CMPSCI 240: Reasoning about Uncertainty

Lecture 21: Game Theory

Andrew McGregor

University of Massachusetts

Last Compiled: April 29, 2017

Outline

Example: Two-finger Morra

- Alice and Bob play a game
- Simultaneously Alice picks $a \in \{1, 2\}$ and Bob picks $b \in \{1, 2\}$
- Bob pays Alice \$$(a + b)$ if $a + b$ is even
- Alice pays Bob \$$(a + b)$ if $a + b$ is odd

Obviously if Bob always plays the same number, Alice can take advantage of this. What if Bob plays different numbers with different probabilities?

Analysis of Two-finger Morra (1/3)

- Suppose Bob plays "1" with prob. q and "2" with prob. $1-q$
- If Alice plays "1" then Bob has expected return

$$-2q + 3(1-q) = -5q + 3$$

- If Alice plays "2" then Bob has expected return

$$3q - 4(1-q) = -4 + 7q$$

- Then no matter what Alice does, Bob expects to get

$$\min(-5q + 3, -4 + 7q)$$

- If Bob sets $q = 7/12$ his expected earning is at least:

$$\min(-5 \times 7/12 + 3, -4 + 7 \times 7/12) = 1/12$$

Analysis of Two-finger Morra (2/3)

- Suppose Alice plays "1" with prob. p and "2" with prob. $1-p$
- If Bob plays "1" then Alice expects to get

$$2p - 3(1-p) = -3 + 5p$$

- If Bob plays "2" then Alice expects to get

$$-3p + 4(1-p) = -7p + 4$$

- Then no matter what Bob does, Alice expects to get

$$\min(-3 + 5p, -7p + 4)$$

- If Alice sets $p = 7/12$ her expected earning is at least:

$$\min(-3 + 5 \times 7/12, -7 \times 7/12 + 4) = -1/12$$

Analysis of Two-finger Morra (3/3)

Conclusion:

- Alice's expected earning plus Bob's expected earnings always sum up to 0.
- Bob can ensure his expected earning is at least $1/12$; so Alice's expected earning is at best $-1/12$
- Alice can ensure his expected earning is at least $-1/12$; so Bob's expected earning is at best $1/12$.
- Hence the strategies that ensured Bob got $1/12$ in expectation and Alice got $-1/12$ in expectation are optimal, i.e., both players is play 1 finger with probability $7/12$ and 2 fingers with probability $5/12$.

Zero-Sum Games

Definition

A *two-player, simultaneous-move, zero-sum game* consists of a set of k options for player A, a set of options ℓ for player B, and a $k \times \ell$ *payoff matrix* P. If A is chooses her ith option and B chooses his jth option then A gets P_{ij} and B gets $-P_{ij}$.

For two-finger Morra, the payoff matrix is

	1 B Finger	2 B Finger
1 A Finger	$+2$	-3
2 A Finger	-3	$+4$

where best strategy was for players to show 1 finger with probability $7/12$ and two fingers with probability $5/12$.

Pure and Mixed Strategies

Definition

If a player picks one of their options, we call it a pure strategy. If they pick a distribution over their options, we call it a mixed strategy. If one option is better than the other no matter what the other player does, we say the first strategy dominates the second.

Example: Two-finger Morra with Slightly Different Pay-off

	1 B Finger	2 B Finger
1 A Finger	$+2$	-3
2 A Finger	-2	$+4$

- If Alice plays "1" with prob. p and "2" with prob. $1-p$ then her expected payoff is at least:

$$\min(2p - 2(1-p), -3p + 4(1-p))$$

  which is at least $2/11$ when $p = 6/11$.

- If Bob plays "1" with prob. q and "2" with prob. $1-q$ then his expected payoff is at least:

$$\min(-2q + 3(1-q), 2q - 4(1-q))$$

  which is at least $-2/11$ when $q = 7/11$.

- Hence, Alice's best strategy is one finger with probability $6/11$ and Bob's best strategy is one finger with probability $7/11$.

Example: Three-finger Morra

- Alice and Bob play a game
- Simultaneously Alice picks $a \in \{1, 2, 3\}$ and Bob picks $b \in \{1, 2, 3\}$
- Bob pays Alice \$$(a + b)$ if $a + b$ is even
- Alice pays Bob \$$(a + b)$ if $a + b$ is odd
- The payoff matrix is

	1 B Finger	2 B Finger	3 B Finger
1 A Finger	$+2$	-3	$+4$
2 A Finger	-3	$+4$	-5
3 A Finger	$+4$	-5	$+6$

Analysis of Three-finger Morra (1/2)

- Suppose A plays "1" with probability r, "2" with probability s, and "3" with probability $1 - r - s$
- If B plays "1" then A's expected reward is

$$2r - 3s + 4(1 - r - s) = 4 - 2r - 7s$$

- If B plays "2" then A's expected reward is

$$-3r + 4s - 5(1 - r - s) = -5 + 2r + 9s$$

- If B plays "3" then A's expected reward is

$$4r - 5s + 6(1 - r - s) = 6 - 2r - 11s$$

- Hence, for $r = 1/4, s = 1/2$, A gets expected return at least 0

Analysis of Three-finger Morra (2/2)

- Suppose B plays "1" with probability r, "2" with probability s, and "3" with probability $1 - r - s$
- If A plays "1" then B's expected reward is

$$-2r + 3s - 4(1 - r - s) = -4 + 2r + 7s$$

- If A plays "2" then B's expected reward is

$$3r - 4s + 5(1 - r - s) = 5 - 2r - 9s$$

- If A plays "3" then B's expected reward is

$$-4r + 5s - 6(1 - r - s) = -6 + 2r + 11s$$

- Hence, for $r = 1/4, s = 1/2$, B gets expected return at least 0

Hence, best strategy for each player is show 1 finger with probability $1/4$ and 2 fingers with probability $1/2$.

Outline

Prisoner's Dilemma

- Two prisoners are being held pending trial for a crime they are alleged to have committed. The prosecutor offers each a deal:

  "Give evidence against your partner and you'll go free, unless your partner also confesses. If both confess, both get 5 year sentences. If neither confess, both get 1 year sentences. If you don't confess but your partner does, you get 10 years!"

- Can represent this as a game but it's not zero-sum:

	B Confesses	B Stays Mute
A Confesses	$-5, -5$	$0, -10$
A Stays Mute	$-10, 0$	$-1, -1$

In each entry, the first number is A's reward and the second number of B's reward.

Nash Equilibrium

Definition

A Nash Equilibrium is a set of strategies for each player where no change by one player alone can improve his outcome.

For the prisoners dilemma the unique Nash Equilibrium is that both prisoners confess.

Theorem (Nash)

Every game where each player has a finite number of options, has at least one Nash equilibrium.