Minmax and Dominance

CPSC 532A Lecture 6

September 28, 2006

Lecture Overview

What are solution concepts?

- Solution concept: a subset of the outcomes in the game that are somehow interesting.
- There is an implicit computational problem of finding these outcomes given a particular game.
- Depending on the concept, existence can be an issue.

Solution concepts we've seen so far:

What are solution concepts?

- Solution concept: a subset of the outcomes in the game that are somehow interesting.
- There is an implicit computational problem of finding these outcomes given a particular game.
- Depending on the concept, existence can be an issue.

Solution concepts we've seen so far:

- Pareto-optimal outcome
- Pure-strategy Nash equilibrium
- Mixed-strategy Nash equilibrium
- Other Nash variants:
  - weak Nash equilibrium
  - strict Nash equilibrium

Mixed Strategies

- It would be a pretty bad idea to play any deterministic strategy in matching pennies
- Idea: confuse the opponent by playing randomly
- Define a strategy s_i for agent i as any probability distribution over the actions A_i.
  - pure strategy: only one action is played with positive probability
  - mixed strategy: more than one action is played with positive probability
    - these actions are called the support of the mixed strategy
- Let the set of all strategies for i be S_i
- Let the set of all strategy profiles be $S = S_1 \times \ldots \times S_n$.

Best Response and Nash Equilibrium

Our definitions of best response and Nash equilibrium generalize from actions to strategies.

- Best response:
  - $s_i^* \in BR(s_{-i})$ iff $\forall s_i \in S_i$, $u_i(s_i^*, s_{-i}) \geq u_i(s_i, s_{-i})$
- Nash equilibrium:
  - $s = \langle s_1, \ldots, s_n \rangle$ is a Nash equilibrium iff $\forall i$, $s_i \in BR(s_{-i})$
- Every finite game has a Nash equilibrium! [Nash, 1950]
  - e.g., matching pennies: both players play heads/tails 50%/50%

Lecture Overview

Recap

Maxmin and Minmax

Linear Programming

Computing

Fun Game

Domination

Max-Min Strategies

- Player i's maxmin strategy is a strategy that maximizes i's worst-case payoff, in the situation where all the other players (whom we denote $-i$) happen to play the strategies which cause the greatest harm to i.
- The maxmin value (or safety level) of the game for player i is that minimum amount of payoff guaranteed by a maxmin strategy.
- Why would i want to play a maxmin strategy?

Max-Min Strategies

- Player i's maxmin strategy is a strategy that maximizes i's worst-case payoff, in the situation where all the other players (whom we denote $-i$) happen to play the strategies which cause the greatest harm to i.
- The maxmin value (or safety level) of the game for player i is that minimum amount of payoff guaranteed by a maxmin strategy.
- Why would i want to play a maxmin strategy?
  - a conservative agent maximizing worst-case payoff
  - a paranoid agent who believes everyone is out to get him

Definition

The maxmin strategy for player i is $\arg\max_{s_i} \min_{s_{-i}} u_i(s_1, s_2)$, and the maxmin value for player i is $\max_{s_i} \min_{s_{-i}} u_i(s_1, s_2)$.

Min-Max Strategies

- Player i's minmax strategy in a 2-player game is a strategy that minimizes the other player $-i$'s best-case payoff.
- The minmax value of the 2-player game for player i is that maximum amount of payoff that $-i$ could achieve under i's minmax strategy.
- Why would i want to play a minmax strategy?

Min-Max Strategies

- Player i's minmax strategy in a 2-player game is a strategy that minimizes the other player $-i$'s best-case payoff.
- The minmax value of the 2-player game for player i is that maximum amount of payoff that $-i$ could achieve under i's minmax strategy.
- Why would i want to play a minmax strategy?
  - to punish the other agent as much as possible

Definition

The maxmin strategy for player i is $\arg\max_{s_i} \min_{s_{-i}} u_i(s_1, s_2)$, and the maxmin value for player i is $\max_{s_i} \min_{s_{-i}} u_i(s_1, s_2)$.

Definition

In a two-player game, the minmax strategy for player i is $\arg\min_{s_i} \max_{s_{-i}} u_{-i}(s_1, s_2)$, and the minmax value for player i is $\min_{s_i} \max_{s_{-i}} u_{-i}(s_1, s_2)$.

Minmax Theorem

Theorem (Minmax theorem (von Neumann, 1928))

In any finite, two-player, zero-sum game it is the case that:

1. *The maxmin value for one player is equal to the minmax value for the other player. By convention, the maxmin value for player 1 is called the value of the game.*
2. *For both players, the set of maxmin strategies coincides with the set of minmax strategies.*
3. *Any maxmin strategy profile (or, equivalently, minmax strategy profile) is a Nash equilibrium. Furthermore, these are all the Nash equilibria. Consequently, all Nash equilibria have the same payoff vector (namely, those in which player 1 gets the value of the game).*

Lecture Overview

Linear Programming

A linear program is defined by:

- a set of real-valued variables
- a linear objective function
  - a weighted sum of the variables
- a set of linear constraints
  - the requirement that a weighted sum of the variables must be greater than or equal to some constant

Linear Programming

$$
\begin{aligned}
\text{maximize} \quad & \sum_i w_i x_i \\
\text{subject to} \quad & \sum_i w_i^c x_i \geq b^c && \forall c \in C \\
& x_i \geq 0 && \forall x_i \in X
\end{aligned}
$$

- These problems can be solved in polynomial time using interior point methods.
  - Interestingly, the (worst-case exponential) simplex method is often faster in practice.

Lecture Overview

Computing equilibria of zero-sum games

$$
\begin{aligned}
\text{minimize} \quad & U_1^* \\
\text{subject to} \quad & \sum_{a_2 \in A_2} u_1(a_1, a_2) \cdot s_2^{a_2} \leq U_1^* && \forall a_1 \in A_1 \\
& \sum_{a_2 \in A_2} s_2^{a_2} = 1 \\
& s_2^{a_2} \geq 0 && \forall a_2 \in A_2
\end{aligned}
$$

- variables:
  - U_1^* is the expected utility for player 1
  - $s_2^{a_2}$ is player 2's probability of playing action a_2 under his mixed strategy
- each $u_1(a_1, a_2)$ is a constant.

Computing equilibria of zero-sum games

$$
\begin{aligned}
\text{minimize} \quad & U_1^* \\
\text{subject to} \quad & \sum_{a_2 \in A_2} u_1(a_1, a_2) \cdot s_2^{a_2} \leq U_1^* && \forall a_1 \in A_1 \\
& \sum_{a_2 \in A_2} s_2^{a_2} = 1 \\
& s_2^{a_2} \geq 0 && \forall a_2 \in A_2
\end{aligned}
$$

- s_2 is a valid probability distribution.

Computing equilibria of zero-sum games

$$
\begin{aligned}
\text{minimize} \quad & U_1^* \\
\text{subject to} \quad & \sum_{a_2 \in A_2} u_1(a_1, a_2) \cdot s_2^{a_2} \leq U_1^* && \forall a_1 \in A_1 \\
& \sum_{a_2 \in A_2} s_2^{a_2} = 1 \\
& s_2^{a_2} \geq 0 && \forall a_2 \in A_2
\end{aligned}
$$

- U_1^* is as small as possible.

Computing equilibria of zero-sum games

$$
\begin{aligned}
\text{minimize} \quad & U_1^* \\
\text{subject to} \quad & \sum_{a_2 \in A_2} u_1(a_1, a_2) \cdot s_2^{a_2} \leq U_1^* && \forall a_1 \in A_1 \\
& \sum_{a_2 \in A_2} s_2^{a_2} = 1 \\
& s_2^{a_2} \geq 0 && \forall a_2 \in A_2
\end{aligned}
$$

- Player 1's expected utility for playing each of his actions under player 2's mixed strategy is no more than U_1^*.
  - Because U_1^* is minimized, this constraint will be tight for some actions: the support of player 1's mixed strategy.

Computing equilibria of zero-sum games

$$
\begin{aligned}
\text{minimize} \quad & U_1^* \\
\text{subject to} \quad & \sum_{a_2 \in A_2} u_1(a_1, a_2) \cdot s_2^{a_2} \leq U_1^* && \forall a_1 \in A_1 \\
& \sum_{a_2 \in A_2} s_2^{a_2} = 1 \\
& s_2^{a_2} \geq 0 && \forall a_2 \in A_2
\end{aligned}
$$

- This formulation gives us the minmax strategy for player 2.
- To get the minmax strategy for player 1, we need to solve a second (analogous) LP.

Computing Maxmin Strategies in General-Sum Games

Let's say we want to compute a maxmin strategy for player 1 in an arbitrary 2-player game G.

Computing Maxmin Strategies in General-Sum Games

Let's say we want to compute a maxmin strategy for player 1 in an arbitrary 2-player game G.

- Create a new game G' where player 2's payoffs are just the negatives of player 1's payoffs.
- The maxmin strategy for player 1 in G does not depend on player 2's payoffs
  - Thus, the maxmin strategy for player 1 in G is the same as the maxmin strategy for player 1 in G'
- By the minmax theorem, equilibrium strategies for player 1 in G' are equivalent to a maxmin strategies
- Thus, to find a maxmin strategy for G, find an equilibrium strategy for G'.

Lecture Overview

Traveler's Dilemma

Two travelers purchase identical African masks while on a tropical vacation. Their luggage is lost on the return trip, and the airline asks them to make independent claims for compensation. In anticipation of excessive claims, the airline representative announces: "We know that the bags have identical contents, and we will entertain any claim between \$180 and \$300, but you will each be reimbursed at an amount that equals the minimum of the two claims submitted. If the two claims differ, we will also pay a reward R *to the person making the smaller claim and we will deduct a penalty* R *from the reimbursement to the person making the larger claim."*

Traveler's Dilemma

- Action: choose an integer between 180 and 300
- If both players pick the same number, they both get that amount as payoff
- If players pick a different number:
  - the low player gets his number (L) plus some constant R
  - the high player gets $L - R$.
- Play this game *once* with a partner; play with as many different partners as you like.
  - $R = 5$.

Traveler's Dilemma

- Action: choose an integer between 180 and 300
- If both players pick the same number, they both get that amount as payoff
- If players pick a different number:
  - the low player gets his number (L) plus some constant R
  - the high player gets $L - R$.
- Play this game *once* with a partner; play with as many different partners as you like.
  - $R = 5$.
  - $R = 180$.

Traveler's Dilemma

- What is the equilibrium?

Traveler's Dilemma

- What is the equilibrium?
  - $(180, 180)$ is the only equilibrium, for all $R \geq 2$.

Traveler's Dilemma

- What is the equilibrium?
  - $(180, 180)$ is the only equilibrium, for all $R \geq 2$.
- What happens?

Traveler's Dilemma

- What is the equilibrium?
  - $(180, 180)$ is the only equilibrium, for all $R \geq 2$.
- What happens?
  - with $R = 5$ most people choose 295–300
  - with $R = 180$ most people choose 180

Lecture Overview

Domination

- Let s_i and s'_i be two strategies for player i, and let S_{-i} be is the set of all possible strategy profiles for the other players

Definition

s_i strictly dominates s'_i if $\forall s_{-i} \in S_{-i}$, $u_i(s_i, s_{-i}) > u_i(s'_i, s_{-i})$

Definition

s_i weakly dominates s'_i if $\forall s_{-i} \in S_{-i}$, $u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i})$ and $\exists s_{-i} \in S_{-i}$, $u_i(s_i, s_{-i}) > u_i(s'_i, s_{-i})$

Definition

s_i very weakly dominates s'_i if $\forall s_{-i} \in S_{-i}$, $u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i})$

Equilibria and dominance

- If one strategy dominates all others, we say it is dominant.
- A strategy profile consisting of dominant strategies for every player must be a Nash equilibrium.
  - An equilibrium in strictly dominant strategies must be unique.

Equilibria and dominance

- If one strategy dominates all others, we say it is dominant.
- A strategy profile consisting of dominant strategies for every player must be a Nash equilibrium.
  - An equilibrium in strictly dominant strategies must be unique.
- Consider Prisoner's Dilemma again
  - not only is the only equilibrium the only non-Pareto-optimal outcome, but it's also an equilibrium in strictly dominant strategies!

Dominated strategies

- No equilibrium can involve a strictly dominated strategy (why?)

Dominated strategies

- No equilibrium can involve a strictly dominated strategy (why?)
  - Thus we can remove it, and end up with a strategically equivalent game
  - This might allow us to remove another strategy that wasn't dominated before
  - Running this process to termination is called iterated removal of dominated strategies.