Section notes for POLSCI 681: Intermediate Game Theory

Samuel Baltz
sbaltz@umich.edu
Winter 2019

Written to complement lectures by Scott Page

Contents

1 Section 1: 2019 January 10

Topics:

- Utility function definition [1]
- Utility function axioms [2]
  ↪ Rationality [3]
- Condorcet voting paradox [4]

After this class, I expect every student to be able to:

- describe what a utility function is and how it relates to preference relations
- list the conditions that a preference relation must satisfy for it to have an associated utility function
- visualise what it means to aggregate individual preferences

BACK TO TABLE OF CONTENTS

① Utility functions:

↪In lecture we considered the idea of a collection of alternatives X, where people are able to choose between any two alternatives in X. Let's use an example to reconstruct the idea of a utility function, which we saw in class.

Example: Say that X is the set $X \equiv \{\text{ducks, squirrels, rabbits}\}$, and suppose that I have the preference ordering ducks $\succ$ rabbits $\succeq$ squirrels. Then one utility function corresponding to this preference ranking would be $u(\text{ducks}) = 2$, $u(\text{rabbits}) = 1$, and $u(\text{squirrels}) = 0$. Another utility function corresponding to this preference ranking would be $u(\text{ducks}) = 7$, $u(\text{rabbits}) = -1$, and $u(\text{squirrels}) = -1$. So, one preference ordering can correspond to many different utility functions.

↪Preferences are the fundamental building block of choice theory, and a utility function is defined with respect to a preference relation. Let's make this idea concrete:

Definition: A function $u : X \to \mathbb{R}$ is a **utility function** corresponding to the preference relation $\succeq$ if, $\forall x, y \in X$, $x \succeq y \iff u(x) \geq u(y)$

② Existence of utility functions:

Theorem: In lecture we saw the **von Neumann-Morgenstern Expected Utility Theorem**, which asserts that a preference relation $\succeq$ corresponds to at least one utility function if and only if that preference relation satisfies all of the following four axioms: completeness, transitivity, independence of irrelevant alternatives, and continuity. We spent some time going over these properties, formally and in plain language, because this fundamental theorem is at the heart of game theory.

↪ **Axiom of completeness:** $x \succeq y$ or $y \succeq x$ $\forall x, y \in X$. So, we can make a choice between any two alternatives. Either I prefer ducks to squirrels, I prefer squirrels to ducks, or I'm indifferent between them.

↪ **Axiom of transitivity:** $x \succeq y,\ y \succeq z \implies x \succeq z$ $\forall x, y, z \in X$. This rules out cycles: I can't prefer ducks to squirrels, squirrels to rabbits and rabbits to ducks. The problem with cycles is that every alternative in the cycle is preferred to every other alternative in the cycle, which means that I cannot name my favourite alternative (I like ducks more than squirrels, and I like squirrels more than rabbits, and I like rabbits more than ducks, and I like ducks more than squirrels, ...).

↪ **Axiom of the independence of irrelevant alternatives (IIA):** Suppose that when choosing between x and y, you will choose x, meaning that $x \succ y$. IIA holds if, when choosing between x, y, and some irrelevant option z, you still prefer x to y, no matter what your opinion is about z. So if z is "irrelevant" to x and y, then our opinion about z does not affect our ranking of x and y.

To understand IIA, it's crucial to notice that some alternatives might not be independent of other alternatives! Amartya Sen gives a remarkable example of this (Binmore, 2009: excellently explained on p. 11, which is where I learned it). A lady is invited to tea and plans to accept, until she is informed that along with tea she will also have the opportunity to snort cocaine. This expansion of her choice set X is definitely *not* an irrelevant alternative, because it gives her some important information about what going to tea will be like! So this lady's preference relation can satisfy IIA while nonetheless allowing the information that there will be cocaine at her friend's house to change her ranking of going to tea versus not going to tea.

Now, what would it look like if the choice had been irrelevant? Let's add this to Sen's example: if she had instead been told that, while at tea, she would have the opportunity to take some cardboard home with her, then in order to satisfy IIA, this information cannot change her preference about whether or not she would like to go to tea. Since the cardboard is an irrelevant alternative, in order to satisfy this axiom, her preference between going to tea and not going to tea is independent of her access to cardboard.

$\hookrightarrow$ **Axiom of continuity:** If $x \succ y \succ z$, then $\exists \alpha \in [0;1]$ which is unique so that $\alpha x + (1-\alpha)z \sim y$. So when we have some alternative y in the middle, we can always weight an alternative that we prefer to y and an alternative that we prefer y to so that this linear combination of those alternatives gives us the same utility as y. Continuity is not nicely interpretable, so we didn't dwell on it.

If a preference relation $\succeq$ satisfies all four of the above axioms, then we can construct at least one utility function which corresponds to $\succeq$.

Definition: If a preference relation $\succeq$ satisfies all four axioms of the von Neumann-Morgenstern Expected Utility Theorem, then we call $\succeq$ **rational**. Any person who uses that preference relation is said to have rational preferences.

Note: We have just defined rationality to be a property of a preference relation, *not* a property of a person! This theoretically crucial distinction is frequently overlooked in political science.

$\hookrightarrow$Let's see for ourselves why the von Neumann-Morgenstern Expected Utility Theorem makes sense.

Exercise: Your goal is to design a perfectly irrational preference relation. Figure out how to define a preference relation over ducks, squirrels, rabbits, and hedgehogs that simultaneously violates completeness, transitivity, and independence of irrelevant alternatives (I leave out continuity because there's nothing interesting involved in violating continuity).

Exercise: Say ducks $\succeq$ squirrels, squirrels $\succeq$ rabbits, and rabbits $\succeq$ ducks. Try to write down a utility function that captures this relationship. Why can't you do it?

Exercise: Say ducks $\succeq$ squirrels and ducks $\succeq$ rabbits, but you simply can't choose between squirrels and rabbits. It's not that you're indifferent; you just have no idea how good they are compared to each other. Again try to write down a utility function that captures this relationship, and explain why it's impossible.

Remark: Evidently, when we talk about "rationality", we're really talking about something very technical that might be better described by the word "coherence" or "consistency".

Exercise: Define a rational preference relation over ducks, squirrels, and rabbits, and then write two different utility functions that correspond to that preference relation.

③ Condorcet voting paradox:

↪To set the stage for Arrow's Theorem and the Gibbard-Satterthwaite Theorem in your homework, let's examine a very simple example in which individual preferences produce a surprising result when we aggregate them up to the level of a community. In particular, Condorcet's Paradox shows that under very straightforward and common-sense rules, a community of individually rational agents can produce group-level decisions that are irrational.

Example: Consider a community in which electors $1, 2, 3$ vote over three alternatives a, b, c by a simple sincere majority vote in pairwise contests. Let's say that these electors have the following preference ranking:

$1 : c \succ b \succ a$
$2 : a \succ c \succ b$
$3 : b \succ a \succ c$

Notice that these electors each individually satisfy our definition of rationality; they have a complete ranking which is transitive (and without loss of generality we can assume that they also satisfy independence of irrelevant alternatives and continuity). Now, consider which alternative wins each possible pairwise contest:

a vs b: b wins with 2 votes.
b vs c: c wins with 2 votes.
c vs a: a wins with 2 votes.

So the community has the following preference ordering: $a \succ c \succ b \succ a$. This is not a transitive preference ordering! So it is very easy for individually rational agents to form an irrational community. Arrow's Theorem is a generalization of this exact type of problem, but it produces a much stronger and more fundamental result.

2 Section 2: 2019 January 17

Topics:

- Showing results on a problem set [1]
- Pivotality and strategic voting [2]
- Limits of Arrow's Theorem [3]
  - ↪ Single-peaked preferences [4]

After this class, I expect every student to be able to:

- explain what pivotality is in the context of elections
- name at least one way to eliminate the contradiction in Arrow's Theorem other than violating the dictator-free assumption
- identify whether or not a utility function is single-peaked

BACK TO TABLE OF CONTENTS

①**Showing results on a problem set**:

↪Let's use question 2a in Lecture Notes 2 as an example of the range of approaches that are okay in this class. This question looks like it is asking for a pretty severe combinatorics proof. But what we're looking for in this class is reasoning. To illustrate this, I included my answer as an example of the sheer variety of answers that I'm happy to accept in this class.

Example: Here's my personal answer, seeking to show that $d(\sigma, \hat{\sigma}) = d(\hat{\sigma}, \sigma)$, which is a pretty informal proof of the result which I present in the form of a procedure to construct $\hat{\sigma}$ from σ, representing each as a concatenation of n digits, as $\sigma = \sigma_1\sigma_2\cdots\sigma_n$ and $\hat{\sigma} = \hat{\sigma}_1\hat{\sigma}_2\cdots\hat{\sigma}_n$:

```
While σ ≠ σ̂ do:
    For i in 1:n do:
        If σ_i matches σ̂_i : σ̂_{i-1} do:
            σ_{i-1} ↔ σ_i
        ElseIf σ_i matches σ_{i+1} : σ_n do:
            σ_{i+1} ↔ σ_i
        EndIf
    EndFor
EndWhile
```

Since σ and $\hat{\sigma}$ are wlog interchangeable in this algorithm, necessarily it takes the same number of iterations to transform σ into $\hat{\sigma}$ as to transform $\hat{\sigma}$ into σ.

Note: Questions in this class will often have many different answers, and what's important is that your approach shows that you fully understood and engaged with the question. What really matters on problem sets in this class is that you are fully responding to the question. When the question asks you to show a general statement, then you need to respond with some sort of general reasoning; an example doesn't suffice. But there is almost never going to be just one way to answer a question, so don't be afraid to take creative approaches.

②**Pivotality and strategic voting**:

↪Let's use the concept of strategic voting to illustrate this idea of pivotality that we've been discussing.

Example: Suppose we have a set of alternatives X, where $X = \{\text{vote}, \text{don't vote}\}$. We have this classical picture in political science, due to Riker and Ordeshook (1968), that says that we can easily argue that for all potential voters, $u(\text{don't vote}) >> u(\text{vote})$. The argument traditionally goes: voting has real tangible costs, in terms of time and even money, for every voter. But whether or not I vote is almost certain to never affect the outcome of an election. So, I should expect a higher utility from voting than from not voting. **Question:** Why does this argument require that we assume that voters are rational?

Definition: In the context of an election, we call an elector **pivotal** if they alone can determine the outcome of the election.

Theorem: The notion of pivotality in the context of the paradox of voter turnout leads to a contradiction. Suppose we accept that every voter gets a bigger utility from not voting than from voting. For one voter, they realize that they have the power of pivotality: by being the only person who votes, they can indeed sway the outcome of the election. And here is the contradiction, because this logic applies equally to every voter, so nobody should expect to be the only person who chooses to vote. We will need game theory to try to resolve this paradox (although it remains a semi-open problem in political science).

Note: We have previously talked about pivotality by another name: a voter who is pivotal is a "dictator" in the sense that we used it while discussing Arrow's Theorem.

(3) **Limits of Arrow's Theorem**:

↪We've discussed Arrow's Theorem as a very serious and fundamental threat to some of the properties that normative democratic theory would like elections to have. So now let's talk about one of the main ways out of Arrow's Theorem.

Recall: Arrow's Theorem tells us that, by assuming 4 simple axioms that describe desirable properties of an election, we arrive at a contradiction, so we are forced to admit one undesirable property. Usually, the violation that's made is the flashiest one: the most explosive result occurs when we resolve the contradiction by violating the axiom that the community is dictator-free. But we could just as easily violate one of the other axioms. Let's remember what those axioms are.

Recall: We reviewed, in plain language (because we had seen a formal treatment in lecture), the four properties of a communal preference ordering which **Arrow's Theorem** claims cannot coexist:

↪ **Dictator-free:** No single voter dictates the preference of the community.

↪ **Universal:** Any possible set of voter preferences corresponds to a unique and complete ordering of alternatives.

↪ **Independent of irrelevant alternatives:** The social preference between two alternatives depends only on the individual preferences between those alternatives. Notice that this is very different from IIA in the context of individual rationality.

↪ **Unanimous:** If every individual prefers some alternative to all other alternatives, then the social preference ranking must also prefer that alternative.

Note: There are plenty of ways out of Arrow's Theorem that are less flashy than allowing a dictator. We could for example allow transitivity in the group preference and violate universality, as we do when constructing the condorcet voting paradox.

Theorem: Another way to get out of both Arrow's Theorem and the condorcet voting paradox is to put a simple constraint on the utility function of every individual in a community, called single-peakedness. By adding an assumption, we can assume away the contradiction which arises when we acecpt only the four Arrow's Theorem assumptions.

Definition: A preference relation $\succeq$ is **single-peaked** over alternatives $\{x_1, x_2, \cdots, x_n\}$ if there is some unique x_i so that, for any $i, j, k \in \{1, 2, \cdots, n\}$,

$$x_j < x_k \leq x_i \implies x_k \succ x_j$$

and

$$x_j > x_k \geq x_i \implies x_k \succ x_j$$

Example: The following utility functions are all single-peaked:

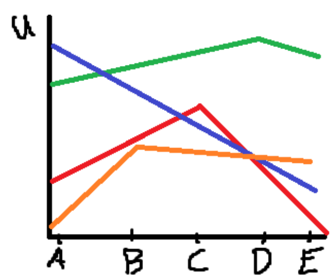

And none of the following utility functions are single-peaked:

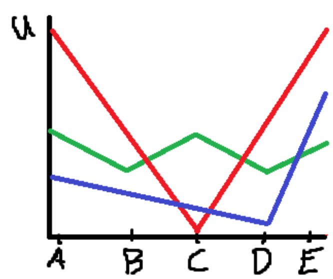

Note: We have to impose some underlying order for this to make sense; in the above example, there has to be some clear way to order the alternatives A, B, C, D, E. This is not too onerous when we're talking about elections; for example, when we study spatial models, we'll talk about the pros and cons of placing candidates on a left-right spectrum.

Theorem: If every individual in the community has single-peaked preferences, then Arrow's Theorem does not apply, since the median position in the community will win a pairwise contest with any other position.

3 Section 3: 2019 January 24

Topics:

- Normal form games [1]
- Nash Equilibria [2]

After this class, I expect every student to be able to:

- read and write a normal form game
- define Nash Equilibrium

BACK TO TABLE OF CONTENTS

① Normal form games

RECALL: A **strategic game** consists of:

1) Players, choosing
2) Strategies, to maximize
3) Payoffs

↪Where did the idea of "maximizing" come from? (And by extension, why did you spend 4 months learning about optimization in POLSCI 598?). This idea is absolutely central to game theory.

Note: We assume that people are rational, so their utilities (∼ aka payoffs) reveal their preferences to us. If people in decision theoretic situations (just 1 player) make decisions to maximize their payoffs, then we know what their preferences are, because if they decide to get x instead of y then we know that $u(x) \geq u(y)$, and so $x \succeq y$. Since preferences were the axiom of the system and the quantity we're ultimately interested in, *this* is why we prefer to focus only on rational decision-makers.

Note: We call this acting **sincerely**; sincere decision-making is closely connected to decision theory. If we make a decision-makers payoffs partly a function of another players' decisions, then they have to start behaving **strategically**, especially if the situation is competitive or adversarial. This is the purview of decision theory.

↪For this reason, "strategic" ≡ "other-consulting"

Note: Another connection is that for all of us, voting is a strategic problem, because sometimes we have to think about what votes other people are casting (like, if the Liberals stand no chance in my riding, but I would like to collude with everyone left-of-center in my riding to defeat the conservatives). For a dictator in the social choice sense of someone who dictates the outcome, voting is a decision-theoretic problem, because they get to dictate the outcome without thinking about anyone else, and therefore they aren't really in a game at all!

② Nash equilibria

↪Now, let's try to get at the idea of an equilibrium.

Definition: In game theory, we use **equilibrium** in the common sense way: equilibrium is a situation in which nothing is changing. This matches other technical definitions of equilibrium; in physics, when a ball rolls down a hill and reaches the bottom, we also say that it reaches equilibrium. In game theory we can imagine lots of **out of equilibrium** situations in which players have incentives to change their strategies. Let's illustrate this with an example game.

Example: To see an example of out of equilibrium behaviour, take the following **coordination game**. My cats have the choice between sleeping on my normal bed in a different room, or on a cat bed which is right where the action is, but is also too small to comfortably contain both cats. Say they have the following payoffs (this is a **normal form game**):

	Big bed′	Small bed′
Big bed	(4,4)	(1,3)
Small bed	(3,1)	(2,2)

↪If Abyzou (vertical) chooses to go to the small bed while Pando (horizontal) chooses to go to the big bed, then each of them has a unlitareal incentive to switch. Pando has 1 but could get 2 by going to the small bed, while Abyzou has 3 but could get 4 by going to the big bed. Nothing stops them from switching forever. We need a way around this.

↪The trick is that each player will plan ahead so that no player has a unilateral incentive to deviate. In other words, they will think: what is my best response to any move? And what is the other players' best response to any move? Then they will pick the **mutual best response**.

Definition: The special type of equilibrium called a **Nash Equilibrium** is the situation in which no player has a unilateral incentive to change their strategy. Both players are playing the **best response** to each other, so we call this strategy a mutual best response.

Exercise: What do the cats' utility functions say about their preferences? Write down their preference ranking. **Answer:** (Big,Big) ≻ (Small, Big) ≻ (Small, Small) ≻ (Big, Small).

Exercise: Replace "big bed" with "don't nuke" and "small bed" with "nuke". Write down the story that the game now tells. **Answer:** $(\neg N, \neg N) \succ (N, \neg N) \succ (N, N) \succ (\neg N, N)$.

Exercise: Is (Nuke, Nuke) a Nash Equilibrium of that game? Why or why not? Write down all the Nash Equilibria of this game. **Answer:** It's a coordination game because coordinating is always an NE and not coordinating never is. This game is actually called "stag hunt": we're deciding whether to hunt a stag or a rabbit.

↪What we've just seen is that Nash Equilibria are not necessarily unique. Games can have more than one Nash Equilibrium. Some games have tremendously large numbers of Nash Equilibria.

↪So far we have only been talking about **pure strategy Nash Equilibria**. These really are (nearly) as simple as I've led you to believe in this section. Where things become arbitrarily messy and complicated is when we permit **mixed strategies**. To see an example of that, consider this fun game:

Example: The following game, called **matching pennies**, has no pure strategy Nash Equilibrium. Try to find the mutual best response and you will be stuck in a loop.

	Heads′	Tails′
Heads	(1,-1)	(-1,1)
Tails	(-1,1)	(1,-1)

Similarly, neither does the Rock Paper Scissors game, which includes a cycle:

	Rock′	Paper′	Scissors′
Rock	(0,0)	(-1,1)	(1,-1)
Paper	(1,-1)	(0,0)	(-1,1)
Scissors	(-1,1)	(1,-1)	(0,0)

Exercise: Thinking about how you actually play Rock Paper Scissors, what do you think the mixed strategy mutual best response should be in the rock paper scissors game? And in the matching pennies game? Try to calculate the utilities that you would get from whatever your intuition is. **Answer:** Randomly select with equal probability, you obtain the sum of the probability of picking each option times the payoff obtained from picking that option.

Example: Today in lecture we saw a game with a very interesting Nash Equilibrium: the 1-dimensional spatial voting game. Recalling my example last week that a community with single-peaked utility functions gets out of Arrow's Theorem because the Nash Equilibrium is acceptable to everyone (a Condorcet Winner: an option that would win a pairwise contest against every other option). This is precisely the observation at the heart of Downsian mode-seeking. Is the median the mutual best response, so can we understand Downs's equilibrium finding as a Nash Equilibrium?

4 Section 4: 2019 January 31

This section was canceled by the University of Michigan due to mild cold.

BACK TO TABLE OF CONTENTS

5 Section 5: 2019 February 7

Topics:

- Concavity and single-peakedness [1]
- Optimizing spatial utilities [2]
- Non-equilibrium results [3]
- Ranking limitations [4]

After this class, I expect every student to be able to:

- identify one major upside and one major downside of spatial voting models
- explain why concave utility functions are desirable
- explain why the spatial voting game in more than one dimension does not generally have an equilibrium

BACK TO TABLE OF CONTENTS

① Concavity and single-peakedness

↪One of the questions in the homework was to show that spatial preferences are quasiconcave. This is very similar to the idea of single-peakedness that we discussed in the second section. Let's explore this relationship a little.

Recall: Single-peaked preferences are those preferences which have one unique maximum, and a strictly monotonic decline on all sides of the maximum. Quick inspection shows how similar that is to the idea of a concave function. There are three reasons that we like this way of defining preferences.

↪ It seems to make sense in politics, as long as can think in terms of a Left-Right issue spectrum.

↪ In one dimension, it can give us a solution to social choice problems like Arrow's Theorem. Remember we said that in single-peaked preferences, the median position defeats Arrow's Theorem.

↪ As you saw in POLSCI 598, strictly[1] concave functions can be easily optimized – they have a unique maximum!

② Optimizing spatial utilities

Let's remind ourself of fun times in POLSCI 598 by convincing ourselves that optimizing on spatial utilities works: namely, that the spatial utility function of a voter positioned at (x_v, y_v) has a unique maximum at exactly their ideal point, (x_v, y_v). Consider a candidate at (x_c, y_c). We said that the voter's spatial utility obtained from voting for that candidate is:

$$u(v, c) = -\sqrt{(x_v - x_c)^2 + (y_v - y_c)^2}$$

$$u(v, c) = -\big((x_v - x_c)^2 + (y_v - y_c)^2\big)^{\frac{1}{2}}$$

First let's find the maximum on the x-dimension. Remember that when we're maximizing, we take the derivative with respect to the variable of interest and we set the function equal to 0:

$$\frac{\partial}{\partial x_c} u(v, c) = -\frac{1}{2}\big((x_v - x_c)^2 + (y_v - y_c)^2\big)^{-\frac{1}{2}} \cdot 2(x_v - x_c)(-1)$$

$$0 = \frac{x_v - x_c}{\sqrt{(x_v - x_c)^2 + (y_v - y_c)^2}}$$

Isolating x_c to figure out where the candidate should locate themselves, we get

[1] I thank Jason Davis who suggested the addition of this word, noting that "**strictly** concave functions have a unique maximum, bruh".

$$0 = x_v - x_c$$

$$x_c = x_v$$

Similarly for the y-dimension,

$$\frac{\partial}{\partial y_c} u(v, c) = \frac{1}{2}\left((x_v - x_c)^2 + (y_v - y_c)^2\right)^{-\frac{1}{2}} \cdot 2(y_v - y_c)(-1)$$

$$0 = \frac{y_v - y_c}{\sqrt{(x_v - x_c)^2 + (y_v - y_c)^2}}$$

$$0 = y_v - y_c$$

$$y_c = y_v$$

And we could repeat this for any number of dimensions. So, spatial utility functions have exactly the property we wanted: we can cleanly optimize them, and the unique maximum of a voter's utility is exactly at that voter's location, which we call their ideal point.[2]

(3) Problem: there's no equilibrium in more than 1 dimension

↪We all know the Downsian picture: parties seek the median in 1 dimension, and then they stay there. But the multidimensional picture is much uglier.

Exercise: Consider a legislative committee that is working on some legislation. Say there are 3 voters on the committee, and they are working on some draft legislation which they are trying to place in the issue space, and for example they are arranged arbitrarily (I drew some example configuration on the board). In the context of this problem, try to write down a reasonable definition for an equilibrium. What should qualify a point to be an equilibrium? Try to identify a few, if you can. **Answer:** Let's call a point A an equilibrium point for the piece of legislation if no majority would agree to move the legislation away from A to any other point in the space.

Exercise: In small groups or by yourself, try to draw an example of 5 voters that **does have** an equilibrium. **One answer:** Plott, 1967: Fig. 2

Exercise: In small groups or by yourself, try to draw an example of 6 voters that **does not have** an equilibrium. **One answer:** Plott, 1967: Fig. 4

[2] The result here contains an extremely ugly contradiction, pointed out by David Chung during section. If $x_c = x_v$ and $y_c = y_v$, then the denominator of the fraction is 0. This is actually a deep hole in the simple spatial optimisation idea. One way that we could work around this is by instead constructing the problem so that, as the candidate's position *approaches* the voter's ideal point, the voter increasingly likes the candidate, so that the denominator is infinitesimal but nonzero. Unfortunately, I think this complication would sap the lesson of its teaching value.

Example: Here I showed output from my spatial elections script, which is available on Canvas.

④ Another (more obscure) limitation of spatial utilities

Note: There is a major problem with the construction of preference rankings using dimensions. Constructing utilities on a number line limits the possible preference rankings *a priori*, for absolutely no good reason.

Exercise: How many possible rankings are there of three political parties, A, B, and C? Write them down. **Answer:** There should be 6 such rankings: $A \succ B \succ C, A \succ C \succ B, B \succ A \succ C, B \succ C \succ A, C \succ A \succ B,$ and $C \succ B \succ A$.

Exercise: Now draw a line (so we're just talking in 1 dimension), label it from (say) -2 to 2, and choose arbitrary positions for parties A, B, C. By thinking of all the different places that you could place a voter v, without moving the parties, how many possible preference rankings can you construct for voter v? What if you had positioned them on a plane instead of on a grid? **Answer:** No matter where you put the candidates on a 1-dimensional line, we will rule out 2 of the 6 rankings.

Note: This means that, for example, a Canadian single-issue voter is *a priori* unable to construct 2 of the 6 possible rankings of the 3 major parties.

Exercise: What happens when you move this to 2 dimensions? Can you make all the possible rankings for an arbitrary placement of candidates? What about if you had 4 parties instead of 3?

6 Section 6: 2019 February 14

Topics: In this section we switched tracks a bit, to make sure that everyone is *really, really* comfortable with the fundamentals!

- Iterated Elimination of Strictly Dominated Strategies and finding Nash Equilibria in Normal Form [1]
- Finding Sequentially Rational Strategies in Extensive Form [2]
- Turning Normal Form into Extensive Form [3]
- Practice! [4]

After this class, I expect every student to be able to:

- find the set of rationalizable strategies using Iterated Elimination of Strictly Dominated Strategies
- find Nash Equilibria in Normal Form
- translate a Normal Form game into an Extensive Form Game
- find sequentially rational strategies using backwards induction

BACK TO TABLE OF CONTENTS

① **Iterated elimination and finding Nash Equilibria**:

Example: Let's use Iterated Elimination of Strictly Dominated Strategies to find the rationalizable strategies for the following normal form game:

	L	M	R
A	(4,3)	(2,0)	(0,4)
B	(3,3)	(8,1)	(-1,2)
C	(2,3)	(1,4)	(1,5)

Do any of player 1's pure strategies strictly dominate? No. For player 2, neither L nor M strictly dominates the other, but notice that R does dominate M:

	L	R
A	(4,3)	(0,4)
B	(3,3)	(-1,2)
C	(2,3)	(1,5)

Now A dominates B:

	L	R
A	(4,3)	(0,4)
C	(2,3)	(1,5)

Then R dominates L:

	R
A	(0,4)
C	(1,5)

Finally, C dominates A, so the only remaining strategy is (C, R).

Definition: Any strategy which is left over after the iterated elimintation of strictly dominated strategies is called **rationalizable**.

↪Next, let's find the Pure Strategy Nash Equilibria of this game.

Recall: To find the Pure Strategy Nash Equilibria of a game in Normal Form, fix on each move by player 2, and find the Best Response that player 1 has to each of player 2's moves. Then, fix on each move by player 1, and find the Best Response that player 2 has to each of these moves. Any strategies which are a best response by both players (a mutual best response) form a Nash Equilibrium.

	L	M	R
A	✓(4,3)	(2,0)	(0,4) ✓
B	(3,3) ✓	✓(8,1)	(-1,2)
C	(2,3)	(1,4)	✓(1,5) ✓

So (C, R) is also the unique Pure Strategy Nash Equilibrium as well as the only rationalizable strategy. We actually already knew this, because of the following theorem:

Theorem: All strategies that are part of a Nash Equilibrium are rationalizable (notice: a Nash Equilibrium is actually a profile – a collection – of strategies played by different players), but not all rationalizable strategies are part of a Nash Equilibrium.

(2) Sequential rationality by backwards induction

↪Now let's turn this normal form game into an extensive form game.

Theorem: Any normal form game corresponds to multiple extensive form games.

Example: By arbitrarily declaring that player 1 goes first, we arrive at the following extensive form game:

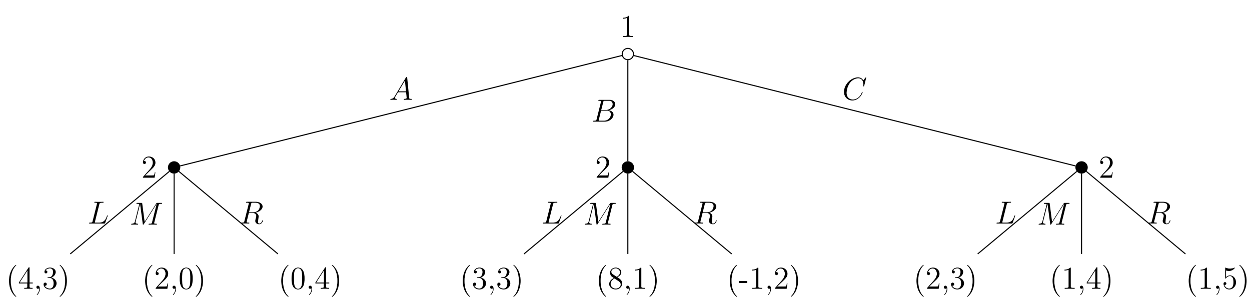

Let's find the set of **sequentially rational strategies**, this time using **backward induction**. We have to start with player 2's decisions, and we may as well go from Left to Right. Clearly, player 2 should pick R in this instance:

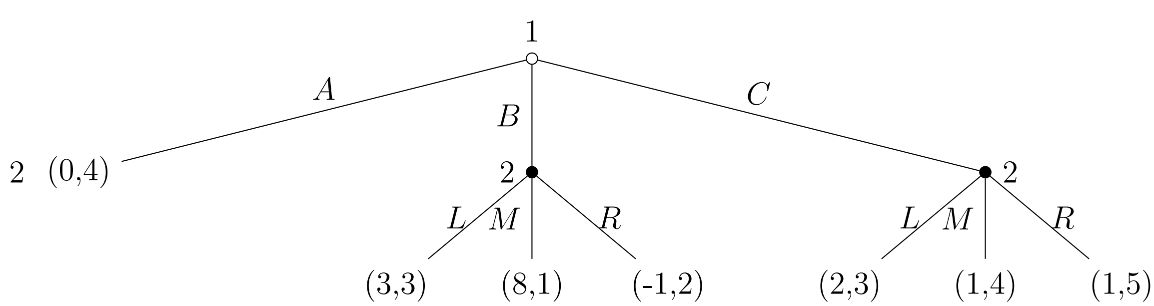

If player 1 plays B, player 2 should play L:

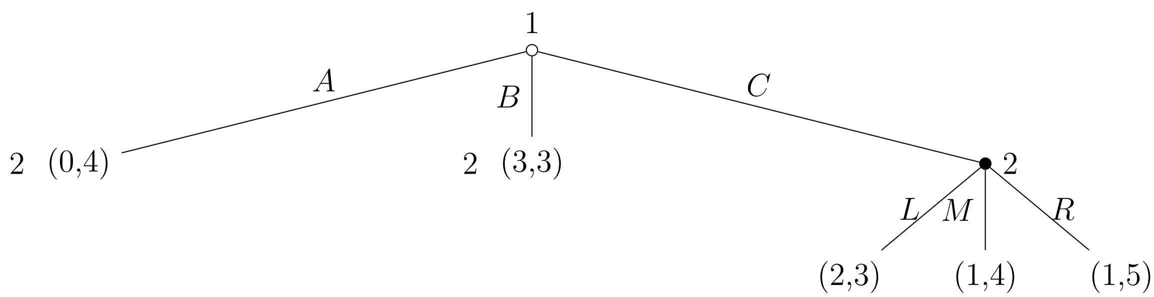

And finally, if player 1 plays C, then player 2 should play R:

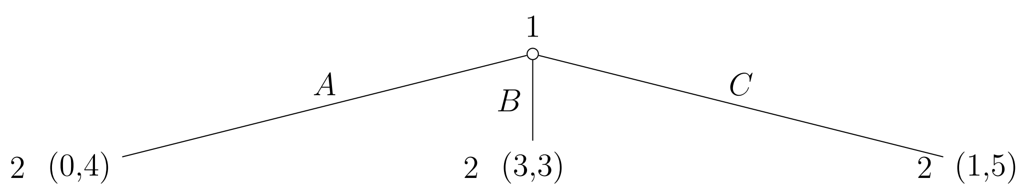

So in the modified game, we found a different sequentially rational strategy: (B, L). This is the power of moving first: player 1 can force a result that they wouldn't have been able to get otherwise.

Exercise: Rewrite the normal form game so that player 2 goes first. What is the result? **Answer:** The game looks like

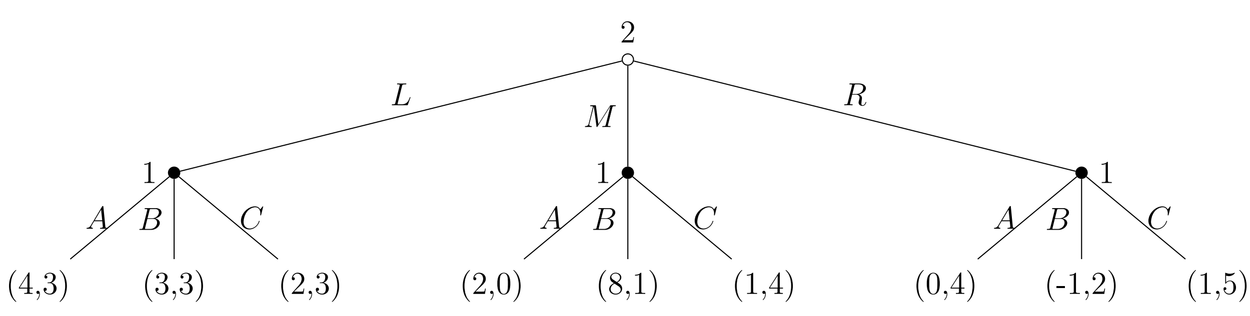

So in this sequence, (R, C) actually is the only sequentially rational strategy.

③ Going from Normal Form to Extensive Form

Theorem: Any game can be equivalently represented in either extensive form or normal form. The only difference is graphical. However, you have to pay careful attention to the information sets. **Question:** How would we write the normal form game we've been discussing in extensive form? To properly translate the normal form game into an extensive form game, we need to write it as:

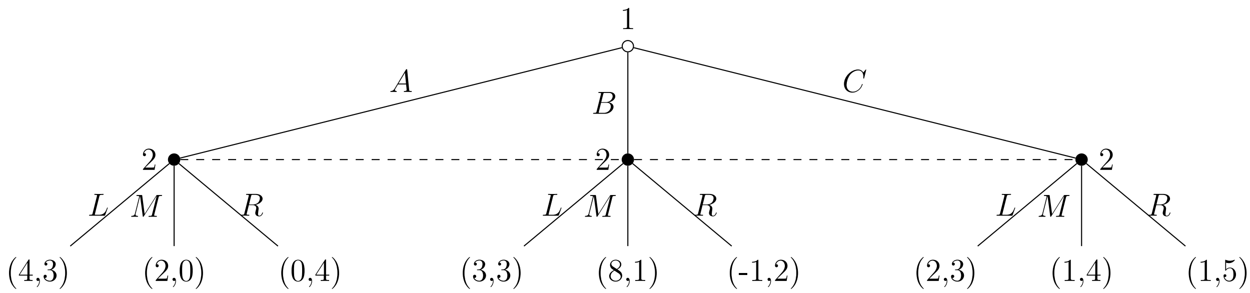

Note: How do we find the Nash Equilibrium of a game in this form? Personally, I would just write it in Normal Form and find the Nash Equilibrium that way! If this was a subgame of a larger extensive form game, you could turn it into Normal Form to find the Subgame Perfect NE.

Remark: Let's step back and consider the implicit process here. We're saying that even if there is a literal time-dimension to the game, then each player will be able to look ahead, check what the other player will do, and act accordingly. So in this way any literal time during the play of the game is irrelevant. In this simple setup, it doesn't matter whether a trade deal takes a day or ten years, rational players will arrive at the same result regardless (although two caveats: 1) game theorists developed the idea of future discounting to handle that, and 2) later in the semester when we study dynamics we will learn modifications of game theory in this direction). This predictive process is one of the central conceits of game theory. The power of rationality is that, if both players know the payoff structures of the other players and know that the other player is payoff-maximizing, then before the game even starts each player can look ahead at what the other player should rationally do, and plan based on that. This planning process eliminates a colossal number of irrelevant options that a rational player would never employ.

Definition: This onerous assumption is called **common knowledge**. Common knowledge is the state in which players are all rational, each player knows that all other players are rational, each player knows that all other players know that they are rational, and so on forever.

(4) Practice with pure strategy rationalizability and NE

Exercise: This question is a modification of question 7.3 in Watson (2013). Find the set of rationalizable strategies of the following game. Also find the NE.

	a	b	c	d
w	(5,4)	(4,4)	(4,5)	(12,2)
x	(3,7)	(8,7)	(5,8)	(10,6)
y	(2,10)	(7,6)	(4,6)	(9,5)
z	(4,4)	(5,9)	(4,10)	(10,9)

Answer: Our approach is to fix on one strategy by one player, and then figure out in what circumstances we would play that strategy. If we do not see any situation in which we would play that strategy, we can cross it out and rewrite the game without it, because it will never be rationally played.

First notice x dominates y:

	a	b	c	d
w	(5,4)	(4,4)	(4,5)	(12,2)
x	(3,7)	(8,7)	(5,8)	(10,6)
z	(4,4)	(5,9)	(4,10)	(10,9)

Next, c dominates d:

	a	b	c
w	(5,4)	(4,4)	(4,5)
x	(3,7)	(8,7)	(5,8)
z	(4,4)	(5,9)	(4,10)

Next, c dominates b:

	a	c
w	(5,4)	(4,5)
x	(3,7)	(5,8)
z	(4,4)	(4,10)

Next, x dominates w:

	a	c
x	(3,7)	(5,8)
z	(4,4)	(4,10)

Next, c dominates a:

	c
x	(5,8)
z	(4,10)

Finally, x dominates z:

	c
x	(5,8)

So the only remaining strategy is (x, c). No other strategy can ever be rationally played.

To find the NE:

	a	b	c	d
w	✓(5,4)	(4,4)	(4,5) ✓	✓(12,2)
x	(3,7)	✓(8,7)	✓(5,8) ✓	(10,6)
y	(2,10) ✓	(7,6)	(4,6)	(9,5)
z	(4,4)	(5,9)	(4,10) ✓	(10,9)

So the unique PSNE is (x, c).

7 Section 7: 2019 February 21

Topics:

- Going from Extensive Form to Normal Form [1]
- Classical 2×2 Normal Form games [2]

After this class, I expect every student to be able to:

- turn an Extensive Form game into a Normal Form game
- recall the payoff tables and properties of a few of the most important 2 by 2 games
- explain why each of the canonical 2-player games has a general representation, and not just one specific set of payoff values

BACK TO TABLE OF CONTENTS

① Going from Extensive Form to Normal Form

Remark: When we translate an Extensive Form game in to the corresponding Normal Form game, each column of the Normal Form game needs to be every *full strategy profile* for each player. Here are some examples.

Example: This example is question 3.3 in Watson (2013). Translate the following Extensive Form game into Normal Form:

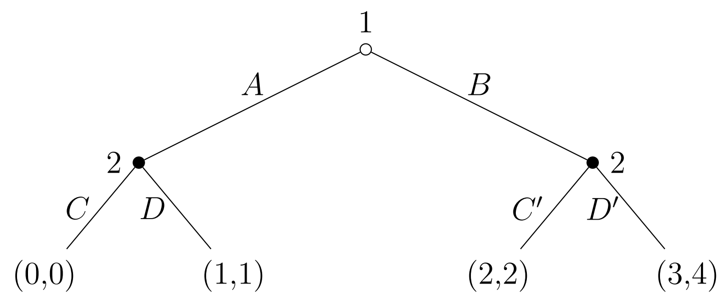

Answer:

	CC'	CD'	DC'	DD'
A	(0,0)	(0,0)	(1,1)	(1,1)
B	(2,2)	(3,4)	(2,2)	(3,4)

Why do we write it in this counter-intuitive way? I'll offer four reasons:

↪ It's closer to the Extensive Form representation.

↪ The purpose of writing out a game is precisely to be able to compare all possible full strategy profiles.

↪ All games should be complete guides to the game, which anybody can follow us no matter what. This is the same rationale that makes assuming rationality useful, and it is also why we care about modeling mistakes as in Trembling Hand Perfect Equilibria.

↪ Games can be repeated, so even if we think we can be sure that we will not see a particular move in the first iteration of the game, that might not be true in repeated games. This means that a thoughtful player needs to commit to a plan for each possibly contingency.

② Classical 2-player games

↪We're going to need some games to practice on. We might as well use the games that you will hear referred to over and over again throughout your career, because they tell useful allegories about social scientific situations: the canonical 2 player simultaneous games. These are part of the core vocabulary of game theory.

Exercise: Find the set of rationalizable strategies and the Pure Strategy Nash Equilibria of all of the following classical 2×2 games.

↪Prisoner's Dilemma:

Example:

	C′	D′
C	(2,2)	(0,3)
D	(3,0)	(1,1)

General:

	C′	D′
C	(a,a)	(b,c)
D	(c,b)	(d,d)

with $c > a > d > b$

Answer: D dominates C and D' dominates C'. So rationalizable (D, D'). PSNE:

	C′	D′
C	(2,2)	(0,3) ✓
D	✓(3,0)	✓(1,1) ✓

↪Battle of the Sexes (notoriously sexist):

Example:

	Opera′	Movie′
Opera	(2,1)	(0,0)
Movie	(0,0)	(1,2)

General:

	C′	D′
C	(a,b)	(c,c)
D	(c,c)	(b,a)

with $a > b > c$

Answer: Every strategy is rationalizable. PSNE:

	Opera′	Movie′
Opera	✓(2,1) ✓	(0,0)
Movie	(0,0)	✓(1,2) ✓

↪Game of Chicken/Hawk-Dove game:

Example:

	Hawk′	Dove′
Hawk	(0,0)	(3,1)
Dove	(3,1)	(2,2)

General:

	C′	D′
C	(a,a)	(b,c)
D	(c,b)	(d,d)

with $b > d > c > a$

Answer: Every strategy is rationalizable. PSNE:

	Hawk′	Dove′
Hawk	(0,0)	✓(3,1) ✓
Dove	✓(3,1) ✓	(2,2)

↪Coordination game:

Example:

	A′	B′
A	(1,1)	(0,0)
B	(0,0)	(1,1)

General:

	C′	D′
C	(a,a)	(b,b)
D	(b,b)	(a,a)

with $a > b$

Answer: Every strategy is rationalizable. PSNE:

	A′	B′
A	✓(1,1) ✓	(0,0)
B	(0,0)	✓(1,1) ✓

↪Pareto coordination game:

Example:

	A′	B′
A	(2,2)	(0,0)
B	(0,0)	(1,1)

General:

	C′	D′
C	(a,a)	(b,b)
D	(b,b)	(c,c)

with $a > c > b$

Answer: Every strategy is rationalizable. PSNE:

	A′	B′
A	✓(2,2) ✓	(0,0)
B	(0,0)	✓(1,1) ✓

8 Section 8: 2019 February 28

Topics:

- Mixed strategy rationalizability [1]
- Mixed strategy Nash Equilibria [2]

After this class, I expect every student to be able to:

- find the set of rationalizable strategies in a game with small strategy sets, accounting for domination by mixed strategies
- understand why we set the opponents' payoffs to be equal in order to find mutually mixed strategy Nash Equilibria
- use this technique to find mixed strategy Nash Equilibria in simple games

BACK TO TABLE OF CONTENTS

① Mixed strategy rationalizability

Recall: Let's notice also that it is possible for one strategy to be dominated by a mixture of other strategies. Consider the following game (Figure 7.3 in Watson):

	X	Y	Z
U	(5,1)	(0,4)	(1,0)
M	(3,1)	(0,0)	(3,5)
D	(3,3)	(4,4)	(2,5)

Consider player 2's mixed strategy $s = (0, \frac{1}{2}, \frac{1}{2})$. **Notation:** This means that X is played with probability 0, Y is played with probability $\frac{1}{2}$, and Z is played with probability $\frac{1}{2}$. The expected utility, which is the sum of the probability of obtaining each payoff times the value of that payoff, is given by:

If player 1 plays U,

$$u_2(s, U) = \frac{1}{2} \cdot 4 + \frac{1}{2} \cdot 0$$

$$u_2(s, U) = 2$$

Since this is greater than 1, player 2 prefers the mix to X in this case. Next, if player 1 plays M,

$$u_2(s, M) = \frac{1}{2} \cdot 0 + \frac{1}{2} \cdot 5$$

$$u_2(s, M) = \frac{5}{2}$$

Since this is greater than 1, player 2 prefers the mix to X in this case also. Finally, if player 1 plays D,

$$u_2(s, D) = \frac{1}{2} \cdot 4 + \frac{1}{2} \cdot 5$$

$$u_2(s, D) = \frac{9}{2}$$

Since this is greater than 3, player 2 prefers the mix to X in this final case. So, X is domated by the mixed strategy which assigns probability $\frac{1}{2}$ to strategy Y and probability $\frac{1}{2}$ to strategy Z. So, the game becomes:

	Y	Z
U	(0,4)	(1,0)
M	(0,0)	(3,5)
D	(4,4)	(2,5)

Next, D dominates U:

	Y	Z
M	(0,0)	(3,5)
D	(4,4)	(2,5)

Next, Z dominates Y:

	Z
M	(3,5)
D	(2,5)

And finally M dominates D, so the only rationalizable strategy is (M, Z). If we had restricted ourselves to pure strategies, we would have been stuck with a 3×3 game. By considering a pure strategy, we realised that only one outcome is rationalizable for all players.

Exercise: How much can we reduce the following game?[3]:

	A′	B′	C′
A	(5,5)	(4,4)	(1,1)
B	(4,4)	(0,0)	(1,1)
C	(0,0)	(4,4)	(1,1)

Answer: The mixed strategy $(\frac{1}{2}, \frac{1}{2}, 0)$ strictly dominates C'. Then, A dominates B.

② Mixed Strategy Nash Equilibria

Note: We can start with one extremely powerful tool for finding Mixed Strategy Nash Equilibria. In order for the best response to be a mixed strategy, it is necessary that the opponent's particular choice of moves will make us completely indifferent between pure strategies. Why? Because if I wasn't completely indifferent between pure strategies, then one of my pure strategies would be dominated! If that were true then I need to remove that

[3]The first half of this example is a worked example taken from Jason Davis's 681 section notes circa 2017.

Pure Strategy from the set of options, because it is not rationalizable, which changes the game that we're considering. So when we're finding a Mixed Strategy Nash Equilibrium, we have the amazing luxury of **setting the payoffs of our pure strategies to all be equal!** This will make our lives way easier.

↪Let's see how we can apply this idea with an example. Let's remind ourselves of how to find Mixed Strategy Nash Equilibria in games with finite strategy sets, using matching pennies as a convenient example:

	H′	T′
H	(1,-1)	(-1,1)
T	(-1,1)	(1,-1)

To complete the matching pennies example, say that player 1 plays H with probability $p_H = p$, and therefore T with probability $p_T = 1 - p$, and similarly say $p_{H'} = q$. Then we say that player 1 has to recieve the same payoff from choosing H and T, and similarly for player 2. So player one can expect the following payoffs:

$$u_1(H) = q + (1 - q) \cdot (-1)$$

$$u_1(H) = q - 1 + q$$

$$u_1(H) = 2q - 1$$

and

$$u_1(T) = -1 \cdot q + (1 - q)$$

$$u_1(T) = -2q + 1$$

Setting these two equal,

$$2q - 1 = -2q + 1$$

$$4q = 2$$

$$q^* = \frac{1}{2}$$

And we know that the same exact process must be true for player 2:

$$u_2(H') = -1 \cdot p + (1 - p)$$

$$u_2(H') = -2p + 1$$

and

$$u_2(T') = p + (1 - p) \cdot (-1)$$

$$u_2(T') = q - 1 + p$$

$$u_2(T') = 2p - 1$$

Setting these two equal,

$$2p - 1 = -2p + 1$$

$$4p = 2$$

$$p^* = \frac{1}{2}$$

So the unique Mixed Strategy Nash Equilibrium of this game is to pick each with equal probability, exactly as we have said when discussing the game in less formal terms: $(\frac{1}{2}, \frac{1}{2})$.

9 Section 9: 2019 March 14

Topics:

- Pure Strategy Nash Equilibria on very large strategy sets [1]
  - ↪ Dollar-splitting game [2]
  - ↪ Classical Cournot Duopoly [3]
- Finding Perfect Bayesian Nash Equilibria in small games [4]

After this class, I expect every student to be able to:

- use the optimization tools from POLSCI 598 to solve for the Pure Strategy Nash Equilibria of 2-player symmetrical games with continuous strategy sets
- have a rough procedure in mind for finding the PBNE of a 2-player signalling game

BACK TO TABLE OF CONTENTS

① Pure Strategy Nash Equilibria on very large strategy sets

At the end of section 8, we discussed the following exercise, to introduce the idea of finding a Pure Strategy Nash Equilibrium on a very large (but still finite) strategy set:

Exercise: Two players are bargaining over how to split a dollar. Simultaneously they pick a share of the dollar they would like to receive, call them shares s_1 for player 1 and s_2 for player 2. If they announce shares such that $s_1 + s_2 \leq 1$, then each player receives the share they named. If players announce shares such that $s1 + s2 > 1$, then both players receive 0. Find and justify the pure-strategy Nash equilibria to this game.[4] **Answer:** We need to ask ourselves: in what situation is there no incentive for any player to deviate? If $s_1 + s_2 > 1$, then any player who picked more than 0 has a unilateral incentive to decrease their quantity. If $s_1 + s_2 < 1$, then both players have a unilateral incentive to increase their choice. There is no incentive for anyone to deviate if $s_1 + s_2 = 1$, so any s_1 and s_2 which satisfy that condition are NE. And finally, there is an extremely strange NE at $s_1 = 1$ and $s_2 = 1$, because it's not possible to deviate upwards, and no player gains by deviating down to 0. While illustrating the logic of how to locate Nash Equilibria, the dollar-splitting example also illustrates some of the very bizarre pathologies of the Nash Equilibrium solution concept.

↪Now let's try finding the Mixed Strategy Nash Equilibria of a game with continuous strategy set. I present this game, even though I detest economics jargon, because it has seeped into political science as standard parlance and is a very typical example of econ jargon that you might encounter in our field:

Example: This is a classical Cournot Duopoly model.[5]

↪Say that q_1 and q_2 are the quantities of some product produced by firms 1 and 2.

↪Set $P(Q) = a - Q$ is the "market-clearing price" (equilibrium price, the price at which quantity supplied exactly equals quantity demanded), for the aggregate quantity on the market $Q \equiv q_1 + q_2$

↪Say also that the cost to firm $i, i \in \{1, 2\}$ of producing quantity q_i is $C_i(q_i) = cq_i$ where c is just some constant.

↪Finally, say that the players move simultaneously.

Let's start by translating this game into a normal form game. We have:

↪Players: firms 1 and 2

[4] Example heavily modified from Tyson (2017).

[5] The meat of this example is entirely taken from Gibbons (1992: p. 15-17), but I have dramatically changed the formatting and wording, and added some commentary.

$\hookrightarrow$Strategies: either firm can choose from the continuous strategy set $S_i = [0; \infty]$.

Notation: We always use S for the set of all strategies and s for an element of that set.

$\hookrightarrow$Payoffs: allow that each firm's payoff is simply its profit. Then the payoff $u_i(s_i, s_j)$ is the price of producing the goods minus the cost of producing them, which is given as follows. **Notation:** π is very frequently used for expected utility.

$$\pi_i(q_i, q_j) = q_i P(q_i + q_j) - cq_i$$

$$\pi_i(q_i, q_j) = q_i(P(q_i + q_j) - c)$$

$$\pi_i(q_i, q_j) = q_i(a - (q_i + q_j) - c)$$

Now, to find the equilibrium, recall that a Nash Equilibrium has to solve $\max\limits_{s_i \in S_i} u_i(s_i, s^*_{-i})$ for all players i.

Notation: the subscript $_{-i}$ is frequently used to mean "every player except player i".

Let's check the first order condition[6] for either player i:

$$0 = \frac{\partial}{\partial q_i} q_i(a - q_i - q_j^* - c)$$

$$0 = \frac{\partial}{\partial q_i} (q_i a - q_i^2 - q_i q_j^* - q_i c)$$

$$0 = a - 2q_i - q_j^* - c$$

$$2q_i = a - q_j^* - c$$

$$q_i^* = \frac{1}{2}(a - q_j^* - c)$$

So, we obtain the following optimal quantity choices for the two firms:

$$q_1^* = \frac{1}{2}(a - q_2^* - c)$$

and

$$q_2^* = \frac{1}{2}(a - q_1^* - c)$$

Now we can just solve the pair of equations by plugging one into the other:

[6]We don't know for sure that it's sufficient because we don't know the shape of the function, but it turns that it that it is sufficient exactly if $q_j^* < a - c$, which is indeed true in this example.

$$q_1^* = \frac{1}{2}(a - \frac{1}{2}(a - q_1^* - c) - c)$$

$$2q_1^* = a - \frac{1}{2}a + \frac{1}{2}q_1^* - \frac{1}{2}c - c$$

$$2q_1^* - \frac{1}{2}q_1^* = \frac{1}{2}a - \frac{1}{2}c$$

$$\frac{3}{2}q_1^* = \frac{1}{2}a - \frac{1}{2}c$$

$$q_1^* = \frac{a-c}{3}$$

And, by the symmetry of the equations, we find similarly

$$q_2^* = \frac{a-c}{3}$$

So this is the unique Pure Strategy Nash Equilibrium of the classical Cournot Duopoly game.

(2) Finding Perfect Bayesian Nash Equilibria in small games

Theorem: We can solve for the Perfect Bayesian Nash Equilibria of any 2-player game using the following (somehwat approximate) procedure:[7]

1. Pick a full strategy profile for player 1.

2. Use Bayes' Rule to calculate updated beliefs if possible for player 2. If Bayes' Rule cannot be used, repeat steps 3 and 4 for all possible arbitrary beliefs.

3. Given optimal beliefs, calculate player 2's optimal action.

4. Check that the strategy we picked for player 1 in step 1 is indeed a best response to player 2's optimal strategy.

5. Repeat steps 1-4 until all of player 1's strategy profiles have been exhausted.

[7] I heavily adapted this procedure from Watson (2013: p. 383)

10 Section 10: 2019 March 21

Topics:

- Perfect Bayesian Nash Equilibria with sequential and uncertain moves [1]
  - ↪ Pooling equilibria [2]
  - ↪ Separating equilibria [3]
- The Gift Game example [4]

After this class, I expect every student to be able to:

- follow our PBNE algorithm to solve for PBNEs in games between few players, where the moves are sequential and players are not certain about previous moves.

BACK TO TABLE OF CONTENTS

① Perfect Bayesian Nash Equilibria with sequential and uncertain moves

Question: How should a payoff-maximizing player approach a game where the players have private information about what moves they play *and* they play sequential moves?

↪Faced with this situation, we claim that a rational player should construct beliefs about what another rational player would do.

Note: Whenever you encounter the world "beliefs" in quantitative political science, the first thing that should come to mind is Bayes Theorem! If we want to put together an equilibrium concept that includes both strategies and beliefs, then we will want it to be motivated by the properties of Nash Equilibria and by the properties of Bayes Theorem.

Recall: For events A and B, Bayes Theorem states that

$$P(A|B) = \frac{P(A)}{P(B)} P(B|A)$$

Which it's often convenient to rephrase as

$$P(A|B) = \frac{P(A)P(B|A)}{P(A)P(B|A) + P(\neg A)P(B|\neg A)}$$

Definition: Consider a strategy profile for all players, and beliefs over the nodes in every information set. These profiles and beliefs form a **Perfect Bayesian Nash Equilibrium (PBNE)** if each player's strategy is optimal given their beliefs about the strategies of other players, and the beliefs are consistent with Bayes' Rule.

Definition: Often our prior beliefs about players come from a move by "Nature" with a known probability, which selects the payoffs that a player obtains from the strategies available to them. In this situation we say that Nature is selecting a "type" of player. For example, a strong country might get a higher payoff from declaring war than a weak country would, and we know roughly how likely it is that a given country is strong.

Definition: A PBNE is **separating** if the possible types of players would choose different strategies.

Definition: A PBNE is **pooling** if the possible types of players would choose the same strategies.

② The Gift Game example

Example: Consider the following game, which we will call the "gift game".[8]

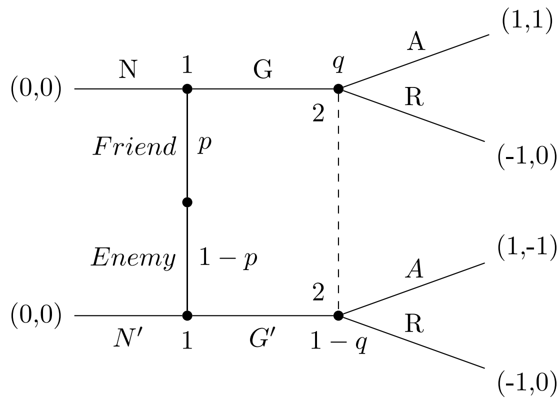

What are the beliefs in the above game? The beliefs are the probabilities that player 2 assigns to being at each of the two indistinguishable nodes, namely $(q, 1-q)$.

We can first compute the expected value of the different moves in terms of q:

$$\pi_2(A) = q + (-1)(1-q)$$

$$\pi_2(A) = 2q - 1$$

While

$$\pi_2(R) = 0$$

So player 2 will select A iff

$$\pi_2(A) > \pi_2(B)$$

$$2q - 1 > 0$$

$$2q > 1$$

$$q > \frac{1}{2}$$

[8]I heavily adapted this example from Watson (2013) figure 28.2 and the surrounding discussion.

So this illustrates how beliefs inform decision-making. But is this pair of strategies and beliefs part of an equilibrium? And what should player 1 do?

To answer these questions, we can use Bayes' Rule. Let α_E represent the probability that an Enemy type would play G, and let α_F represent the probability that a Friend type would play G. Then we have:

Top node with probability: $p \cdot \alpha_F$

Bottom node with probability: $(1-p) \cdot \alpha_E$

Now, notice that with $q \equiv P(F|G)$, by Bayes' Rule we have

$$P(A|B) = \frac{P(A)P(B|A)}{P(A)P(B|A) + P(\neg A)P(B|\neg A)}$$

$$P(F|G) = \frac{P(F)P(G|F)}{P(F)P(G|F) + P(E)P(G|E)}$$

$$q = \frac{P(F)P(G|F)}{P(F)P(G|F) + P(E)P(G|E)}$$

Substituting in the values we found for each of those events, we then have

$$q = \frac{p\alpha_F}{p\alpha_F + (1-p)\alpha_E}$$

The game is in Nash Equilibrium for any combination of strategies by player 1 and probabilities by Nature which satisfy Bayes' Theorem.

Let's check for those strategies one by one:

Question: Could NG' be part of a (separating) equilibrium? Then $q = 0$, so player 2 should play R. But if they are expecting R, then player 1 would strictly prefer not to play G' if they are an Enemy. So NG' cannot be part of a PBNE.

Exercise: Check if there is an equilibrium with GN'. Is this equilibrium separating or pooling? **Answer:** Pooling. Here $q = 1$, so player 2 should play A. But then player 1 would strictly prefer to play G' over N' if they are an Enemy. So NG' cannot be part of a PBNE.

Exercise: Could GG' be part of an equilibrium, and of what type? **Answer:** Pooling. Then $q = p$ exactly, so recalling our expected utility calculations from earlier, player 2 should play A iff $q \geq \frac{1}{2}$. So if the exogenous $p \geq \frac{1}{2}$, then there is a PBNE with the belief $q = p$ together with the strategy profile (GG', A).

Question: What about NN'? Then let's try to find q using Bayes' Rule:

$$P(Friend|Give) = \frac{P(Friend)}{P(Give)} P(Give|Friend)$$

$$P(Friend|Give) = \frac{p}{0} P(Give|Friend)$$

Uh oh! So Bayes' Rule does absolutely nothing for us here; in fact, it completely breaks. This is the problem that Scott mentioned in class: in out-of-equilibrium paths, Bayes' Rule gives us no directions. However, standard payoff-maximization can give us an answer: R will continue to dominate so long as $q \leq \frac{1}{2}$. So any $q \leq \frac{1}{2}$ supports a PBNE where (NN', R) is played together with the belief q.

11 Section 11: 2019 March 28

Topics:

- Mixed Strategy Nash Equilibria and Bayesian Nash [1]
- PBNE with simultaneous moves and uncertain payoffs [2]

After this class, I expect every student to be able to:

- have a rough procedure in mind to solve for PBNEs in games between few players, where the moves are simulateneous and players have sufficiently little uncertainty about the payoffs.
- understand that Mixed Strategy Nash Equilibria and Perfect Bayesian Nash Equilibria nearly always converge to the same solution as uncertainty in the payoffs approaches zero.

BACK TO TABLE OF CONTENTS

① Mixed Strategy Nash Equilibria and Bayesian Nash

↪How should we interpret the fact that we've discussed Mixed Strategy Nash Equilibria as the correct way to handle uncertainty in some situations, and Perfect Bayesian Nash Equilibria as the correct solution concept for uncertainty in other situations?

Theorem: Mixed Strategy Nash Equilibria can (almost always) be understood as the Perfect Bayesian Nash Equilibria of a closely related game with a sufficiently small amount of missing information.

① PBNE with simultaneous moves and uncertain payoffs

Example: Let's illustrate this idea with the Battle of the Sexes. We'll study this particular game so that we can see a Bayesian Nash in a simultaneous game with uncertainty about the payoffs, since so far we've only really seen Bayesian Nash Equilibria in sequential games with uncertainty about the moves.

Exercise: Find the set of rationalizable strategies, the Pure Strategy Nash Equilibria, and any Mixed Strategy Nash Equilibria of the following game:

	Opera′	Movie′
Opera	(2,1)	(0,0)
Movie	(0,0)	(1,2)

Answer: Every strategy is rationalizable. PSNE:

	Opera′	Movie′
Opera	✓(2,1) ✓	(0,0)
Movie	(0,0)	✓(1,2) ✓

MSNE: Say player 1 plays (O, M) with probability $(p, 1-p)$ and player 2 plays (O', M') with probability $(q, 1-q)$. Then at MSNE, player 1's utility is given by:

$$u_1(O) = 2q^*$$

And

$$u_1(M) = 1 - q^*$$

Setting these utilities equal so that player 1 will choose to mix,

$$2q^* = 1 - q^*$$

$$3q^* = 1$$

$$q^* = \frac{1}{3}$$

And similarly,

$$u_2(O') = p^*$$

And

$$u_2(M') = 2 - 2p^*$$

Setting these utilities equal so that player 1 will choose to mix,

$$p^* = 2 - 2p^*$$

$$3p^* = 2$$

$$p^* = \frac{2}{3}$$

So player 1 plays (O, M) with probabilities $(\frac{2}{3}, \frac{1}{3})$, and player 2 plays (O', M') with probabilities $(\frac{1}{3}, \frac{2}{3})$.

Now, let's modify the game so that it has some notion of uncertainty built into it (heavily expanded from Gibbons 153). Say that we randomly draw modifications to each players' preferred payoffs, which are only known to their owners when the strategies are selected. Then the game becomes:

	Opera′	Movie′
Opera	(2 + t_1,1)	(0,0)
Movie	(0,0)	(1,2 + t_2)

Let's say that t_1, t_2 are drawn from a uniform distribution from $[0; x]$, for sufficiently small x. Then we will follow a hunch to seek equilibria where player 1 plays O if t_1 exceeds some critical value c_1, so with probability $\frac{x-c_1}{x}$, and player 2 plays M if t_2 exceeds some critical value c_2, so with probability $\frac{x-c_2}{x}$.

Suppose that the players do indeed select these strategies. Holding x constant, what values of c_1 and c_2 make this strategy profile a Bayesian Nash Equilibrium?

First let's calculate player 1's expected utility:

$$u_1(O) = \big(1 - \frac{x - c_2}{x}\big) \cdot (2 + t_1) + \frac{x - c_2}{x} \cdot 0$$

$$u_1(O) = \big(\frac{x}{x} - \frac{x - c_2}{x}\big) \cdot (2 + t_1)$$

$$u_1(O) = \frac{c_2}{x} \cdot (2 + t_1)$$

And

$$u_1(M) = \left(1 - \frac{x - c_2}{x}\right) \cdot 0 + \frac{x - c_2}{x} \cdot 1$$

$$u_1(M) = 1 - \frac{c_2}{x}$$

So playing O is a Best Response for player 1 iff

$$u_1(O) \geq u_1(M)$$

$$\frac{c_2}{x} \cdot (2 + t_1) \geq 1 - \frac{c_2}{x}$$

$$2 + t_1 \geq \frac{x}{c_2} - 1$$

$$t_1 \geq \frac{x}{c_2} - 3$$

This was the threshold for player 1 that we set out to solve for: in other words, $c_1 = \frac{x}{c_2} - 3$. Then solving similarly for player 2,

$$u_2(O') = \left(1 - \frac{c_1}{x}\right) \cdot 1 + \frac{c_1}{x} \cdot 0$$

$$u_2(O') = 1 - \frac{c_1}{x}$$

And

$$u_2(M') = \left(1 - \frac{c_1}{x}\right) \cdot 0 + \frac{c_1}{x}(2 + t_2)$$

$$u_2(M') = \frac{c_1}{x}(2 + t_2)$$

So playing O' is a Best Response for player 2 iff

$$\frac{c_1}{x}(2 + t_2) \geq 1 - \frac{c_1}{x}$$

$$2 + t_2 \geq \frac{x}{c_1} - 1$$

$$t_2 \geq \frac{x}{c_1} - 3$$

So, $c_2 = \frac{x}{c_1} - 3$.

Now let's solve for the threshold values. We have

$$\frac{x}{c_1} - 3 = \frac{x}{c_2} - 3$$

$$c_1 = c_2$$

And it is also true that, using the above equality,

$$\frac{x}{c_1} - 3 = c_2$$

$$x = c_2 c_1 + 3c_1$$

$$0 = c_1^2 + 3c_1 - x$$

Solving this quadratic (I omit the steps) shows that $\frac{x-c_1}{x}$, which is the probability that player 1 plays O, and $\frac{x-c_2}{x}$, which is the probability that player 2 plays M', both equal the expression

$$1 - \frac{-3 + \sqrt{9+4x}}{2x}$$

Let's look at the properties of this big ugly probability p as $x \to 0$, so as the noise in the payoffs approaches 0 also:

$$p = \lim_{x \to 0} \left(1 - \frac{-3 + \sqrt{9+4x}}{2x} \right)$$

$$p = 1 - \frac{1}{2} \lim_{x \to 0} \left(\frac{-3 + \sqrt{9+4x}}{x} \right)$$

Using

$$\frac{\sqrt{4x+9} - 3}{x} = \frac{(\sqrt{4x+9} - 3)(\sqrt{4x+9} + 3)}{x(\sqrt{4x+9} + 3)}$$

$$\frac{\sqrt{4x+9} - 3}{x} = \frac{4}{\sqrt{4x+9} + 3}$$

Then we have

$$p = 1 - \frac{1}{2} \lim_{x \to 0} \left(\frac{4}{3 + \sqrt{9+4x}} \right)$$

$$p = 1 - \frac{1}{2} \cdot \frac{4}{6}$$

$$p = 1 - \frac{1}{3}$$

$$p = \frac{2}{3}$$

Is there a lesson behind all this pain and suffering? Yes, there is, and it's a very pleasant one:

Remark: In a simple simultaneous game with uncertainty about the payoff values, as that uncertainty decreases, the Bayesian Nash Equilibrium almost always collapses into exactly the Mixed Strategy Nash Equilibrium. This is a wonderfully deep connection between our two solution concepts that involve uncertainty.

12 Section 12: 2019 April 3

Topics:

- The Setting of Evolutionary Game Theory [1]
- Payoff Matrices and Mixed Strategies as Frequencies [2]
- Evolutionarily Stable Strategies [3]

After this class, I expect every student to be able to:

- explain the setting of Evolutionary Game Theory
- write down the payoff matrix for a symmetric game
- have a rough procedure in mind to find the Evolutionarily Stable Strategies of a small game

BACK TO TABLE OF CONTENTS

① The Setting of Evolutionary Game Theory

Remark: We want to work up to the idea of an Evolutionarily Stable Strategy (ESS). Part of the reason that we care about this is that, for your homework, you can use the following incredibly useful theorem:

Theorem: The Evolutionarily Stable States of a system are all also rest points of the replicator dynamics. If we solve for the ESS, we have found rest points of the replicator dynamics for the game under consideration.

Remark: So, let us quickly discuss the setup in which the idea of an ESS makes sense. This setup is called Evolutionary Game Theory, which is the field of study that uses game theory to study hypothetical populations of animals. We have to modify a few of the ideas we've developed in this class before we can make this substantive change.

Recall: We said that a Mixed Strategy is a probability distribution over Pure Strategies. In EGT, we'll imagine an arbitrarily large population where each individual chooses a (pure or mixed) strategy to play, and pairs of individuals are constantly being chosen to play against each other. Without assuming rationality, we can now make players still be payoff-seeking by declaring that payoffs represent the probability of reproduction, called "reproductive fitness". Rather than assuming rationality, all we need is the Darwinian conclusion that some latent process (evolution) maximizes reproductive fitness.

Note: The idea of a population all playing the same game with the same payoff structure is only possible for a very special type of game: symmetric games. Let's remind ourselves what that means.

Recall: A **symmetric** 2-player game is a game which remains identical when player 1 and player 2 are swapped. Prisoner's Dilemma is a symmetric game because the Prisoners could be switched and nothing changes, but Matching Pennies is not symmetric because there has to be a Matching player and an Anti-matching player.

② Payoff Matrices and Mixed Strategies as Frequencies

Note: Because EGT *only* applies to symmetric games, they get to use a really wonderful notational trick. Encode one list of mixed strategy probabilities in a vector x and another in a vector y. Then we can rewrite the Normal Form game as a payoff matrix U which fully encodes the game, and xUy will be the expected payoff of playing strategy x against strategy y.

Example: Let's prepare the Hawk-Dove game for the EGT setup. The traditional Hawk-Dove game looks like:

	Hawk	Dove
Hawk	$\left(\frac{G-C}{2}, \frac{G-C}{2}\right)$	$\left(G, 0\right)$
Dove	$\left(0, G\right)$	$\left(\frac{G}{2}, \frac{G}{2}\right)$

(1)

Consider the strategy vector $x = \begin{bmatrix} x_H \\ x_D \end{bmatrix}$, $x_H + x_D = 1$, where x_H, x_D represent the chance of playing Hawk or Dove.

We can write the Hawk-Dove payoff matrix as

$$U = \begin{bmatrix} \frac{G-C}{2} & G \\ 0 & \frac{G}{2} \end{bmatrix}$$

Then we can find the expected payoff of playing x against x as

$$x^T U x = \begin{bmatrix} x_H \, x_D \end{bmatrix} \begin{bmatrix} \frac{G-C}{2} & G \\ 0 & \frac{G}{2} \end{bmatrix} \begin{bmatrix} x_H \\ x_D \end{bmatrix}$$

$$x^T U x = x_H x_H \tfrac{G-C}{2} + x_H x_D G + x_D x_H 0 + x_D x_D \tfrac{G}{2}$$

Now notice that each term has a clear interpretation: if you are a randomly selected player in a population of $x-$players, the first term is the probability that you play a Hawk times the probability that your opponent plays a Hawk, times the payoff that a Hawk obtains when it plays against a Hawk. The second is the probability that you play a Hawk and your opponent plays a Dove, times the payoff obtained by a Hawk for facing a Dove. Etc.

Note: In other words, why did we get to say "x against x"? Suppose for a second that x is the optimal strategy, so we expect the whole population to tend to pick strategy x. Then at the precise moment when we choose two players to compete, there is precisely probability x_H of the opponent picking Hawk and precisely probability x_D of the opponent picking Dove.

Definition: In this setup, we can write a **Nash Equilibrium** as follows: A strategy $\hat{p}$ is a Nash Equilibrium strategy if $p^T U \hat{p} \leq \hat{p}^T U \hat{p} \quad \forall p$

Remark: In words: for any vector of frequencies p, the expected payoff for playing p in response to $\hat{p}$ is lower than the expected payoff of paying $\hat{p}$ against $\hat{p}$. This is exactly what we have meant throughout this class by a "mutual best response".

③ **Evolutionarily Stable Strategies**

Definition: A strategy $\hat{p}$ is an **Evolutionarily Stable Strategy** of the game with payoff matrix U if:

a) $\hat{p}$ is a Nash Equilibrium, so $p^T U \hat{p} \leq \hat{p}^T U \hat{p} \quad \forall p$, and

b) $\hat{p}$ is non-invadable, so $p \neq \hat{p}$ and $p^T U \hat{p} = \hat{p}^T U \hat{p} \quad \implies \quad p^T U p < \hat{p}^T U p$

Remark: In words: a NE strategy $\hat{p}$ is non-invadable if it plays better against its alternative best responses than they play against themselves (so, if any other strategy p which is also a BR against $\hat{p}$ obtains a lower payoff by playing against itself than $\hat{p}$ obtains by playing against p).

Theorem: Strict Nash Equilibrium $\implies$ ESS $\implies$ Nash Equilibrium

Imagine p as a mutation that could evolve in a population playing exclusively $\hat{p}$. If $\hat{p}$ is Nash but not an ESS, then there are mutations p which could arise which would get a better payoff against itself than $\hat{p}$ gets against it, and also get as good a payoff against $\hat{p}$ as $\hat{p}$ gets against itself. In this sense $\hat{p}$ is unstable, because you can imagine such a mutation taking over the population.

Exercise: Translate the Rock Paper Scissors Normal Form game below into Payoff Matrix form:

	Rock′	Scissors′	Paper′
Rock	(0,0)	(1,-1)	(-1,1)
Scissors	(-1,1)	(0,0)	(1,-1)
Paper	(1,-1)	(-1,1)	(0,0)

Answer:

$$U = \begin{bmatrix} 0 & 1 & -1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \end{bmatrix}$$

Simple – The payoff matrix just contains player 1's payoffs!

Example: Let's show that the game with this payoff matrix has no Evolutionarily Stable Strategies.

Answer: Any Evolutionarily Stable Strategy must be a Nash Equilibrium, so let's check all of the Nash Equilibria and see if they are non-invadable.

We've talked about the Nash Equilibria of the Rock Paper Scissors game before, so I will skip finding the Nash Equilibria (you can do this as an Exercise if you like). We know that there are no PSNE, and that the unique MSNE is $\hat{p} = \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)$.

Note: Non-invadability requires that this strategy must get a strictly higher payoff against any alternative best response p than p gets against itself. **Exercise:** Try to think of a counterexample.

Answer: Here's mine. Consider the alternative mixed strategy $p = \left(0, \frac{1}{2}, \frac{1}{2}\right)$. First we have to check that $p \neq \hat{p}$, which holds for this example. Second check that $p^T U \hat{p} = \hat{p}^T U \hat{p}$. Here we have

$$p^T U \hat{p} = \begin{bmatrix} 0 \, \frac{1}{2} \, \frac{1}{2} \end{bmatrix} \begin{bmatrix} 0 & 1 & -1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \end{bmatrix} \begin{bmatrix} \frac{1}{3} \\ \frac{1}{3} \\ \frac{1}{3} \end{bmatrix}$$

$$p^T U \hat{p} = \begin{bmatrix} 0 - \frac{1}{2} \, \frac{1}{2} \end{bmatrix} \begin{bmatrix} \frac{1}{3} \\ \frac{1}{3} \\ \frac{1}{3} \end{bmatrix}$$

$$p^T U \hat{p} = 0$$

While also

$$\hat{p}^T U \hat{p} = \begin{bmatrix} \frac{1}{3} \, \frac{1}{3} \, \frac{1}{3} \end{bmatrix} \begin{bmatrix} 0 & 1 & -1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \end{bmatrix} \begin{bmatrix} \frac{1}{3} \\ \frac{1}{3} \\ \frac{1}{3} \end{bmatrix}$$

$$\hat{p}^T U \hat{p} = \begin{bmatrix} 0 \, 0 \, 0 \end{bmatrix} \begin{bmatrix} \frac{1}{3} \\ \frac{1}{3} \\ \frac{1}{3} \end{bmatrix}$$

$$\hat{p}^T U \hat{p} = 0$$

So let's look back at condition b) in the definition of an ESS: In order for $\hat{p} = \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)$ to be non-invadable, we will require that $p^T U p < \hat{p}^T U p$. We have just found that $\hat{p}^T U p = 0$, so let's check the value of $p^T U p$:

$$p^T U p = \begin{bmatrix} 0 \, \frac{1}{2} \, \frac{1}{2} \end{bmatrix} \begin{bmatrix} 0 & 1 & -1 \\ -1 & 0 & 1 \\ 1 & -1 & 0 \end{bmatrix} \begin{bmatrix} 0 \\ \frac{1}{2} \\ \frac{1}{2} \end{bmatrix}$$

$$p^T U p = \begin{bmatrix} 0 \, -\frac{1}{2} \, \frac{1}{2} \end{bmatrix} \begin{bmatrix} 0 \\ \frac{1}{2} \\ \frac{1}{2} \end{bmatrix}$$

$$p^T U p = 0$$

Since $p^T U p = \hat{p}^T U p$ we have found a counterexample: we found a p that satisfies $p \neq \hat{p}$ and $p^T U \hat{p} = \hat{p}^T U \hat{p}$ but for which it is not true that $p^T U p < \hat{p}^T U p$. So, $\hat{p}$ cannot be an ESS. Since all ESS must also be NE, and the only NE is not an ESS, therefore there are no ESS of the Rock Paper Scissors Game.

Remark: What did we just see, in words? Consider a population where every single individual is randomizing between Rock, Paper, and Scissors. On average over the long run, they will get 0 payoff per contest. Now imagine that one of these individuals suddenly decides to stop playing Scissors completely, and only play Rock or Paper with equal probability. This individual will continue to obtain 0 payoff. That behaviour could spread with absolutely no punishment; playing Rock/Paper against Rock/Paper/Scissors has an expected payoff of 0, playing Rock/Paper against Rock/Paper has an expected payoff of 0, and playing Rock/-Paper/Scissors against Rock/Paper has an expected payoff of 0. So Rock/Paper can spread arbitrarily much in the population without penalty.

Exercise: (NOT COVERED IN CLASS) How many ESS can you find in the game with the following payoff matrix?

$$U = \begin{bmatrix} 1\,0\,0 \\ 0\,1\,0 \\ 0\,0\,1 \end{bmatrix}$$

Answer: When we're trying to find ESS, a good first step is to ask: can we identify any Strict Nash Equilibria, which will necessarily be ESS?

We need some labels: call the first column/row of the matrix R_1, the second R_2, and the third R_3. Then we can immediately by inspection notice that $BR(R_1) = \{R_1\}$, $BR(R_2) = \{R_2\}$, and $BR(R_3) = \{R_3\}$. Since none of the Nash Equilibrium strategies has any alternative best responses (they are Strict Nash Equilibria), each pure strategy is an ESS.

In words: a population of players which plays any of these strategies with certainty cannot be invaded by a population which plays another of these strategies, since any deviation from the dominant strategy will be met with 0 payoff when played against the dominant strategy. This means that there are three Evolutionarily stable strategies: $R_1, R_2,$ and R_3.

13 Section 13: 2019 April 11

Topics:

- Some subtle problems with rationality [1]
  - ↪ Too many calculations [2]
  - ↪ Incompleteness of games [3]
  - ↪ Systematic irrationality [4]

After this class, I expect every student to be able to:

- name a few criticisms of rationality that aren't half a century old and extremely boring

BACK TO TABLE OF CONTENTS

① Some subtle problems with rationality

Remark: Bounded rationality is motivated by a similar concern to the one that we hear all the time from political scientists: in lots of situations that we care about, there's no reason to expect rational play. Let's first quickly remind ourselves what rationality *is*.

Recall: A function $u : X \to \mathbb{R}$ is a utility function corresponding to the preference relation $\succeq$ if, $\forall x, y \in X$, $x \succeq y \iff u(x) \geq u(y)$. If a preference relation satisfies the following four axioms for all alternatives x, y, z in the set of all alternatives X, then there exists at least one utility function corresponding to it:

↪ *Complete:* $x \succeq y$ or $y \succeq x$

↪ *Transitive:* $x \succeq y,\, y \succeq z \implies x \succeq z$

↪ *Independent of Irrelevant Alternatives:* $x \succeq y \in X \setminus z \implies x \succeq y \in X$ when z irrelevant to x, y

↪ *Continuous:* $x \succ y \succ z \implies \exists \alpha \in [0; 1]$ which is unique so that $\alpha x + (1-\alpha) z \sim y$

Question: We hear criticisms of rationality all the time. Frankly, most of them are tired (the ones that get tricked by the terrible word "rationality" and forget that we're really talking about something closer to "consistency"). Instead of talking about the tired problems, let's focus on some wired problems. What are some plausible situations where this idea might not apply? Especially since this is Scott Page's class after all, we're hardly orthodox game theory evangelists here ...

Problem 1 (Incredible volume of calculations): The decision theory places an immense pressure on people to be extremely sophisticated calculating machines. Prisoner's Dilemma is easy enough to solve, but try applying game theory to a game of chess (not to mention finding the Nash Equilibrium of how to vote in a Russian election).

Proposed solution 1: Simple bounded rationality and satisficing. Say that I have a cognitive budget for how many calculations I'm willing to do, and I use the common budget-saving decide that I will play any **epsilon-optimizing strategy** s, that is, any s for which

$$u(s) \geq u^* - \epsilon$$

for sufficiently small ϵ. This is exactly the idea of satisficing. Now suppose that I am the auctioneer in a Dutch Auction (where I start high and lower the bid until someone bids on the item), and I know that someone in the audience is a satisficer. The price of the item is publicly known to be τ (it's a boring auction), so everybody wants to bid arbitrarily close to τ from below. For example, nobody would have a unilateral incentive to defect if everyone bid τ, and let's suppose the community lands on this silly equilibrium.

Now, suppose that the satisficer's ϵ is perceptibly large, say $\epsilon = 0.01$. Then $u^* = 0$, and the set of acceptable options is the set $u(s) \geq -0.01$. Then the satisficer will be content to bid $\tau + 0.01$. This means that as the auctioneer, if I put up infinite goods for sale, the satisficer would be perfectly happy to pay me infinite money!

So what does this tell us? It means that if you are rational, and you think that someone is satisficing, and the system is simple enough that there's *any way* that you can do the full rational calculation, you might be able to exploit the system to gain a huge advantage. This is why Scott says that rationality holds when the stakes are high and especially when play is repeated; people have time to learn the rational moveset, and they have the incentive to behave rationally.

Proposed solution 2: Extremely informal heuristics. Imagine that we have an election where I calculate pivotality as follows:

P(pivotal) = Probability of creating a first place tie between two parties I like + Probability of breaking a first-place tie in favour of a party I like

Suppose there are 3 other independent electors ν_1, ν_2, ν_3 and 3 candidates c_1, c_2, c_3, I have $c_1 \succ c_2 \succ c_3$, and abstention A is not allowed.[9] Then, assuming that only pairwise ties are possible (just for tractability):

$$\begin{aligned} P(\text{pivotal}) = \; & \text{Probability of creating a first place tie between } c_1 \text{ and } c_2 + \\ & \text{Probability of creating a first place tie between } c_1 \text{ and } c_3 + \\ & \text{Probability of creating a first place tie between } c_2 \text{ and } c_3 + \\ & \text{Probability of breaking a first place tie between } c_1 \text{ and } c_2 + \\ & \text{Probability of breaking a first place tie between } c_1 \text{ and } c_3 + \\ & \text{Probability of breaking a first place tie between } c_2 \text{ and } c_3 \end{aligned}$$

$$\begin{aligned} P(\text{pivotal}) = \; & P(\nu_1, c_1)P(\nu_2, c_2)P(\nu_3, c_2) + \\ & P(\nu_1, c_1)P(\nu_2, c_3)P(\nu_3, c_3) + \\ & P(\nu_1, c_2)P(\nu_2, c_3)P(\nu_3, c_3) + \\ & P(\nu_1, c_1)P(\nu_2, c_2)P(\nu_3, c_3) \end{aligned}$$

How can I possibly make conjectures about *all* of those probabilities? With 3 electors it's hard enough, imagine 80,000,000 electors. Just as awful, how can I account for the fact that this is not a decision theoretic problem, it's a game theoretic problem, because everybody else in the system should also be doing the exact same calculations! One answer is heuristics.

Example: Anything we say about pivotality calculations is controversial, but one

[9] Because of the classical paradox of voting, it might seem that the only reason to focus on pivotality is if abstention is possible, but notice that you also don't want to waste your vote by voting for a candidate who has an extremely small chance of winning. The only situation in which pivotality is truly irrelevant is if abstention is not permitted and there are only two candidates. I don't think that describes any democratic system.

result is that when people have low or even negative costs, they should end up voting for a small number of relatively popular parties that are not their least-favourite party. Here's an example of three simple rules that probably get you close to the result of the actual full pivotality calculations:

Simplification 1: I can use public information like polls and past results to estimate the chance of ties using aggregate information.

Simplification 2: What the specific electors do is irrelevant; actually all that matters is how many electors choose each strategy.

Simplification 3: It doesn't matter what the probabilities are, just the relative likelihood of each outcome, together with how much I care about those outcomes.

Simplification 4: Because polls are public information, most peoples' guesses about their pivotal probabilities are similar to my own guesses, so the rule they adopt will be similar to the one I adopt.

Then we can employ those simplifications as follows:

By simplification 1: Check the polls and decide what the most likely tie is

By simplifications 2 and 3: Just take the most likely tie that I care about and vote as if that tie were going to happen

By simplification 4: Each person will strategize to boost their own party just like I am, so they'll strategize in proportion to how popular each party is in the polls, so the strategic behaviour will come out in the wash.

The result should be similar-ish to the pivotality logic result: people will pick the parties that are most likely to win, and vote for whichever they like better among those. This is the Cox 1994 result, but by extremely informal simplifications.

Even if we believe that the machinery of pivotal voting logic is the strictly correct way of fleshing out the actual process, the actual process doesn't involve the undoable game theoretic calculations.

The challenge: So what was all of this for then? Didn't we know that hand-wavy stuff already?

My answer: The simple approach didn't actually buy us any insight; it was only a useful story to tell once we already had the actual result. However, ideas of heuristics that are much closer to the formalized system, like satisficing, can be more useful.

Problem 2 (Gödel's Incompleteness Theorem): This one is out there, but it's amazing. Ken Binmore in many books and articles argued that Gödel's Incompleteness Theorem implies that there is no such thing as perfect rationality (Binmore, 2009). Why would that be?

One of **Gödel's incompleteness results** is (very roughly) that any consistent formal theory which is adequate to contain number theory is incomplete. Gödel showed how to construct a sentence S of number theory so that neither S nor $\neg S$ is a theorem, and yet because they are statements about number theory, one of them must be true (Stoll, 1961: p. 166). Gödel did this using a procedure of self-reference, and Binmore (1990) points out that this should make practicioners of game theory quite nervous about the totality of their theory, since the common knowledge assumption ("I know that you know that I know that ...") is just about the most self-referential thing you can possibly imagine!

Gödel's Incompleteness Theorem is an incredibly abused and misused result, so we need to be careful about accepting arguments that refer to it. On the face of it, we might just as easily accept that because quantum tunneling exists, perfect play is impossible and trembles are inevitable with some nonzero probability (I guess that thought is true, but it's definitely not useful). Luckily, this has been formalized in one of my all-time favourite papers, in which Tsuji et al. (1998) apply this idea. To set up the problem they first quote the economist Eswar Prasad, who wrote

"For [finite] games (i.e., with finite numbers of players and strategies) it is easy to describe a computational procedure for finding Nash equilibria. Since the problem is inherently finite, one approach would be to consider exhaustively all possible strategy combinations and to check if any player can gain from unilateral deviation."

This argument asserts that Nash Equilibria can be computed in any finite game.

Note: Here's a way of talking about games that we haven't adopted in this class, to avoid over-formalizing, but this is a good moment to introduce it. A non-cooperative game Γ consists of a **von Neumann Triplet** $\langle N, S_i, u_i \rangle$ for $i \in \{1, 2, \ldots, N\}$ where N is the number of players, S_i is the strategy set of player i, and u_i is a utility function which maps each combination of strategies onto a real number.

Now, consider the predicate $P(x, y)$ which says that y is a Nash Equilibrium strategy profile which corresponds to the von Neumann Triplet x. Tsuji et al. (1998) use Gödel Incompleteness to show that there are Nash Equilibrium profiles y corresponding to some von Neumann Triplet x which are not computible within the logical system that supports the definition of the game.

So even if the game is finite, and it feels like we should be able to line up all the possible strategy profiles and check them against every possible deviation, it's actually certain that there exist games that have Nash Equilibria that are not possible to compute!

Never forget that the truth hurts.

Problem 3 (Systematic departures from rationality): We talked about trembles, but we never talked about when people are systematically irrational. (Binmore, 1990) argues that if someone is irrational at node x, we should be more likely to expect that they'll be irrational at node y: we should update. But that's not what the theories we've talked about so far this semester have said; they've dealt only with purely random, non-systematic departures from rationality.

Solution 1: Binmore (1990) suggests explicitly modeling the players' decision-making processes to handle this, and specifically he wants to incorporate a learning process into the model itself. So, for example, if I learn that you're irrationally cooperating in the Prisoner's Dilemma over and over again, then I should start cooperating too.

This is extremely reminiscent of the *Evolution of Cooperation* (Axelrod, 1984), which has a very similar finding. Recall that in Axelrod's tournament TFT wins specifically because it scored OK against mean strategies and extremely well against nice strategies. In fact, this latter capacity – taking advantage of nice strategies being nice – is the key to TFT's success that Axelrod intensely focuses on in his book.

14 Section 14: 2019 April 18

This section was spent entirely on review.

BACK TO TABLE OF CONTENTS

References

Axelrod, R. (1984). Evolution of cooperation.

Binmore, K. (1990). *Essays on the foundations of game theory*. Basil Blackwell.

Binmore, K. (2009). *Rational decisions*. Princeton University Press.

Gibbons, R. (1992). *Game theory for applied economists*. Princeton University Press.

Plott, C. R. (1967). A notion of equilibrium and its possibility under majority rule. *American Economic Review*, 57(4):787–806.

Riker, W. H. and Ordeshook, P. C. (1968). A theory of the calculus of voting. *The American Political Science Review*, 62(1):25–42.

Stoll, R. R. (1961). *Sets, logic and axiomatic theories*. W. H. Freeman and Company.

Tsuji, M., Costa, N. C. A. D., and Doria, F. A. (1998). The incompleteness of theories of games. *Journal of Philosophic Logic*, 27:553–568.

Tyson, S. (2017). Notes on a game theory course.

Watson, J. (2013). *Strategy: An introduction to game theory*. W. W. Norton & Company, 3 edition.