = 0$

The Complementary Slackness Property

P: $\begin{array}{ll} \textbf{Max} & \sum_{j=1}^{n} c_j x_j \\ \textbf{s.t.} & \sum_{j=1}^{n} a_{ij} x_j \leq b_i \\ & x_j \geq 0 \end{array}$

D: $\begin{array}{ll} \textbf{Min} & \sum_{i=1}^{n} b_i y_i \\ \textbf{s.t.} & \sum_{i=1}^{n} a_{ij} y_i \geq c_j \\ & y_i \geq 0 \end{array}$

A pair of feasible solutions x and y are said to satisfy *Complementary Slackness* if and only if

For all $i = 1, \ldots, m$ (at least) one of the following holds:

- $\sum_{j=1}^{n} a_{ij} x_j = b_i$ (The ith primal constraint has zero slack)
- $y_i = 0$ (The ith dual variable is zero)

and for all $j = 1, \ldots, n$ (at least) one of the following holds:

- $\sum_{i=1}^{m} a_{ij} y_i = c_j$ (The jth dual constraint has zero slack)
- $x_j = 0$ (The jth primal variable is zero)

Complementary Slackness Theorem

P: **Maximize** $c^T x$
subject to $Ax \leq b, x \geq 0$

D: **Minimize** $b^T y$
subject to $A^T y \geq c$, $y \geq 0$

Theorem

Let x and y be feasible solutions to P and D. Then

x and y both optimal

if and only if

x and y satisfy complementary slackness.

A diet problem

Buy food items in order to satisfy daily intake of energy, protein and calcium, minimizing the cost.

Food	Serving size	Energy	Protein	Calcium	Price
Oatmeal	28	110	4	2	3
Chicken	100	205	32	12	24
Eggs	2	160	13	54	13
Whole milk	237	160	8	285	9
Cherry pie	170	420	4	22	20
Pork with beans	260	260	14	80	19

Linear Programming Formulation and dual

Minimize $3x_1 + 24x_2 + 13x_3 + 9x_4 + 20x_5 + 19x_6$ **subject to**

$$\begin{array}{rcr} 110x_1+205x_2+160x_3+160x_4+420x_5+260x_6 & \geq & 2000 \\ 4x_1+\ 32x_2+\ 13x_3+\ \ 8x_4+\ \ 4x_5+\ 14x_6 & \geq & 55 \\ 2x_1+\ 12x_2+\ 54x_3+285x_4+\ 22x_5+\ 80x_6 & \geq & 800 \end{array}$$

$x_1 \geq 0, \ldots, x_6 \geq 0$

Maximize $2000y_1 + 55y_2 + 800y_3$ **subject to**

$$\begin{array}{rcrcrcr} 110y_1 & + & 4y_2 & + & 2y_3 & \leq & 3 \\ 205y_1 & + & 32y_2 & + & 12y_3 & \leq & 24 \\ 160y_1 & + & 13y_2 & + & 54y_3 & \leq & 13 \\ 160y_1 & + & 13y_2 & + & 54y_3 & \leq & 9 \\ 420y_1 & + & 4y_2 & + & 22y_3 & \leq & 20 \\ 260y_1 & + & 14y_2 & + & 80y_3 & \leq & 19 \end{array}$$

$y_1 \geq 0, y_2 \geq 0, y_3 \geq 0$

Interpreting the Dual

Maximize $2000y_1 + 55y_2 + 800y_3$ **subject to**

$$\begin{array}{rcrcrcr}
110y_1 & + & 4y_2 & + & 2y_3 & \leq & 3 \\
205y_1 & + & 32y_2 & + & 12y_3 & \leq & 24 \\
160y_1 & + & 13y_2 & + & 54y_3 & \leq & 13 \\
160y_1 & + & 13y_2 & + & 54y_3 & \leq & 9 \\
420y_1 & + & 4y_2 & + & 22y_3 & \leq & 20 \\
260y_1 & + & 14y_2 & + & 80y_3 & \leq & 19
\end{array}$$

$y_1 \geq 0, y_2 \geq 0, y_3 \geq 0$

- How much is a consumer willing to pay for energy, protein and calcium tablets.

Understanding optimality using Complementary Slackness

In optimal solutions we have the

(at least) one of the following holds

- No oatmeal is purchased.
- The price of oatmeal equals the cost of equivalent energy, protein and calcium tablets.

etc...

and dually, (at least) one the following holds

- The price of calcium is zero.
- The amount of calcium purchased equals the daily requirement.

etc...

Bill matching game

Max and Minnie play the following game:

They each, in secret, hide either a one-dollar bill or a hundred-dollar bill (of their own money).

Then the bills are revealed. If they differ, Max gets both. If they are the same, Minnie gets both.

Would you rather be Max or Minnie?

Bill matching game, observations

Many (but not all) people choose to be Max - he only has to bet 1 dollar to possibly win 100 dollars. On the other hand, if he chooses the strategy of betting 1 dollar, a simple counter strategy of Minnie is to also bet 1 dollar.

So who has the advantage and how should the game be played?

How Max should play the game depends on how Minnie is going to play the game. But suppose Max has no clue about that!

Cautious strategy (for Max)

Dont't try to second guess what Minnie might do. Play the game so that the loss is as small as possible *assuming worst case behavior of Minnie* (with negative loss = gain).

Cautious strategy (for Max)

Dont't try to second guess what Minnie might do. Play the game so that the loss is as small as possible *assuming worst case behavior of Minnie* (with negative loss = gain).

This leads Max to bet 1 dollar...

Cautious strategy (for Max)

Dont't try to second guess what Minnie might do. Play the game so that the loss is as small as possible *assuming worst case behavior of Minnie* (with negative loss = gain).

This leads Max to bet 1 dollar...

The cautious strategy for Minnie is also to bet 1 dollar

Cautious strategy (for Max)

Dont't try to second guess what Minnie might do. Play the game so that the loss is as small as possible *assuming worst case behavior of Minnie* (with negative loss = gain).

This leads Max to bet 1 dollar...

The cautious strategy for Minnie is also to bet 1 dollar

... but then Max loses 1 dollar every time he plays with Minnie!

Randomized cautious strategy (for Max)

Play the game in a *randomized* way so that the *expected* loss is as small as possible, *assuming worst case behavior of Minnie*.

A randomized strategy is also called a *Mixed Strategy*

A deterministic strategy is also called a *Pure Strategy*

Randomized cautious strategy (for Max)

Play the game in a *randomized* way so that the *expected* loss is as small as possible, *assuming worst case behavior of Minnie*.

A randomized strategy is also called a *Mixed Strategy*

A deterministic strategy is also called a *Pure Strategy*

Strategy: Bet 1 dollar with probability p and 100 dollars with probability $1 - p$.

How to choose p?

Analysis

If Minnie bets 1 dollar, Max's expected gain is

$$g = p \cdot (-1) + (1 - p) \cdot 1 = 1 - 2p$$

If Minnie bets 100 dollars, Max's expected gain is

$$g = p \cdot 100 + (1 - p) \cdot (-100) = 200p - 100$$

Choose p so that g is maximized, where $g = \min(1 - 2p, 200p - 100)$.

Analysis

If Minnie bets 1 dollar, Max's expected gain is

$$g = p \cdot (-1) + (1-p) \cdot 1 = 1 - 2p$$

If Minnie bets 100 dollars, Max's expected gain is

$$g = p \cdot 100 + (1-p) \cdot (-100) = 200p - 100$$

Choose p so that g is maximized, where $g = \min(1-2p, 200p-100)$.

Solution: $p = \frac{1}{2}$, $g = 0$.

Linear Programming

Finding Max's cautious mixed strategy can be formulated as a linear program.

Find (p, g) maximizing g so that

$$\begin{aligned} p &\geq 0 \\ p &\leq 1 \\ g &\leq 1 - 2p \\ g &\leq 200p - 100 \end{aligned}$$

Are cautious strategies too cautious?

In real life, if you are very timid you may get exploited by bullies. (?)

Suppose both players play cautiously.

Suppose Max *learns* for sure that Minnie play cautiously. Can he then exploit her by deviating from his cautious strategy?

Randomized cautious strategy (for Minnie)

Play game in a *randomized* way so that the *expected* loss is as small as possible, *assuming worst-case behavior of Max*.

Bet 1 dollar with probability q and 100 dollars with probability $1 - q$.

If Max bets 1 dollar, Minnie's expected loss is

$$l = q \cdot (-1) + (1 - q) \cdot 100 = 100 - 101q$$

If Max bets 100 dollars, Minnie's expected loss is

$$l = q \cdot 1 + (1 - q) \cdot (-100) = 101q - 100$$

Choose q so that l is minimized, where $l = max(100 - 101q, 101q - 100)$.

Solution: $q = \frac{100}{101}$, $l = 0$.

Linear Programming

Finding Minnie's cautious mixed strategy can be formulated as a linear program.

Find (q, l) minimizing l so that

$$\begin{aligned} q &\geq 0 \\ q &\leq 1 \\ l &\geq 100 - 101q \\ l &\geq 101q - 100 \end{aligned}$$

0=0

Max's guaranteed lower bound on his expected gain when he plays his cautious mixed strategy is *equal* to Minnie's guaranteed upper bound on her expected loss when she plays *her* cautious mixed strategy.

Thus Max cannot exploit Minnie if he learns that she will play by the cautious strategy. **A priori, this is not obvious - intuitively, the cautious strategies are very timid and "pessimistic".**

Since the bounds are the same, both Max and Minnie can announce their strategies before playing without making the other player wish to change strategy as a result.

The two cautious strategies are together called a **Nash Equilibrium** for the game.

Matrix Games

A Matrix Game is given by *payoff matrix* $A = \{a_{ij}\} \in \mathbf{R}^{n \times m}$.

Row player (Max) chooses $i \in \{1, \ldots, n\}$ (without seeing Minnie's move).

Column player (Minnie) chooses $j \in \{1, \ldots, m\}$ (without seeing Max's move).

Max gains a_{ij} dollars.

Minnie loses a_{ij} dollars.

Matrix games are Zero-Sum as Max gains exactly what Minnie loses.

Bill matching as matrix game

	hide \$1	hide \$100
hide \$1	-1	100
hide \$100	1	-100

Optimal Mixed Strategy (for Max)

Play the game in a *randomized* way so that the *expected* gain is as big as possible, *assuming worst case behavior of Minnie*.

Row Players (Max's) Optimal Mixed Strategy

Optimal mixed strategy and guaranteed lower bound on expected gain for Max is $(p_1, p_2, \ldots, p_n, g)$ which is a solution to the LP:

Find $(p_1, p_2, \ldots, p_n, g)$ maximizing g so that

$$\begin{aligned} p_1, \ldots, p_n &\geq 0 \\ p_1 + p_2 + \cdots + p_n &= 1 \\ \sum_i p_i a_{i1} &\geq g \\ \sum_i p_i a_{i2} &\geq g \\ &\ldots \\ \sum_i p_i a_{im} &\geq g \end{aligned}$$

Optimal g is called the *value* of the game.

Column Players (Minnie's) Optimal Mixed Strategy

Optimal strategy and guaranteed upper bound on expected loss for Minnie is $(q_1, q_2, \ldots, q_m, l)$ which is a solution to the LP:

Find $(q_1, q_2, \ldots, q_m, l)$ minimizing l so that

$$
\begin{aligned}
q_1, \ldots, q_m \geq 0 & \\
q_1 + q_2 + \cdots + q_m &= 1 \\
\sum_j q_j a_{1j} &\leq l \\
\sum_j q_j a_{2j} &\leq l \\
\cdots & \\
\sum_j q_j a_{nj} &\leq l
\end{aligned}
$$

Crucial Observation

Max's program and Minnie's program are each others duals!

Consequence

For any Matrix game, Max's guaranteed lower bound on his expected gain when he plays *his* optimal mixed strategy is *equal* to Minnie's guaranteed upper bound on her expected loss when she plays *her* optimal mixed strategy. They are equal to the value of the game.

A priori, this is not obvious - intuitively, the "optimal" strategies are very timid and "pessimistic", as they are optimal only when assuming a worst case opponent.

Since the bounds are the same, both Max and Minnie can *announce* their strategies before playing *without making the other player wish to change strategy as a result*: The two optimal strategies form a **Nash Equilibrium** for the game.

Remarks

Note: If Max plays by (not necessarily optimal) mixed strategy $p = (p_1, p_2, \dots, p_n)$ and Minnie plays by the mixed strategy $q = (q_1, q_2, \dots, q_m)$ the expected gain of Max is $p^T Aq$.

If the value of the game for Max is v, we have that

$$v = \max_p \min_q p^T Aq$$

Reason: v is Max's best possible expected gain, assuming a worst case move of Minnie and hence also assuming a worst case mixed strategy of Minnie.

von Neuman Min-Max Theorem

For any real Matrix A,

$$\max_p \min_q p^T Aq = \min_q \max_p p^T Aq$$

where p (resp. q) are arbitrary probability distributions on rows (resp. columns) of A.

Fair and symmetric games

A matrix game is called *symmetric* if its matrix is antisymmetric.

Fair and symmetric games

A matrix game is called *symmetric* if its matrix is antisymmetric.

Example: Rock/Scissors/Paper.

Fair and symmetric games

A matrix game is called *symmetric* if its matrix is antisymmetric.

Example: Rock/Scissors/Paper.

A matrix game is called *fair* if its value is 0.

Fair and symmetric games

A matrix game is called *symmetric* if its matrix is antisymmetric.

Example: Rock/Scissors/Paper.

A matrix game is called *fair* if its value is 0.

Theorem: All symmetric games are fair.

Fair and symmetric games

A matrix game is called *symmetric* if its matrix is antisymmetric.

Example: Rock/Scissors/Paper.

A matrix game is called *fair* if its value is 0.

Theorem: All symmetric games are fair.

The opposite is not necessarily true. Example?

Fair and symmetric games

A matrix game is called *symmetric* if its matrix is antisymmetric.

Example: Rock/Scissors/Paper.

A matrix game is called *fair* if its value is 0.

Theorem: All symmetric games are fair.

The opposite is not necessarily true. Example?

The bill matching game.

“Should” one use the optimal strategy?

The optimal mixed strategy for Max is the optimal strategy *for the situation where he has no information about how Minnie will play and assume the worst possible behavior*.

Is this a realistic assumption?

"Should" one use the optimal strategy?

The optimal mixed strategy for Max is the optimal strategy *for the situation where he has no information about how Minnie will play and assume the worst possible behavior*.

Is this a realistic assumption?

In actual situations, one may have such (reliable) information and one can then modify the play. *But because of the duality theorem, this will only be useful for situations where Minnie will be doing worse than she could by her optimal mixed strategy*. The only Nash Equilibria for the game are the ones where both players play an optimal mixed strategy.

How about exploiting non-optimal opponents

Can we play an optimal strategy and still exploit opponents that do *not* play by the optimal strategy? (Exploit = Achieve a better expectation that the value of the game)

Can we do this in the Bill matching game?

Bill matching game

Value of the game is 0

Optimal strategy for Max: Bet 1 with probability $\frac{1}{2}$, bet 100 with probability $\frac{1}{2}$.

Optimal strategy for Minnie: Bet 1 with probability $\frac{100}{101}$, bet 100 with probability $\frac{1}{101}$.

How about exploiting non-optimal opponents

Can we play an optimal strategy and still exploit opponents that do *not* play by the optimal strategy? (Exploit = Achieve a better expectation that the value of the game)

Can we do this in the Bill matching game?

No!

How about exploiting non-optimal opponents

Can we play an optimal strategy and still exploit opponents that do *not* play by the optimal strategy? (Exploit = Achieve a better expectation that the value of the game)

Can we do this in the Bill matching game?

No!

In general, if *and only if* our opponent plays a pure strategy that should be played with probability 0 in *all* of his optimal mixed strategies, we *can* exploit him without deviating from an optimal strategy ourselves.

How about exploiting non-optimal opponents

Can we play an optimal strategy and still exploit opponents that do *not* play by the optimal strategy? (Exploit = Achieve a better expectation that the value of the game)

Can we do this in the Bill matching game?

No!

In general, if *and only if* our opponent plays a pure strategy that should be played with probability 0 in *all* of his optimal mixed strategies, we *can* exploit him without deviating from an optimal strategy ourselves.

Problem 3 of the March exam of 2005.

Beyond one move games - Kuhn's three-card poker

Two players, Max and Minnie, each pays a dollar to get dealt a card from a 3-card deck with an Ace, a King and a Queen.

Two players, Max and Minnie, each pays a dollar to get dealt a card from a 3-card deck with an Ace, a King and a Queen.

Max either *checks* or *bets* an additional dollar.

Beyond one move games - Kuhn's three-card poker

Two players, Max and Minnie, each pays a dollar to get dealt a card from a 3-card deck with an Ace, a King and a Queen.

Max either *checks* or *bets* an additional dollar. If Max checks, Minnie may either check or bet. If Max bets, Minnie may either *fold* or *call*.

Beyond one move games - Kuhn's three-card poker

Two players, Max and Minnie, each pays a dollar to get dealt a card from a 3-card deck with an Ace, a King and a Queen.

Max either *checks* or *bets* an additional dollar. If Max checks, Minnie may either check or bet. If Max bets, Minnie may either *fold* or *call*. If Minnie bets, Max may either fold or call.

Two players, Max and Minnie, each pays a dollar to get dealt a card from a 3-card deck with an Ace, a King and a Queen.

Max either *checks* or *bets* an additional dollar. If Max checks, Minnie may either check or bet. If Max bets, Minnie may either *fold* or *call*. If Minnie bets, Max may either fold or call.

If the last player to act folds, the player betting wins the pot. If the last player to act calls, the player with the highest card wins the pot.

Analyzing three-card poker

In three-card poker, the players do *not* act simultaneously and a random card is dealt to them.

How does the theory apply to this kind of game?

Analyzing three-card poker

In three-card poker, the players do *not* act simultaneously and a random card is dealt to them.

How does the theory apply to this kind of game?

A pure strategy may be described by an instruction for what to do in every possible situation that may arise during the game.

A pure strategy for Max

If I get an Ace, I shall bet.

If I get a King, I shall check. If Minnie then bets, I shall call.

If I get a Queen, I shall check. If Minnie then bets I shall fold.

A pure strategy for Max

If I get an Ace, I shall bet.

If I get a King, I shall check. If Minnie then bets, I shall call.

If I get a Queen, I shall check. If Minnie then bets I shall fold.

A pure strategy for Max (or Minnie) is a *tree structure.*

Matrix game formulation of three-card poker

For each pure strategy s for Max and pure strategy t for Minnie, define a_{st} as the *expected* (under the random cards dealt) *gain for Max when these two pure strategies are used*.

A mixed strategy is a probability distribution over all pure strategies.

Game matrix: 27×64 matrix.

Unique optimal mixed strategy for Max

If I get the ace, with probability $\frac{1}{2}$ I shall bet. With probability $\frac{1}{2}$, I shall check and call if Minnie bets.

Unique optimal mixed strategy for Max

If I get the ace, with probability $\frac{1}{2}$ I shall bet. With probability $\frac{1}{2}$, I shall check and call if Minnie bets. **Slowplaying!**

Unique optimal mixed strategy for Max

If I get the ace, with probability $\frac{1}{2}$ I shall bet. With probability $\frac{1}{2}$, I shall check and call if Minnie bets. **Slowplaying!**

If I get the king, I shall check. If Minnie bets, I shall call with probability $\frac{1}{2}$ and fold with probability $\frac{1}{2}$.

Unique optimal mixed strategy for Max

If I get the ace, with probability $\frac{1}{2}$ I shall bet. With probability $\frac{1}{2}$, I shall check and call if Minnie bets. **Slowplaying!**

If I get the king, I shall check. If Minnie bets, I shall call with probability $\frac{1}{2}$ and fold with probability $\frac{1}{2}$.

If I get a queen, I shall check with probability $\frac{5}{6}$. If I check and Minnie bets, I fold. But with probability $\frac{1}{6}$ I shall start by betting.

Unique optimal mixed strategy for Max

If I get the ace, with probability $\frac{1}{2}$ I shall bet. With probability $\frac{1}{2}$, I shall check and call if Minnie bets. **Slowplaying!**

If I get the king, I shall check. If Minnie bets, I shall call with probability $\frac{1}{2}$ and fold with probability $\frac{1}{2}$.

If I get a queen, I shall check with probability $\frac{5}{6}$. If I check and Minnie bets, I fold. But with probability $\frac{1}{6}$ I shall start by betting. **Bluffing!**

Unique optimal mixed strategy for Max

If I get the ace, with probability $\frac{1}{2}$ I shall bet. With probability $\frac{1}{2}$, I shall check and call if Minnie bets. **Slowplaying!**

If I get the king, I shall check. If Minnie bets, I shall call with probability $\frac{1}{2}$ and fold with probability $\frac{1}{2}$.

If I get a queen, I shall check with probability $\frac{5}{6}$. If I check and Minnie bets, I fold. But with probability $\frac{1}{6}$ I shall start by betting. **Bluffing!**

Value of the game for Max: $-\frac{1}{18}$ (Intuitive reason: Minnie has positional advantage.)

Conclusions for poker

Concealing your hand by bluffing and slowplaying is necessary rational behavior - even against completely rational opponents (who knows that you are rational).

General conclusions

The theory of matrix games can be used to analyze games where several rounds of actions are made, information is leaked "bit by bit" and external randomness enters the game.

General conclusions

The theory of matrix games can be used to analyze games where several rounds of actions are made, information is leaked “bit by bit” and external randomness enters the game.

(... but the game matrix may become very large. The idea described gives a matrix with size *exponential* in the size of the *game tree*.)

General conclusions

The theory of matrix games can be used to analyze games where several rounds of actions are made, information is leaked “bit by bit” and external randomness enters the game.

(... but the game matrix may become very large. The idea described gives a matrix with size *exponential* in the size of the *game tree*.)

Much more on this and a remedy in the course “Computational Game Theory”!

Analysis of Algorithms and von Neuman Min-Max Theorem

Consider a *randomized* algorithm as follows:

- Allowed to “flip coins” and continue execution based on the outcome.
- Must always provide a correct answer when terminating (Las Vegas algorithm).

Reasonable measure of complexity?

Analysis of Algorithms and von Neuman Min-Max Theorem

Consider a *randomized* algorithm as follows:

- Allowed to "flip coins" and continue execution based on the outcome.
- Must always provide a correct answer when terminating (Las Vegas algorithm).

Reasonable measure of complexity?

Expected execution time, on a worst case instance.

Analysis of Algorithms and von Neuman Min-Max Theorem

Consider a *randomized* algorithm as follows:

- Allowed to "flip coins" and continue execution based on the outcome.
- Must always provide a correct answer when terminating (Las Vegas algorithm).

Reasonable measure of complexity?

Expected execution time, on a worst case instance.

How to analyse this?

Formulation as a Game

We may view a randomized algorithm as a probability distribution on deterministic algorithms.

Assumption:

- Only a finite number of possible deterministic algorithms.
- Only a finite number of possible inputs.

Formulation as a Game

We may view a randomized algorithm as a probability distribution on deterministic algorithms.

Assumption:

- Only a finite number of possible deterministic algorithms.
- Only a finite number of possible inputs.

Consider the following game:

- Max chooses an input.
- Minnie chooses an algorithm.
- Minnie pays Max for every unit of execution time.

von Neuman Min-Max Theorem

For any real Matrix A,

$$\max_p \min_q p^T A q = \min_q \max_p p^T A q$$

where p (resp. q) are arbitrary probability distributions on rows (resp. columns) of A.

von Neuman Min-Max Theorem

For any real Matrix A,

$$\max_p \min_q p^T Aq = \min_q \max_p p^T Aq$$

where p (resp. q) are arbitrary probability distributions on rows (resp. columns) of A.

- Rows of A are indexed by inputs; columns of A are indexed by algorithms; entries are complexities.
- Max and Minnie have optimal strategies.

Yao's Principle

von Neuman Min-Max Theorem

For any real Matrix A,

$$\max_p \min_q p^T Aq = \min_q \max_p p^T Aq$$

where p (resp. q) are arbitrary probability distributions on rows (resp. columns) of A.

- Rows of A are indexed by inputs; columns of A are indexed by algorithms; entries are complexities.
- Max and Minnie have optimal strategies.

Yao's Principle: The expected complexity of a the best randomized algorithm on a worst-case input *equals* the expected complexity for a worst-case distribution of inputs using the best deterministic algorithm.

Using Yao's Principle

- Suppose we wish to prove a lower bound on the complexity of a randomized algorithm.

Using Yao's Principle

- Suppose we wish to prove a lower bound on the complexity of a randomized algorithm.
- By Yao's principle we may instead construct a "hard" input distribution and analyse a deterministic algorithm on this (A seamingly easier task).

Using Yao's Principle

- Suppose we wish to prove a lower bound on the complexity of a randomized algorithm.
- By Yao's principle we may instead construct a "hard" input distribution and analyse a deterministic algorithm on this (A seamingly easier task).
- The best lower bound possible *can* be obtained this way!

Using Yao's Principle

- Suppose we wish to prove a lower bound on the complexity of a randomized algorithm.
- By Yao's principle we may instead construct a "hard" input distribution and analyse a deterministic algorithm on this (A seamingly easier task).
- The best lower bound possible *can* be obtained this way!

Much more on this in the course "Randomized Algorithms"!

Example: Sorting

How many comparisons must be used to sort n elements?

Example: Sorting

How many comparisons must be used to sort n elements? $\Omega(n \log n)$

What about a randomized algorithm?

Example: Sorting

How many comparisons must be used to sort n elements? $\Omega(n \log n)$

What about a randomized algorithm?

By Yao's priciple we may analyse a best deterministic algorithm on a chosen input distribution.

Example: Sorting

How many comparisons must be used to sort n elements? $\Omega(n \log n)$

What about a randomized algorithm?

By Yao's priciple we may analyse a best deterministic algorithm on a chosen input distribution.

We know that a deterministic sorting algorithm must use $\Omega(n \log n)$ comparisons in the worst case.

The same proof shows also that $\Omega(n \log n)$ comparisons are needed on expectation to sort a permutation selected uniformly at random.