CSC304: Algorithmic Game Theory and Mechanism Design
Fall 2016

Allan Borodin (instructor)
Tyrone Strangway and Young Wu (TAs)

September 14, 2016

Lecture 3

Announcements

- The start of Assignment 1 has been posted. Note that this includes (question 3c) the use of an LP solver to compute the value of a zero-sum game as defined in questions 3a and 3b.
- **Talk of possible interetst**: This coming Tuesday, September 20, there will be a seminar “Preferences and Manipulative Actions in Elections” by Gabor Erdelyi. The talk will take place in Pratt 266 at 11AM.

Agenda for today

We ended the second class stating the “principle of indifference “ for a two person game and gave one example showing how to find a mixed NE for the stag-hare game.

- We begin today with a few more simple 2 by 2 two person examples; namely,
  1. The cheetah-antelope game
  2. The driver-inspector game (that we mentioned last class)
  3. The prisoner’s dilemma
- Next we consider the “tragedy of the commons”, a many player game where each player has an infinite number of strategies.
- We conclude with a brief critique of Nash equilibria

The cheetah and antelope game

In the chettah and anteople game, two cheetahs are both chasing a large and small antelope. Similar to the stag and hare game (where a stag is worth more than a hare), the large antelope is worth more than the small antelope. But unlike the stag and hare game, a cheetah can catch any anteope on its own but if both cheetahs go for the same antelope they must share the value. This is a symmetric game where each cheetah has two strategies, L (catch the large antelope) and S (catch the small antelop).

The cheetah and antelope game

In the chettah and anteople game, two cheetahs are both chasing a large and small antelope. Similar to the stag and hare game (where a stag is worth more than a hare), the large antelope is worth more than the small antelope. But unlike the stag and hare game, a cheetah can catch any anteope on its own but if both cheetahs go for the same antelope they must share the value. This is a symmetric game where each cheetah has two strategies, L (catch the large antelope) and S (catch the small antelop).

Let ℓ (resp. s) be the payoff of the large (resp. small) antelope where obviously we can assume $s \leq \ell$. When $\ell \geq 2s$, strategy L is a dominant strategy for both cheetahs. When $s \leq \ell < 2s$ (the text does not consider $s = \ell$), (S, L) and (L, S) are pure NE. But (as the text says), how would they agree on who gets which antelope?

The symmetric mixed strategy for the cheetah and antelope game

Using the principle of indifference and letting x_1 being the probability of of a cheetah choosing the “greedy strategy” L we have :

$$\frac{\ell}{2}x_1 + (1 - x_1)\ell = s \cdot x_1 + (1 - x_1)\frac{s}{2}$$

which implies that for the symmetric mixed NE, $x_1 = \frac{2\ell - s}{\ell + s}$. For example, when $\ell = \frac{4}{3}$ and $s = 1$, $x_1 = \frac{5}{7}$.

The symmetric mixed strategy for the cheetah and antelope game

Using the principle of indifference and letting x_1 being the probability of of a cheetah choosing the “greedy strategy” L we have :

$$\frac{\ell}{2}x_1 + (1 - x_1)\ell = s \cdot x_1 + (1 - x_1)\frac{s}{2}$$

which implies that for the symmetric mixed NE, $x_1 = \frac{2\ell - s}{\ell + s}$. For example, when $\ell = \frac{4}{3}$ and $s = 1$, $x_1 = \frac{5}{7}$.

Interestingly, the text points out that if we think of a cheetah playing this game reperatedly against a random cheetah, one can argue why the cheetah population would evolve to the mixed NE strategy.

The symmetric mixed strategy for the cheetah and antelope game

Using the principle of indifference and letting x_1 being the probability of of a cheetah choosing the "greedy strategy" L we have :

$$\frac{\ell}{2}x_1 + (1 - x_1)\ell = s \cdot x_1 + (1 - x_1)\frac{s}{2}$$

which implies that for the symmetric mixed NE, $x_1 = \frac{2\ell - s}{\ell + s}$. For example, when $\ell = \frac{4}{3}$ and $s = 1$, $x_1 = \frac{5}{7}$.

Interestingly, the text points out that if we think of a cheetah playing this game reperatedly against a random cheetah, one can argue why the cheetah population would evolve to the mixed NE strategy.
Think about who has an advantage when the cheetah plays more or less greedily. The claim is that such mixed NE have been observed in real evolutionary biological populations.

The game of "chicken"

While it may seem like we are obsessed with animals, this game is about human behaviour. The game of "chicken" is depicted in two rather iconic moves, "Rebel without a cause" and "Footloose". What happens when two drivers (men in both movies trying to prove something) are driving in opposite directions in a single lane? Disaster (i.e. some large negative payoff M) is incurred if neither "chickens out" whereas if only one chickens out, he only suffers a minor humiliation (say payoff = -1) while the other gains approval (say payoff = 2) for his macho performance.

player I \ player II	Swerve (S)	Drive (D)
Swerve (S)	$(1, 1)$	$(-1, 2)$
Drive (D)	$(2, -1)$	$(-M, -M)$

Figure: The chicken game matrix in KP

Observations about the chicken game

- There are two pure NE, namely (S,D) and (D,S). But (like the cheetah and antelope game), who is going to swerve (chicken out) and who is going to drive?
- The mixed NE (not surprisingly) depends on the value of M. Using the principle of indifference, the symmetric mixed strategy is for each driver to swerve with probability $x = 1 - \frac{1}{M}$ so that the expected payoff for each driver is $1 - \frac{2}{M}$.
- Perhaps surprisingly, as the text notes, even though the "matrix is decreasing in value" as M increases,, the equilibrium payoff for each player is increasing. The payoff is lower than the payoff to the "winner" in the pure NE. '
- Perhaps surprisingly, as M increases it is limiting to (S,S) which is not a Nash equilibrium.
- Finally, as the text notes, there is a great benefit to the first driver who sends a binding commitment (tearing out the steering wheel) that he will not swerve.

The mixed strategy for the driver and parking inspector game

We observed last time that the driver and inspector game did not possess a pure NE. But from Nash's theorem, we know it has a mixed NE.

Driver \ Inspector	Don't Inspect	Inspect
Legal	(0, 0)	(0, -1)
Illegal	(10, -10)	(-90, -6)

Figure: Driver and parking inspector game; Example 4.1.4 in KP

Driver-inspector mixed strategy

To find a mixed NE, lets say that the driver plays the strategy (x_1, x_2) and the inspector playes strategy (y_1, y_2)

Note: In question 2a of the assignment, we ask for a proof that in any 2 by 2 two person game that does not have a dominant strategy, any NE is either pure or fully mixed.

From the principle of indifference:

- We must have that the expected payoff to the inspector is the same whether he inspects or doesn't inspect. That is,

$$0 \cdot x_1 + (-10) \cdot (1 - x_1) = (-1) \cdot x_1 + (-6)(1 - x_1)$$

  which implies $x_1 = .8$.
- Similarly, $0 = 10 \cdot y_1 + (-90)(1 - y_1)$ so that $y_1 = .9$.

Driver-inspector mixed strategy

To find a mixed NE, lets say that the driver plays the strategy (x_1, x_2) and the inspector playes strategy (y_1, y_2)

Note: In question 2a of the assignment, we ask for a proof that in any 2 by 2 two person game that does not have a dominant strategy, any NE is either pure or fully mixed.

From the principle of indifference:

- We must have that the expected payoff to the inspector is the same whether he inspects or doesn't inspect. That is,

$$0 \cdot x_1 + (-10) \cdot (1 - x_1) = (-1) \cdot x_1 + (-6)(1 - x_1)$$

  which implies $x_1 = .8$.
- Similarly, $0 = 10 \cdot y_1 + (-90)(1 - y_1)$ so that $y_1 = .9$.

Note: I claim that for a 2 by 2 two person game we do not have to verify that a mixed strategy computed by the principle of indifference is a mixed NE. Why? But in general, one does have to verify whether or not any proposed mixed is an NE.

The prisoners dilemma

One of the classic examples in game theory is the well-known prisoners dilemma game. There can be many instances of this game

The Prisoner's Dilemma narrative: page 67 of KP text)

Two prisoners are suspected of and are being interogated for a serious crime. But they need a confession in order to convict. If they get a single confession and use that to convict the other prisoner, they set free the person who confesses and the convicted prisoner will serve serious time. If they both confess, there is a reduced sentence and if neither confess, the police have to settle for conviction of a minor crime.

Cartoon from Smbc-comics.com: family in a prisoners dilemma

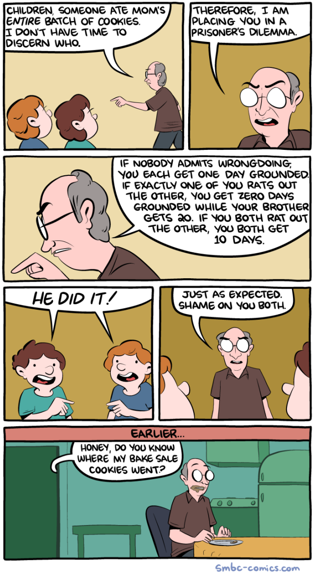

Another prisoners dilemma cartoon

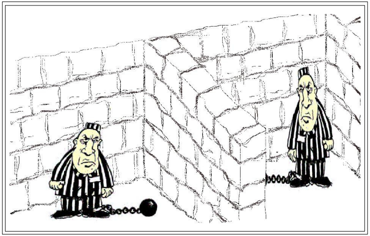

FIGURE 4.1. Two prisoners considering whether to confess or remain silent.

Figure: A cartoon for the prisoners dilemma; Figure 4.1 in KP text. Note that the prisoners cannot communicate and must decide what to do individually

Once we indicate the actual jail times, we can again use a matrix to precisely define the game.

A prisoners dilemma matrix

prisoner I	prisoner II: silent	prisoner II: confess
silent	$(-1,-1)$	$(-10,0)$
confess	$(0,-10)$	$(-8,-8)$

Figure: The prisoners dilemma matrix from KP text. Following the KP text, negative numbers represent the sentence in years. The goal of a prisoner is to *maximize* their payoff = minimize their prison time. Equivalently, one can use positive costs and then the prisoner's goal is to *minimize* their cost.

Observations about the prisoners dilemma game

- This is a symmetric game; we cold define a non symmetric game where say one prisoner (with a bigger criminal record) receives a harsher sentence.
- Confessing is a (strongly) dominant strategy (for each prisoner).

Observations about the prisoners dilemma game

- This is a symmetric game; we cold define a non symmetric game where say one prisoner (with a bigger criminal record) receives a harsher sentence.
- Confessing is a (strongly) dominant strategy (for each prisoner).
- It follows that (confess,confess) is (the only) Nash equilibrium Why?

Observations about the prisoners dilemma game

- This is a symmetric game; we cold define a non symmetric game where say one prisoner (with a bigger criminal record) receives a harsher sentence.
- Confessing is a (strongly) dominant strategy (for each prisoner).
- It follows that (confess,confess) is (the only) Nash equilibrium Why?
- But clearly (wrt to the prisonsers) the "social welfare" optimum is (silent,silent).

Observations about the prisoners dilemma game

- This is a symmetric game; we cold define a non symmetric game where say one prisoner (with a bigger criminal record) receives a harsher sentence.
- Confessing is a (strongly) dominant strategy (for each prisoner).
- It follows that (confess,confess) is (the only) Nash equilibrium Why?
- But clearly (wrt to the prisonsers) the "social welfare" optimum is (silent,silent).
- Of course achieving optimality here (and in general) usually requires coordination (either amogst the players or orchestrated by a central authority as we will see in mechansim design). It may not be possible to achieve optimality in a Nash equilibirum.
- A weaker (but still useful) solution concept is *Pareto optimality*. A solution is Pareto optimal if there is no other solution for which at least one player has higher payoff and no other player has a lower payoff. The pairs (confess,silent)) and (silent,confess) are both Pareto optimal.

The tragedy of the commons

There are (unfortunately) many situations where selfish behaviour can lead to bad outcomes for all players (and then clearly also for the social welfare). One tragedy of the commons example is given in Example 4.5.1 of KP attributed to Example 1.4 in the Roughgarden et al AGT text.

A tragedy of the commons narrative

A shared unit capacity communication channel is being shared by n players each wanting to send information along the channel. Each player i can send $x_i \in [0, 1]$ units of flow as long as the capacity is not exceeded. Because of the impact of congestion, the defined utility for i when sending i units of flow is $x_i(1 - \sum_{j \neq i} x_j)$. There is no utility to any player if the capacity is exceeded.

Note that this is a game where each agent has a continuously infinite choice of strategies.

So what is the tragedy?

- The *best response* for each player i to players $j \neq i$ is to choose $x_i = (1 - \sum_{j \neq i} x_j)/2$ resulting in the equilibrium $x_i = \frac{1}{n+1}$ and hence utility $\frac{1}{(n+1)^2}$ for all i with social welfare $\frac{n}{(n+1)^2}$.
- On the other hand each player could have decided to choose $x_i = \frac{1}{2n}$ resulting in $\frac{1}{2}$ unused capacity and a non-equilibrium solution. However, each player's utility is now $\frac{1}{4n}$ and the social welfare is $\frac{1}{4}$.

So what is the tragedy?

- The *best response* for each player i to players $j \neq i$ is to choose $x_i = (1 - \sum_{j \neq i} x_j)/2$ resulting in the equilibrium $x_i = \frac{1}{n+1}$ and hence utility $\frac{1}{(n+1)^2}$ for all i with social welfare $\frac{n}{(n+1)^2}$.
- On the other hand each player could have decided to choose $x_i = \frac{1}{2n}$ resulting in $\frac{1}{2}$ unused capacity and a non-equilibrium solution. However, each player's utility is now $\frac{1}{4n}$ and the social welfare is $\frac{1}{4}$.
- Computer scientists like to refer to this loss of social welfare as the *price of anarchy POA* which is defined as the worst case ratio between an optimal solution and an equilibrium solution. In this example, we have seen that the POA is at least $\frac{1/4}{\frac{n}{(n+1)^2}} = \frac{(n+1)^2}{4n} \approx \frac{n}{4}$.

So what is the tragedy?

- The *best response* for each player i to players $j \neq i$ is to choose $x_i = (1 - \sum_{j \neq i} x_j)/2$ resulting in the equilibrium $x_i = \frac{1}{n+1}$ and hence utility $\frac{1}{(n+1)^2}$ for all i with social welfare $\frac{n}{(n+1)^2}$.
- On the other hand each player could have decided to choose $x_i = \frac{1}{2n}$ resulting in $\frac{1}{2}$ unused capacity and a non-equilibrium solution. However, each player's utility is now $\frac{1}{4n}$ and the social welfare is $\frac{1}{4}$.
- Computer scientists like to refer to this loss of social welfare as the *price of anarchy POA* which is defined as the worst case ratio between an optimal solution and an equilibrium solution. In this example, we have seen that the POA is at least $\frac{1/4}{\frac{n}{(n+1)^2}} = \frac{(n+1)^2}{4n} \approx \frac{n}{4}$.
- The price of anarchy for our prisoners dilemma game was $\frac{-16}{-2} = 8$.

Some (for now) concluding comments on general-sum games; the importance and limitations of Nash equilibria

We have been assuming a full information game; that is, all players know the payoff matrix. We are also assuming players are strategic and that their only goal is to maximize their individual payoff.

What is the importance of Nash equilibria?
Suppose the game is being played repeatedly. IF we are not in a NE, we would expect some player to change their strategy so as to improve their (expected) payoff. So an NE is a more stable solution. In fact, the KP text states: "**we expect the outcome of a game to be a Nash equilibrium**"

As stated in the AGT text, ...the notion of a Nash equilibrium, which, despite its shortcomings ...has emerged as *the central concept in game theory*.

Brief critique of Nash equilibrium

The KP chapter notes and AGT text point out criticisms regarding the relevance of Nash equilibria. There are many critiques (both positive and negative) about Nash equlibrium. Some of the arguments against NE are:

- The development thus far assumes each player has perfect an comolete information about the the entire matrix.
- Players are rational and only care about their own payoffs.

Brief critique of Nash equilibrium

The KP chapter notes and AGT text point out criticisms regarding the relevance of Nash equilibria. There are many critiques (both positive and negative) about Nash equlibrium. Some of the arguments against NE are:

- The development thus far assumes each player has perfect an comolete information about the the entire matrix.
- Players are rational and only care about their own payoffs.
- There can be multiple NE and hence no reason to believe we will end up in a favourable NE even if the game is repeated and we find an NE.

Brief critique of Nash equilibrium

The KP chapter notes and AGT text point out criticisms regarding the relevance of Nash equilibria. There are many critiques (both positive and negative) about Nash equlibrium. Some of the arguments against NE are:

- The development thus far assumes each player has perfect an comolete information about the the entire matrix.
- Players are rational and only care about their own payoffs.
- There can be multiple NE and hence no reason to believe we will end up in a favourable NE even if the game is repeated and we find an NE.
- Individuals have difficulty thinking about mixed strategies.

Brief critique of Nash equilibrium

The KP chapter notes and AGT text point out criticisms regarding the relevance of Nash equilibria. There are many critiques (both positive and negative) about Nash equlibrium. Some of the arguments against NE are:

- The development thus far assumes each player has perfect an comolete information about the the entire matrix.
- Players are rational and only care about their own payoffs.
- There can be multiple NE and hence no reason to believe we will end up in a favourable NE even if the game is repeated and we find an NE.
- Individuals have difficulty thinking about mixed strategies.
- It is now known (and this is a result in AGT not previously considered in GT), based on reasonable complexity assumptions, that it can be computationally difficult to find an NE. As stated by Kamal Jain "If your laptop can't find an equilibrium neither can the market".

Brief critique of Nash equilibrium

The KP chapter notes and AGT text point out criticisms regarding the relevance of Nash equilibria. There are many critiques (both positive and negative) about Nash equlibrium. Some of the arguments against NE are:

- The development thus far assumes each player has perfect an comolete information about the the entire matrix.
- Players are rational and only care about their own payoffs.
- There can be multiple NE and hence no reason to believe we will end up in a favourable NE even if the game is repeated and we find an NE.
- Individuals have difficulty thinking about mixed strategies.
- It is now known (and this is a result in AGT not previously considered in GT), based on reasonable complexity assumptions, that it can be computationally difficult to find an NE. As stated by Kamal Jain "If your laptop can't find an equilibrium neither can the market".
- Beyond these criticisms, there is an entire field called **behavioural economics** including **prospect theory** giving alternative explanations as to how individuals will make decisions.

In defense of Nash equilibrium

All these criticsms can be offset to some extent.

- The concept of perfect information games and Nash equilibrium can be extended to one of probabilistic information. More specifically in Bayesian settings, each player draws their playoff values from some distribution (and these distributions can be independent or correlated).
- Agents are now sometimes algorithms which (perhaps arguably) are rational agents that can implement mixed strategies.
- The complexity arguments are worst case arguments and do not necessarily tell us what to expect "in practice".
- Even if individuals may not think in terms of probabilities and/or have trouble learning good mixed strategies, there is some evidence that populations (e.g. as in the cheetah-antelope game and the soccer penalty shot game) as a whole will evolve to learn mixed strategies.
- Some other factors (beyond pure self-interest as we assume in rationality) such as risk aversion (or attraction), or interest in the social welfare or welfare of others can often be accomodated by modifying the payoffs.

The emphasis of this course

Our emphasis on rational decision making

This course is a more typical game theory/mechanism design course with its emphasis on rational decision making (but with some attention to the computational and informational asepcts of AGT) and concepts such as Nash equilibrium for which John Nash received the 1994 Nobel Prize in Economics. It should be noted that Daniel Kahneman received the 2002 Nobel Prize in Economics for his work with Amos Tversky (who died in 1996) for their work on alternative explanations of human decision making.

The emphasis of this course

Our emphasis on rational decision making

This course is a more typical game theory/mechanism design course with its emphasis on rational decision making (but with some attention to the computational and informational asepcts of AGT) and concepts such as Nash equilibrium for which John Nash received the 1994 Nobel Prize in Economics. It should be noted that Daniel Kahneman received the 2002 Nobel Prize in Economics for his work with Amos Tversky (who died in 1996) for their work on alternative explanations of human decision making.

Do you make decisions rationally?

The emphasis of this course

Our emphasis on rational decision making

This course is a more typical game theory/mechanism design course with its emphasis on rational decision making (but with some attention to the computational and informational asepcts of AGT) and concepts such as Nash equilibrium for which John Nash received the 1994 Nobel Prize in Economics. It should be noted that Daniel Kahneman received the 2002 Nobel Prize in Economics for his work with Amos Tversky (who died in 1996) for their work on alternative explanations of human decision making.

Do you make decisions rationally?

To what extent does intuition, experience, or emotion determine your important decisions vs more time consumptive and demanding rational analysis?