15-780 – Numerical Optimization

J. Zico Kolter

April 23, 2014

Outline

Introduction: Some (possibly) familiar games

Basic game theory

Computing equilibria

Special cases and extensions

Prisoner's dilemma

- Two prisoners being interrogated, can either stay silent or implicate the other one

- If both stay silent, each sentenced to a year in jail; if only one implicates another, he goes free and other gets 5 years in jail; if both implicate each other, both get 3 years

	Silent	**Implicate**
Silent	-1,-1	-5,0
Implicate	0,-5	-3,-3

- Even though Silent/Silent is best for both, each one strictly benefits from implicating the other, regardless of other's actions

Guess 2/3 the mean

- All of you will play the game
- Pick a number between 1 and 10 (inclusive)
- The student whose number is *closest to 2/3* of the mean of all the guesses wins (and breaking ties randomly)

General ideas

- Both these games differ slightly from what we have seen so far in class: in order to decide our action we need to account for other agents that are also acting (and trying to account for our actions, ad infinitum)

- We focus here on the special cases of *noncooperative* game theory and games in *normal form*

  - Non-cooperative doesn't mean that agents don't cooperate, just that they are *self-interested*

  - Normal form here just means "one-shot" games, as opposed to turn-based games

Outline

Introduction: Some (possibly) familiar games

Basic game theory

Computing equilibria

Special cases and extensions

Games in normal form

- A *normal form game* is defined by (N, A, u), where
  - N is a number of players, each indexed by i
  - $A = A_1 \times A_2 \times \ldots \times A_n$ is a set of actions, where each A_i is a *finite* set of actions available to player i
  - $u : A \to \mathbb{R}^n$ is a utility function that maps each set of actions $a \in A$ to a set of N utilities, one for each agent; i.e., $u_i(a)$ denotes the utility of agent i for the actions a

- Example: Prisoner's dilemma

	Silent	**Implicate**
Silent	-1,-1	-5,0
Implicate	0,-5	-3,-3

- $N = 2$
- $A = \{\text{Silent}, \text{Implicate}\} \times \{\text{Silent}, \text{Implicate}\}$
- $u(a) = \begin{cases} \begin{bmatrix} -1 \\ -1 \end{bmatrix} & \text{if } a = (\text{Silent}, \text{Silent}) \\ \begin{bmatrix} -5 \\ 0 \end{bmatrix} & \text{if } a = (\text{Silent}, \text{Implicate}) \\ \begin{bmatrix} 0 \\ -5 \end{bmatrix} & \text{if } a = (\text{Implicate}, \text{Silent}) \\ \begin{bmatrix} -3 \\ -3 \end{bmatrix} & \text{if } a = (\text{Implicate}, \text{Implicate}) \end{cases}$

- Example: Guess 2/3 the mean
  - N arbitrary
  - $A = \{1, 2, \ldots, 10\}^N$
  - $u_i(a) = \begin{cases} \frac{1}{\sum_{k=1}^{m} \mathbf{1}\left\{a_k = \left[\frac{2}{3}\mathrm{mean}(a)\right]\right\}} & \text{if } a_i = \left[\frac{2}{3}\mathrm{mean}(a)\right] \\ 0 & \text{otherwise} \end{cases}$
  - Note that utilities here refer to expected utilities: although we are not guaranteed to win if we pick the average, we still have a chance proportional to one over the number of others who pick the same mean

- Battle of the sexes
  - Husband and wife planning movie for the evening: husband wants to see Wondrous Love (WL) wife wants to see Lethal Weapon (LW)
  - Different utilities for each movie, but both equally unhappy if they end up seeing different movies

	WL	**LW**
WL	2,1	0,0
LW	0,0	1,2

  - Like prisoner's dilemma, points where there is no incentive for either player to deviate, but here there are two such points

- Rock, paper, scissors
  - Rock beats scissors, scissors beats paper, paper beats rock

	Rock	**Paper**	**Scissors**
Rock	0,0	-1,1	1,-1
Paper	1,-1	0,0	-1,1
Scissors	-1,1	1,-1	0,0

  - If we play a fixed strategy, other player will always be able to beat us

- Some special cases:
  - Zero-sum game: two player game where $u_1(a) = -u_2(a)$, $\forall a \in A$ (e.g., rock paper scissors)
  - Coordination game: payoffs for all players are the same $u_i(a) = u_j(a)$, $\forall i, j \in \{1, \ldots, N\}$, $a \in A$

	L	**R**
L	1,1	-1,-1
R	-1,-1	1,1

Pure and mixed strategies

- A strategy for player i, denoted $s_i : A_i \to [0,1]$ is a probability distribution over actions: $s_i(a_i)$ denotes probability that player i takes action action a_i (think of s_i as a vector in $[0,1]^{|A_i|}$ that must sum to one)

- A strategy profile s is a set of strategies for each player $s = (s_1, \ldots, s_N)$

- The *support* of a strategy s_i is the set of actions that have non-zero probability

- Strategy with a support of size one is called *pure* strategy, otherwise *mixed* strategy

- The probability of set of actions a under strategy profile s is

$$s(a) = \prod_{i=1}^{N} s_i(a_i)$$

i.e., the actions are all chosen independently

- The expected utility for a strategy profile s is given by

$$u(s) = \sum_{a \in A} u(a) s(a)$$

- This can also be written elementwise, for example

$$u_i(s) = \sum_{a \in A} u_i(a) \prod_{j=1}^{N} s_j(a_j)$$

Best response

- Best response refers to a best (potentially mixed) strategy that a player can play *given the strategies of all opponents*

- Define s_{-i} be be strategy profile s omitting the strategy of the ith player

- Formally, best response for player i given strategy profile s_{-i} is a strategy $s_i^\star$ such that $u(s_i^\star, s_{-i}) \geq u(s_i, s_{-i})$ for all strategies s_i

- Of course, in general we don't know the strategies of the other opponents

Nash equilibrium

- Key definition: a strategy profile s is a Nash equilibrium if s_i is a best response to s_{-i} for all $i = 1, \ldots, N$

- Intuitively corresponds to a “rational” set of strategies for each agent: no agent gains an advantage by switching their strategy

- Can be one or more Nash equilibria for a game

- *Strict Nash* if for all i and $s_i' \neq s_i$

  $$u(s_i, s_{-i}) > u(s_i', s_{-i})$$

  i.e., s_i is strictly preferable to all other strategies

- *Weak Nash* otherwise, i.e., can have s_i' such that $u(s_i, s_{-i}) = u(s_i', s_{-i})$

- Prisoner's dilemma

	Silent	**Implicate**
Silent	-1,-1	-5,0
Implicate	0,-5	-3,-3

What are NE? Are they strict?

- Rock paper scissors

	Rock	**Paper**	**Scissors**
Rock	0,0	-1,1	1,-1
Paper	1,-1	0,0	-1,1
Scissors	-1,1	1,-1	0,0

What are NE? Are they strict?

- Battle of sexes

	WL	**LW**
WL	2,1	0,0
LW	0,0	1,2

What are NE? Are they strict?

- In 1950, John Nash proved that every game has at least one equilibrium point (important, requires mixed strategies)
- 27 pages, typeset like on the right (probably about 5 pages in dense latex, the same as your class project writeup ... hmmm); two references, one to his own paper
- Work won the Nobel prize in economics

Existence of Equilibrium Points

I have previously published [Proc. N. A. S. 36 (1950) 48–49] a proof of the result below based on Kakutani's generalized fixed point theorem. The proof given here uses the Brouwer theorem.

The method is to set up a sequence of continuous mappings:

$s \to s'(s,1)\,;\; s \to s'(s,2)\,;\; \cdots$ whose fixed points have an equilibrium point as limit point. A limit mapping exists, but is discontinuous, and need not have any fixed points.

THEO. 1: Every finite game has an equilibrium point.

Proof: Using our standard notation, let s be an n-tuple of mixed strategies, and $p_{i\alpha}(s)$ the pay-off to player i if he uses his pure strategy $\pi_{i\alpha}$ and the others use their respective mixed strategies in s. For each integer λ we define the following continuous functions of s:

$$q_i(s) = \max_{\alpha} p_{i\alpha}(s),$$

$$\phi_{i\alpha}(s,\lambda) = p_{i\alpha}(s) - q_i(s) + 1/\lambda, \text{ and}$$

$$\phi^{+}_{i\alpha}(s,\lambda) = \max\left[0, \phi_{i\alpha}(s,\lambda)\right].$$

Now $\sum_{\alpha} \phi^{+}_{i\alpha}(s,\lambda) \geq \max_{\alpha} \phi^{+}_{i\alpha}(s,\lambda) = 1/\lambda > 0$ so that

$$c'_{i\alpha}(s,\lambda) \doteq \frac{\phi^{+}_{i\alpha}(s,\lambda)}{\sum_{\beta} \phi^{+}_{i\beta}(s,\lambda)}$$ is continuous.

Define $s'_i(s,\lambda) = \sum_{\alpha} \pi_{i\alpha} c'_{i\alpha}(s,\lambda)$ and $s'(s,\lambda) = (s'_1, s'_2, \ldots s'_n)$. Since all the operations have preserved continuity, the mapping $s \to s'(s,\lambda)$ is con-

Outline

Introduction: Some (possibly) familiar games

Basic game theory

Computing equilibria

Special cases and extensions

Can we compute Nash equilibria?

- Since this is a computer science course after all...
- How do we actually compute the Nash equilibria of a game (for now, let's just consider two-player games)?
- In 2005, shown to be a PPAD-complete problem (not quite like NP, since every game has a Nash equilibrium, but main intuition is similar, thought to require solvable exponential time in game size in the worst case)
- But "hard" problems don't faze us in this course (see search, mixed integer programming, etc)

Computing an equilibrium with known support

- If we just want to look at pure strategies (again,in two player case) this is easy: just check all $|A_1| \times |A_2|$ possible strategies
- But, a game may not have a pure strategy equilibrium
- Key idea: For a *given* support, we can compute NE (if one exists) by solving a set of linear equations
- Thus, problem really becomes one of searching for the correct support

- Battle of sexes, let's guess that the support for a mixed strategy contains both WL and LW for husband and wife, and say that husband chooses WL with probability p

	WL	**LW**
WL	2,1	0,0
LW	0,0	1,2

- Key idea: In order for strategy to be a NE, wife must be *indifferent* between alternatives

$$u_2(\mathrm{WL}) = u_2(\mathrm{LW})$$
$$p \cdot 1 + (1-p) \cdot 0 = p \cdot 0 + (1-p) \cdot 2$$
$$p = 2/3$$

- So $s_1 = (2/3, 1/3), s_2 = (1/3, 2/3)$ is mixed strategy NE

- The general case
  - Hypothesize some supports $\mathcal{A}_1 \subseteq A_1, \mathcal{A}_2 \subseteq A_2$ for players
  - Utilities for all actions in support must be equal for both players

$$u_1(a) = u_1(a'), \forall a, a' \in \mathcal{A}_1 \Longrightarrow |\mathcal{A}_1| - 1 \text{ linear equations}$$

$$u_2(a) = u_2(a'), \forall a, a' \in \mathcal{A}_2 \Longrightarrow |\mathcal{A}_2| - 1 \text{ linear equations}$$

$$\sum_{a \in \mathcal{A}_1} s_1(a) = 1, \sum_{a \in \mathcal{A}_2} s_2(a) = 1 \Longrightarrow |\mathcal{A}_2| - 1 \text{ linear equations}$$

Variables $s_1(a), \forall a \in \mathcal{A}_1, s_2(a), \forall a \in \mathcal{A}_2, \Longrightarrow |\mathcal{A}_1| + |\mathcal{A}_2|$ unknowns

- What happens when we try a support that does not produce a NE?

- Prisoner's dilemma, mixed strategy with full support (S,I)

	Silent	**Implicate**
Silent	-1,-1	-5,0
Implicate	0,-5	-3,-3

Finding the support

- But, $2^{|A_i|} - 1$ possible supports for each player
- Do we have to try them all? In the worst case, yes (unless PPAD = P)
- But, many times we will find a solution much faster (c.f. search, mixed integer programming, etc)
- In fact, a procedure that looks a lot like local hill-climbing search is guaranteed to find a solution for the two-player case

- Lemke-Howson algorithm (stated very imprecisely)
- Start with some initial support $\mathcal{A}_1, \mathcal{A}_2$ and repeat:
    1. Choose (according to a specific rule), to add, drop, or swap action from support
    2. Solve resulting linear systems, if they are consistent with mixed strategy, we have found a NE
    3. Otherwise, continue
- Essentially the same procedure as the simplex algorithm for linear programming, for those who may be familiar with that

Outline

Introduction: Some (possibly) familiar games

Basic game theory

Computing equilibria

Special cases and extensions

N-player games

- For more than two players, precise analogue of Lemke-Howson algorithm doesn't exist, but local search can still be effective

- Can also formulate as optimization problem

$$\begin{array}{ll} \underset{s}{\text{minimize}} & \sum_{i=1}^{N} \sum_{a_i \in A_i} \max\{u_i(a_i, s_{-i}) - u_i(s), 0\}^2 \\ \text{subject to} & 1^T s_i = 1, \;\; s \geq 0 \end{array}$$

- At any NE, objective value will be zero (no incentive to any other pure strategy)

- Of course, a non-convex problem, with potential local optima

Special case: zero-sum games

- Two-player zero-sum games *can* be solved efficiently (in polynomial time) by formulating it as a linear program

$$\begin{array}{ll} \underset{s_1}{\text{minimize}}\,\underset{s_2}{\text{maximize}} & s_1^T C s_2 \\ \text{subject to} & s_1 \geq 0, 1^T s_1 = 1, s_2 \geq 0, 1^T s_2 = 1 \end{array}$$

- Requiring player 2 to play a pure strategy, equivalent to

$$\begin{array}{ll} \underset{s_1}{\text{minimize}} \max_{i=1,\ldots,|A_2|} & (C^T s_1)_i \\ \text{subject to} & s_1 \geq 0, 1^T s_1 = 1 \end{array}$$

which is equivalent to linear program

$$\begin{array}{ll} \underset{s_1,t}{\text{minimize}} & t \\ \text{subject to} & s_1 \geq 0, 1^T s_1 = 1, C^T u \leq t1 \end{array}$$

- Somewhat surprisingly, this is actually the optimal strategy for player 1, even if player 2 can play mixed strategies (proof involves an optimization concept called duality)
- Key aspect of zero-sum game is that we could express game as minimization and maximization over the same objective terms by the two agents, can't do this in general case