Game Theory and Algorithms*
Lecture 1: Introduction

February 22, 2011

Summary: We give an overview of the course and some motivation. We give the most basic model of a mathematical game, namely *strategic games*, which underlie everything else in the course. We introduce notation that will be used throughout the course.

1 All About This Course

A *game* in mathematics is a situation where multiple distinct *players* make choices, and each player gets some sort of a *utility/payoff* depending on their own choice and the choices of the other players. All analysis of games is done under the assumption that each player wants to maximize their own utility; thus game theory with just one player is, at least superficially, the theory of optimization. But when there are multiple players whose choices affect one another, a much richer set of possibilities opens up. It becomes possible to model competition and cooperation.

Some prototypical examples of situations which can be modelled with games are elections, auctions, price-setting of competing companies, negotiation, evolution, sharing common goods (e.g. traffic congestion), or building common infrastructure (e.g. network formation). We will see several of these examples (some, multiple times from different angles). Traditional games like tic-tac-toe, chess, checkers, poker, and backgammon also can be cast as mathematical games.

Q: Two casino games that are not suitable for this course are roulette and blackjack. Why?

A: These are essentially one-player games; we thus consider them just as optimization problems. (For example in roulette, the payoffs to a gambler depend only on his choices and on a fixed randomized procedure of the casino; different gamblers at the same table have no effect on one another.)

A common theme in game theory literature is that there is more than one way to model a given situation. For example, we will start with *strategic games*, which are the most fundamental class of games since it is in principle possible to cast any game as a strategic game. Then we will come to the class of *extensive games*, where a notion of time/history comes into play and the players can make decisions dependent on past actions. Although we can "typecast" any extensive game into a strategic game, we will find out that some natural solutions (Nash equilibria) of the resulting strategic game don't really make sense (leading to *subgame perfect equilibria*).

Another example of multiple different models is price-setting competition by two firms (*duopoly*), where we give two different strategic games to model the situation.

In short, the game-theoretic predictions made by a model are only accurate if we accurately modelled the game in the first place.

Since game theory and algorithmic game theory are huge disciplines, we try to stick to the parts that are simplest, most useful, and under least dispute. The intent is that a student completing the course should be well-equipped to apply game theory in practice and to understand modern literature from (algorithmic) game theory.

* Lecture Notes for a course given by David Pritchard at EPFL, Lausanne.

1.1 Course Contents Overview

I give a point-by-point overview on the course website `http://disopt.epfl.ch/gta` but here is a more informal summary.

First, we will study strategic games, where everyone makes one choice, all at the same time. This simple model already contains most of the richness of game theory. We will look at one type of "prediction" which is very natural, (iterated) elimination of strictly dominated strategies, illustrated at first by the prisoner's dilemma.

We move on to (pure) Nash equilbria, which has more predictive value than iterated elimination. We give many examples of games where the Nash equilibria seem to give useful predictions, including duopoly and elections.

Not all strategic games have a pure Nash equilibria, but all (finite) ones have a *mixed* Nash equilibrium, where we allow players to randomize. This is Nash's theorem and one of the cornerstones of the field; we will give some examples and a sketch of the proof, which is topological. For 2 players we give direct algorithmic proofs, including the *Lemke-Howson algorithm* which has connections to so-called *PPAD-completeness* and deep results about the complexity of finding mixed Nash equilbria.

These topics (strategic games) should take up roughly 1/3 of the lectures. We then move on to:

- games with explicit histories — *extensive games* — and (infinitely) repeated games
- auctions, an example of *mechanism design*, where we want to design a game with specific properties, rather than analyze the properties of a game
- elections and Arrow's impossibility theorem
- the Price of Anarchy — can we quantify how badly the greedy nature of players hurts them?
- impartial combinatorial games: a family of games where we can compute optimal strategies
- finally, student presentations.

Some topics from game theory which we *won't* cover are: games with incomplete information, such as poker, and the related *Bayesian games*; coalitional games; and *cost-sharing* from algorithmic game theory.

1.2 Format of the Course and Requirements

We meet in 10:15-12:00 Tuesdays in MA A1 10 (except Mar 22) and Thursdays in a different room, MA A3 30. There are three requirements in the course.

1. I will present several exercises every week. Approximately once per two weeks, we will replace a lecture with an exercise session to discuss the solutions as a group. To get credit, you must participate in this activity. I will schedule office hours on demand for anyone who wishes to get hints and advice with exercises. You can also contact me by e-mail any time. You are expected to work alone on the exercises; ask me for permission in advance if you want to collaborate with anyone else.

2. I will ask each student to act as a scribe (at least once, more if people are interested). There is a LaTeX template on the course website for this purpose. The main reason for scribing is that, since there is no textbook, I don't want everyone to spend their time copying everything down in lecture. However, since there is a delay in getting the scribe notes online, and since we might not scribe every lecture, I will upload scans of my hand-written lecture notes when needed.

3. Finally, each student will present a paper of their choice at the end of the term. This serves as the final exam for the course. Each presentation will be about half of a lecture period (approx. 40 minutes long).

2 Strategic Games

Definition 1. A *strategic game* consists of three things:

- A finite set $\{1, 2, \ldots, n\}$ of players
- For each player i, a set A_i of *actions* (sometimes called *strategies*)
- For each player i, a *payoff function* $u_i : \prod_i A_i \to \mathbb{R}$ (sometimes called *utility functions*)

It is worthwhile to unpack the term $\prod_i A_i$ in this definition. Each element $a \in \prod_i A_i$ is a list of strategies, one per player, and such an a is called an *action profile* (or an *outcome*). So $a = (a_1, \ldots, a_n)$ where each a_i is some element of A_i. The meaning if the payoff function is that it represents, for each player, their preferences over the outcomes. When $u_i(a) > u_i(a')$ for two outcomes a and a', this means that player i prefers outcome a to outcome a'.

In analyzing strategic games, we will assume for now that players make their choices simultaneously, and that every player knows each A_i and u_i (*complete information*). Initially we assume nothing else but later definitions (e.g. Nash equilibria) will model stronger assumptions by players on the actions of the other players.

Example 2. We will model a situation where player 1 is a taxpayer and player 2 is the Swiss Revenue Agency. The action set A_1 for player 1 is {lie (L), don't lie (DL)} and the action set A_2 for player 2 is {audit (A), don't audit (DA)}. (Although in real life player 1 moves before player 2, we can equivalently model their moves as simultaneous since player 2 needs to act without knowing what player 1 chose.) We choose preferences to model the fact that *there is a cost to the SRA for auditing, but it is compensated by a penalty assessed when the taxpayer is caught lying*:

$$u_1(L, DA) > u_1(DL, DA) > u_1(DL, A) > u_1(L, A)$$

$$u_2(L, A) = u_2(DL, DA) > u_2(DL, A) > u_2(L, DA).$$

This is a little arbitrary but not hard to imagine, e.g. $u_1(DL, DA) > u_1(DL, A)$ represents that being audited is annoying.

By giving the u_i specific numeric values, we can represent the game by a table.

p1 \p2	A	DA
L	1,3	4,1
DL	2,2	3,3

Here and later the convention is that the entry in row a_1 and column a_2 is $u_1(a), u_2(a)$ for $a = (a_1, a_2)$.

Note: these values don't accurately represent dollar amounts. We'll come back to a related issue when we discuss how probabilistic preferences are modelled.

2.1 All Games are Strategic

In principle, any game can be converted to a strategic form. We will return to this later, but sketch the general idea on an example game. At the start of a football game, team 1 decides to kick (K) or receive (R), and then team 2 decides to play westwards (W) or eastwards (E). Converting this game to strategic form doesn't seem possible on face value, since in a strategic game each player makes their choices simultaneously, while in this football scenario we want to allow team 2 to make different choices depending on the choice of team 1.

The fix this, we define the set of *strategies* for player 2 in such a way that each strategy indicates a *pair* of choices: first the choice of what team 2 would do when team 1 picks K, and then the choice of what team 2 would do when team 1 picks R. In other words, take $A_1 = \{K, R\}$ and $A_2 = \{(W, W), (W, E), (E, W), (E, E)\}$. This is indeed a strategic game: for any action profile (a_1, a_2) the payoffs are well-defined. We will return to this later.

2.2 The Prisoners' Dilemma, and Dominated Strategies

Two suspects in a major crime have been arrested and are being held in separate cells. There is enough evidence to convict each of them of a minor offense, but not enough evidence to convict either one of a major crime unless the other one confesses against the other ("finks"). They will be interrogated separately and both will be given the option of keeping quiet (Q) or finking (F).

If they both stay quiet, they will both be convicted of a minor crime. If one player finks and the other stays quiet, the finker is set free and the quiet one will be convicted of a major crime. But if both of them fink, they will share the penalty for the major crime. This leads to the following payoffs:

p1 \p2	Q	F
Q	2,2	0,3
F	3,0	1,1

3: set free
2: minor conviction
1: share major conviction
0: major conviction alone
e.g. $u_2(F, F) = 1$

What will happen when the prisoners play this game, if we assume they are interested in maximizing their utility? Next we will give a logical argument showing that the outcome (F, F) is most likely. This gives each player payoff 1; this is counterintuitive because if they both picked Q instead, they would *both* get a better payoff of 2. This notation helps explain the argument and a more general approach.

Definition 3. A *partial action profile* a_{-i} is an $(n-1)$-tuple consisting of one action for each player *except* player i. That is to say each a_{-i} is an element of $\prod_{j:j\neq i} A_j$. If a_{-i} is such a partial action profile, and $a_i \in A_i$ is an action for player i, we write (a_{-i}, a_i) for the (complete) action profile obtained by combining them. I.e.,

$$(a_i, b_{-i}) = (b_1, b_2, \ldots, b_{i-1}, a_i, b_{i+1}, \ldots, b_n).$$

Definition 4. For player i, we say action a_i *strictly dominates* action a_i' if, for all partial action profiles a_{-i} of the other players,

$$u_i(a_i, a_{-i}) > u_i(a_i', a_{-i}).$$

In other words, a_i *strictly dominates* a_i' if a_i is always better than a_i' for player i.

So assuming that a player i is rational and seeks to maximize their own utility, if a_i strictly dominates a_i', they will never choose the strictly dominated action a_i', rather in such a situation they should always play a_i.

Now look at the prisoners' dilemma: for player 1, F strictly dominates Q, since 3¿2 and 1¿0. So we would expect player 1 never to pick Q, and always to pick F. The same holds for player 2 since it is a symmetric game. So, as claimed, the outcome (F, F) is the one to logically expect.

2.2.1 Discussion

Is there any logical way for the players to pick (Q, Q)? No: even if they know about the upcoming situation and "agree" to both stay quiet, the only rational action is to double-cross one another.

Q: What if they are friends and don't want to see the other prisoner get punished?

A: Applying this sort of reasoning is equivalent to modifying their utility functions and thereby changing the game. As usual, the result of the model is only as good as its accuracy.

Later we will see a modified "infinite" version of the game where cooperation becomes a plausible outcome.

3 Iterated Elimination of Strictly Dominated Strategies

We can come up with a simple algorithm to predict the outcome of certain games:

Procedure ITERATEDELIMINATION

while any player has a strictly dominated action a_i, **do**
    delete a_i from A_i
end while

The argument in the last section amounts to saying the following: if the information of the game is completely known, if each player seeks to maximize their utility, and all players know that fact, and all players know that all players know that fact, etc., then ITERATEDELIMINATION preserves all rational outcomes of the game. And thus if ITERATEDELIMINATION comes up with only one outcome, it is the unique rational outcome of the game, and the one we should predict under the given assumptions.

Exercise. Compute the result of applying ITERATEDELIMINATION to the following game.

p1 \p2	L	M	R
U	2,3	0,2	1,1
D	1,1	5,0	0,4

There is also a more organic (and informal) way to describe iterated elimination. Suppose the players play the game many times. After a while, player 2 will learn that M is a bad choice, and stop playing it. Later, although player 1 might continue to use D for a while hoping that player 2 would pick M, eventually player 1 will stop choosing D. Once player 2 learns this, they will only pick L. (Note, in this setting, each player does not even need to know anything about their opponents' payoff functions!)

Exercise. Compute the result of applying ITERATEDELIMINATION to the "income tax" game.

So, iterated elimination does not produce a single outcome for all games. The next topic we introduce is a more specific solution concept for games, the *Nash equilibrium*. Eventually this leads to one kind of "solution concept" that exists in all games.

Exercise. (A little subtle) In ITERATEDELIMINATION, it may happen at a given point in time that there are several possible choices of dominated strategies to eliminate. Prove that the final result of the algorithm is independent of the choice made in each iteration.

3.1 Challenge Exercise

Here is a challenge exercise which uses the conversion of arbitrary games to strategic form.

Exercise. A *linear inequality system* $\{x \in \mathbb{R}^n : 0 \leq x_i \leq 1 \text{ for } i = 1, \ldots, n \text{ and } Ax \geq b\}$ consists of a matrix $A \in \mathbb{R}^{m \times n}$, a vector x of n real variables, and a vector $b \in \mathbb{R}^n$. (Each row A_i, b_i gives a linear constraint $\sum_i A_{ij} x_j \geq b_j$.) The simplex algorithm can be used to determine whether a given linear inequality system has any solution. Now, consider the *quantified* linear inequality system

$$\exists x_1 \in [0,1] \forall x_2 \in [0,1] \exists x_3 \in [0,1] \forall x_4 \in [0,1] \cdots : (Ax \geq b).$$

Give an algorithm to determine whether a quantified linear inequality expression of this form is true or false. (Hint 1: it won't be a polynomial-time algorithm. Hint 2: convert it to a game where the "$\exists$ player" tries to make the statement true and the "$\forall$ player" tries to make the statement false. Hint 3: show that the $\exists$ player can assume the $\forall$ player only ever chooses the values $\{0, 1\}$.)