Game Theory

Adapted from Mihai Manea's Lecture Notes (MIT Opencourseware)

What is Game Theory?

Game Theory is the formal study of strategic interaction.

In a strategic setting the actions of several agents are interdependent.

Each agent's outcome depends not only on his actions, but also on the actions of other agents.

How to predict opponents' play and respond optimally?

Everything is a game...

- poker, chess, soccer, driving, dating, stock market
- advertising, setting prices, entering new markets, building a reputation
- bargaining, partnerships, job market search and screening
- designing contracts, auctions, insurance, environmental regulations
- international relations, trade agreements, electoral campaigns

Most modern economic research includes game theoretical elements.

Eleven game theorists have won the economics Nobel Prize so far.

Brief History

- Cournot (1838): quantity setting duopoly
- Zermelo (1913): backward induction
- von Neumann (1928), Borel (1938), von Neumann and Morgenstern (1944): zero-sum games
- Flood and Dresher (1950): experiments
- Nash (1950): equilibrium
- Selten (1965): dynamic games
- Harsanyi (1967): incomplete information
- Akerlof (1970), Spence (1973): first applications

- 1980s boom, continuing nowadays: repeated games, bargaining, reputation, equilibrium refinements, industrial organization, contract theory, mechanism/market design
- 1990s: parallel development of behavioral economics
- more recently: applications to computer science, political science, psychology, evolutionary biology

Key Elements of a Game

- Players: Who is interacting?
- Strategies: What are the options of each player? In what order do players act?
- Payoffs: How do strategies translate into outcomes? What are players' preferences over possible outcomes?
- Information/Beliefs: What do players know/believe about the situation and about one another? What actions do they observe before making decisions?
- Rationality: How do players think?

Normal-Form Games

A normal (or strategic) form game is a triplet (N, S, u) with the following properties:

- $N = \{1,2, \dots, n\}$: finite set of players
- $s_i \in S_i$: s_i is a pure strategy for player i; S_i is the set of pure strategies of player i
- $s = (s_1, \dots, s_n) \in S_1 \times \cdots \times S_n = S$: set of pure strategy profiles
- $s_{-i} \in \prod_{j \neq i} S_j = S_{-i}$: pure strategy profiles of i 's opponents
- $u_i : S \to \mathbb{R}$: payoff function of player i; $u = (u_1, \dots, u_n)$.

A pure strategy is deterministic.

Outcomes are interdependent.

> Player $i \in N$ receives payoff $u_i(s_1, \dots, s_n)$ when $s = (s_1, \dots, s_n) \in S$ is played.

The structure of the game is common knowledge

> All players know (N, S, u), and know that their opponents know it, and know that their opponents know that everyone knows, and so on.

The game is finite if S is finite.

Rock-Paper-Scissors

	R	P	S
R	0,0	$-1,1$	$1,-1$
P	$1,-1$	0,0	$-1,1$
S	$-1,1$	$1,-1$	0,0

Mixed and Correlated Strategies

- $\Delta(X)$: set of probability measures (or distributions) over the measurable space X (usually, X is either finite or a subset of a Euclidean space)
- $\sigma_i \in \Delta(S_i)$: mixed strategies of player i
- $\sigma \in \Delta(S_1) \times \cdots \times \Delta(S_n)$: mixed strategy profile, specifies a mixed strategy for each player
- $\sigma \in \Delta(S)$: correlated strategy profiles
- $\sigma_{-i} \in \Delta(S_{-i})$: correlated belief for player i
- $\prod_{j \neq i} \Delta(S_j)$: set of independent beliefs for i

Player i has von Neumann-Morgenstern preferences-expected utility—over $\Delta(S)$, i.e., u_i extends to $\Delta(S)$ as follows:

$$u_i(\sigma) = \sum_{s \in S} \sigma(s) u_i(s)$$

Flight plan:

First thing we did: describe the environment (setup) we want to study.

In our case: agents, physical/institutional environment, and the way they interact.

Assumptions describe a specific environment.

Next: make **predictions** about what will happen.

In our case: what agents will do (or at least what we think they will not do), and what will follow from their decisions, considering their interaction.

We may talk about a prediction, or an equilibrium.

We will need to define equilibrium concepts, and consider their tradeoffs when making predictions.

Typically: if a prediction (ie., an equilibrium concept) is more precise, it has more restrictive assumptions (ie, applies to more restricted environments), and hence loses generality.

This is the basic tradeoff in theoretical work, mathematical or otherwise: generality x precision.

(All models make assumptions, mathematical or not…)

Dominated Strategies

Are there obvious predictions about how a game should be played?

Basic idea: a player should not choose something that is clearly bad for him.

This goes back to a basic, individual choice problem.

This may help us develop an equilibrium concept that does not rely on many assumptions.

Advertising War: Coke vs. Pepsi

- Let's begin describing the environment.
- Without any advertising, each company earns \$5b/year from Cola consumers.
- Each company can choose to spend \$2b/year on advertising.
- Advertising does not increase total sales for Cola, but if one company advertises while the other does not, it captures \$3b from the competitor.

		Pepsi	
		No Ad	Ad
Coke	No Ad	\$5b, \$5b	\$2b, \$6b
	Ad	\$6b, \$2b	\$3b, \$3b*

- Now, let's try to make predictions, and study them.
- What will the Cola companies do? **In other words: what's the equilibrium?**

- Is there a better feasible outcome then the actual result of their interaction? **In other words: is the equilibrium efficient, or has some other desirable property?**

Prisoners' Dilemma (PD)

Flood and Dresher (1950): RAND corporation's investigations into game theory for possible applications to global nuclear strategy

- Two persons are arrested for a crime.
- There is not enough evidence to convict either.
- Different cells, no communication.
- If a suspect testifies against the other ("Defect") and the other does not ("Cooperate"), the former is released and the latter gets a harsh punishment.
- If both prisoners testify, they share the punishment.
- If neither testifies, both serve time for a smaller offense.

	C	D
C	2, 2	0, 3
D	3, 0	1, 1*

- Each prisoner is better off defecting regardless of what the other does.

  We say D strictly dominates C for each prisoner.
- The resulting outcome is (D, D), which is worse than (C, C).
- This prediction does not assume any knowledge by player i about player j.
- It just assumes a player knows his own payoffs associated to each action, and can actually make a choice: bare minimum.
- So it's very general.
- But there is always the usual tradeoff: it is not very precise, as in many games it does not yield a prediction, nor rule out any action by any players.
- **In short: our first prediction, or equilibrium concept, is based on the idea that players do not play dominated strategies.**

Modified Prisoners' Dilemma

Consider the game obtained from the prisoners' dilemma by changing player 1's payoff for (C, D) from 0 to 2.

	C	D
C	2, 2	2, 3*
D	3, 0	1, 1

- No matter what player 1 does, player 2 still prefers D to C.
- This is helpful as it rules out some outcomes, but is less than precise than the previous example.
- **If player 1 knows that 2 never plays C**, then he prefers C to D, and we get to the final prediction: (C, D).
- Unlike in the prisoners' dilemma example, we use an **additional assumption** to reach our final prediction in this case: player 1 needs to deduce that player 2 never plays a dominated strategy.
- Less generality (due to additional assumption), but more precision (since we cannot make a unique prediction without this additional assumption).

Strictly Dominated Strategies

Definition 1

A strategy $s_i \in S_i$ is strictly (s.) dominated by $\sigma_i \in \Delta(S_i)$ if

$$u_i(\sigma_i, s_{-i}) > u_i(s_i, s_{-i}), \forall s_{-i} \in S_{-i}$$

That is, σ_i is strictly better than s_i for player i whatever the other players do.

Pure Strategies May Be Dominated by Mixed Strategies

	L	R
T	3, x	0, x
M	0, x	3, x
B	1, x	1, x

B is s . dominated by $1/2T + 1/2M$.

The Beauty Contest

- Players: everyone in the class
- Strategy space: any number in $\{1,2, \dots ,100\}$
- The person whose number is closest to $2/3$ of the class average wins the game.
- Payoffs: one randomly selected winner receives $1.

Why is this game called a beauty contest?

Keynesian Beauty Contest

Keynes described the action of rational actors in a market using an analogy based on a newspaper contest.

> Entrants are asked to choose a set of 6 faces from photographs that they find "most beautiful."
>
> Those who picked the most popular face are eligible for a prize.

A naive strategy would be to choose the 6 faces that, in the opinion of the entrant, are most beautiful.

A more sophisticated contest entrant, wishing to maximize the chances of winning against naive opponents, would guess which faces the majority finds attractive, and then make a selection based on this inference.

This can be carried one step further to account for the fact that other entrants would each have their own opinions of what public perceptions of beauty are.

What does everyone believe about what everyone else believes about whom others find attractive?

The Beauty Contest and the Stock Market

"It is not a case of choosing those faces that, to the best of one's judgment, are really the prettiest, nor even those that average opinion genuinely thinks the prettiest. We have reached the third degree where we devote our intelligences to anticipating what average opinion expects the average opinion to be. And there are some, I believe, who practice the fourth, fifth and higher degrees."

(John Maynard Keynes, General Theory of Employment, Interest and Money, 1936)

Keynes suggested that similar behavior is observed in the stock market.

Shares are not priced based on what people think their fundamental value is, but rather on what they think everyone else thinks the value is and what they think about these beliefs, and so on.

Iterated Deletion of Strictly Dominated Strategies

Our second equilibrium concept, less general but more precise than the previous one, is based on the idea that a player knows that other players do not play dominated strategies.

> Notice the additional assumption: now I player needs to know something about other players.

We can iteratively eliminate dominated strategies, under the assumption that "I know that you know that other players know. . . that everyone knows the payoffs and that no one would ever use a dominated strategy."

Definition 2

For all $i \in N$, set $S_i^0 = S_i$ and define S_i^k recursively by

$$S_i^k = \left\{ s_i \in S_i^{k-1} | \nexists | \sigma_i \in \Delta\left(S_i^{k-1}\right), u_i(\sigma_i, s_{-i}) > u_i(s_i, s_{-i}), \forall s_{-i} \in S_{-i}^{k-1} \right\}$$

The set of pure strategies of player i that survive iterated deletion of s. dominated strategies is

$$S_i^\infty = \cap_{k \geq 0} S_i^k$$

The set of surviving mixed strategies is

$$\{\sigma_i \in \Delta(S_i^\infty) \mid \nexists \sigma_i' \in \Delta(S_i^\infty), u_i(\sigma_i', s_{-i}) > u_i(\sigma_i, s_{-i}), \forall s_{-i} \in S_{-i}^\infty\}$$

Remarks

In a finite game, the elimination procedure ends in a finite number of steps, so S^∞ is simply the set of strategies left at the final stage.

In an infinite game, if S is a compact metric space and u is continuous, then one can use Cantor's theorem (a decreasing nested sequence of non-empty compact sets has nonempty intersection) to show that $S^\infty \neq \emptyset$.

Definition assumes that at each iteration all dominated strategies of every player are deleted simultaneously.

> In a finite game, the limit set S^∞ does not depend on the particular order in which deletion proceeds.

Outcome does not change if we eliminate s. dominated mixed strategies at every step.

> A strategy dominated against all pure strategies of the opponents iff it is dominated against all their mixed strategies.
>
> Eliminating mixed strategies for player i at any stage does not affect the set of s. dominated pure strategies for any $j \neq i$ at the next stage.

Detour on Common Knowledge

- Is common knowledge a sensible assumption? What does the definition of S_i^{100} entail?
- Higher order beliefs, common knowledge of rationality...
- Why did the strategy of choosing 1 not win in the beauty contest?

The Case of the Dishonest Business Partners

In a town, there are 100 business partnerships, each consisting of two partners, and a respected judge who ensures fairness in the community.

Before entering into any partnership, each partner had to pass a rigorous exam in logic and ethics.

It is common knowledge that the judge is always truthful in her statements.

Each partner in the town is aware of dishonest behavior in other partnerships but does not know if their own partner has been dishonest.

No one ever informs a partner if their own colleague has been dishonest.

The judge knows that some partners have been engaging in unethical practices and decides that such behavior should no longer be tolerated.

She gathers all the partners in the courthouse and announces that the integrity of the town's business community has been compromised—there is at least one dishonest partner among them.

She also reminds them that while no one knows if their own partner is dishonest, each person is aware of dishonest behavior in other partnerships.

The judge orders that any partner certain of their colleague's dishonesty should immediately dissolve the partnership at midnight, **and announce it to everyone immediately**.

39 silent nights passed, and then, on the 40th night, papers of dissolution were filed by several partners.

Questions:

- How many partnerships were dissolved?
- Were all the dishonest partners revealed?
- How did the partners learn of their colleague's dishonesty after 39 nights of silence?

Rationalizability

- Solution concept introduced independently by Bernheim (1984) and Pearce (1984).
- Like iterated dominance, rationalizability derives restrictions on play from common knowledge of payoffs and the fact that players are "reasonable."
- **Dominance**: unreasonable to use a strategy that performs worse than another (fixed) one in every scenario.
- **Rationalizability**: irrational for a player to choose a strategy that is not a best response to some *beliefs* about opponents' strategies.

What is a "Belief"?

- Bernheim & Pearce: every player i's beliefs σ_{-i} about the play of $j \neq i$ must be independent, i.e., $\sigma_{-i} \in \prod_{j \neq i} \Delta(S_j)$.
- Alternatively, allow player i to believe that the actions of opponents are correlated, i.e., any $\sigma_{-i} \in \Delta(S_{-i})$ is a possibility.
- The two definitions have different implications for $n \geq 3$.

Focus on case with correlated beliefs.

> Such beliefs represent a player's uncertainty about his opponents' actions, not necessarily his theory about their deliberate randomization and coordination.
>
> Player i may place equal probability on two scenarios: either both j and k pick action A or they both play B.
>
> If i is not sure which theory is true, then his beliefs are correlated even though he knows that j and k are acting independently.

Best Responses

Definition 3

A strategy $\sigma_i \in S_i$ is a **best response** to a belief $\sigma_{-i} \in \Delta(S_{-i})$ if

$$u_i(\sigma_i, \sigma_{-i}) \geq u_i(s_i, \sigma_{-i}), \forall s_i \in S_i$$

That is: σ_i is a best response if there is something feasible for other players I can think of (σ_{-i}) such that σ_i maximizes my payoff if they play σ_{-i}.

Rationalizabile Strategies

Common knowledge of payoffs and rationality imposes restrictions on play...

We can apply the same iterated deletion as before.

Definition 4

Set $S^0 = S$ and let S^k be given recursively by

$$S_i^k = \{s_i \in S_i^{k-1} \mid \exists \sigma_{-i} \in \Delta(S_{-i}^{k-1}), u_i(s_i, \sigma_{-i}) \geq u_i(s_i', \sigma_{-i}), \forall s_i' \in S_i^{k-1}\}$$

The set of correlated rationalizable strategies for player i is $S_i^\infty = \bigcap_{k \geq 0} S_i^k$.

A mixed strategy $\sigma_i \in \Delta(S_i)$ is rationalizable if there is a belief $\sigma_{-i} \in \Delta(S_{-i}^\infty)$ s.t. $u_i(\sigma_i, \sigma_{-i}) \geq u_i(s_i, \sigma_{-i})$ for all $s_i \in S_i^\infty$.

The definition of independent rationalizability replaces $\Delta(S_{-i}^{k-1})$ and $\Delta(S_{-i}^\infty)$ with $\prod_{j \neq i} \Delta(S_j^{k-1})$ and $\prod_{j \neq i} \Delta(S_j^\infty)$, resp.

Rationalizability in Cournot Duopoly

Two firms compete on the market for a divisible homogeneous good.

- Each firm $i = 1,2$ has zero marginal cost and simultaneously decides to produce an amount of output $q_i \geq 0$.
- The resulting price is $p = 1 - q_1 - q_2$.
- Profit of firm i is $q_i(1 - q_1 - q_2)$.

Best response of one firm if the other produces q is

$$B(q) = \max(0, (1 - q)/2)(j = 3 - i)$$

B is decreasing

If $q \lesseqgtr r$ then $B(q) \gtreqless (1 - r)/2$.

- Since $q \geq q^0 := 0$, only strategies $q \leq q^1 := B(q^0) = (1 - q^0)/2$ are best responses, $S_i^1 = [q^0, q^1]$.
- Then only $q \geq q^2 := B(q^1) = (1 - q^1)/2$ survives the second round of elimination, $S_i^2 = [q^2, q^1]$ ...
- We obtain a sequence

$$q^0 \leq q^2 \leq \cdots \leq q^{2k} \leq \cdots \leq q^{2k+1} \leq \cdots \leq q^1$$

where $q^{2k} = \sum_{l=1}^{k} 1/4^l = \left(1 - 1/4^k\right)/3$ and $q^{2k+1} = \left(1 - q^{2k}\right)/2$ s.t. $S_i^{2k+1} = [q^{2k}, q^{2k+1}]$ and $S_i^{2k} = [q^{2k}, q^{2k-1}]$ for all $k \geq 0$.

$\lim_{k\to\infty} q^k = 1/3$, so the only rationalizable strategy for firm i is $q_i = 1/3$ (the Nash equilibrium).

What strategies are rationalizable with more than two firms?

Never Best Responses

A strategy $\sigma_i \in \Delta(S_i)$ is **never a best response** for player i if it is not a best response to any correlated belief $\sigma_{-i} \in \Delta(S_{-i})$.

Recall that $\sigma_i \in \Delta(S_i)$ is s. dominated if $\exists \sigma_i' \in \Delta(S_i)$ s.t.

$$u_i(\sigma_i', s_{-i}) > u_i(\sigma_i, S_{-i}), \forall s_{-i} \in S_{-i}$$

Theorem 1

In a finite game, a strategy is never a best response iff it is s. dominated.

Corollary 1

Correlated rationalizability and iterated strict dominance coincide.

If σ_i is s . dominated by σ_i', then σ_i is not a best response for any belief $\sigma_{-i} \in \Delta(S_{-i})$: σ_i' yields a higher payoff than σ_i for player i against any σ_{-i}.

Left to prove that a strategy that is not s . dominated is a best response for some beliefs.

Proof

Suppose $\tilde{\sigma}_i$ is not s. dominated for player i.

- Define set of "dominated payoffs" for i by

$$D = \{x \in \mathbb{R}^{S_{-i}} \mid \exists \sigma_i \in \Delta(S_i), x \leq u_i(\sigma_i, \cdot)\}$$

D is non-empty, closed, and convex.

- $u_i(\tilde{\sigma}_i, \cdot)$ does not belong to the interior of D because it is not s. dominated by any $\sigma_i \in \Delta(S_i)$.
- By the supporting hyperplane theorem, $\exists \alpha \in \mathbb{R}^{S_{-i}} \setminus \{\mathbf{0}\}$ s.t.

$$\alpha \cdot u_i(\tilde{\sigma}_i, \cdot) \geq \alpha \cdot x, \forall x \in D$$

In particular, $\alpha \cdot u_i(\tilde{\sigma}_i, \cdot) \geq \alpha \cdot u_i(\sigma_i, \cdot), \forall \sigma_i \in \Delta(S_i)$.

- Since D is not bounded from below, α has non-negative components.
- Normalize α so that its components sum to 1; α interpreted as a belief in $\Delta(S_{-i})$ with the property that

$$u_i(\tilde{\sigma}_i, \alpha) \geq u_i(\sigma_i, \alpha), \forall \sigma_i \in \Delta(S_i)$$

Thus $\tilde{\sigma}_i$ is a best response to belief α.

Iteration and Best Responses

Theorem 2

For every $k \geq 0$, each $s_i \in S_i^k$ is a best response (within S_i) to a belief in $\Delta\left(S_{-i}^{k-1}\right)$.

Proof.

Fix $s_i \in S_i^k$; s_i is a best response within S_i^{k-1} to some $\sigma_{-i} \in \Delta\left(S_{-i}^{k-1}\right)$. If s_i were not a best response within S_i to σ_{-i}, let s_i' be a best response.

Since s_i is a best response within S_i^{k-1} to σ_{-i} and s_i' is a better response than s_i to σ_{-i}, we need $s_i' \notin S_i^{k-1}$.

Then s_i' was deleted at an earlier stage, say $s_i' \in S_i^{I-1}$ but $s_i' \notin S_i^I$ for some $I \leq k-1$. This contradicts the fact that s_i' is a best response in $S_i \supseteq S_i^{I-1}$ to $\sigma_{-i} \in \Delta\left(S_{-i}^{k-1}\right) \subseteq \Delta\left(S_{-i}^{l-1}\right)$.

Corollary 2

If the game is finite, then each $s_i \in S_i^\infty$ is a best response (within S_i) to a belief in $\Delta(S_{-i}^\infty)$.

Closed under Rational Behavior

Definition 5

A set $Z = Z_1 \times \ldots \times Z_n$ with $Z_i \subseteq S_i$ for $i \in N$ is closed under rational behavior if, for all i, every strategy in Z_i is a best response to a belief in $\Delta(Z_{-i})$.

Theorem 3

If the game is finite (or if S is a compact metric space and u is continuous), then S^∞ is the largest set closed under rational behavior.

Proof.

S^∞ is closed under rational behavior by Corollary 2. Suppose that there exists $Z_1 \times \ldots \times Z_n \not\subset S^\infty$ that is closed under rational behavior.

Consider the smallest k for which there is an i such that $Z_i \not\subseteq S_i^k$. It must be that $k \geq 1$ and $Z_{-i} \subseteq S_{-i}^{k-1}$.

By assumption, every element of Z_i is a best response to an element of $\Delta(Z_{-i}) \subset \Delta\left(S_{-i}^{k-1}\right)$, contradicting $Z_i \not\subseteq / S_i^k$.

Epistemic Foundations of Rationalizability

Formalize the idea of common knowledge and show that rationalizability captures the idea of common knowledge of rationality (and payoffs).

Definition 6 (Information Structure)

An information (or belief) structure is a list $(\Omega, (I_i)_{i \in N}, (p_i)_{i \in N})$

Ω is a finite state space

$I_i: \Omega \to 2^\Omega$ is a partition of Ω for each $i \in N$ s.t. $l_i(\omega)$ is the set of states that i thinks are possible when the true state is ω;

$\omega' \in I_i(\omega) \Leftrightarrow \omega \in I_i(\omega')$

$p_{i,l_i(\omega)} \in \Delta(l_i(\omega))$: i's belief at ω

Interpretation

- State ω summarizes all the relevant facts about the world. Only one of the states is true; all others are hypothetical states needed to encode players' beliefs.
- In state ω, player i is informed that the state is in $l_i(\omega)$ and gets no other information.
- Such an information structure arises if each player observes a state-dependent signal and $l_i(\omega)$ is the set of states for which player i's signal is identical to the signal at ω.

Knowledge and Common Knowledge

Definition 7

For any event $F \subseteq \Omega$, player i knows at ω that F obtains if $l_i(\omega) \subseteq F$. The event that i knows F is

$$K_i(F) = \{\omega \mid l_i(\omega) \subseteq F\}$$

The event that everyone knows F is defined by

$$K(F) = \cap_{i \in N} K_i(F)$$

Let $K^0(F) = F$ and $K^{t+1}(F) = K(K^t(F))$ for $t \geq 0$. Set $K^\infty(F) = \bigcap_{t \geq 0} K^t(F)$. $K^\infty(F)$ is the set of states where F is common knowledge.

Public Events

$K(K^\infty(F)) = K^\infty(F) \to$ alternative definition of common knowledge Event F' is public if $F' = \cup_{\omega' \in F'} l_i(\omega')$ for all i. If F' is public, then $K(F') = F'$, so $K^\infty(F') = F'$.

Lemma 1

An event F is common knowledge at ω iff there exists a public event F' with $\omega \in F' \subseteq F$.

If F is common knowledge at ω, there exists a submodel that includes state ω and respects the information structure where F is true state by state.

Strategies

Fix a finite game (N, S, u). To give strategic meaning to information states, introduce a strategy profile s : $\Omega \to S$.

Definition 8

A strategy profile $\mathbf{s}\colon \Omega \to S$ is adapted with respect to $(\Omega, (l_i)_{i\in N}, (p_i)_{i\in N})$ if $\mathbf{s}_i(\omega) = \mathbf{s}_i(\omega')$ whenever $I_i(\omega) = I_i(\omega')$.

Players must choose a constant action at all states in each information set since they cannot distinguish states in the same information set.

Definition 9

An epistemic model $(\Omega, (I_i)_{i\in N}, (p_i)_{i\in N}, \mathbf{s})$ consists of an information structure and an adapted strategy profile.

Common Knowledge of Rationality

Definition 10

For an epistemic model $(\Omega, (I_i)_{i\in N}, (p_i)_{i\in N}, \mathbf{s})$, player i is rational at $\omega \in \Omega$ if

$$\mathbf{s}_i(\omega) \in \arg\max_{s_i \in S_i} \sum_{\omega' \in l_i(\omega)} u_i\left(s_i, \mathbf{s}_{-j}(\omega')\right) p_{i,l_i(\omega)}(\omega')$$

Definition 11

A strategy $s_i \in S_i$ consistent with common knowledge of rationality if there exists a model $\left(\Omega, (l_j)_{j\in N}, (p_j)_{j\in N}, \mathbf{s}\right)$ and state $\omega^* \in \Omega$ with $\mathbf{s}_i(\omega^*) = s_i$ at which it is common knowledge that all players are rational.

Equivalently, $\exists \left(\Omega, (l_j)_{j\in N}, (p_j)_{j\in N}, \mathbf{s}\right)$ s.t. $\mathbf{s}_j(\omega)$ is a best response to $\mathbf{s}_{-j}$ at each $\omega \in \Omega$ for every player $j \in N$.

Result

Theorem 4

For any $i \in N$ and $s_i \in S_i$, s_i is consistent with common knowledge of rationality iff $s_i \in S_i^\infty$.

Can extend result to allow for payoff uncertainty (adding the hypothesis that payoffs are common knowledge at the relevant state).

Proof

($\Rightarrow$) Fix s_i consistent with common knowledge of rationality. $\exists \left(\Omega, \left(l_j\right)_{j\in N}, \left(p_j\right)_{j\in N}, \mathbf{s}\right)$ with $\omega^* \in \Omega$ s.t. $\mathbf{s}_i(\omega^*) = s_i$ and

$$\mathbf{s}_j(\omega) \in \arg\max_{s_j\in S_j} \sum_{\omega'\in l_j(\omega)} u_j\left(s_j, \mathbf{s}_{-j}(\omega')\right) p_{j,l_j(\omega)}(\omega'), \forall j \in N, \omega \in \Omega$$

Define $Z_j = \mathbf{s}_j(\Omega)$. Note that $s_i = \mathbf{s}_i(\omega^*) \in s_i(\Omega) = Z_i$. By Theorem 3, to show that $s_i \in S_i^\infty$, it suffices to prove that Z is closed under rational behavior.

$\forall z_j \in Z_j, \exists \omega \in \Omega$ s.t. $z_j = \mathbf{s}_j(\omega)$. Define $\mu_{j,\omega} \in \Delta\left(Z_{-j}\right)$ by

$$\mu_{j,\omega}\left(s_{-j}\right) = \sum_{\omega'\in l_j(\omega), \mathbf{s}_{-j}(\omega')=s_{-j}} p_{j,l_j(\omega)}(\omega')$$

Then

$$z_j = \mathbf{s}_j(\omega) \in \arg\max_{s_j\in S_j} \sum_{\omega'\in \ell_j(\omega)} u_j\left(s_j, \mathbf{s}_{-j}(\omega')\right) p_{i,l_j(\omega)}(\omega')$$

$$= \arg\max_{s_j\in S_j} \sum_{s_{-j}\in Z_{-j}} \mu_{j,\omega}\left(s_{-j}\right) u_j\left(s_j, s_{-j}\right)$$

Proof

($\Leftarrow$) Since S^∞ is closed under rational behavior, for every $s_i \in S_i^\infty$, there exists $\mu_{i,s_i} \in \Delta(S_{-i}^\infty)$ for which s_i is a best response. Define the model $(S^\infty, (l_i)_{i\in N}, (p_i)_{i\in N}, \mathbf{s})$:

$$l_i(s) = \{s_i\} \times S_{-i}^\infty$$

$$p_{i,s}(s') = \mu_{i,s_i}(s'_{-i})$$

$$\mathbf{s}(s) = s$$

In this model, it is common knowledge that every player is rational:

$$\forall s \in S^\infty, s_i(s) = s_i \in \arg\max_{s_i'\in S_i} \sum_{s_{-i}\in S_{-i}^\infty} u_i(s_i', s_{-i})\mu_{i,s}(s'_{-i})$$

$$= \arg\max_{s_i'\in S_i} \sum_{s'\in l_i(s)} u_i(s_i', s_{-i}) p_{i,s}(s')$$

For every $s_i \in S_i^\infty$, there exists $s = (s_i, s_{-i}) \in S^\infty$ s.t. $s_i(s) = s_i$, showing that s_i is consistent with common knowledge of rationality.

Nash Equilibrium

Many games are not dominance solvable.

Nevertheless, the involved parties find a solution.

	H	T
H	1, −1	−1, 1
T	−1, 1	1, −1

Matching Pennies has no Nash equilibrium (in pure strategies…).

A Nash equilibrium is a strategy profile with the property that no player can benefit by deviating from his corresponding strategy.

Definition 12 (Nash 1950)

A mixed-strategy profile σ^* is a Nash equilibrium if for every $i \in N$,

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(s_i, \sigma_{-i}^*), \forall s_i \in S_i$$

Interpretation:

It's an intersection of best responses.

When everyone plays Nash, I better play Nash.

There is no incentive for <u>individual</u> deviation.

2,3	5,1	3,2
1,4	4,5	0,6
0,1	2,2	4,1

Remarks

- If in equilibrium a player uses a mixed strategy that places positive probability on several pure strategies, he must be indifferent between all pure strategies in its support.
- The fact that there is no profitable deviation in pure strategies implies there is no profitable deviation in mixed strategies either.
- Strategies that do not survive iterated strict dominance (or are not rationalizable) cannot be played with positive probability in a Nash equilibrium.

Let's prove the first of these points. (The others are left as exercises.)

Statement:

$S_i^+ \subseteq S\_i$ is the set of pure strategies played with strictly positive probability in a mixed strategy $\sigma = (\sigma_1, \dots, \sigma_I)$. Strategy profile σ is a Nash equilibrium if and only if for all players:

$u_i(s_i, \sigma_{-i}) = u_i(s_i', \sigma_{-i})$ for all $s_i, s_i' \in S_i^+$

$u_i(s_i, \sigma_{-i}) \geq u_i(s_i', \sigma_{-i})$ for all $s_i \in S_i^+, s_i' \notin S_i^+$

Proof:

Necessity: if σ is a Nash equilibrium, then i and ii must hold.

Assume first i does not hold. Hence there are strategies $s_i, s_i' \in S_i^+$ such that $u_i(s_i, \sigma_{-i}) \neq u_i(s_i', \sigma_{-i})$. Without loss of generality, assume $u_i(s_i, \sigma_{-i}) > u_i(s_i', \sigma_{-i})$. Hence player i may increase his payoff by playing s_i whenever σ_i indicates s_i', and increase his payoff. This means he increases the probability of s_i and decreases of s_i' given that both are played with strictly positive probability under σ. But then σ_i is not optimal, and hence there is a profitable individual deviation: σ is not a Nash equilibrium.

Assume now that ii does not hold. Then there are $s_i \in S_i^+, s_i' \notin S_i^+$ such that $u_i(s_i, \sigma_{-i}) < u_i(s_i', \sigma_{-i})$. Hence player i may increase his payoff by playing s_i' whenever σ_i indicates s_i, and increase his payoff. Since s_i is played with strictly positive probability under σ, we have again a profitable individual deviation, and again σ is not a Nash equilibrium.

In short, we have proven that if i and ii do not hold simultaneously, then σ is not a Nash equilibrium. This is equivalent to stating that if σ is a Nash equilibrium, then i and ii must hold, proving necessity.

Sufficiency: if i and ii hold, then σ is a Nash equilibrium.

Suppose i and ii hold but σ is not a Nash equilibrium. We will show that this leads to an absurd conclusion, and hence establish that i and ii imply that σ is a Nash equilibrium.

If σ is not a Nash equilibrium, then by definition there must be some profitable individual deviation for some player i: that is, some strategy σ_i' such that $u_i(\sigma_i', \sigma_{-i}) > u_i(\sigma_i, \sigma_{-i})$. If this holds, then there must be at least one pure strategy s_i' played with strictly positive probability under σ_i' such that $u_i(s_i', \sigma_{-i}) > u_i(\sigma_i, \sigma_{-i})$. But we know that $u_i(\sigma_i, \sigma_{-i}) = u_i(s_i, \sigma_{-i})$ for all $s_i \in S_i^+$. Hence $u_i(s_i', \sigma_{-i}) > u_i(s_i, \sigma_{-i})$. Hence i does not hold, which is an absurd because we assumed that it holds.

Back to Matching Pennies

Matching Pennies does have a Nash equilibrium in mixed strategies.

How to find it?

Two ways:

1. Solve the complete problem
2. Use Remark 1above, noticing its limitations

What Are the Assumptions behind NE?

- Nash equilibria are "consistent" predictions (or "stable" conventions) of how the game will be played.
- If all players expect that a specific Nash equilibrium will arise, then no player has incentives to play differently.
- Each player must have correct conjectures about the strategies of his opponents and play a best response to his conjecture.
- We interpret mixed strategies as beliefs regarding opponents' play, not necessarily as deliberate randomization.
- Assumes knowledge of strategies (beliefs) and rationality.

Do Soccer Players Flip Coins?

Penalty kicks

- Kicker's strategy space: $\{L, M, R\}$
- Goalie's strategy space: $\{L, M, R\}$

-What are the payoffs?

-What's the Nash equilibrium?

- Simultaneous move game? (125mph, 0.2 seconds reaction time)
- What do players do in reality?

Chiappori, Levitt, and Groseclose (2002):

- 459 kicks in French and Italian first leagues
- 162 kickers, 88 goalies

Matrix of Shots Taken

	Kicker			
Goalie	Left	Middle	Right	Total

Left	117	48	95	260
Middle	4	3	4	11
Right	85	28	75	188
Total	206	79	174	459

Notes: The sample includes all French first-league penalty kicks from 1997-1999 and all Italian first-league kicks (1997-2000). For shots involving left-footed kickers, the directions have been reversed so that shooting left corresponds to the "natural" side for all kickers.

Tennis Service Game

Player 1 chooses whether to serve to player 2's forehand, center or backhand side, and player 2 chooses which side to favor for the return.

Unique mixed strategy equilibrium, which puts positive probability only on strategies C and B for either player.

	F	C	B
F	0, 5	2, 3	2, 3
C	2, 3	0, 5	3, 2
B	5, 0	3, 2	2, 3

- For player 1 , playing C with probability ϵ and B with probability $1-\epsilon$s. dominates F.
- If player 1 never chooses F, then C s. dominates F for player 2.
- In the remaining 2×2 game, there is a unique equilibrium, in which both players place probability $1/4$ on C and $3/4$ on B.

Multiple Equilibria

Nash equilibrium needs not be unique.

	L	R
L	1, 1	0, 0
R	0, 0	1, 1

	T	S
T	3, 2	1, 1
S	0, 0	2, 3

The concept of Nash equilibrium is silent about the selection of equilibria.

Stag Hunt

Each player can choose to hunt hare by himself or hunt stag with the other.

Stag offers a higher payoff, but only if players team up.

	S	H
S	9, 9	0, 8
H	8, 0	7, 7

The game has two pure strategy Nash equilibria: (S, S) and (H, H)

And also a mixed strategy Nash equilibrium: $(7/8S + 1/8H, 7/8S + 1/8H)$.

Which Equilibrium is More Plausible?

- We may expect (S, S) to be played because it is Pareto dominant. However, if one player expects the other to hunt hare, he is much better off hunting hare himself; and the potential downside of choosing stag is bigger than the upside-hare is the safer choice.
- Harsanyi and Selten (1988): H is the risk-dominant action-if each player expects the other to choose either action with probability $1/2$, then H has a higher expected payoff (7.5) than S (4.5).
- For a player to optimally choose stag, he should expect the other to play stag with probability $\geq 7/8$.
- Coordination problem may persist even if players communicate: regardless of what i intends to do, he would prefer j to play stag, so attempts to convince j to play stag are cheap talk.

Epistemic Foundations

- Aumann and Brandenburger (1995): a framework that can be used to examine the epistemic foundations of Nash equilibrium.
- The primitive of their model is an interactive belief system in which

  there is a possible set of types for each player

  each type has associated to it a payoff for every action profile, a choice of which action to play, and a belief about the types of the other players.
- In a 2-player game, if the game being played, the rationality of the players, and their conjectures are all mutually known, then the conjectures constitute a Nash equilibrium.
- For games with more than 2 players, we need to assume additionally that players have a common prior and that conjectures are commonly known.

  This ensures that any two players have identical and separable (independent) conjectures about other players, consistent with a mixed strategy profile.

Evolutionary Foundations

- Solution concepts motivated by presuming that players make predictions about their opponents' play by introspection and deduction, using knowledge of their opponents' payoffs, rationality...
- Alternatively, assume players extrapolate from past observations of play in "similar" games and best respond to expectations based on past observations.
- Cournot (1838) suggested that players take turns setting their outputs in the duopoly game, best responding to the opponent's last-period action.
- Simultaneous action updating, best responding to average play, populations of players anonymously matched (another way to think about mixed strategies), etc.
- If the process converges to a particular steady state, then the steady state is a Nash equilibrium.

Convergence

How sensitive is the convergence to the initial state? If convergence obtains for all initial strategy profiles sufficiently close to the steady state, we say that the steady state is asymptotically stable.

Shapley (1964) Cycling

	L	M	*R*
U	0,0	4,5	5,4
M	5,4	0,0,	4,5

D	4,5	5,4	0,0

Remarks

- Evolutionary processes are myopic and do not offer a compelling description of behavior.
- Such processes do not provide good predictions for behavior in the actual repeated game, if players care about play in future periods and realize that their current actions can affect opponents' future play.

Existence of Nash Equilibrium

Prove the existence of Nash equilibria in a more general setting.

- Continuity and compactness assumptions are indispensable, usually needed for the existence of solutions to optimization problems.
- Convexity is usually required for fixed-point theorems.

Topology prerequisites

Correspondences

Topological vector spaces S and Y

- A correspondence $B: S \rightrightarrows Y$ is a set valued function taking elements $s \in S$ into subsets $B(s) \subseteq Y$.
- $G(B) = \{(s, y) \mid y \in B(x)\}$: graph of B
- $s \in S$ is a fixed point of B if $s \in B(s)$
- B is non-empty/closed-valued/convex-valued if $B(s)$ is non-empty/closed/convex for all $s \in S$.

Closed Graph

- A correspondence B has closed graph if $G(B)$ is a closed subset of $S \times Y$.

- If S and Y are first-countable spaces (such as metric spaces), then B has closed graph iff for any sequence $(s_m, y_m)_{m\geq 0}$ with $y_m \in B(s_m)$ for all $m \geq 0$, which converges to a pair (s, y), we have $y \in B(s)$.
- Correspondences with closed graph are closed-valued.

The Maximum Theorem

Theorem 7 (Berge's Maximum Theorem)

Suppose that $u: S \times Y \to \mathbb{R}$ is a continuous function, where S and Y are metric spaces and Y is compact.

(1) The function $V: S \to \mathbb{R}$, defined by

$$V(x) = \max_{y\in Y} u(s, y)$$

is continuous.

(2) The correspondence $B: S \rightrightarrows Y$,

$$B(x) = \arg\max_{y\in Y} u(s, y)$$

is nonempty valued and has a closed graph.

A Fixed-Point Theorem

Theorem 8 (Kakutani's Fixed-Point Theorem)

Let S be a (i) non-empty, (ii) compact, and (iii) convex subset of a Euclidean space
Let the correspondence $B: S \rightrightarrows S$ have (iv) closed graph and (v) non-empty (vi) convex values.
Then the set of fixed points of B is non-empty and compact.

In game theoretic applications, S is usually the strategy space, assumed to be compact and convex when we include mixed strategies – that is, if we consider $\Delta(S)$.

B is typically the best response correspondence, which is non-empty valued and has a closed graph by the Maximum Theorem.

Hence we need to check that S has properties i to iii, and B has properties iv to vi.

Convexity

To ensure that B is convex-valued, assume that payoff functions are quasi-concave.

Definition 13

If S is a convex subset of a real vector space, then the function $u: S \to \mathbb{R}$ is quasi-concave if

$$u(ts + (1-t)y) \geq \min(u(s), f(y)), \forall t \in [0,1], s, y \in S$$

Quasi-concavity implies convex upper contour sets and convex arg max.

Now we can turn to the proof of existence of Nash equilibria.

Existence of Nash Equilibrium

Theorem 9

Consider a game (N, S, u) such that S_i is a convex, non-empty and compact subset of a Euclidean space and that u_i is continuous in s and quasi-concave in s_i for all $i \in N$. Then there exists a pure strategy Nash equilibrium.

The result implies the existence of pure strategy Nash equilibria in generalizations of the Cournot competition game.

Importantly, theorem 9 also implies the existence of mixed strategy Nash equilibria in finite games.

> **S is not necessarily convex nor compact in many applications, but $\Delta(S)$ is always convex and compact.**

Proof

Let $B_i(s_{-i}) := \arg\max_{s_i' \in S_i} u_i(s_i', s_{-i})$ and define $B: S \rightrightarrows S$,

$$B(s) = \{(s_1^*, \dots, s_n^*) \mid s_i^* \in B_i(s_{-i}), \forall i \in N\} = \prod_{i \in N} B_i(s_{-i}), \forall s \in S$$

Since S is compact and the utility functions are continuous, the Maximum Theorem implies that B_i and B are non-empty valued and have closed graphs.

As u_i is quasi-concave in s_i, the set $B_i(s_{-i})$ is convex for all i and s_{-i}, so B is convex-valued.

Kakutani's fixed-point theorem $\Rightarrow$ F has a fixed point:

$$s^* \in F(s^*)$$

$s_i^* \in B_i(s_{-i}^*), \forall i \in N \Rightarrow s^*$ is a Nash equilibrium.