The main objective of the class is to give an overview of important topics in game theory. We will start the course by reviewing some basic concepts: what is a simultaneous-move game, what is a Nash equilibrium, how to prove that it exists, and give some important examples of games. After that, instead of learning from a textbook, we will start learning from original papers. Papers will be assigned to read, and I will present papers in class. Also, there will be student presentations during some of the classes. Student presentations will not be graded, and they are not required, but they are very good for you if you volunteer to present.

There are clearly many advantages to using a textbook as opposed to learning the material from papers, and many teachers would argue for using a textbook. Here are two advantages for a textbook:

- Over time, as concepts were understood by many people and taught in classes, they become crystallized in a simpler and cleaner form. Therefore, the picture comes across cleaner from a textbook than from a paper
- It is more difficult to read papers and understand what they are referring to. It is more difficult to understand what is more important and what is less important in a paper.

Here are the advantages of learning material from papers.

- Ultimately what we need as researchers is to take a topic, which is unstructured, and identify structure behind it and get results. By learning from textbook, we learn to think in a structured way about standard well-know concepts, and we become close-minded this way. Then it becomes more difficult to identify structure in chaos, because one thinks that there is no structure beyond the standard material that one already knows.
- The second advantage is that by reading papers themselves, we get a better understanding of how to write papers. A paper conveys many things that a textbook does not, which researchers inevitably have to deal with. When a fundamental paper appears for the first time, it is not always welcomed immediately. Some people realize that it is important, but many people are skeptical. There are many established opinions, and alternative ways of how something should be done. Authors should deal with all that. Thus, papers reflect the feedback that authors got, comments that they got in the seminars or from referees. While you may be entering this class hoping to simply learn the standard material, ultimately it is important to know how to answer people’s questions in the seminars, how to talk to people on the job market, etc. I think reading papers will help that.
- Finally, by including papers, one can learn slightly different material than textbook material, which is nevertheless important. Textbook material is something that very many people have read, learned and thought about. Therefore, it is highly unlikely that you will find an attractive topic for a paper in connection with something that you learn from a textbook. Even if you think of something, it is likely that this has been done. By contrast, by reading a paper

you get exposed to a different set of open issues, and it is much more likely that you find a topic for new research.

Of course, there is still the challenge that papers are harder to understand than textbook. This is okay, because as I present papers in class, I try to isolate what is important. Also, problem sets will be based on the papers, and problems will guide you to understand the results of the papers better.

So, the bottom line is, there will be difficult moments as you read papers, and you will wish that things would be spelled out clearly in a textbook. However, this is all good for you, and in the end, the skills that you learn will be useful not only for the field exam in theory, but also in research, and in convincing people that your research is important.

In addition, let me tell you a little bit about my teaching style, because some of you may become concerned that my teaching is too informal and not sufficiently rigorous. Before this happens, let me explain that I may sound informal because I try use informal language, avoid jargon and give colorful explanations to formal concepts. However, with informal words the meaning is the same as a dry formal explanation, except it takes less time and effort to understand. Let me give an example. Here is a *formal and rigorous* definition of a simultaneous-move game with private information from Osborne and Rubinstein:

Definition 25.1 A **Bayesian game** consists of

- a finite set N (the set of **players**)
- a finite set Ω (the set of **states**)

and for each player $i \in N$

- a set A_i (the set of **actions** available to player i)
- a finite set Y_i (the set of **signals** that may be observed by player i) and a function $y_i : \Omega \rightarrow Y_i$ (the signal function of player i)
- a probability measure p_i on Ω (the **prior belief** of player i) for which $p_i(\tau_i^{-1}(y_i)) > 0$ for all $y_i \in Y_i$
- a preference relation $\geq_i$ on the set of probability measures over $A \times \Omega$ (the **preference relation** of player i), where $A = \times_{j \in N} A_j$.

To me, this definition is extremely hard to understand. It makes me feel like a game is a bunch of formal sets and functions. I remember studying these definitions as an undergraduate. I would read this definition, think hard about it, put things together, and at the end figure out what a Bayesian game is. After that I would get a feeling of satisfaction and accomplishment, feeling like I have learned and achieved something. After that, I would do problems, read an informal setting, identify all sets and functions involved in the game, work things out mathematically. Sometimes I would see the big picture, but most of the time I was just concerned with getting the homework done on time. Now, let me explain how I would teach it.

Definition. N players participate in the game. First, the nature chooses the state of the world ω from the probability distribution p. Each player privately sees a random signal about the state of the world $y_i \in Y_i$ with probability $q(y_i | \omega)$. After seeing his signal, each player privately takes an action $a_i \in A_i$. The payoff of player i is given by $g_i(a, \omega)$, where $a = (a_1, a_2, \dots a_N)$ is the vector of actions that the players took.

If you compare these two definitions, of course they are not exactly equivalent. The definition from Osborne and Rubinstein allows players to have different priors, whereas here I have a common prior p for all players. The reason I do this is because this definition applies to most games of interest, and if we wanted to extend it to games with non-common prior, this would be very easy to do.

The second distinction is redundancy: I allow signals to depend stochastically on the state of the world. Why is this redundancy? Because one can always change the model by making the set of states of the world richer, so the state of the world also contains information about each player's signal. However, this redundancy has practical value. If we think about an auction, where different companies bid for an oil tract. We think that the state of the world is the total amount of oil in the tract, and each company performs an estimate based on its research. It is confusing to think that the state of the world contains all possible information: the amount of oil, and the content of research that each company has performed.

Lastly, by language I try to better convey the flavor of the game that is played. From this definition one can imagine each player, the information that he gets, and the agony about reaching the decision about which action to take.

Normal-form games.

First, I begin by defining a game in normal form. It is useful to think of it as a simultaneous-move game.

Definition. A normal form game is defined as follows. There are N players. Each player chooses an action $a_i \in A_i$. The payoff of player i is given by $g_i(a)$, where $a = (a_1, a_2, \dots a_N)$ is the vector of actions that the players took.

Definition. A pure-strategy *Nash equilibrium* is an action profile $a = (a_1, a_2, \dots a_N)$, such that each player's action is a best response to the actions of his opponent. Action a_i is a *best response* to a_{-i} if a_i maximizes $g_i(a_i, a_{-i})$.

A mixed strategy μ_i in a normal form game is simply a probability distribution over the players actions. Basically, player i randomizes when he plays a mixed strategy, choosing action a_i with probability $\mu_i(a_i)$. Payoffs for mixed strategies are defined are defined as follows

$$g(\mu_1, \dots \mu_N) = E_{\mu 1 \dots \mu N} [g(a)]$$

where the expectation is under the probability distribution over the action profile, induced by the mixture of the players.

At this point, let me gives some examples of games, and we will do some exercises before we move on.

Examples of normal form games:

Prisoners' Dilemma:

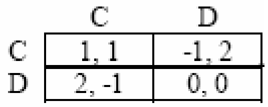

	C	D
C	1, 1	-1, 2
D	2, -1	0, 0

Battle of the Sexes.

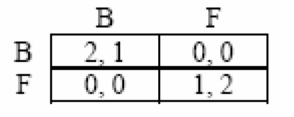

	B	F
B	2, 1	0, 0
F	0, 0	1, 2

Find the Nash equilibrium in mixed strategies.

Matching pennies:

		Player 2	
		H	T
Player 1	H	-1, 1	1, -1
	T	1, -1	-1, 1

Cournot Duopoly:

There are two sellers with zero marginal cost. Each chooses a quantity $q_i \in [0, \infty)$. The market price is $p = 1 - q_1 - q_2$ and firm i's payoff is pq_i.

Find the Nash equilibrium.

Existence of the Nash equilibrium.

Generally, the Nash equilibrium does not have to exist. The existence of the Nash equilibrium under some conditions is standard material in a game theory course. The proof that the Nash equilibrium exists is purely technical and there is no economic intuition behind it whatsoever. It is generally considered good for something to exist. Many people who teach a game theory course spend several lectures developing the mathematical machinery and proving existence. I will not go into the proof in great detail, but I will explain to you what is the backbone of the argument. The main mathematical result behind the proof of existence is Kakutani's fixed point theorem:

Kakutani's fixed point Theorem. Let X be a compact convex subset of R^n and let f : X → X be a set-valued function for which

- for all $x \in X$, the set f(x) is nonempty and convex
- the graph of f is closed

Then there exists $x^* \in X$ such that $x^* \in f(x^*)$

(people also use the terminology *upper hemi continuous* correspondence, which basically means that the graph is closed, and images of all compact sets are compact)

Existence Theorem. The Nash equilibrium exists if the sets A_i is compact and convex, and the payoff functions g_i are continuous on A and quasi-concave on A_i.

Corollary: Every normal form game with finitely many actions has a mixed strategy Nash equilibrium.

(Indeed, the set of mixed actions M_i is a simplex, so it is compact and convex, the payoff functions g_i are linear in the mixture of each player, thus continuous and quasi-concave).

Proof of the Existence Theorem. Let B : A → A be a best response correspondence (for each $a \in A$, B(a) gives the set of $(a_1', \ldots a_N')$, such that a_i' is a best response to a_{-i}). Under the assumption of the theorem, this correspondence satisfies the assumptions of Kakutani's theorem (in particular, it has closed graph). Therefore, the fixed point exists, and it must be a Nash equilibrium.

At this point, let me give some intuition behind Kakutani's fixed point theorem and the existence theorem. For Kakutani's fixed point theorem, the set X is compact and convex. Let's take the interval [0, 1] of the real line as an example, and draw different correspondences with a closed graph. Which of the following correspondences satisfy the assumptions of Kakutani's fixed point theorem?

Groups:
X compact
X convex
f(x) nonempty
f(x) convex
graph of f is closed

Now, let me explain the assumptions of the Existence Theorem.
A_i compact convex
g_i continuous
g_i quasi-concave

<u>Enough on existence.</u>

Example: War of attrition. Two players are bargaining over how to split \$3. Each player initially makes a demand of \$2 to himself and \$1 to his opponent. After that, they play a game. Each player chooses a time $t_i \in [0, \infty]$ to concede to his opponent. If player i concedes to player j first at time t, then player j gets his demand and the payoffs are $2e^{-rt}$ to player j and e^{-rt} to player i, where $r > 0$ is a discount factor. If the players concession times simultaneously, then the player who gets his demand is chosen randomly.

<u>Exercise</u>: find all pure-strategy Nash equilibria.

<u>Exercise</u>: does this game satisfy the conditions of the Existence Theorem?

Now, this game also has mixed-strategy equilibrium. In particular, there is one where each player has a density of concession characterized by the function F, which has full support and satisfies F(0) = 0. The intuition behind this equilibrium is the following: each player waits and may concede at any moment. For any time t, he must be indifferent between conceding at t or a moment later. The trade-off between conceding now or a moment later is as follows: by conceding a moment later, one gets less value due to discounting, but at the same time one has a chance to get a concession from one's opponent.

<u>Exercise:</u> Find the mixed-strategy equilibrium, in which function F_i for each players i = 1,2 has full support and F(0) = 0.

Answer: the payoff from conceding at time t is:

$$\int_0^t 2e^{-rs}\,dF(s) + e^{-rt}(1 - F(t))$$

It must be constant. Differentiating with respect to t, we get

$$2e^{-rt} f(t) - re^{-rt}(1 - F(t)) - e^{-rt} f(t) = 0$$

$$F'(t) = -r(1 - F(t)) \qquad (*)$$

$$F(t) = 1 - e^{-rt}$$

Are there other Nash equilibria?

- Can function F have an atom at t > 0?
- Can concessions stop at a time T > 0?
- Can there be an interval where nobody concedes?
- Can there be an interval where one player does not concede?

We conclude that there are concessions at all intervals. Because there are no atoms, a player's payoff function as a function of time is continuous. Because it is constant on a dense set (at all concession points), it must be constant. Therefore, F must satisfy ODE (*). This ODE has a number of solutions, which correspond to an atomic concession at time 0 and a density of concessions with a constant hazard rate thereafter.

- Is it possible that one player concedes with a positive probability at time 0? How about both players?

Simultaneous-move game with imperfect information.

Definition. N players participate in the game. First, the nature chooses the state of the world ω from the probability distribution p. Each player privately sees a random signal about the state of the world $y_i \in Y_i$ with probability $q(y_i \mid \omega)$. After seeing his signal, each player privately takes an action $a_i \in A_i$. The payoff of player i is given by $g_i(a, \omega)$, where $a = (a_1, a_2, \ldots a_N)$ is the vector of actions that the players took.

Definition. A strategy in such a game assigns an action for each signal that a player may receive. In a Bayesian Nash equilibrium, each player's strategy must maximize his payoff given the strategies of other players.