PSCI552 - Formal Theory (Graduate)

Lecture Notes

Dr. Jason S. Davis

Spring 2020

Contents

1 Preliminaries

1.1 What is the role of theory in the social sciences?

- We can't escape it!
- First off, what is "the" scientific method? How would we characterize it?
- Is it just a collection of hypotheses and tests?
- Once we have some empirical result, if we want to generalize from it, or understand how it fits with other results, then we need some kind of theory about what generated the result.
- Even in physical sciences, theory often takes us far beyond the individual results. Special relativity and the constant speed of light (see light clocks). String theory.
- Hypothetico-deductive model. Create hypothesis/model, see if you can falsify it.
- By this model, can never prove something, only disprove it. Whatever best withstands falsification is what you should believe.
- Even "atheoretical" empirical work usually involves some understanding of how, say, measurement works.
- Measurement assumptions are often *highly* controversial. Consider measurements of poverty, inequality, etc.
- Still, some work/approaches are less theory dependent than others (e.g. machine learning, drug trials, etc.)

1.2 So why formalize that theory? And why take this course?

- If you want to learn a particular "lens" for looking at the world (learning game theory has transformed the way I think about the world more than learning any particular set of information).
- Disciplining one's thinking by adopting what Skip Lupia (Michigan/NSF) calls a "commitment to precision" (even if you're not writing a model, thinking precisely about the logical structure of arguments and about what one is trying to prove can be important).
  - Lead people from where they are to where you are. What do things you believe imply?
  - You may not believe my conclusion, but you can see how I came to it. So where have I gone wrong?
  - **Note:** *All* political scientists make assumptions about players, actions, strategies, information, beliefs, outcomes, payoffs, methods of inference, etc.
  - Conclusions depend on these assumptions, and the link between those assumptions and conclusions.
  - Formal theory in general is just a way of making that chain explicit, while game theory is more specifically about a particular collection of approaches to strategic interaction.
- All of this is useful even if you are primarily interested in doing empirical work.
- Can also use the math to generate new intuitions you might not have started off with (i.e. as an "intuition pump").
- Finally, there is value to learning to thinking in terms of abstraction.

- Abstraction "new ideas or conceptions are formed by considering several objects or ideas and omitting the features that distinguish them." Rosen (2014), quoted in the Kieran Healy piece.

- "Furniture", "honor killing", "social-democratic state", "white privilege"

- Temptation to criticize theory by simply *identifying* things that have been left out.

- All theories are simplifications of the world. Mathematical models are "reductionist", but not always more than any other theory.

- Models are "maps" that are designed to provide insight into some aspect of reality that is of interest. Often we want the minimum level of detail required to illustrate the relationship or trade-off we are interested in.

- Adding in more detail can be useful if it allows us to get leverage on more trade-offs, comparative statics, conditionalities, etc. Can allow for better assessment against data, more "rich" conclusions. But detail for detail's sake is not the point.

- Relationship between formal theory and data can be subtle: given that all models are literally false, you don't "test" whether models are "true".

- Not necessarily problematic to "fit" a model to data; oftentimes models have been driven to explain something that seems like a puzzle. For example, there has been a significant amount of effort spent on generating models that explain high, positive turnout in the face of the "irrationality of voting" paradox.

2 Proofs

2.1 Introduction

- Early problem sets in this course will require writing some simple proofs.

- Moreover, formal modeling in general is about proofs, insofar as you are showing how a certain set of modeling assumptions leads to a set of conclusions. This is what a proof sets out to do.

- Two books worth referencing for proofs: Velleman's 'How to Prove it' and Cupillari's 'Nuts and Bolts of Proofs'.

- Velleman is a good structured introduction on how to write proofs, while Cupillari's proofs are heavier on the math, so maybe a little closer to what you'll actually see.

2.2 Logic

- Will sometimes use logical notation. I.E. $A \to B$, $A \to \neg B$, $A \leftrightarrow B$. I don't intend to spend too much time on logic in this course, but it's useful for exposition.

- $A \to B$. What if we have $\neg B$? What does this imply about A?

- Contrapositive: $(A \to B) \leftrightarrow (\neg B \to \neg A)$

- What if we have B? Do we know anything about A? Nope! Fallacy of affirming the consequent.

- AND: both elements must be true, represented $\wedge$. OR: one element must be true, represented $\vee$.

- Say we have $A \vee B$. We have A. What do we know about B?

- Say we have $\neg B$. What do we know about the statement $A \wedge B$?

- Proofs are based on using a series of logic statements to show that some B is implied by A. Formal models are thus just this: series of logical statements.
- Conditionals versus biconditions.
- Review of some laws:
  - DeMorgan's Laws: $\neg(P \wedge Q) \leftrightarrow \neg P \vee \neg Q$ and $\neg(P \vee Q) \leftrightarrow \neg P \wedge \neg Q$
  - Distributive $P \vee (Q \wedge R) \leftrightarrow (P \vee Q) \wedge (P \vee R)$
  - $\neg\neg P \leftrightarrow P$
- Example: Simplify $P \vee (Q \wedge \neg P)$
- Composite statements. e.g. $A \rightarrow (B \rightarrow C)$
- Quantifiers. $\forall x,\ P(x)$ or $\exists x,\ s.t.\ P(x)$
- May need to convert natural language statement into formal logical structure.
- Everyone in the class thinks Jason is great: $\forall x \in C, P(x)$ where $P(x)$ represents "Jason is great"
- Equivalent formulation: $\nexists x \in C\ s.t. \neg P(X)$

2.3 What are we proving?

- Equality of numbers/variables.
- A implies B.
- **Aside:** A increases the probability of B?
  - Oftentimes formal models are framed in deterministic terms (i.e. A is the unique equilibrium if B). Only in some instances do we have mixed strategies or stochastic components that would allow for randomness within the model. How do we reconcile this with a world in which there seem to be counterexamples to every theory?
  - Particularly important if we want to estimate a model with data. "Zero likelihood" problem; likelihood function would be zero with one counterexample.
  - Can incorporate stochastic elements into the theoretical model. Then it's "deterministic" given the resolution of the stochastic component. Sometimes this can be done "usefully"; maybe we think people make errors stochastically, or some state of the world is realized stochastically, and incorporating that allows the model to provide greater insight.
  - Sometimes we're adding complications that aren't providing any real insight. Focus should be on whether the model is "useful" not whether it is realistic.
  - Can incorporate stochastic elements into the statistical model (e.g. Signorino's stuff). Can conceptualize this as actor "errors" or as "stuff not in the model which also matters but which we can't measure or account for systematically" or as measurement error.
  - Signorino's work builds off a Quantal Response Equilibrium (QRE, McKelvey and Palfrey) approach even if the original model wasn't QRE.
- A if and only if B.
- Proofs involving sets (equality, subsets, etc.).
- Existence.

- Uniqueness.
- "For all" statements.
- Any other complicated statements. Just be sure to prove all the components.

2.4 Different proof strategies

2.4.1 Direct

- Simply show the different steps. Directly.
- Proving $A = B$, can do the old $LS = RS$.
- For a statement like $A \to B$, think about what it is that you're proving. IF you have A THEN you have B. So you assume the antecedent (A) and then see if you can show that B must also be true.
- Example: prove if m is even, then m^2 is even.

  Proof. If m is even, it can be written as $m = 2r$. Then we have $m^2 = 2^2r^2 = 2(2r^2)$. Thus m^2 is divisible by 2, and thus is even by definition. □

- Alternatively, can prove by contrapositive: Assume $\neg B$ and show $\neg A$ as $\neg B \to \neg A \leftrightarrow A \to B$.
- Keep track of what you've been given, and then see if you can combine that information logically to get to the "goal".
- E.g. do both directions to get a biconditional, i.e. $A \leftrightarrow B$
- Sometimes the statements we're proving include antecedents that are themselves logical relationships. For instance, transitivity of R is equivalent to saying that $xRy \wedge yRz \to xRz$. If we want to prove implications of transitivity, we need to assume that relationship as the antecedent.
- Existence: Just need to show it's true for some example. To prove $\exists x \in [0, 1]$ *s.t.* $x > 0.5$, $x = 0.75$ is sufficient.

2.4.2 Indirect/contradiction

- Any proof by contrapositive can be formulated as a proof by contradiction. Do what makes you happy/comes easiest (not independent things I'm sure).
- Proof that contradiction proofs work.

  Proof. Want to prove that $P \to Q$. We start by assuming P and $\neg Q$ and generate a contradiction.

$$\begin{aligned}\neg(P \wedge \neg Q) &\leftrightarrow \neg P \vee Q \\ &\leftrightarrow \neg Q \to \neg P \\ &\leftrightarrow P \to Q\end{aligned}$$

□

- Assume that what you're trying to prove is not true, and then show that this leads to a contradiction. If it couldn't not be true, it must be true!
- Uniqueness: after showing an element exists, assume that there are two elements and find a contradiction.

- Example: Proof that there is an infinite number of prime numbers.

Proof. Assume that this was not the case, i.e. that there is a *finite* number n of prime numbers. We can thus write the set of prime numbers as:

$$\{p_1, p_2, p_3, \ldots, p_n\}$$

If this is *all* the prime numbers, then we shouldn't be able to find another one. The trick here is to prove that there has to be another prime number which isn't included in this set; this way, for any finite set of prime numbers, there has to be at least one more, and we would generate a contradiction. To start, we take the product of all the prime numbers in our finite set and add one, and assume that this must be equal to an integer c times a prime number that is included in our finite set, i.e.:

$$p_1 p_2 p_3 \ldots p_n + 1 = c p_k$$

If we divide both sides by p_k we get the following:

$$\frac{p_1 p_2 p_3 \cdots p_n}{p_k} + \frac{1}{p_k} = c$$

Which rearranging, gets us:

$$\underbrace{c}_{\text{integer}} - \underbrace{\frac{p_1 p_2 p_3 \cdots p_n}{p_k}}_{\text{integer}} = \underbrace{\frac{1}{p_k}}_{\text{not an integer!}}$$

Now if p_k is in the finite set of prime numbers, then it must be the case that both terms on the left-hand side are integers (because the product of all the prime numbers must be divisible by that prime number). However, the right hand side is not an integer! So $LHS \neq RHS$ and we get a contradiction. □

2.4.3 Induction

- Most complicated in terms of logical structure.
  1. Prove base case.
  2. Inductive step: Prove that if relationship is true for n it is also true for $n+1$.
- This demonstrates that it's true for all natural numbers. A domino analogy is sometimes used; you've shown it for the first case, and if it being true for n means it's true for $n+1$, then it being true for $n = 1$, means it's true for $n + 1 = m = 2$, which means it's true for $m + 1 = 3$ and so on.
- Example: $0 + 1 + 2 + ... + n = \frac{n(n+1)}{2}$
- Base case: $0(0 + 1)/2 = 0$
- Inductive step:
  - Assume antecedent, i.e. that it is true for n, so $0 + 1... + n = \frac{n(n+1)}{2}$.
  - Now let's check $n+1$. If antecedent is true, $0+1...+n+n+1 = \frac{n(n+1)}{2} + n+1 = \frac{n(n+1)}{2} + \frac{2(n+1)}{2} = \frac{(n+1)(n+2)}{2} = \frac{(n+1)((n+1)+1)}{2}$. So we're done!
- Note: These can be used for finite sets, and *certain kinds of infinite sets*. The distinction is it has to be countable for induction to be used. You cannot use induction for uncountably infinite sets (which happens when you have continuity).

3 Decision Theory

3.1 Preference Relations and Rationality

- Q: What does it mean for someone to be rational?
  - Economics/game theory doesn't "own" the word, so it's clear there are multiple definitions.
  - More substantive: not being emotional, not caring about certain things, focusing on material consequences, etc.
  - "Rational choice" formal theory takes a different tack, though.
- Rationality in formal theory is usually defined with respect to a **preference relation**.
- A preference relation, or preference order, ranks the different **outcomes** of a **decision problem**.
- For instance, if the outcome space is $S = \{Beilers, Federal, Dunkin\}$, my preference relation would order these.
- To define some notation, if x, y, are outcomes, we have $xRy \equiv x \succeq y$ if x is "at least" as good as y. This is sometimes referred to as "weakly preferring" x over y.
- We have $xIy \equiv x \sim y$ iff both xRy and yRx. This is the *indifference* relation.
- We have $xPy \equiv x \succ y$ iff we have both xRy and $\neg yRx$. This is having a *strict* preference for x over y.
- A preference relation is rational iff it is:
  1. Complete (provides an ordering for every pair of outcomes)
     Formally: for any x, y in outcome set X, $xRy \vee yRx$
  2. Transitive. Formally, transitivity is:

$$xRy \wedge yRz \rightarrow xRz$$

- These do not impose dramatic substantive assumptions about preferences, though transitivity is the most controversial.
- A rational preference relation assumes that someone exhibits *goal-oriented* behavior, but does not judge those goals.
- One's main goal could be to make someone else worse off, or it could be to lose as much money as possible: this is all fine within our formal framework of rationality
- Because of this, some complain that formal theory's definition of rationality is almost unfalsifiable, since most behavior can be rationalized.
- A game theorist would say: sure! That's a virtue, not a vice.
- In fact, the goal here is to show conclusions about goal-oriented behavior that have to be true given *minimal* assumptions. So the *less substantive the better*, since this increases the scope conditions.
- To emphasize, the goal here is *not* to provide a substantive understanding of rationality - though this may be useful for other purposes! Instead, it is to characterize things that are true about the broader subsets of human behavior that can be described as goal-oriented.
- Given formal theorists' obsession with generality, we sometimes find cases where results will be demonstrated with *even more general* assumptions on preferences. Two examples are:

- Quasi-transitivity, i.e.:
  $xPy \wedge yPz \to xPz$
- Acyclicity: R is acyclic on the outcome space X iff:
  For any $S = \{x_1, x_2, x_3, \dots, x_n\} \subseteq X$:
  $(x_1 P x_2 \wedge x_2 P x_3 \dots \wedge x_{n-1} P x_n) \to x_1 R x_n$

- Why can't acyclicity be defined in terms of triples, like transitivity and quasi-transitivity??
  Ans: Take quasi-transitivity and imagine you have four elements $\{x, y, z, a\}$. If you have xPy, yPz then you also have xPz. So now we can look at xPz, zPa, breaking everything up into triples. We can't do this for acyclity because xPy, yPz only implies xRz. So imagine if we had xPy, yPz, zPa, aPx. This is clear a cycle. But if we tried to break it up into triples, we'd get $xPy, yPz \to xRz$ and then xRz, zPa, aPx which is not a cycle.

- We get to *no* cycles via the fact that in order for R to be acyclic, the relationship needs to hold *for any subset* of the choice space X.

- E.g. imagine you have choice set $X = \{x, y, z, a\}$ and xPy, yPz, zPa, zPx, xRa. For the subset including all the elements, we have xPy, yPz, zPa, xRa which doesn't violate acylicity. However, for the subset $X' = \{x, y, z\}$ we have xPy, yPz, zPx which is a cycle and violates the acyclicity of R assumption.

- Note I've only spoken of acyclicity of R because P is defined in terms of R (i.e. $xPy \leftrightarrow xRy \wedge \neg yRx$)

- Next let's talk about maxima/maximal sets as defined by a preference relation.

- An element x of a set S is a **maximum** iff, for all $y \in S$, xRy.

- In Assignment 3, you're asked to prove that a rational preference releation (i.e. an ordering with completeness and transitivity) implies the existence of at least one maximal element in any finite set. Let's do the same thing, but using a weaker assumption on preferences: acyclicity.

  Proof. Take any finite set $S \subseteq X$. Now take an arbitrary element $x \in S$. Either xRy for all $y \in S$, so x is a maximal element, or there exists a y such that *not* xRy, which by completeness means we have yRx and not xRy, i.e. yPx.

  So now let's look at y. We know that yRx. We can now check every other element in the set (i.e. in $S \backslash \{x, y\}$) and see if there exists a z such that *not* yRz. If we find such an element, we know for that z it must be the case that zPy. Thus we have zRy and by acyclicity, zRx, and can focus on z now.

  We then check all elements $a \in S \backslash \{x, y, z\}$, looking to see if an a exists such that $\neg zRa$. If it doesn't, a is a maximal element. If it does, we now have collected aPz, zPy, and yPx, which via acyclicity gets us aRz, aRy, and aRx. We keep repeating this process, and because the set is finite, we eventually exhaust the set and *have* to find a maximal element. $\square$

- Why is proving existence of maximal elements so important? Basically all of game theory and decision theory is about (1) showing what optimal choices are, given the structure of the game; (2) showing how these choices change as features of their environment change.

- However, the step of showing that there *is* an optimal choice is not always obvious! And we can't talk about behavior of optimal choices without an optimal choice.

- In this sense, the proof you're doing in Assignment 3 Q4 is *foundational* to formal theory, because it "buys" you a lot in terms of scope conditions: any case where the outcome set is finite, and you have a rational preference relation, requires at least one "optimal" choice.

- As we move to non-finite spaces, we have to do more to demonstrate the existence of maximal elements before we can really do anything else. This is the crux of Nash's proof of equilibrium existence: we can talk more about behavior of equilibria once we show they exist.
- We won't get into the details of these existence proofs, because I don't want to bore you all too much with real analysis/topology.
- But to give you some flavor: many of the existence proofs are loosely derived from a generalization of "extreme value theorem" known as Weierstrass theorem.
- Weierstrass Theorem: if you have a continuous function on a compact set, you have a maximal element.
- Compactness: closed and bounded.
- Continuous function: think of intuition, won't get into details.
  1. Counterexample 1: continuous function $f(x) = x$ on non-compact set $[0, 1)$. Not closed, doesn't have a maximum.
  2. Counterexample 2: same function, on set of integers $\mathbb{Z}$. This set is actually closed, but not bounded! And also doesn't have a maximum.
  3. Counterexample 3: discontinuous function $f(x) = x$ when $x \in [0, 1)$, $f(x) = 0$ when $x = 1$, on compact set [0,1]. No maximum.
- So you can see how both *continuity* and *compactness* apply when we're talking about functions and intervals on $\mathbb{R}$. Suffice to say that Weierstrass generalizes this to multiple dimensions (and, indeed, arbitrary metric and topological spaces).

3.2 Imposing Structure on Preferences

- Transitivity can fail! Two instances demonstrated "in the wild" are:
  - "Just perceptible" differences. E.g. I may be indifferent between each additional crystal of sugar in my coffee, I would not be indifferent between a cup of sugar or no sugar in my coffee. This violates transitivity.
  - Framing differences (see Kahenman and Tversky).
    Example: you are about to buy a stereo for \$125 and a calculator for \$15, and are told that you can drive to another store where the calculator is \$5 off. In surveys, people are more willing to drive in that case than when they are told it's \$5 off the stereo; the difference in outcomes is the same, but people rank them differently based on framing.
    * This has led to a whole subset of behavioral science talking about "nudges" and their effect on behavior.
- However, for the most part, social science theory is rarely oriented around behavior that violates transitivity.
- Indeed, in practice, most models end up imposing more structure on preferences than "there exists a transitive ordering".
- 2x2 matrices don't need to, because there is a very restrictive set of outcomes. However, what if our outcome space is, for instance, all real numbers?
- At this point, it is helpful to talk about what kinds of structure we might place on preferences, which for the most part, *also* requires placing restrictions on *outcomes*.

- So let's for now think of outcomes as being represented by **vectors** in Euclidean space, *i.e.* $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^n = X$
- In economics, these vectors often represent commodity bundles that a consumer might purchase. Each entry in the vector in another thing one could buy.
- Having defined *outcomes* in this way we can now think about what kinds of structure we might place on *preferences* over these outcomes.
- An example: **convex** preferences are a very common assumption.
- What is a convex combination? Sort of like a weighted average.
- More precisely: it's a linear combination of points in $\mathbb{R}^n$ such that all the coefficients add to one, so for some set of vectors $\{\boldsymbol{x_i}\}_{i=1}^n$
$$\sum_i^n \lambda_i \boldsymbol{x_i} \; s.t. \; \sum_i^n \lambda_i = 1$$
- Convex sets: a set which includes all the convex combinations of its points.
- Convex hull: the smallest set which contains all the convex combinations of a set of points.
- Convex preferences: if xRz and yRz then $(\lambda x + (1-\lambda)y)Rz$, for $\lambda \in (0,1)$
- Strict convexity: if xRz and yRz, and $x \neq y$, then $(\lambda x + (1-\lambda)y)Pz$
- Alternative definition, which you will prove follows in problem set:
Strict convexity: if xIy, $x \neq y$, then $(\lambda x + (1-\lambda)y)Px$
- In words, the above is: a weighted average of two options that one is indifferent between is preferred to either option.
- Strict convexity also implies that if there exists a maximal element, it is unique. This will also be in the problem set.
- Convex preferences as a concept are closely related to the idea of "decreasing returns" or "diminishing marginal utility". We will talk more about this soon!
- Weakly monotonic: if $x > y$ then xPy.
- Strict monotonicity: if $x \geq y$ and $x \neq y$ then xPy.
- Above, vector inequalities have a natural interpretation: strict $x > y$ means that every component of x is higher than the corresponding component of y.
- $x \geq y$ means, straightforwardly, that each component is "at least as" high. Furthermore, imposing this AND $x \neq y$ *implies* that at least one component must be strictly higher (can you see why?).
- Local non-satiation (LNS): $\forall x \in X$ and $\forall \epsilon > 0$ there exists a $y \in X$ such that $d(x,y) < \epsilon$ and yPx
  - Intuition: in any ϵ-ball around a commodity bundle, there exists a commodity bundle which is strictly preferred.
  - Above allows for "bads", i.e. commodities that are disliked by the consumer.
  - However, it still ensures that a consumer's preferences don't reach a point where they don't want more or less of anything.

3.3 Utility Functions

- Up until now, we have been dealing with preference relations, not utility functions.
- Utility functions *map* from the set of alternatives to the real numbers. Formally, for a utility function u and set of alternatives/outcomes X:
$$u : X \to \mathbb{R}$$
- A utility function **represents** a preference relation $\succeq$ iff:
$$x \succeq y \leftrightarrow u(x) \geq u(y), \ \forall x, y \in X$$
- *Only* rational preference relations can be represented by a utility function.
  - Quick Q: Does this mean all rational preference relations can be represented by a utility function? No: example is lexicographic preferences.
- Any rational preference relation defined on a **finite** set of outcomes can be represented by a utility function.
  - Proof: just order the elements and assign numbers ("proof by construction").
- Once we have a utility function representation of *ordinal* preferences, we have an *infinite* set of such functions, as the original ranking is *invariant to monotonic transformations.*
- Example of monotonic transformation: $log(u(x)) = v(x)$, where $v(x)$ represents the same ordinal preferences.
- For now, a utility function simply "functions" (pun intended) as a more tractable (i.e. easy to work with) representation of a preference relation.
- We also have characteristics of utility functions that correspond to characteristics of preferences:
  - Monotonic utility functions imply monotonic preferences.
  - *Concave* utility functions imply *convex* preferences.
    - To see this, note that the definition of a concave function is that:
$$f(\lambda x + (1-\lambda)y) \geq \lambda f(x) + (1-\lambda)f(y)$$
    - So the value of the convex combination of outcomes is better than just the weighted average of utilities.
    - **Q:** From what we know about strict convexity: what does strict concavity imply about maxima of a utility function?
    - **A:** They must be unique!
  - Quasiconcavity? $f(\lambda x + (1-\lambda)y) \geq min\{f(x), f(y)\}$
    - This is actually the more direct analogue to convex preferences, but it's a little less convenient mathematically, so we often end up assuming or looking for concavity instead of quasiconcavity.
- Other properties:
  - Continuity.
  - Differentiability.
  - Additive separability.

- Characteristics like the above have mathematical characteristics that can make using them more useful.
- However! They also imply certain substantive things about the preferences.
- We want to keep track of these things so we know what exactly we're assuming when we derive conclusions.
- Imposing a particular functional form on a utility function (e.g. $u(x) = \sqrt{x}$) "bakes in" a lot of different assumptions.
- Sometimes we need to do this anyway in order to get anywhere, and sometimes we can justify a particular choice as being a reasonable representation of some general form of utility function that's a bit easier to work with.
- But always need to be careful; how much "generality" are we losing with a particular assumption?
  - As an aside, you may have seen the phrase "without loss of generality". This is how it comes into play: an assumption or simplification is made, do we lose generality in the conclusions?
- This is why formal theorists often prefer to work at pretty high levels of abstraction; e.g. generate results about existence of equilibria and comparative statics that rely on seemingly arcane fixed point theorems and invocations of upper hemicontinuity.

3.4 Uncertainty, Expected Utility

- **Review Appendix 5.16-5.18 before continuing**
- While we have defined a rational preference relation as one that is complete, reflexive, and transitive, we can go a bit further than that in defining rationality overall in the context of decision/game theory.
- Specifically, we want to impose that one's preferences are rational over outcomes, AND:
  - They are aware of the **action** space.
  - They know how actions lead to **outcomes** (in some cases, it is important to differentiate actions and outcomes).
- In many cases, the mapping between actions and outcomes is straightforward: if I choose to order a hamburger, I get a hamburger.
- When there is uncertainty, this starts to break down. You choose the *action* of taking this class, but there are multiple outcomes: it could be a good courses, an OK course, or a terrible course.
- You may have a sense of the probability of each, and you have preferences over each, but your actions don't neatly map to outcomes.
- We can thus represent the outcome from an action as a "lottery".
- This lottery can be **simple** or **compound** without it mattering very much, since either induces a probability distribution over outcomes.
- In this case, we need some kind of tool for connecting the choice of *actions* to the preferences over *outcomes.*
- The common tool for this is the von Neumann-Morgenstern expected utility function.

- In the discrete case, the approach here is fairly straightforward and intuitive: we have some set of outcomes, we know the probability of each, and we have a utility function that represents $u(\cdot)$ over those outcomes. So we just multiply the utility of each outcome with its probability and sum it up, i.e.:

$$\sum_{x \in X} Pr(x)u(x)$$

- We have to use integrals in cases where the outcome space is continuous/uncountably infinite. However, the structure is basically identical. If outcomes $x \in X$ are distributed according to probability density function $f(x)$, we get:

$$\int_{x \in X} u(x)f(x)dx$$

- For now, we're going to punt (am I using the sports analogy correctly here?) on integrals though, since they aren't *integral* (hold for laughter) to the the earlier content in this course. We may come back to it later when we talk about Bayesian games, time permitting.
- Q: When we take an expected utility function, do we preserve the property that utility functions are invariant to monotonic transformations? Why or why not?
- We don't: now utility functions need to represent **cardinal** utility as well as ordinal.
- "Intensity" of preferences matters now, because we are combining outcomes via a weighting. So, for instance, for $u(x) = x$, we get very different expected utilities from the two lotteries with probability distribution $\boldsymbol{p} = (p_1, p_2, p_3, p_4) = (0.25, 0.25, 0.25, 0.25)$ and $\boldsymbol{x_1} = (5, 4, 3, 2)$ versus $\boldsymbol{x_2} = (10, 4, 3, 2)$, even though they represent the same ordering over outcomes.
- So something to keep in mind when we consider what we're "baking in" to a model.

3.5 The Many Faces of Decreasing Returns

- As mentioned earlier, convex preferences are, essentially, decreasing returns/marginal utility.
- Risk aversion is also essentially equivalent to decreasing returns. Creates incentives for insurance.
  - To see this, consider two "gambles":

    A: \$0.50 with certainty.
    B: \$1.00 with probability 0.5, \$0.00 with probability 0.5.

    Clearly these have equal expected values. A risk averse person will, however, prefer the choice with less variance, i.e. the option with certainty, while the risk loving person will prefer the opposite. Risk aversion is an implication of decreasing returns: the intuition can be seen from imagining someone who's starting off with the \$0.50 with certainty and considering whether to trade it for the gamble laid out in B. If they switch to B, they have a 50% chance of gaining \$0.50, but an equal chance of losing \$0.50, relative to where they started with A. If there are decreasing returns, that "extra" \$0.50 is worth less to them than the initial \$0.50, so switching to B won't be worth it to them if there is an equal probability of gaining and losing the same amount. This is what it means to say that they are risk averse!

    Consider instead if there were increasing returns. Then the second \$0.50 is worth MORE to them than the first \$0.50, and they are happy to take on a gamble that gives them a chance of earning the extra amount, even with equal chance of losing the first \$0.50. This corresponds to being risk-loving.

Consider that in any case the above example's expected utility is:

$EU(A) = (1)u(0.5)$
$EU(B) = 0.5u(1) + 0.5u(0)$

If $u(x) = x$ (i.e. constant returns/risk neutrality) we have:

$EU(A) = 1(0.5) = 0.5$
$EU(B) = 0.5(1) + 0 = 0.5$

As we can see, risk neutrality implies that one is indifferent between all gambles that produce the same expected value, so they don't care which of A or B they get.

If $u(x) = \sqrt{x}$ (example of decreasing returns/risk aversion) we have:

$EU(A) = (1)\sqrt{0.5} = 0.707$
$EU(B) = (0.5)\sqrt{1} + (0.5)\sqrt{0} = (0.5)(1) = 0.5$

As we can see, risk aversion implies they prefer the gamble with less variance, so they prefer A to B.

If $u(x) = x^2$ (example of increasing returns/risk loving) we have:

$EU(A) = (1)(0.5)^2 = 0.25$
$EU(B) = (0.5)1^2 + (0.5)0^2 = 0.5$

Risk loving implies that one prefers the gamble with more variance, so they prefer B to A.

- Some other important definitions:
  - Certainty equivalence.
    - This is the amount you'd accept with certainty instead of taking the gamble. For the gambles discussed above for concavity and risk aversion, consider when $u(x) = \sqrt{x}$. A is already with certainty, so the certainty equivalent is just the same value. For B, the expected utility is 0.5, and the certainty equivalent is found by solving $u(x) = 0.5$, so $\sqrt{x} = 0.5 \leftrightarrow x = 0.5^2 = 0.25$. So the certainty equivalent is 0.25.
  - Mean-preserving spread.
    - You keep the same expected value, but move more "weight" to the tails of the distribution, such that you "preserve" the same mean but increase the variance. Going from N(0,1) to N(0,2) is an example of a mean-preserving spread.
  - First-order stochastic dominance.
    - First order stochastic dominance: Intuitively, first order stochastic dominance basically implies that throughout the distribution of x, the dominating gamble is producing higher returns. For CDFs this can be stated as:$F_A(x) \leq F_B(x)\ \forall x \wedge \exists x\ s.t.\ F_A(x) < F_B(x)$. The reason this makes sense is because what of what it implies about the pdf: in order for F_B to be higher than F_A at all x (so strict first order stochastic dominance) it would have to be the case that the lower xs are disproportionately weighted. In an extreme case, imagine comparing gamble A in which x is uniformly distributed between 0 and 1, and gamble B that returns $x = 0$ with

probability 1. On $x \in [0,1]$, the CDF of A is x, while the CDF of B is 1. So A first order stochastic dominates B, as $F_A = x \leq 1 = F_B \; \forall x \in [0,1]$ and $F_A = x < 1 = F_B \; \forall x \in [0,1)$.

- Second-order stochastic dominance.
  * Second order stochastic dominance can be phrased in terms of mean preserving spreads, i.e. a gamble A is second-order stochastic dominated by another gamble B if A is a mean-preserving spread of B . Also: Imagine comparing an asset x, uniformly distributed from $[-1,1]$ with another asset y which is uniformly distributed $[-2,2]$. These obviously have the same mean, while y has higher variance (it's spread over a wider range). So x second order stochastically dominates y; this constitutes more uncertainty, because there is a wider range of possible outcomes (the probability mass is more spread out). A risk averse person prefers the second order stochastic dominant gamble because risk aversion essentially implies decreasing returns; the new probability of getting a value in $[1,2]$ is exactly the same probability as getting a value in $[-2,-1]$, but with decreasing returns, the $[1,2]$ part isn't valued as highly.

- What do decreasing returns imply about dynamics?
- Question: How many of you having savings? Why?
- How much should you save if your expected income in the future is higher, and the discount factor is $\delta = 1$ (assuming no risk)?
- Problem: Imagine you are going to make \$20 thousand for the next five years, and \$100 thousand every year following that for the next 45 years, after which "the game ends". Your utility function function is concave (i.e. exhibits decreasing returns). Borrowing from future you is costless, there is no inflation, and zero interest rate on investment (capital yields zero returns), and no uncertainty about outcomes.. How much should you consume this year? How much should you save/borrow this year? **Ans:** $\frac{20*5+45*100}{50} = \frac{4600}{50} = 92$. Which implies you should consume \$92 thousand dollars (more precisely, should consume the goods that \$92 thousand dollars can purchase... probably don't eat the money), which means that if your current income is \$20,000, you should be borrowing *\$72,000 a year*.
- With decreasing returns, savings can enable *consumption smoothing*, or insuring against risk.
- Thus, borrowing (for consumption) and saving (for later consumption) and the buying of insurance can all be explained by decreasing marginal utility.
- Given all this, should you be saving?
- "A PhD student is someone who forgoes current income in order to forgo future income." (From Shit Academics Say)

3.5.1 Miscellaneous stuff

- When should we have discounting?
- Do preferences change over time? (Usually model things such that preferences are the only thing that *don't* change)

3.6 Optimization

- **Review Appendix 5.7-5.10 before continuing**
- At this point, we have considered optimization in a *decision theoretic* context. But many of the same tools will be useful when we move to a *strategic* context, i.e. when we actually start talking about game theory.

- Decision theory means there is one actor making a choice, while game theory looks at cases where choices are interdependent: my evaluation of different options depends on choices that you make.
- Note that *most* of economic theory is formal, but a small minority of it is *strategic*. Indeed, economists often use clever modeling choices (e.g. monopolistic competition) to try to avoid the strategic dimensions generated by firm competition, whatnot.
- Political science is so dominated by game theory that it sometimes ends up being used interchangably with all formal theory (as with this course's title).
- Most of the math on this we'll cover is in the appendix, but let's briefly also discuss **constrained optimization**.
- Constrained optimization, as the name implies, is about making choices when one of your choice variables cannot vary freely.
- This could take the form of inequality constraints (e.g. $x < 10$) or equality constraints ($x + y = 10$).
- In practice, we can sometimes structure the decision problem to turn inequality constraints, which are a little trickier to deal with, into equality constraints.
- To see this, consider a classic problem from consumer theory in economics: choosing how much to purchase of various goods, subject to a budget constraint. i.e.
  for two goods x, y:
$$\max_{x,y} u(x, y) \; s.t. \; p_1 x + p_2 y \leq B$$
- If preferences are monotonic in x, y then this *has* to be satisfied with equality, i.e. the consumer must expend their budget.
- Intuition of this is pretty clear: if you like more of everything (and assuming savings has no value), you're not going to not use all your money.
- Techniques of constrained optimization can get a little complicated, but if we have an *equality* constraint there is a straightforward procedure that works and doesn't require us to learn any more calculus: substitute in for one of the variables.
- So if $p_1 x + p_2 y = B$, then we have $y = \frac{B - p_1 x}{p_2}$. Which transforms $u(x, y)$ from a two-variable optimization problem to a *one variable* optimization problem:
$$\max_{x} u\left(x; y = \frac{B - p_1 x}{p_2}\right)$$

3.7 Comparative Statics

- Having established how to find the optima of decision theoretic problems, we may be further interested in how those decisions vary as you change certain parameters.
- For instance, how do I change my purchasing choices as the price of a good increases? How does my vote change as a candidate moves further away from me ideologically?
- This is the subject of the analysis of **comparative statics**.
- They are "statics" because they represent "stable" equilibrium values, and "comparative" because they compare equilibrium outcomes with different parameters.
- The simplest way of finding comparative statics: solve for the optimal x^* and then take a partial derivative w.r.t. the parameter.

- Example: Let's consider a **Cobb-Douglas** utility function, $u(x, y) = x^\alpha y^\beta$. Now to make things easier, let's take a natural logarithmic transformation of this function (which, if you'll recall, should give us a new utility function that represents the same preferences, so: $v(x, y) = ln(u(x, y)) = \alpha ln(x) + \beta ln(y)$.

  Let's consider a consumer choice problem where someone is choosing how much of x and y to buy, subject to their budget constraint W, where p_x, p_y are the prices of x and y respectively. Thus, the budget constraint is:

  $$p_x x + p_y y \leq W$$

  Which will be satisfied with equality, given that the utility function is strictly monotonic. Thus, we can "build in" the constraint into the problem by substituting:

  $$y = \frac{W - p_x x}{p_y}$$

  Giving us:

  $$v(x) = \alpha ln(x) + \beta ln\left(\frac{W - p_x x}{p_y}\right)$$

  We solve this in the usual way, by taking a derivative and setting it equal to zero:

  $$\begin{aligned}
  \frac{\partial v(x)}{\partial x} &= \frac{\alpha}{x} + \frac{\beta}{\left(\frac{W - p_x x}{p_y}\right)}\left(\frac{-p_x}{p_y}\right) = 0 \\
  &= \frac{\alpha}{x} - \frac{\beta p_x}{W - p_x x} = 0 \\
  \leftrightarrow \alpha(W - p_x x) &= \beta p_x x \\
  \leftrightarrow x(\beta p_x + p_x \alpha) &= W\alpha \\
  \leftrightarrow x^* &= \frac{W\alpha}{p_x(\alpha + \beta)}
  \end{aligned}$$

  Which if we substitute into the expression for y gives us:

  $$y^* = \frac{W - p_x x^*}{p_y} = \frac{W}{p_y}\left(\frac{\beta}{\alpha + \beta}\right)$$

  This has the nice interpretation that you spend fixed shares of your budget on x and y depending on how fast your utility "goes up" in each. Once we have these **explicit** expressions for the optimal choices of x and y, we can derive comparative statics very simply: just take partial derivatives of these expressions with respect to the parameter we are interested in finding comparative statics. So, for instance:

  $$\frac{\partial y^*}{\partial p_y} = \frac{W}{p_y}\left(\frac{\beta}{\alpha + \beta}\right)(-1)\frac{1}{p_y^2} < 0$$

  So, since this is negative, we know that the amount consumed of y is *decreasing* in the price of y (which, of course, makes intuitive sense. So much so that economists call this the "law of demand").

  We can also evaluate what happens with an increase in budget:

  $$\frac{\partial y^*}{\partial W} = \frac{1}{p_y}\left(\frac{\beta}{\alpha + \beta}\right) > 0$$

  Since this is greater than zero, we know an increase in budget leads one to consume more of y. So y is a normal good in this case. Thus, we have directly derived some comparative statics for these choice variables.

- We may also be interested in how the objective function (in this case, the utility function) changes as these parameters change. For instance, does your utility increase as your budget increases?
- The most natural way to find this is to substitute in x^* and y^* in your expression for the utility function, then take a derivative with respect to the parameter you want. A somewhat simpler way to find the answer to this is to invoke something called **envelope theorem**.
- Envelope theorem states that if you take a derivative of the original objective function with respect to the parameter, and then *evaluate* that derivative at the optimal values, this will give you an expression for how the objective function varies with the parameter.
- So, for instance:

$$\frac{\partial v(x)}{\partial W} = \frac{\beta}{\frac{W-p_x x^*}{p_y}} \frac{1}{p_y} = \frac{\beta}{W - p_x x^*} = \frac{\beta}{W - p_x \frac{W\alpha}{p_x(\alpha+\beta)}} = \frac{\beta}{W\left(\frac{\alpha+\beta-\alpha}{\alpha+\beta}\right)} = \frac{\beta}{\frac{W\beta}{\alpha+\beta}} = \frac{\alpha+\beta}{W} > 0$$

  So this is positive! Which is what we'd expect; your utility goes up as your budget increases.
- This approach relies on solving **explicitly** for the solutions, i.e. finding the optimal values and then using those to find comparative statics.
- However, in many cases we **may not be able to solve explicitly for the solution**. This is especially true in cases where we have decided to favor generality in our specification of the model, i.e. have tried not to bake in too many assumptions about structure into preferences, outcomes, etc.
- Under these circumstances, the standard approach is to use something called "implicit function theorem". Given time constraints and the math required, we won't be talking about this any further, but suffice to say if you wanted to pursue this kind of work you should probably learn it.
- In recent years (the last 15 years or so) a new approach has been becoming popular: monotone comparative statics.[1]
- While the details behind how monotone comparative statics works can get a little complicated, the intuition is actually very straightforward.
- It relies on assumptions about *complementarity* between variables.
- In fact, complementarity is a common theme in formal theory: oftentimes models reveal that certain choices or parameters are complements in ways that might otherwise not be obvious.
- Complementarity here means: an increase in one variable increases the value of another.
- If the function is fully differentiable[2], we can represent the complementarity of x and y as:

$$\frac{\partial^2 u}{\partial x \partial y} > 0$$

- Substitutability is the opposite of complementarity.
- If all the choice variables are complementary, then we know they all "move the same direction" in some sense. This makes it possible to find clear "monotone" comparative statics in a way that would be difficult if they did not.

[1]See Milgrom and Shannon 1994 and Edlin and Shannon 1998, or in political science, Ashworth and Bueno de Mesquita 2006

[2]One of the big advantages of this approach is that we don't *need* to assume full differentiability, but to avoid a fair degree of complication we will assume this in this course.

- There is a natural analogy here to linear regression and omitted variable bias. Without knowing the full covariance structure of the other variables, it is difficult to determine the impact of leaving a variable out of the statistical model. It will depend on the magnitude of different effects.
- From the world of economic theory: consider a case where there are two goods that are substitutes, say, twinkies and plain scones.
- The price of sugar has a **direct** effect on both of these goods, because they both use sugar as an input. So as the price of sugar decreases, you might expect demand for both goods to go up.
- However! Twinkies are obviously *much* more dependent on sugar. So it will effect the price of twinkies more.
- Meanwhile, because twinkies and scones are substitutes, the impact on demand for scones will depend on the comparative magnitude of two effects: the **direct effect** on scones, as the price of scones falls, and the **indirect effect** of substitution between scones and twinkies.
- Another example: income effects from consumer theory. If two goods are complements, then an increase in income will increase demand for both. If two goods are substitutes, who knows? **Complementarity** is a key concept in formal theory.
- "Moving all in the same direction" here is represented by a condition known as **supermodularity**.
- For instance, if we have a three variable function $u(x, y, z)$, the simplest supermodularity condition would be represented by:
$$\frac{\partial^2 u}{\partial x \partial y} \geq 0, \frac{\partial^2 u}{\partial x \partial z} \geq 0, \frac{\partial^2 u}{\partial y \partial z} \geq 0$$
If the above holds, we would have that $u(\cdot)$ is supermodular in (x, y, z). We also *may* be able to establish a supermodularity condition with a certain degree of substitutability between some subset of the variables, so long as we can establish a certain consistency in the substitutability.
- What I mean is: can we perform a change in variables such that we end up with a supermodularity condition? Specifically, changing a variable for its negative.
- So we can get supermodularity with the following:
$$\frac{\partial^2 u}{\partial x \partial y} > 0, \frac{\partial^2 u}{\partial x \partial z} < 0, \frac{\partial^2 u}{\partial y \partial z} < 0$$
- As this function is now supermodular in $(x, y, -z)$. But it only works if *both* x and y are substitutes with z.
- Once we've established supermodularity in the choice variables, we then can look to see the impact of a **parameter** on the equilibrium quantities.
- Parameters don't need to be supermodular with each other (they are not choices, so there can't be indirect effects between them), but they do need to impact all the choice variables in the same direction.
- Otherwise we end up in the same world of "direct versus indirect effects".
- So, to use a consumer theory example like what we were talking about before, consider a case where one is deciding how many pens (f) or mechanical pencils (c) they want. These are substitutes, and there is a parameter p_I, the price of ink, makes pens less attractive as it goes up, but has no effect on pencils. So we have a utility function $u(e, c)$ with the following properties:
$$\frac{\partial^2 u}{\partial f \partial c} < 0, \frac{\partial^2 u}{\partial f \partial p_I} < 0, \frac{\partial^2 u}{\partial c \partial p_I} = 0$$

- Here, note that $\frac{\partial^2 u}{\partial c \partial p_I} \geq 0$, since it's equal to zero. So we have supermodulairty in $(-e, c)$, and the parameter p_I then effects both choice variables (i.e. $-f$ and c) in the same direction (weakly). So we can establish that an increase in p_I increase both $-f$ and c; i.e. it decreases the number of pens bought and increases the number of pencils bought.

- Now, alternatively, imagine we considered the price of plastic p_L. Assume that both pens and mechanical pencils are made with plastic, so both are made less attractive with an increase in the price of plastic. So now we have:
$$\frac{\partial^2 u}{\partial\text{-}f \partial p_L} > 0, \frac{\partial^2 u}{\partial c \partial p_L} < 0$$
And thus we can't sign the comparative static.

- Now consider if we added another thing you could buy: erasers e. These are clearly complementary with pencils, and have no relationship with pens. So we have:
$$\frac{\partial^2 u}{\partial f \partial c} < 0, \frac{\partial^2 u}{\partial c \partial e} > 0, \frac{\partial^2 u}{\partial f \partial e} = 0$$

- This allows us to establish supermodularity in $(-f, c, e)$! So we can still derive monotone comparative statics, if a parameter effects everything in the same direction.

- So let's modify things: now the pen is made out of metal (it's a fancy pen). And let's consider the price of petroleum p_M, since both plastics (and thus mechanical pencils) and erasers are petroleum-based. So now, in addition to supermodularity of $(-f, c, e)$ we have:
$$\frac{\partial^2 u}{\partial\text{-}f \partial p_M} = 0, \frac{\partial^2 u}{\partial c \partial p_M} < 0, \frac{\partial^2 u}{\partial e \partial p_M} < 0$$

- So the price of petroleum affects all these parameters in the same direction (negative), allowing us to establish monotone comparative statics. Specifically, an increase in the price of petroleum increases the number of pens purchased, decreases the number of pencils purchased, and decreases the number of erasers purchased. Formally, by invoking monotone comparative statics we can derive:
$$\frac{\partial f^*}{\partial p_M} > 0, \frac{\partial c^*}{\partial p_M} < 0, \frac{\partial e^*}{\partial p_M} < 0$$

- And we've done all this without saying *anything else* about the functional forms here; no assumptions about strict concavity, etc. A few cross-partial derivatives and we're done.

4 Game Theory

4.1 Preliminaries

- The content up until this point in the course has focused on **decision theory**, i.e. the study of how a single actor/agent would make decisions in isolation, when their choices do not affect the utility functions of other agents/actors.

- We now move to **game theory** (the somewhat erroneous title of the course), which is the study of **interdependent** decision-making, or **strategic behavior**.

- These are the mathematical techniques that are most commonly used by formal theorists to study situations in which utility functions are interdependent, i.e. my best choice depends on decisions you make, and vice versa.

- Essentially all of the mathematical and substantive content that we've learned up until now about decision theory carries over into this topic - e.g. optimization in game theory still often involves taking derivatives and setting them equal to zero! Assumptions on utility functions still mean the same kinds of things (e.g. concave utility functions implying risk aversion).
- But now, we see what happens when we build upon this to consider more complicated, interdependent situations.
- Notably, because the tools required here are more complicated, and are less likely to produce unique solutions, economists will often go through great lengths to develop decision theoretic representations of strategic situations (see, for instance, monopolistic competition models instead of models of oligopoly) to avoid having to use game theory. So it is probably the case that *most* formal theory is *not* game theoretic.
- However, most formal theory work in political science is game theoretic, as politics is pretty inherently strategic. Hence the disproportionate focus in this course.
- **Example game:** 2/3 of the average. See uploaded boardwork or textbook for more discussion.

4.2 Static games of complete information

- The first class of games we will discuss are **static** games of **complete information**.
- Static:
  - Simultaneous choices. This doesn't mean that choices are made *literally* simultaneously, but that no player observes another player's choice before making their decision, and thus they cannot **condition** on that choice.
  - Payoffs are determined as a function of the outcomes, which are determined by strategy profiles.
- Complete information. Must know:
  - All possible actions of all players
  - All possible outcomes
  - How strategy profiles $\rightarrow$ outcomes
  - Preferences of **each** player over outcomes.
- We also sometimes refer to things that are "common knowledge". Some statement E is common knowledge iff:
  - Everyone knows E.
  - Everyone knows that everyone knows E
  - ...continue ad infinitum.
- **Actions** are the things that each player can do. **Strategies** are "plans of action" for each player, that say what they will do in every circumstance where they get a chance to make a choice.
- In static games of complete information, **pure strategies** (i.e. strategies that are deterministic) are trivially equivalent to actions, since they involve simply choosing one action. But as we talk about **mixed strategies** or talking about **extensive-form** games, there will be more clear differences between strategies and actions.
- Each player's set of strategies is denoted S_i for every player i in the set of all players N.

- We can thus define the strategy space as:

$$S_1 \times S_2 \times \dots S_n \equiv S$$

Any element $x \in S$ is called a **strategy profile**, and it entails a strategy choice for each player of the game.

- Payoffs to each player are utility functions u_i, where:

$$u_i : S \to \mathbb{R}$$

- So they map elements of the strategy space (strategy profiles) to real numbers.
- In words: they give a payoff for everyone possible strategy profile, where in the case of pure strategy profiles, we get one outcome for everyone strategy profile. So pretty straightforward.
- A normal form game G is thus defined by the triple:

$$G = \langle N, S, \{u_i\}_{i=1}^n \rangle$$

- Where N is the set of players, S is the strategy space, and $\{u_i\}_{i=1}^n$ is the utility function for each player.
- A *matrix representation* of a normal form game is what many of you may be most familiar with if you've seen game theory in some introductory class, or some other substantive class.
- The standard matrix form can be used to represent any two-player finite game. More complicated games (for instance, those with continuous strategy spaces) will not usually be representable in this way.
- So, for instance, let's consider the classic Prisoner's dilemma, represented in common matrix form:

	C	D
C	$-3, -3$	$-5, 0$
D	$0, -5$	$-4, -4$

- We could formally write the normal form of this game as follows:

$$N = \{1, 2\}$$

$$S_1 = \{C, D\}, S_2 = \{C, D\}, S = S_1 \times S_2 = \{CC, CD, DC, DD\}$$

$$u_1(CC) = -3, u_1(DC) = 0, u_1(CD) = -5, u_1(DD) = -1, u_2(CC) = -3, u_2(CD) = 0, u_2(DC) = -5, u_2(DD) = -4$$

- So as we can see, the matrix representation is a more compact and easy to read representation of this normal form game, though in many cases matrix representations will not be possible.

4.3 The Concept of the Solution Concept

- Any strategy profile is a potential candidate for an equilibrium.
- Solution concepts are ways of restricting our attention to a set of strategy profiles in a systematic fashion. They are a "filter" defined by a set of rules.
- We can contrast this to the ad hoc removal of strategy profiles that we think are "unreasonable". Being systematic allows us to be clear and transparent about why we're removing certain strategies.
- Each solution concept has differing advantages and weaknesses. We can consider evaluating them based on such criteria as:

  - Uniqueness. Does it give us a precise prediction about behavior, or does it only narrow things down to a set of plausible options?
  - Existence. Nash's dissertation's "existence proof" of mixed strategy Nash equilibria in all finite games is one big reason why Nash equilibrium has become so common-place; it has high levels of *generality* because we know that a wide class of games have at least one Nash equilibrium.
  - Inherent plausibility/desirability of the rules? Some solution concepts are attractive because they mimic features of behavior that we expect to be prominent or salient.

- **Rationalizability** is a particularly weak solution concept (weaker than Nash. Every Nash is rationalizable but not vice versa). Involves iterated elimination of strictly dominated strategies.
- What is strict domination? Why would such a strategy not be "rationalizable"?
- If there is no world where making that choice would not make that player strictly worse off, then a rational person would never pick it.
- We can be more precise about strict domination. If a strategy profile S^i is an element of the strategy space S, we can denote S_{-i} (read: "not i") as:

$$S_{-i} \equiv S_1 \times S_2 \times S_3 \times \ldots S_{i-1} \times S_{i+1} \times \ldots S_N$$

- In words: we take the strategy space, but remove player i.
- In general, we with use $-i$ to refer to "everyone but i"
- Consider two strategies for player i, s_i and s_i'. s_i strictly dominates s_i' iff:

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

- In words: a strategy strict dominates another if, no matter what choices everyone else makes, it's still better.
- *Iterated elimination of strictly dominated strategies* (**IESDS**) follows directly from *common knowledge of rationality.* How so?
  1. Rational players will never play a strictly dominated strategy.
  2. But if *I know* you're rational, I know you won't play a strictly dominated strategy. So I can eliminate these, and then consider what is strictly dominated in what's left.
  3. But if *you know I know* you're rational, you know that I know you won't play a strictly dominated strategy, and have thus eliminated those in making a decision. So you can eliminate from consideration any strategies that would be strictly dominated for me once I eliminate from consideration any strategies that are strictly dominated for you.
  4. ...ad infinitum.
- We will say that the strategy profiles that survive this process of iterated elimination are **rationalizable**, which is our first solution concept.
- Note: Tadelis will tell you that these are different, and defines rationalizability using a "never best response" concept. Suffice to say that as long as we have mixed strategies (we'll talk about these in a bit) these approaches are virtually always identical, to the point where some other textbook treatments don't get into the weeds of what diferentiates them. In keeping with that, we will not differentiate between IESDS and rationalizability.
- Because all Nash are rationalizable, you can also use iterated elimination of strictly dominated strategies to simplify the process of looking for Nash equilibrium.

- So, for instance, in the Prisoner's Dilemma, the unique Nash equilibrium is also the unique rationalizable strategy profile! Note from below that D dominates C for each player:

	C	D
C	$-3,-3$	$-5,0$
D	$0,-5$	$-4,-4$

- An important point of confusion to avoid: keep in mind that the equilibrium here is DD **not** $(-4,-4)$. This is because $(-4,-4)$ is the **payoff** to the equilibrium, but an equilibrium is a characterization of each player's behavior. So for instance, you could have two strategy profiles that produce the same payoffs, but represent different behavior. As an example, consider a game where two cars driving in opposite directions are choosing which side of the road to drive on. We could represent the payoffs to this game as:

	Left	Right
Left	$0,0$	$-1000,-1000$
Right	$-1000,-1000$	$0,0$

- Obviously, both drives are way worse off if they drive on different sides of the road (since then they crash...). But the equilibria here -$(Left, Left)$ and $(Right, Right)$ - represent different behavior (e.g. the difference between the US and the UK).

- As a general matter, Nash equilibrium tends to be the starting point for a lot of analysis.

- Then if Nash doesn't get us all the way to where we want, we can impose further structure (subgame perfection, Markov perfection, etc.) to restrict our interest to a subset of Nash equilibria.

- Example:

	A	B	C
A	-1,-1	1,2	0,0
B	-1,-1	0,0	2,1
C	-2,-2	-2,-2	-2,-2

- Note that no pure strategy initially strictly dominates for player 2. However, for player 1, B strictly dominates C, so we can restrict our attention to the game where we eliminate that strategy:

	A	B	C
A	-1,-1	1,2	0,0
B	-1,-1	0,0	2,1

- Now if player 2 knows that player 1 is rational, they know that player 1 will not choose C. Given this, player 2 looks at the restricted game above, and eliminates the strictly dominated strategy A, to get the following reduced game. Note: A strategy that becomes strictly dominated in an iterated process should not be described as dominated in the original game unless it was dominated before the iteration took place.

	B	C
A	1,2	0,0
B	0,0	2,1

- No more strategies are strictly dominated, so the remaining strategy profiles $S = \{(A,B),(A,C),(B,B),(B,C)\}$ are the set of *rationalizable* strategy profiles.

- Jumping a bit ahead, it's clear that not all four of these are Nash equilibria. Nash equilibria is a filter that gets us to $S' = \{(A,B),(B,C)\}$.

- Jumping even further ahead, imagine player 1 gets to play first. We have the same Nash equilibria, but if we apply the more restrictive filter of Subgame Perfect Nash Equilibrium, we are left with only $S'' = \{(B,C)\}$.

4.3.1 Quick side note: Weak Dominance

4.3.2 Quick side note: Weak Dominance

- It is important for the solution of rationalizability that we deal with *strict* dominance, not *weak* dominance. Weak dominance has the same definition, but with a weak inequality, so s_i weakly dominates s_i' iff:

$$u_i(s_i, s_{-i}) \geq u_i(s_i', s_{-i}) \; \forall s_{-i} \in S_{-i}$$

- If we eliminate all weakly dominated strategies, we will often remove Nash equilibria. So this would not be a weaker solution concept than NE, and thus doesn't come up as much in formal theory discussions compared to IESDS.

- For instance, in the game:

	A	B
A	$1,1$	$0,0$
B	$0,0$	$0,0$

- A weakly dominates B for both players, but (B, B) is also a Nash equlibrium (as we will discuss more later).

4.4 Evaluating Outcomes/Normative Theory

- While this will not be a huge focus of this course, we may be interested in the question of how we would evaluate outcomes, rather than simply predict them.

- This is the domain of **normative theory** (as contrasted with "positive" theory).

- There are lots of different ways we could do this, and many different ethical models of distributional effects, etc.

- John Roemer at Yale has done interesting work on this using formal theory.

- Some possible normative models:
  - Utilitarianism/aggregate utility
  - Equality of opportunity
  - "Veil of ignorance"/Rawlsian.
  - Rights-based frameworks.

- Political economists usually focus on positive rather than normative models, but not exclusively, and sometimes normative considerations end up in the mix either explicitly or implicitly.

- To the extent that formal theorists deal with welfare effects, it usually doesn't go much further than utilitarianism, because even aggregate utility calculations involve a lot of complications.

- Consider: as we moved from utility to expected utility, we had to abandon simple ordinal utility functions for "cardinal" utility functions in which the numbers represented intensity of preferences of different bundles. So now, we had to specify if, say, an additional apple was worth two additional oranges.

- This gets even more complicated with aggregate utility calculations, since now we need to make *interpersonal utility comparisons*. I.e. is an apple to you worth more than an apple to someone else? How do you like them apples?[3]

[3]Sorry.

- To avoid these complications, formal theorists often focus on especially weak ethical criteria, of which the most common is **Pareto optimality**.
- Something is **Pareto optimal** if you cannot make anyone else better off without making someone worse off.
- This leaves in place many prima facie undesirable outcomes: e.g. one person having all the wealth while everyone else has none.
- However, if an outcome is Pareto-dominated, i.e. if there exists an outcome where at least one person is better off without making anyone worse-off, it may be relatively uncontroversial to conclude that this Pareto-dominated outcome is undesirable relative to the outcome that Pareto-dominates it.
- Pareto optimality comes up all the time in formal theory; consider for instance, the Prisoner's Dilemma, in which one of the key features in that the only rationalizable outcome is Pareto-dominated by cooperation; thus, a considerable amount of attention has been directed towards determining when the "bad" uncooperative outcome can be avoided (see Axelrod's Evolution of Cooperation, for instance).

4.5 Best Response Correspondences

- A **correspondence** is like a function, but it produces a **set** of outcomes rather than a single outcome.
- So you can think of it as a set-valued function.
- So, for instance, while a utility function maps elements of the strategy space to real numbers representing the utility from that outcome, i.e.:

$$u : S \to \mathbb{R}$$

- A **best response correspondence** maps the strategy choices of others (i.e. $x \in S_{-i}$) to **all** the optimal choices of player i. If there is only one, then this is a best response function, but there may be more than one. So:

$$BR : S_{-i} \to x \subseteq S_i$$

- Where $s_i \in BR(s_{-i})$ iff:

$$u_i(s_i, s_{-i}) \geq u(s_i', s_{-i}) \; \forall s_i' \in S_i$$

4.6 Nash Equilibrium

- We can think of actors as having **beliefs** about what other players are going to play s_{-i}.
- Rationalizability as a solution concept involves strategy profiles in which the strategy for each player but be a best response to "some" belief.
- Nash equilibrium can be thought of as requiring that these beliefs be "correct"; i.e. they must be consistent with the strategy profile in question.
- So, for instance, with this classic coordination game:

	O	F
O	$2, 1$	$0, 0$
F	$0, 0$	$1, 2$

- OF and FO are both rationalizable, because they are best responses to a player thinking that the other player is going to choose O or F and being wrong. But only OO and FF are Nash equilibria, because they are best responses to the correct belief about what the other player is doing.

- Indeed, NE can be defined as *mutual best responses.* Formally, a few equivalent definitions of NE are: $s^* \in S$ is a NE iff:

  1. $u_i(s_i^*, s_{-i}^*) \geq u_i(s_i', s_{-i}^*)\ \forall s_i' \in S_{-i}, \forall i \in N$
  2. $s_i^* \in BR_i(s_{-i}^*) \forall i \in N$
  3. Define the vector of best response correspondences $BR_1, BR_2, \ldots, BR_n = BR$. Then $s^* \in BR(s^*)$, i.e. s^* is a **fixed point** of this vector correspondence.

- Often people think of this in terms of the outcome of each player being "rational". This can be misleading in many circumstances.

- Consider, for example:

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

- Reasonably, one may say that (R, R) seems like something a "rational" person would deviate from, because they can't do worse off by choosing L. However, Nash equilibrium tells us nothing about this.

- We could ad hoc say "well, (R,R) doesn't make sense, so I won't consider it", but this is pretty loose. More restrictive solution concepts are a way to make these kinds of restrictions more systematic.

- For instance, "trembling hand" Nash equilibrium. In the presence of "trembles", one would want to choose L instead of R.

- Generally, a better way to think of Nash equilibria is that they are in some sense "stable".

- For instance, in the US, when people pass each other on the sidewalk, they generally pass on the right. This is an equilibrium, but it's not the outcome of each player being individually rational.

- The core concept is: stability is generated when **no individual has an incentive to unilaterally deviate**.

- Otherwise, if someone knows they can do better without having to coordinate with anyone else, they will find their way to do so, either just through trial and error or reasoning.

- How do we go about finding pure strategy NE? Well, in the case of finite two player games (i.e. those that are representable by matrices), we can just mechanically go through and put stars next to the best responses for each player, and then see which strategy profiles have two stars next to them, indicating mutual best response.

	A	B	C
A	$7,7$	$2,2$	$1,8^*$
B	$8^*,4$	$3^*,5^*$	$2^*,3$
C	$8^*,1^*$	$3^*,1^*$	$0,0$

- So, NE are $\{CA, CB, BB\}$

- Not every finite game has a *pure strategy* NE. For instance, Rock Paper Scissors can be represented as:

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

- So as you can see, there are no double-starred strategy profiles. Mixed strategies (which will talk about later) will be important for getting us to existence of NE in a wider class of games.

- Indeed: existence proofs are essential to why NE is such a backbone of game theoretic analysis, because it ensures the solution concept has generality.
- So finding pure strategy NE is fairly straightforward for two player finite games that can be resprented by matrices. What about more general strategy spaces? Non-finite games? More than two players?
- Sometimes you can't write a matrix representation of a game, so you need to think through possible strategy profiles as candidates, and then determine whether or not there's an incentive to deviate.
- **Example:** k-threshold public goods game. Public good is provided in k people contribute. Benefit exceeds costs of contributing ($B > c$). Equilibria are: (1) exactly k people contribute; (2) no-one contributes. If you think through the strategy profiles, these are the only ones that don't induce at least one person to want to change their strategy.
- Median voter models, similarly: both politicians located at the median, because otherwise they have an incentive to move *just* closer than their opponent towards the median to capture a clear majority of votes. Only stable outcome is at the median.
- Non-finite strategy spaces? The same logic of mutual best response applies, but we end up using tools of calculus to determine best response functions, and then determine where these best response functions intersect.
- Finding the intersection means: solving the system of equations.
- Consider the following example question:

Two allied countries $i \in \{1, 2\}$ are deciding how much to invest in their military. Because they are allies, one country's military improves the security of the other, but not by as much as that country's own military, as alliance reliability is not 100%. (Leeds et al. 2000) For military investments $f_i, f_j \in \mathbb{R}^+$, we have:

$$u_i(f_1, f_2) = ln\left(f_i + \frac{1}{2}f_j\right) - cf_i$$

(a) What is the strategy set for each player?
Ans: $f_i \in \mathbb{R}^+$

(b) What is the strategy space (i.e. space of strategy profiles) for the game?
Ans: $f_1 \times f_2 \in \mathbb{R}^+ \times \mathbb{R}^+$

(c) Find the best response functions for each player as a function of f_{-i}.
Ans:

$$\begin{aligned}\frac{\partial u_1}{\partial f_1} &= \frac{1}{f_1 + 0.5f_2} - c = 0 \\ &\leftrightarrow c(f_1 + 0.5f_2) = 1 \\ &\leftrightarrow f_1 = \frac{1}{c} - 0.5f_2\end{aligned}$$

Then, since the utility functions are symmetric, we have:

$$\leftrightarrow f_2 = \frac{1}{c} - 0.5f_1$$

(d) Find the Nash Equilibrium of this game.
Ans:

$$f_1 = \frac{1}{c} - 0.5\left(\frac{1}{c} - 0.5f_1\right)$$
$$f_1 = \frac{1}{c} - 0.5\frac{1}{c} + 0.25f_1$$
$$\leftrightarrow 0.75f_1 = 0.5\frac{1}{c}$$
$$\leftrightarrow f_1^* = \frac{2}{3c}$$

And once again by symmetry

$$\leftrightarrow f_2^* = \frac{2}{3c}$$

(e) Comparative statics: how will the equilibrium f_1^* change with a change in c (the cost of investments into military capacity)? **Ans:**

$$\frac{\partial f_1^*}{\partial c} = -\frac{2}{3c^2}$$

This is negative, so the equilibrium f_1^* declines as c increases. This is pretty intuitive.

(f) How will country 1's utility change with a change in c? (hint: remember envelope theorem)
Ans: Applying envelope theorem:

$$\frac{\partial u_1^*}{\partial c} = \frac{\partial u_1}{\partial c}(f_1^*, f_2^*) = -f_1^* = -\frac{2}{3c}$$

4.7 Mixed Strategies

- What is the strategy space? Convex hull or simplex of pure strategies.
- For each player, we denote this ΔS_i
- Strategy space is continuous because mixed strategies smooth things out. Convexifies the choice set.
- In future, may hear of "feasible set" of payoffs. These are the technically feasible payoffs achievable from a combination of strategies.
- In this part of the course, we will focus on mixed strategies over finite pure strategy spaces.
- Mixed strategies over continuous strategy spaces require integrals, as they involve probability distributions of non-finite spaces. Since we're mostly avoiding those (at least for now), we'll avoid these.
- Mixed strategies are the simplex on pure strategies because they are all possible probability distributions over the pure strategies.
- Formally, a mixed strategy σ_i on a finite pure strategy space is defined by the following properties:
  1. $\sigma_i(s_i) \geq 0, \forall s_i \in S_i$
  2. $\sum_{s_i \in S_i} \sigma_i(s_i) = 1$
- In other words, a mixed strategy gives a positive probability or zero to each pure strategy, and the probabilities of all strategies must sum to one.

- In this sense, the pure strategy space is now the sample space, and clearly the sum of the probabilities of all elements in the sample space must equal zero, since *something* must happen by definition.
- A pure strategy s_i' is described as "in the support" of σ_i if it has a strictly positive probability, i.e. $\sigma_i(s_i') > 0$
- *Beliefs*, i.e. each player's understanding of what the other player will do, are now *probability distributions* over pure strategies.
- To deal with the concept of best responses, choices, etc. in this context, we need to rely on our standard tool for dealing with uncertainty: **von-Neumann Morgenstern expected utility functions**.
- In this case, the uncertainty is generated by the randomness in each player's strategy.
- The expected utility to any player i of a pure strategy, given that the other players are playing mixed strategies, is defined as:

$$u_i(s_i, \sigma_{-1}) = \sum_{s_{-i} \in S_{-i}} \sigma_{-i}(s_i) u_i(s_i, s_{-i}) = \sum_{s_{-i} \in S_{-i}} \underbrace{\sigma_{-i}(s_i)}_{\text{(Prob. of strat.)}} \underbrace{u_i(s_i, s_{-i})}_{\text{(payoff of strat. profile)}}$$

- To emphasize, this is for when player i is *not* randomizing, but is a useful step towards defining their payoff where both randomize.
- If player i is playing a *mixed* strategy, there are now two nested sources of randomness, so we get the kind of compound lottery structure we talked about earlier in the course. Specifically:

$$u_i(\sigma_i, \sigma_{-i}) = \sum_{s_i \in S_i} \sigma_i(s_i) u_i(s_i, \sigma_{-i})$$

  Subbing in the expression from above into this one, we get:

$$u_i(\sigma_i, \sigma_{-i}) = \sum_{s_i \in S_i} \left(\sum_{s_{-i} \in S_{-i}} \sigma_i(s_i) \sigma_{-i}(s_i) u_i(s_i, s_{-i}) \right)$$

- Note the compound structure: each combination of $\sigma_i(s_i)$ and $\sigma_{-i}(s_{-i})$ gives a particular probability of a pure strategy profile (s_i, s_{-i}), which weights the payoff of that pure strategy profile to player i.
- Having defined these payoffs to mixed strategy profiles, we can thus provide the natural analog to our earlier definition of pure strategy Nash equilibrium, thus generalizing the concept to the world of mixed strategies, i.e σ^* is a NE if:

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(\sigma_i', \sigma_{-i}^*) \; \forall \sigma_i' \in \Delta S_i, \forall i \in N$$

- So, the same idea of mutual best response, but now with mixed strategies and expected utilities.
- So let's try some of this out with an example, the classic game: matching pennies.

	H	T
H	$1, -1$	$-1, 1$
T	$-1, 1$	$1, -1$

- Clearly no pure strategy NE. So we want to figure out mixed strategies.
- Strategy space for each player: $\Delta S_1 = (P_1(H), P_1(T)), \Delta S_2 = (P_2(H), P_2(T))$

- Applying the definition of $u_i(s_i, \sigma_{-i})$ for player 1 we get:

$$u(H, \sigma_2) = u_1(H, H)P_2(H) + u_1(H, T)P_2(T)$$

$$u(T, \sigma_2) = u_1(T, H)P_2(H) + u_1(T, T)P_2(T)$$

- And then using $u(\sigma_i, \sigma_{-i})$ to consider a possible mixed strategy by player 1, we sub in and get:

$$u(\sigma_1, \sigma_2) = P_1(H)(u_1(H, H)P_2(H) + u_1(H, T)P_2(T)) + P_1(T)(u_1(T, H)P_2(H) + u_1(T, T)P_2(T))$$

- At this point, useful to consider an important proposition:
If σ^* is a NE, and s_i and s'_i are in the support of σ^*_i then:

$$u_i(s_i, \sigma^*_{-i}) = u_i(s'_i, \sigma^*_{-i}) = u_i(\sigma^*_i, \sigma^*_{-i})$$

- Consider: a player will *not* randomize if they get a higher expected payoff from any particular strategy. If you have two strategies s_1 and s_2 and your expected utility $EU(s_1) = x > EU(s_2) = y$, and you probability of palyer each in a mixed strategy is p and $1 - p$ respectively, then your utility from any mixed strategy is:

$$px + (1 - p)y$$

which is maximized by choosing the higher possible p, i.e. $p = 1$. So, no mixed strategy.

- So if the expected utility to each pure strategy in a mixed strategy profile has to be the same, then we know that the expected utility of the mixed strategy is just equal to the payoff to any individual pure strategy.

- To see this, consider if the expected utility to each strategy s_i in the support of σ_i is equal to x. Then for probabilities $p_1, p_2, p_3, \ldots, p_n$ for each of the n strategies with support, where $p_1 + p_2 + \ldots p_n = 1$ by definition, we have the expected utility of the mixed strategy as:

$$p_1x + p_2x + p_3x \ldots p_nx = (p_1 + p_2 + \ldots p_n)x = 1x = x$$

- So: returning to matching pennies.

	H	T
H	$1, -1$	$-1, 1$
T	$-1, 1$	$1, -1$

- We can see that each pure strategy has a pure strategy best response, and there are no pure strategy equilibria. And we are armed with this proposition from above, and know that in a mixed strategy each player has to be indifferent between all the strategies with support.

- So, defining $P_2(H) = p$ and $P_2(T) = 1 - p$, we get:

$$u_1(H, \sigma_2) = 1p + (-1)(1 - p) = 2p - 1$$

$$u_1(T, \sigma_2) = (-1)p + (1)(1 - p) = 1 - 2p$$

- And in any NE, $u_1(H, \sigma_2) = u_1(T, \sigma_2) = u_1(\sigma_1, \sigma_2)$, so:

$$2p - 1 = 1 - 2p \leftrightarrow p = \frac{1}{2}$$

- This gives us a mixed strategy for **player 2** of:

$$\sigma_2 = (P_2(H) = 0.5, P_2(T) = 0.5)$$

- Which since the model is symmetric, if we do the same calculations using player 2's utility function, we would get the same mixed strategy for player 1:

$$\sigma_1 = (P_1(H) = 0.5, P_2(T) = 0.5)$$

- This gives us a mixed strategy NE in which the players play σ_1, σ_2 as defined above.
- In each case, the best responses are truly **correspondences**, because they imply an infinite set of best responses to the other player's strategy:

$$BR_1(\sigma_2) = \{P_1(H) \in [0,1]\}$$

$$BR_2(\sigma_1) = \{P_2(H) \in [0,1]\}$$

- But to be a Nash equilibrium, each strategy only has to **a** best response to the other strategy. This makes it *stable* in some sense, but to emphasize again, it's not necessarily the most natural to think of this as arising from just each player playing 3D chess or whatever.
- Another unintuitive thing: each person's strategies end up being determined by the **other** player's utility function. Indeed, they are indifferent between any combination of probabilities, and only end up playing that particular combination because anything else induces changes.
- Consider that in this example, these are the **only** values that don't lead some other player to want to unilaterally deviate. So if $P_2(H) = p = 0.75$ we would get:

$$u_1(H, \sigma_2) = 2p - 1 = 2(0.75) - 1 = 0.5$$

$$u_1(T, \sigma_2) = 1 - 2p = -0.5$$

- So player 1 would now just choose H with certainty instead. In which case σ_2 would no longer be a best response. So this isn't stable.
- Do mixed strategies exist in the real world? Computing a mixed strategy can be complicated, so some have asked under what conditions is it reasonable to believe that people are actually playing mixed strategies.
- Consider again the classical example of rock paper scissors:

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

- If we wanted to compute the mixed strategy Nash equilibrium, we would need to solve a system of three equations:
  Defining $P_2(R) = p, P_2(P) = q, P_2(S) = 1 - p - q$

$$u_1(R, \sigma_2) = 0p - 1q + 1(1 - p - q)$$

$$u_1(P, \sigma_2) = p + 0q - (1 - p - q)$$

$$u_1(S, \sigma_2) = -p + q + 0(1 - p - q)$$

  Need to solve for p, q where $u_1(R, \sigma_2) = u_1(P, \sigma_2) = u_1(S, \sigma_2)$

- This is a little tricky maybe to do, but it's also not that surprising that people might arrive on the equilibrium strategy over time of playing each with equal probability, so $p = q = 1 - p - q = 1/3$.

- Perhaps through some kind of intuition about ways to improve by, say, starting to favor rock against a player who seems to play a lot of scissors, it's not that crazy to imagine moving around until "settling" on the equilibrium, at which point no-one has an incentive to change.

- Plenty of examples from sports, certainly, where people randomize.

- There is also a way of interpreting mixed strategies as being about *beliefs*. In this interpretation, people only need to play in such a fashion where they believe that the other player is playing the mixed strategy as defined by the NE.

- Can a mixed strategy strictly dominate a pure strategy?

- What about in cases where neither pure strategy in the mixture is strictly dominant? Consider the following.

	A	B	C
A	$5,5$	$4,4$	$1,1$
B	$4,4$	$0,0$	$1,1$
C	$0,0$	$4,4$	$1,1$

- Consider for player 2. Neither A or B strictly dominates C. However, imagine playing $\sigma_2 = (Pr(A) = 0.5, Pr(B) = 0.5, Pr(C) = 0)$. Keeping in mind that $EU_2(s_2 = C) = 1$ (i.e. payoff of 1 irrespective of player 1's strategy), $EU_2(s_2 = \sigma_2 | s_1 = A) = (0.5)(5) + (0.5)(4) = 4.5 > 1$, $EU_2(s_2 = \sigma_2 | s_1 = B) = (0.5)(4) + (0.5)(0) = 2 > 1$, $EU_2(s_2 = \sigma_2 | s_1 = A) = (0.5)(0) + (0.5)(4) = 2 > 1$. So σ_2 strictly dominates C for player 2, even though neither A nor B strictly dominates C.

- What about the following?:

	A	B	C
A	$5,5$	$4,4$	$1,1$
B	$10,10$	$0,0$	$1,1$
C	$0,0$	$2,2$	$1,1$

- In this case, $0.5A, 0.5B$ does not strictly dominate. But we can determine a strategy that does systematically in the following fashion:

$$u_2(\sigma_2, s_1 = A) = 5p + 4(1-p) = 4 + p > 1 \; \forall p$$

$$u_2(\sigma_2, s_1 = B) = 10p + 0(1-p) = 10p > 1 \leftrightarrow p > \frac{1}{10}$$

$$u_2(\sigma_2, s_1 = C) = 0p + 2(1-p) = 2 - 2p > 1 \leftrightarrow p < \frac{1}{2}$$

- So all $p \in \left(\frac{1}{10}, \frac{1}{2}\right)$ satisfy all three of these inequalities, insuring that any such strategy σ_2 will strictly dominate C for player 2.

- Strict domination by mixed strategies of pure strategies is absolutely *key* to the general equivalence of IESDS and rationalizability.

- Consider the "2/3 of the average" game we did in class. Are there any strategies that are dominated by a pure strategy? Nope. But are there any strategies that are dominated by a *mixed* strategy? (This is a problem set exercise...)

- Rationalizability is sometimes defined in terms of eliminating strategies that are "never a best response". But in two player games at least, we can use mixed strategies to get to equivalence.

Brief Note on NE Existence and Fixed Points

- I don't want to get too deep into the weeds on this, but given that I've emphasized that NE is used to widely because of its generality, which is supported by these existence proofs that were central to John Nash's dissertation, it's worth taking a moment to sketch some of this.
- As discussed before, a fixed point of a function is something like $f(x) = x$. The output is equal to the input. So, for $f(x) = x$, every point is a fixed point.
- For correspondences, where the output is a set, it only needs to be the case that $x \in f(x)$ for a correspondence $f(\cdot)$.
- Brouwer's fixed point theorem: for a continuous function $f : [0,1] \to [0,1]$, it will achieve a fixed point. **Proof:** Try it! Draw an x and y axis that go from 0 to 1. Now try to draw a curve that stays on that plane without touching the 45 degree line. Do it! I dare you!
- The fixed point theorems that become relevant to proving the existence of mixed strategy of NE are basically more complicated generalizations of htis kind of idea.
- Recall that we can construct a vector correspondence which just combines the best response correspondences for each player:

$$BR(\sigma) = BR_1(\sigma), BR_2(\sigma), BR_3(\sigma), \ldots, BR_n(\sigma)$$

- Where a NE is now just a fixed point of this best response correspondence, i.e. $\sigma \in BR(\sigma)$ implies by definition that everyone is best responding to everyone else's strategy.
- John Nash's equilibrium existence proof is just a simple application of Kakutani's fixed point theorem (as von Neumann apparently told him, snidely).
- Kakutani's fixed point theorem: a correspondence $C : X \rightrightarrows X$ has a fixed point if:
  1. X is non-empty, compact, convex subset of $\mathbb{R}^n$ (this is satisfied for BR with mixed strategies, which convexify the space, and include the boundaries)
  2. $C(x)$ is non-empty for all x (we've shown before that there are always best responses)
  3. $C(x)$ is convex for all x (this also follows from the mathematical setup here)
  4. C has "closed graph" (going a little beyond scope of this part of the course here)
- Without worrying too much about the details, hopefully this gives you some flavor for how with mixed strategies, we can use the mathematical properties of the best response correspondence to prove the existence of a fixed point, which thus proves the existence of a Nash equilibrium in mixed strategies for any game with a finite section of pure strategies.
- Condition (1) is basically the correspondence generalization of the $[0,1] \to [0,1]$ property from Brouwer's.
- Conditions (2),(3),(4) are basically correspondence continuity assumptions, so give us the mathematical properties we want. And they follow from the mixed strategy setup.

4.8 Extensive Form Games

4.8.1 Introduction

- Up until now, we've been dealing with **static** games of **complete** information.
- For this purpose, we defined the "normal form", which defines a game Γ in terms of:

1. A set of players N
2. A strategy space S
3. A set of utility functions u_i for each player $i \in N$, defined on the outcome space

- Now we are interested in bringing **dynamics** into our framework, as well as **incomplete** or **imperfect** information.
- **Dynamics** are when choices are made over time, and most importantly, made sequentially.
- To allow our framework to incorporate these kinds of characteristics, we define something called the "extensive form" of a game, which expands on the set of characteristics that are defined for any given game Γ. Specifically, we now have:

1. A set of players N
2. A set of payoffs defined on the outcomes of the games.
3. An ordering of moves.
4. A strategy space S, where a strategy is defined as a **complete contingent plan**.
5. What knowledge each players have when they move.
6. The distributions of each random, exogenous outcome.
7. **All of 1-6 is common knowledge**

- This is the **formal** definition of an extensive form game, but as with matrix representations of normal form games, we will often use convenient shorthand to represent these characteristics.
- Central amongst these is the **game tree**. A game tree consists of:

1. Nodes = points of decision of players.
2. A partitioning of those nodes between players, identifying who is making those decisions, where one player can be "Nature", i.e. a random player representing exogenous events.
3. A **partial** ordering of those nodes, which identifies which nodes precede other nodes. Some nodes it will not be possible to order (e.g. if they end up on different branches of some complicated tree).
4. Payoffs defined on **terminal nodes**. Terminal nodes are nodes that do not precede another node.

- At this point, it's worth looking at an example tree to see how these features are represented:

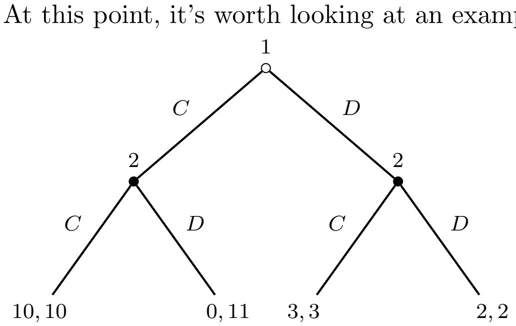

- In the above, you can see the terminal nodes are where the payoffs are located.
- Moreover, we can see the node for player 1 and the nodes for player 2 are ordered, since player 2 plays second.
- Meanwhile, the two nodes for player 2 are not ordered, i.e. neither precedes the other. They are just in different branches.

- There is no Nature player in this game, and we'll get to explaining that in a bit.
- So, from the definition of the extensive form game above, we can see how this game tree captures:
  - The set of players $N = \{1, 2\}$
  - The set of payoffs defined over outcomes.
  - The ordering of moves.
  - The strategy space, where here strategies are $S_1 = \{C, D\}$ and $S_2 = \{CC, CD, DC, DD\}$
- So 1-4 of the definition.
- 6 (probability distributions over exogenous events) is captured by the possibility of a random player Nature, who just chooses branches according to some probability distribution (i.e. non-strategically). This is just to include the possibility of exogenous uncertainty, created by random events like, for instance, weather, flipping a coin, rolling a die, etc.
- 5 is captured by the concept of **information sets**, which we can now introduce:
- Each player i's nodes are partitioned into a collection of information sets $h_i \in H_i$ with the following properties:
  1. If only one node is in an information set h_i, then player i knows exactly where they are in the game if they reach that node.
  2. If these is more than one node in the information set, player i does not know which of those nodes they are at when they reach any of those nodes.
  3. Given that they don't know which of the nodes they are at, they have to choose one action for all of the nodes in any given information set; they cannot condition on the particular node, because they don't know which one they're in when they make a choice.
- You may recall that we said earlier than simultaneous games could be represented as sequential games in which the players don't know what the other player has done before they make a choice.
- So we can use information sets to represent these simultaneous games in extensive form. For instance, the prisoner's dilemma, which in normal form is:

	C	D
C	3,3	0,5
D	5,0	2,2

- Is in extensive form:

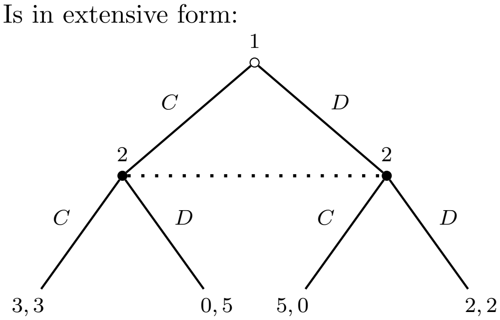

- The dotted line here indicates that when player 2 makes a choice, they don't know what player 1 does. This is mathematically identical to the idea of "simultaneous" decisions.
- So information sets capture the idea of "what does each player know at each point in the game".

- In general, we can always describe an extensive-form game in terms of its normal form, but doing so involves 'throwing out" all the additional characteristics/information we had when we defined a normal form game.

- So for instance, if we have the extensive-form game here:

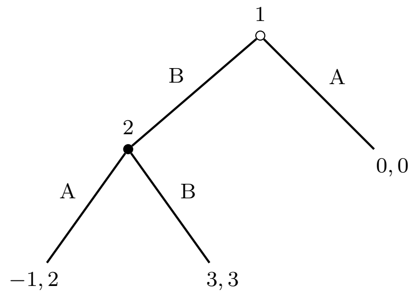

- The normal form representation would be, in matrix form:

	A	B
A	0,0	0,0
B	-1,2	3,3

- You can see that this matrix captures the key normal-form properties of: (1) the players; (2) the strategies; (3) the payoffs to outcomes.

- However, we lose the sequencing!

- In this case, we have (A, A) and (B, B) are Nash equilibria, but only (B, B) seems **credible** in some sense, given that it doesn't seem reasonable that P2 would choose B if the game arrives at their decision node.

- Indeed, only (B, B) will survive our next solution concept designed to deal with choices in sequence: **subgame perfect Nash equilibrium**.

4.8.2 Imperfect Information

- In a game of **perfect** information, each player knows exactly where they are at every part of the game, and there is no exogenous uncertainty. In other words, each information set is a singleton, and there is no Nature player.

- **Imperfect** information introduces: non-singleton information sets, and Nature players.

- **Incomplete** information is above a breakdown in the common knowledge of the game. Because this is somewhat intractable mathematically, we will introduce later something called the **Harsanyi tranformation** to represent a class of these situations as games of **complete but imperfect** information. But for now, just keep in mind the distinction between these two concepts of imperfect and incomplete.

- So what about these Nature players? How do they work?

- Ultimately, it's pretty straightforward. It's just a convention for making people deal with expected utilities when outcomes are non-deterministic.

- So consider this example, where someone is choosing whether to invest or not in a project, and the project can be either good or bad. Here, Nature chooses whether the project is good or bad, but according to a probability distribution, not strategically.

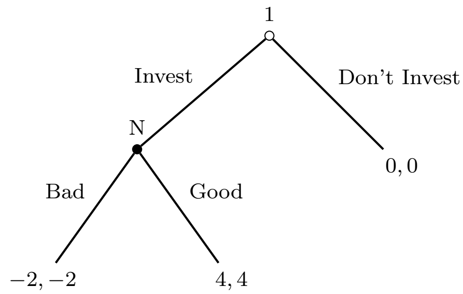

- So let's say the probability Nature chooses Good is p and thus the probability Nature chooses bad is $1-p$. Player 1 is thus choosing to invest in the following way:

$$EU(invest) = p(4) + (1-p)(-2) = 6p - 2$$

$$EU(\neg invest) = 0$$

- So they'd invest if $p > \frac{1}{3}$
- To emphasize; so far, we've been talking about the **pure strategies** of players in extensive form games. But just like with normal form games, it will also be possible to have **mixed strategies**.
- Mixed strategies induce **endogenous uncertainty** and do not require a Nature player. Nature is essential for **exogenous** uncertainty.
- If a **pure strategy** of an extensive-form game is a complete contingent plan, i.e. a choice at every information set for each player i, then a mixed strategy is a randomization over those plans.
- This differentiates **mixed strategies** from **behavioral strategies**, as behavioral strategies randomize actions at particular nodes. But in almost every case behavioral strategies are representable as mixed strategies, so we will generally not worry about this distinction.

4.8.3 Subgame Perfection

- Having defined the characteristics of an extensive-form game, we can now think about ways in which we can leverage the sequential structure of the game to limit the set of strategy profiles we consider.
- Key to this analysis is the idea of **sequential rationality**. We want to define a solution concept that rules out people making decisions that would not be **credible**, i.e. would not make sense for them once they actually arrive at that node.
- Before going any further, let's try out one such game: the centipede game!

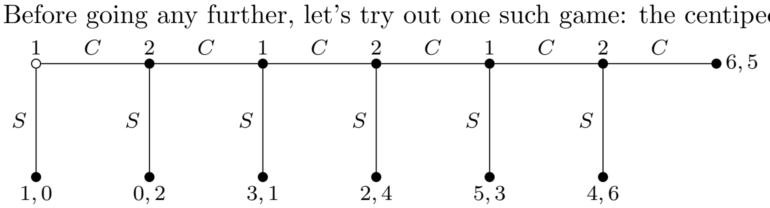

- The **subgame perfect Nash equilibrium** (SGPE) of this game is that Player 1 immediately chooses S at the beginning! Or more precisely, both players choose strategies of SSS, which leads to that outcome.
- This is the only thing that survives sequential rationality in the way we will define it. Why? Well let's find out.

- **Sequential rationality** is the idea that players must be best responding at every one of their information sets (i.e. the places where they get to make decisions).
- Implementing this idea more systematically requires a solution concept, of which the most key one is SGPE.
- What is a subgame?
  - Single initial node that is only member of that node's information set.
  - All successor nodes are in subgame.
  - All nodes in the information set of any node in the subgame must also be in the subgame.[4]
- What is a strategy in the context of an extensive game?
- A strategy in an extensive game specifies a strategy *for every subgame*; this includes strategies for nodes that are never reached.
- Question: Why must a strategy specify an action for *every* subgame of the game, and not just the actions taken in equilibrium?
- Answer: Because the equilibrium *depends* on the strategies **off the equilibrium path.**
- Here, a node is **on the equilibrium path** if it is reached with positive probability give the strategies played by each player.
- You may have heard of **backwards induction** in some previous class. You may even have gotten the impression that this is "the" solution concept for extensive form games.
- Backwards induction is a **procedure** for finding subgame perfect Nash equilibria. It works great in these finite extensive form games of complete information, and is helpful in other circumstances as well.
- So, to give an example, let's go back to the game we looked at earlier:

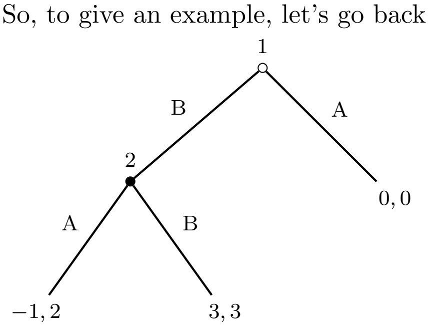

- The normal form representation would be, in matrix form:

	A	B
A	0,0	0,0
B	-1,2	3,3

- There are two subgames, and this is a finite game of perfect information.
- (A, A) does not survive SGPE because in the second subgame, it is not Nash. It requires P2 not to be best responding "off the equilibrium path".
- We can thus figure out the unique SGPE by using backwards induction. Furthermore, the following results are true for these finite games of perfect information:

[4] Morrow 1994

1. Each has at least one SGPE.
2. If no two terminal nodes produce identical payoffs, that SGPE is **unique**.
3. Backwards induction will find this unique equilibrium.

- However, as we shall see, backwards induction is not *identical* to SGPE.
- Why does backwards induction lead to to subgame perfection?
- Backwards induction is a technique used to ensure that strategies are Nash at every node, including those which are not reached.
- Important to note is that these strategies off the equilibrium path are often absolutely essential to the equilibrium.
- Example:

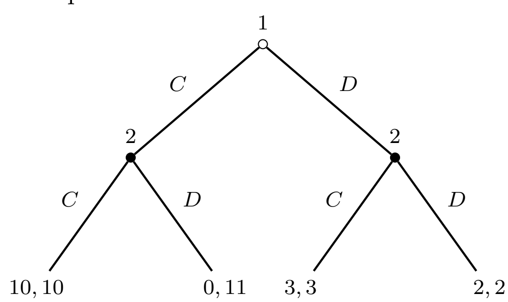

- In the above case, the only subgame perfect Nash equilibrium is $(\sigma_1, \sigma_2) = (D, DC)$.
- Note that although C is not played by player 1 in equilibrium, it is important to specify that Player 2 would play D if Player 1 played C in order for this equilibrium to hold. If we instead had the strategy profile (D, CC), Player 1 would have an incentive to deviate to playing C, in which case we would now arrive at a subgame $(C-)$ where playing D is not incentive compatible for Player 2.
- We can also construct the normal form of this game:

	CC	CD	DC	DD
C	10,10	10,10	0,11	0,11
D	3,3	2,2	3,3	2,2

- Where we can see that there is also only one *Nash* equilibrium, i.e. (D, DC). So these off-eq-path strategies matter here even if we're not using our SGPE solution concept.
- Sequential battle of the sexes?

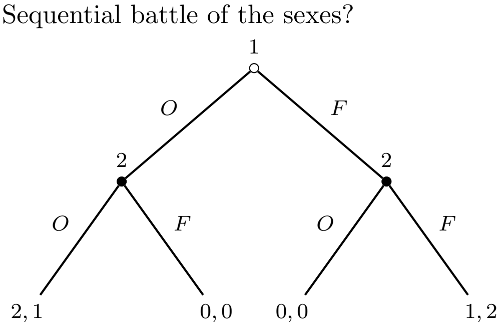

- With normal form:

	OO	OF	FO	FF
O	2,1	2,1	0,0	0,0
F	0,0	1,2	0,0	1,2

- Similar setup, but now the SGPE filter becomes important. NE retains the (F, FF) equilibrium and the (O, OO) equilibrium, but SGPE only retains (O, OF), taking into account that F is not a NE of that subgame.

4.8.4 Example Questions

- k threshold public goods game in sequential form. (**Ans:** Last k contribute)
- Now let's consider a "coffee buying" game with five people. Need \$6 to buy coffee, each person values having coffee at the meeting as \$10. What's the SGPE? (**Ans:** Last person contributes all \$6. Everyone else contributes nothing.)
- Ultimatum game. Generalize to any take-it-or-leave-it bargaining game. (**Ans:** Key intuition: make an offer that makes the other player exactly indifferent. This structure of game allocates the first-mover all of the surplus.)
- Three-round alternating-offer bargaining model. Fixed cost of delay. (**Ans:** Work backwards. P2 will be made indifferent in the last round.)
- Finitely repeated Prisoner's dilemma. (**Ans:** Will defect in last round. This unravels all the way back for any finite number of repetitions. Infinitely repeated games can have different equilibria, since there's no last round in which defection dominates.
- Here's an interesting example. Imagine that you play a two stage game where in the first stage you play simultaneous game:

	A	B
A	3,3	0,5
B	5,0	2,2

  and in the second stage you play:

	C	D
C	7,7	0,0
D	0,0	1,1

- What are the subgame perfect Nash equilibria of this game?
- We can condition on the history (i.e. choices made in the first stage), although the second stage needs to be an equilibrium. So consider the following strategy profiles of the form (σ_1, σ_2):
  $\sigma = ((B, D|AA, D|AB, D|BA, D|BB), (B, D|AA, D|AB, D|BA, D|BB))$
  $\sigma' = ((A, C|AA, D|AB, D|BA, D|BB), (A, C|AA, D|AB, D|BA, D|BB))$
- Second stage *always* has to be a Nash equilibrium. If there's only one, then it will always be that Nash equilibrium.
- However, if there's more than one, we can now start to condition on histories.
- Of the above, both σ, σ' are subgame perfect Nash!
- This assumes $\delta = 1$ discounting. What's the cutpoint on δ below which the σ' SGPE is no longer sustaintable?
- Consider: you get an extra payoff of 6 from the second stage by getting the "better" equilibrium. You would get an extra payoff of 2 by deviating from A to B in the first stage. So when is $2 > \delta 6$?
- When $\delta < \frac{1}{3}$, you cannot sustain the the SGPE, as each player will have an incentive to cheat in the first round.

- Imagine we change the second stage to:

	C	D
C	7,7	2,2
D	2,2	1,1

- Is σ' still a SGPE? **Ans:** No! Because now there's no "punishment equilibrum".
- Interestingly, the second stage game has higher payoffs to every strategy combination, but leads to a lower OVERALL payoff over both stages relative to σ' before the change in payoffs.
- As with NE, we can also apply the concept of SGPE to continuous strategy spaces.
- Here, the key characteristic is that subsequent players observe specific choices made before they make their own choices. And earlier players know that they will be "committed" to a choice after they make it, and that subsequent players will be respond to it.
- So, we find Best Response Correspondences for the players furthest down the game tree, and then "plug in" these to the preceding player's utility function.
- We then figure out preceding player's optimal response considering what the player will do after them.
- **Key to this:** we don't find the preceding player's best response function until after the later player's, because the earlier players will take subsequent choices into account when determining their optimal choice.
- Furthermore, **they know** that subsequent choices will take into account their choice and response optimally.
- So, using the security investments from before, we get the following:

Now, however, Country 1 has to announce how much they're investing in the military before Country 2 does (maybe they have to propose the budget in the legislature earlier). Country 2 then observes this decision before determining how much to invest.

(a) Find all the Subgame Perfect Nash Equilibria of this game.
Ans: To do this, note that not only does Country 2 observe Country 1's choice, but Country 1 *knows* that Country 2 will observe Country 1's choice, and knows Country 2's objective function. Thus, backwards induction involves taking into account what Country 2's response will be before committing to a particular choice of f_1.

$$\begin{aligned} u_1(f_1) &= ln\left(f_1 + 0.5\left(\frac{1}{c} - 0.5f_1\right)\right) - cf_1 \\ &= ln\left(0.75f_1 + \frac{1}{2c}\right) - cf_1 \end{aligned}$$

Taking first order conditions for this new expression,

$$\begin{aligned} \leftrightarrow \frac{\partial u_1}{\partial f_1} &= \frac{0.75}{0.75f_1 + \frac{1}{2c}} - c = 0 \\ \leftrightarrow 0.75 &= c\left(0.75f_1 + \frac{1}{2c}\right) \\ \leftrightarrow 0.75f_1 &= \frac{3}{4c} - \frac{1}{2c} \\ \leftrightarrow f_1^* &= \frac{12}{12c} - \frac{4}{6c} = \frac{1}{3c} = \frac{1}{3}\frac{1}{c} \end{aligned}$$

Note that we can compare this to what we obtained in Problem Set 4 ($f_1^* = \frac{2}{3}\frac{1}{c}$). In the SGPE of this new sequential game, Country 1 invests far less (half as much) in their military! This is because they can commit to investing less first and know what the reaction of country 2 will be, and can exploit this.

(b) Is the Nash equilibrium from the normal form representation still a Nash equilibrium? Explain your answer.
Ans: Yes. It is always the case that the NE of a game will stay NE as you start revealing more information sequentially (subject to the game having common priors and whatnot). Consider that Country 2 can adopt a strategy whereby they always invest $\frac{2}{3c}$ in the military (the equilibrium quantity from before) *irrespective* of what Country 1 chooses. This is non-credible, because this implies not best responding to some potential choices by Country 1, but for NE we don't need to worry about credibility, and only need to check if players are best responding on the equilibrium path. If Country 2 commits to $\frac{2}{3c}$ no matter what, then country 1 will also choose $\frac{2}{3c}$, and we have the NE from before. Note however that this is NOT a SPNE, because of the fact that it is not commitment credible; if Country 1 invests an amount different than $\frac{2}{3c}$, it will no longer be incentive compatible for Country 2 to invest $\frac{2}{3c}$.

4.9 Static Games of Incomplete Information

4.9.1 Learning

- Before we talk about Bayesian Nash equilibrium, let's take a moment to talk about **learning** more broadly.
- What is our formal framework for learning? It tends to involve a particular procedure for incorporating new information: **Bayesian** learning (or "rational learning").
- In general, learning involves:
  1. Starting with some prior belief about X.
  2. Being provided new information.
  3. Updating your belief about X in response to (2).
- Bayesian learning models the procedure from prior $\rightarrow$ posterior as occurring "correctly" as a matter of computing conditional probabilities. Thus it is "Bayesian" because it follows Bayes' rule.
- However, this is a reasonably contentious substantive assumption. While we can sometimes make even weird beliefs fit Bayes' rule with some mullarkey with priors, clearly people aren't information processing machines, and everything from motivated reasoning to cognitive shortcuts can affect this learning process.
- However, Bayesian learning is a pretty decent first approximation of how learning can be incorproated into games, and can be a good benchmark which alternative, pscyhologically-grounded models of learning can be compared.
- If you're building a model that's focused on, say, psychological biases, you can do so using an alternative learning procedure, and compare it to a model that use Bayes.
- Anyways, what is Bayes' rule? How do we compute conditional probabilities?
- A further discussion of this is in the Appendix, but for a quick review, for any event A the conditional probability of it occuring given some other even B is:

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

- Where the denominator here is a consequence of the law of total probability, i.e. if the set $\{B_i\}_{i=1}^n$ are disjoint events that partition the sample space, then:

$$P(A) = \sum_{i=1}^{N} P(A|B_i)P(B_i)$$

- In summary: in the rest of this course, we will tend to assume people learn/update their beliefs according to Bayes' rule, and this is the formal procedure for learning that is built in to the core solution concepts that deal with incomplete and imperfect information.

4.9.2 Bayesian Nash Equilibrium

- With this solution concept, we are dealing with uncertainty, but it should be noted we are not yet in a world where "signaling" occurs, as we are talking about simultaneous games instead of talking about sequential/dynamic games.
- Indeed, moving to this topic from subgame perfection, we are getting rid of the sequencing/subgames, and introducing player types.
- Whereas in the previous section on extensive form games we would sometimes have **imperfect** information (there may have been exogenous or endogenous uncertainty, and people may not have been aware of what specific nodes they were at at any given point), here we deal with games of **incomplete** information.
- **Incomplete** information, to put it simply, is when people don't know what other people's payoffs are.
- With the extensive form, we assume that all the features of the game (e.g. the utility functions for each player) were common knowledge. But obviously, there are many substantive situations where this isn't the case.
- If I'm asked to work with you on a group project, for instance, I don't know if you're a hard worker or lazy, but I have to make strategic decisions anyways.
- The trick to managing this kind of situation is called the **Harsanyi transformation**, which won John Harsanyi a Nobel prize (he shared it).
- It transforms a game of **incomplete** information into a game of **complete but imperfect** information.
- With this transformation, players may have different types, but Nature assigns the type of a player according to a distribution that is common knowledge.
- These types reflect *private information* of the player; for instance, I may know whether I'm a hard worker or not, but you may not observe this.
- A Bayesian Nash Equilibrium (BNE) occurs when no player has an incentive to deviate from their strategies, given the strategies of the other player, and given their beliefs about types. In this case, a player for which there exists unresolved uncertainty will be comparing expected utilities of different strategies.
- Formally, a Bayesian game is:

$$\langle N, \{A_i\}_{i-1}^n, \{\Theta_i\}_{i=1}^n, \{u_i(\cdot|\theta_i, \theta_{-i}), \theta_i \in \Theta_i, \theta_{-i} \in \Theta_{-i}\}_{i=1}^n, \{\phi_i\}_{i=1}^n\rangle$$

- So that's a mouthful! But let's go through it:

- N is just the set of players.
- $\{A_i\}_{i-1}^n$ is the set of actions for the players.
- $\{\Theta_i\}_{i=1}^n$ is the type space for the players.
- $\{u_i(\cdot|\theta_i, \theta_{-i}), \theta_i \in \Theta_i, \theta_{-i} \in \Theta_{-i}\}_{i=1}^n$ is the set of utility functions for the players. So this is long, as it's worth flagging a couple of this. Utility is dependent both on the type of player i (i.e. θ_i), and on the types of other players (θ_{-i}). The book focuses on the **private values** case where my utility is unaffected by your type, except indirectly through how your type affects your action choice. I have specified this in a more complete way, that allows for **common values**, which are some of the most interesting applications.
- $\{\phi_i\}_{i=1}^n$ is the set of beliefs for each player about everyone else's type.

- Important to note here is that like with SGPE, where a strategy was a complete contingent plan for every information set in a dynamic game, here a strategy is a complete contingent plan **for every possible type**.
- Although you will only end up being one type in actuality, the strategies you would play if you were different types are essential for determining what other players' best response will be, and are thus key parts of the equilibrium.
- A Bayesian Nash equilibrium is just a Nash equilibrium in which everyone's type-dependent strategies are best responses, incorporating their beliefs about what types other players are via expected utility calculations.
- Meanwhile, those beliefs are determined by Bayes' rule whenever appropriate.
- So let's consider an example. Let's say Player 1 is unsure if Player 2 is "cooperative" or "uncooperative", and thus doesn't know whether or not P2 will treat the game more like a prisoner's dilemma, or more like a coordination game.
- Player 1's has only one type, which is commonly known. Player 2 has two types, $\Theta_2 = \{T_1, T_2\}$.
- The sequence goes like this:
  - Nature determines whether Player 2 is T_1 or T_2. Probabilities of each are π and $1 - \pi$.
  - Player 2 (the column player) learns their type, and thus knows whether they are playing Game 1 or Game 2, but player 1 does not.
  - Player 1 chooses either T or B and Player 2 chooses either L or R, but Player 2 can condition on their type, because they observe it. So the strategy sets are $S_1 = \{T, B\}$, $S_2 = \{TT, TB, BT, BB\}$.
  - Game 1:

	L	R
T	5,5	2,4
B	4,2	3,3

    Game 2:

	L	R
T	5,4	2,5
B	4,2	3,3

- Something that's key here: I have kept Player 1's payoffs the same to each action combination in both games. This means the game has **private values**; Player 2's type will impact Player 1, but not **directly**, only via how it affects Player 2's choices.
- However, because Player 2's payoffs are affected by their type, they view the first game as essentially a cooperative game (they want to choose L if P1 chooses T, and choose R if P1 chooses B), but the second game as something akin to a prisoner's dilemma (R dominates).

- In other words, T_1 players are cooperative, and T_2 players are uncooperative. But P1 doesn't know which of them they're facing when they play the game.

- In extensive form, the game looks like this:

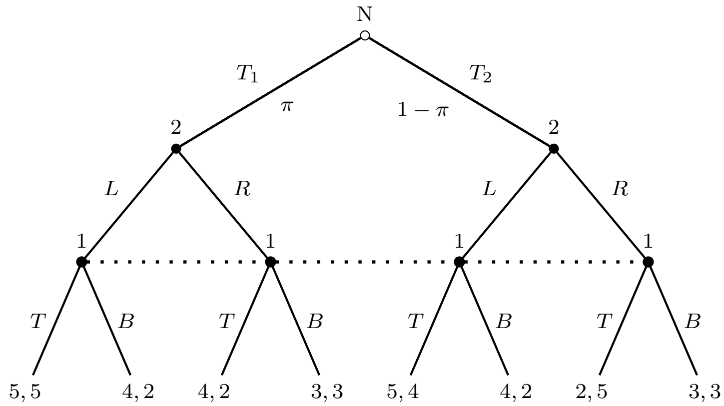

- If you look carefully at this tree, you'll note that $P1$ observes their type, but $P2$ observes neither $P1$'s type, nor what choice $P2$ makes before they move.

- To find equilibria, we need to check possible strategy profiles, where a strategy must specify an action at every information set.

- Player 2 has two information sets (see above) while player 1 has one information set.

- Each of Player 2's information sets here reflects a different "type". So again, they will only end up being one of these kinds of players, but what they *would* do if they were the other type of player is essential for characterizing the equilibrium.

- Rather than check *every* possible strategy profile to see if it's an equilibrium, we can note a few features to simplify the process.

- In this game, $P2$ always chooses R if they're the uncooperative type. So we know that $R|T_2$ is part of their strategy.

- We also know, from above, that if they are T_1 they will choose L if $P1$ chooses T and R if $P1$ chooses B. This actually only leaves us with two possible strategy profiles:

$$s^1 = (T, LR), s^2 = (B, RR)$$

- For each of these, we've already checked that LR and RR are a best response for $P2$. So now we just need to check $P1$'s expected utility from T versus B, conditional on $P2$'s strategies as defined in those profiles, to see if sticking with each is incentive compatible.

- So, consider the case of $\pi = 0, 25, 1 - \pi = 0.75$. We would have the following.

- For (T,LR):
  $EU_1(T|LR) = (0.25)(5) + (0.75)(2) = 2.75$ $EU_2(B|LR) = (0.25)(4) + (0.75)(3) = 3.25$ So there *is* an incentive to deviate, and (T, LR) is not a Bayesian Nash equilibrium.

- For (B,RR):
  $EU_1(B|RR) = (0.75)(3) + (0.25)(3) = 3$
  $EU_1(T|RR) = (0.75)(2) + (0.25)(2) = 2$ Therefore, (B,RR) is a Bayesian Nash equilibrium! There is no incentive to deviate.

- Indeed, this equilibrium is unique.
- The following example has a similar structure, but with Player 1's type being revealed, and Player 2 responding.
- It is also different in that it implies **common values**, because Player 2's payoffs also depend on Player 1's type **directly**, and not just through how it affects Player 1's strategies.
  - Game 1:

	L	R
T	12,6	5,9
B	7,12	5,9

Game 2:

	L	R
T	10,5	8,6
B	12,4	3,6

- Which in extensive form looks like:

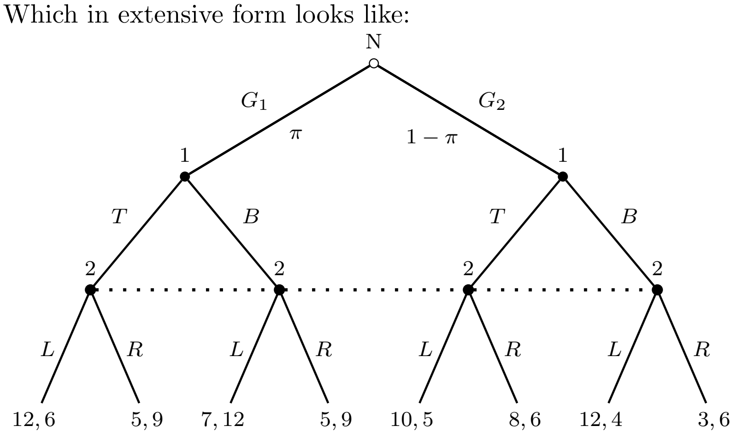

- To find equilibria, we need to check possible strategy profiles.
- Now, Player 1 has two information sets while player 2 has one information set.
- Let's assume $\pi = 1 - \pi = 0.5$ and solve the problem.
- I start by reducing the set of strategy profiles we need to look at, by assuming a particular strategy by Player 2 and figuring out what best responses would be for Player 1. Because Player 1 knows their type, and can condition on it, this just means picking their pure strategy best responses in each game to any single action choice by Player 2.
- If Player 2 chooses L: $\{(TB, L)\}$
- If Player 2 chooses R: $\{(TT, R), (BT, R)\}$
- Then, we just need to compute expected utilities to figure out whether these are Bayesian Nash equilibria.
- For (TB,L):
  $EU_2(L) = (0.5)(6) + (0.5)(4) = 5$
  $EU_2(R) = (0.5)(9) + (0.5)(6) = 7.5$
  Therefore, (TB,L) is not a Bayesian Nash equilibrium.

- For (TT,R):
  $EU_2(L) = (0.5)(6) + (0.5)(5) = 5.5$
  $EU_2(R) = (0.5)(9) + (0.5)(6) = 7.5$
  Therefore, (TT,R) is a Bayesian Nash equilibrium!
- For (BT,R):
  $EU_2(L) = (0.5)(12) + (0.5)(5) = 8.5$
  $EU_2(R) = (0.5)(9) + (0.5)(6) = 7.5$
  Therefore, (BT,R) is not a Bayesian Nash equilibrium.
- So we are left with only one Bayesian Nash equilibrium: (TT,R)
- As you might expect, you can have significantly more complicated games that we'd want to compute Bayesian Nash equilibria for, but this is the basic structure: Nature defines a type, strategies are actions for each possible type, and then we check to see if everyone is best responding using expected utility calculations when necessary.
- In the next part of the class, we will combine these type-dependent strategies with the time/sequencing of the extensive form games we covered earlier. This will give us *dynamic games of incomplete information*, which allows for interesting situations like learning from another player's actions and making choices accordingly.

4.10 Dynamic Games with Incomplete Information

4.10.1 Perfect Bayesian Nash Equilibrium

- Up until this point we've dealt with the following solution concepts, designed for different game structures:
  - Nash Equilibrium - the basic framework upon which we build most of the other solution concepts.
  - Subgame Perfect Nash Equilibrium: introduces a particular instantiation of sequential ration (i.e. NE in every subgame) to help deal with the challenges of credibility in dynamic games.
  - Bayesian Nash Equilibrium: introduced types as a way of dealing with incomplete information via the Harsanyi transformation, and brought beliefs into our framework.
- Now, we are interested in combining characteristics of essentially all these frameworks to study **dynamic** games of **incomplete** information.
- For these kinds of games, we want to come up with a solution concept that will give us some traction on strategic situations that evolve over time/in sequence (the dynamic part), but also involve uncertainty about other players' utility functions (incomplete information).
- This kind of structure will also allow us to address questions of "strategic information transfer" - e.g. when is telling the truth an equilibrium, under what conditions can we learn about other players' types from their actions.
- The core solution concept for these kinds of games is Perfect Bayesian Nash Equilibrium. It is defined as a Bayesian Nash Equilibrium with the following properties:
  1. A well-defined system of beliefs: players must have a probability distribution over the nodes in each of their information sets.
  2. On the path beliefs must be consistent with Bayes' rule.
  3. Off the equilibrium path beliefs, however, can be whatever.
  4. We need *sequential rationality.* Players must be best responding with respect to their beliefs.

- A key insight here is that subgame perfection doesn't really help us with (4), because in these games of incomplete information we often want to impose sequential rationality on future decisions starting at non-singleton information sets.
- Given the non-singleton start, it's not a proper subgame, and subgame perfection doesn't tell us anything one way or another.
- In general, there are two types of Perfect Bayesian Nash Equilibria (PBNE): **separating** and **pooling**.[5]
- Separating: types do different things.
  - This tends to imply *perfect revelation of information.* If only one type does x, then if you observe x you know it's that type.
  - This is a trivial application of Bayes' rule. If Type 1 chooses A and Type 2 choose B, for instance:

$$Pr(t_1|A) = \frac{Pr(A|t_1)Pr(t_1)}{Pr(A|t_1)Pr(t_1) + Pr(A|t_2)Pr(t_2)} = \frac{(1)Pr(t_1)}{(1)Pr(t_1) + (0)Pr(t_2)} = \frac{Pr(t_1)}{Pr(t_1)} = 1$$

- Pooling: types do the same thing.
  - This implies *no revelation of information.* If all types do the same thing, then your posterior beliefs should be equal to your prior beliefs.
  - Again, formally (but trivially) with Bayes' rule, if Type 1 chooses A and Type 2 chooses A, and $Pr(t_1) = p$ and $Pr(t_2) = 1 - p$

$$Pr(t_1|A) = \frac{Pr(A|t_1)Pr(t_1)}{Pr(A|t_1)Pr(t_1) + Pr(A|t_2)Pr(t_2)} = \frac{(1)p}{(1)p + (1)(1-p)} = \frac{p}{1} = p$$

- To clarify these concepts, let's consider an entry-deterrence game, drawn from Tadelis.
- In this game, a potential entrant (Player 1) into a market is either a competitive (C) or uncompetitive (U). They can choose whether to enter the market (E) or stay out (O).
- The existing market participant (Player 2) has the option of fighting them after they enter (F), or accepting their entry (A). The game tree thus looks as follows:

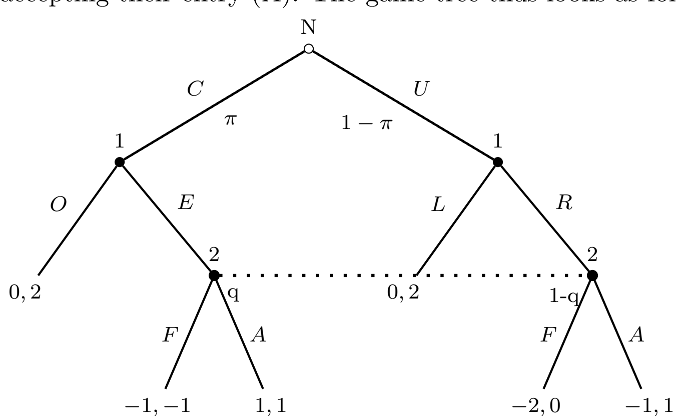

- This game has two Bayesian Nash Equilibria: (OO, F) and (EO, A).
- However, it should be reasonably clear that F is not especially credible from Player 2: they receive a strictly lower payoff than from A **regardless** of Player 1's type.

[5]There are also hybrid "semi-separating" or "semi-pooling" equilibria, which are where one or more types play mixed strategies, but these won't be a focus of this course.

- This would thus seem not to be sequentially rational. But because we don't have a proper subgame that starts with Player 2's decision (because it's not a singleton information set), subgame perfection doesn't get us anywhere.

- Perfect Bayesian Nash Equilibrium, however, does. Simply by having *any* consistent set of beliefs (condition 1 of PBNE, labeled here as q and $1 - q$) and requiring that decisions be consistent with those beliefs (condition 4 of PBNE) we rule out F as a response off-the-equilibrium-path to entry. This is because:

$$EU(F) = (-1)q + (0)(1-q) = -q$$
$$EU(A) = (1)q + (1)(1-q) = 1$$

  And therefore $EU(F) > EU(A)$ for all Q

- Thus, this allows for a new instantiation of **sequential rationality**.

- This particular application of sequential rationality rules out the (OO, F) equilibrium, leaving only (EO, A) as a Perfect Bayesian Nash Equilibrium. Notably, (EO, A) is an example of a **separating equilibrium**, since the types are choosing different actions.

- We can tweak the payoffs to the game slightly to illustrate the importance of off-equilibrium-path beliefs:

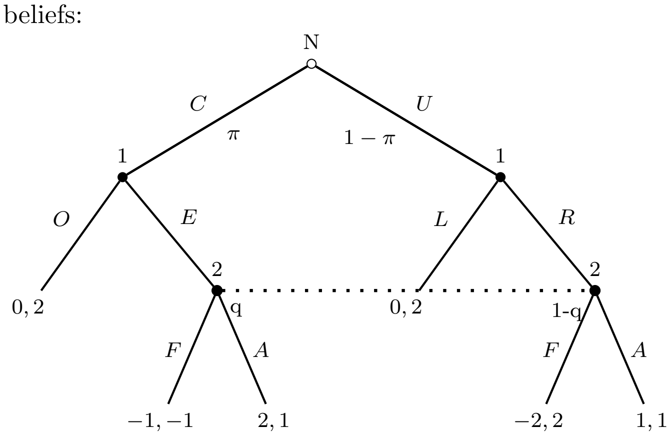

- Now, Player 2 prefers to fight uncompetitive entrants, but prefers to accept competitive entrants.

- Can we sustain the (OO, F) equilibrium now? Yes, but it will depend on Player 2's beliefs if they observe E, even though this occurs with zero probability if the strategies dictated by the equilibrium are played.

- Specifically, it must be the case that $EU_2(F|E) > EU_2(A|E)$, which with this new setup means:

$$(-1)q + (2)(1-q) = 2 - 3q \geq 1$$

$$\leftrightarrow q \leq \frac{1}{3}$$

- Intuitively, they must believe that the type is competitive with sufficiently low probability to make them want to fight upon observing entry.

- Note as well: this is an example of a **pooling equilibrium**, since both types are playing the same strategy.

4.10.2 Signaling Games

- Below are two game trees that represent the same basic signaling game.
- Note that the structure of the first looks a whole lot like the example given for Bayesian Nash equilibrium given done before. What's changed?
- Difference is that now Player 2 observes Player 1's choice. Allows for possibility of strategic information transfer.

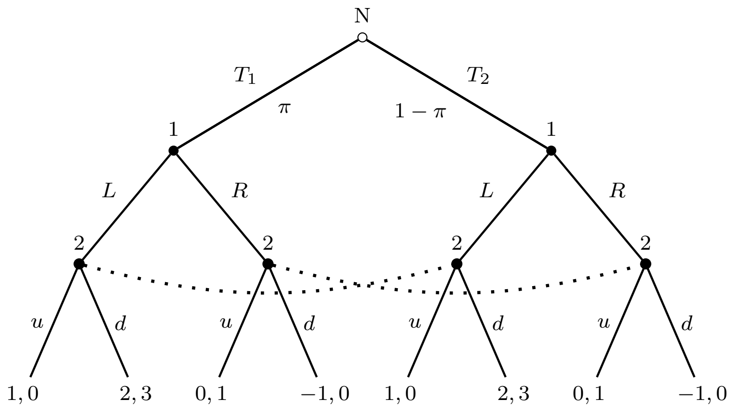

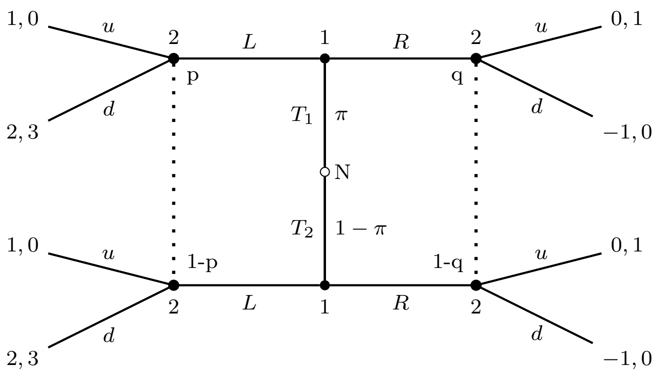

- What are Perfect Bayesian Nash Equilibria (PBNE) to these above game?
- If L is chosen, d dominates u for P2 irrespective of P1's type. Similarly, u dominates d if R is chosen. Thus, the equilibrium is is $\sigma = (L|T_1, L|T_2, u|L, d|R, p = \pi)$. Not super interesting. Let's do another, randomly selecting numbers.

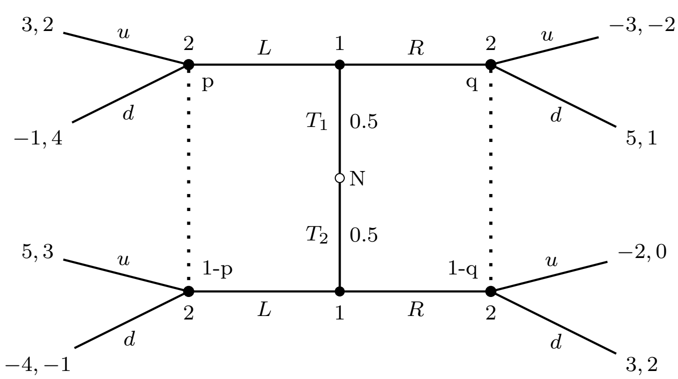

- If R is chosen, d dominates u. If L is chosen, there is no dominant strategy.
- To limit the set of strategy profiles we need to check, I first choose R and L (to determine separating or pooling) and then choose $P2$'s strategies such that they are a best response. Then I examine whether $P1$ would have an incentive to
- Test $(R, L, u|L, d|R, p = 0, q = 1)$. $P1$ has no incentive to deviate; this is a separating equilibrum.
- Test $(L, R, d|L, d|R, p = 1, q = 0)$. $P1$ has incentive to deviate when T_1, because $5 > -1$.
- Test $(R, R, d|L, D|R, p =?, q = 0.5)$. To determine if $P1$ has incentive to deviate, need to find off the equilibrum path beliefs that would make $P2$ choose d when observing L, otherwise $P1$ will deviate to L when T_2.

$$\begin{aligned}
EU(u) &\leq EU(d) \\
\leftrightarrow 2p + 3(1-p) &\leq 4p + -1(1-p) \\
\leftrightarrow 2p + 3 - 3p &\leq 4p - 1 + p \\
\leftrightarrow -6p &\leq -4 \\
\leftrightarrow p &\geq 2/3
\end{aligned}$$

- So $(R, R, d|L, d|R, p \geq 2/3, q = 0.5)$ is a pooling equilibrum.
- Test $(L, L, u|L, d|R, p = 0.5, q =?)$. $P1$ will always deviate when T_1, so this is not an equilibrium.
- So we have found one separating and one pooling equilibrium.
- We might wonder whether the off the equilibrium path beliefs in the pooling equilibrium, i.e. $p \geq 2/3$, are reasonable. After all, it would seem $P1$ would only even conceivably have an incentive to deviate when they are T_2, as it is in these cases that there's even a possibility of a higher payoff to them.
- Some authors have proposed systematic ways of restricting beliefs. This may allow us to focus our attention on the more reasonable equilibria. Morrow 1994 p.244 has a good introduction to this.
- We'll introduce one such example in the next section: the intuitive criterion.

4.10.3 Restrictions on Beliefs

- To explore the topic of restrictions on beliefs, let's first introduce the "Beer-Quiche" game, which was used in the paper (Cho and Kreps 1987) that initially developed the particular belief refinement we will focus on (the intuitive criterion).
- In this game, the story is that Player 1 is settling down at a bar to have breakfast, and is choosing between Beer and Quiche as their breakfast choice.
- There are two types of Player 1: surly and wimpy, i.e. $\Theta_1 = \{S, W\}$, and surly types prefer Beer for breakfast, and wimpy types prefer Quiche.
- However, both types care even more about avoiding a bar fight.
- Player 2 is an aggressive bar patron, who wants to pick a fight.
- However, they only want to fight wimpy types; the surly types would probably beat them.
- So their preference ranking is: fight wimpy > no fight > fight surly.

- This leads to the following game tree:

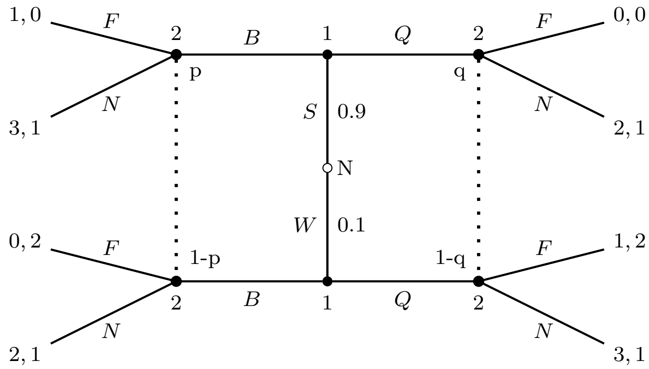

- A quick inspection should make it clear that there isn't a separating equilibrium to this game: if types separate so that $P2$ knows which type they are facing, they will fight W types and not fight S types. But this gives W types an incentive to deviate to imitate the W types and avoid getting beat up.
- What about pooling equilibria? Well it turns out there are two types of pooling equilibria, which are sustained by particular off-equilibrium-path beliefs.

$$(Beer, Beer, N|Beer, F|Quiche, p = 0.9, q \leq 0.5)$$

$$(Quiche, Quiche, F|Beer, N|Quiche, p \leq 0.5, q = 0.9)$$

- These both ensure that on the path, not fighting is the best response for Player 2, but prevents deviations from the type getting their less-preferred breakfast option by setting off-eq-path beliefs such that fighting is chosen upon observing the off-the-path breakfast.
- But is the second of these actually reasonable? Player 2 has to believe, upon observing Beer, that the Beer-choosing player is *less* likely to be a Surly type, but it is only the Surly types that could conceivably benefit from such a deviation.
- Indeed, wimpy types are getting their best possible outcome in the second pooling equilibrium.
- Intuitive criterion restricts off-eq-path beliefs to only put positive weight on types that could conceivably benefit from a deviation.
- The "intuition" behind this is that such a type could say to Player 2: "look, only a Surly type would ever do this, so it should be clear that I'm surly".
- Thus, applying the Intuitive Criterion restriction on beliefs rules out the second equilibrium, leaving us only with:

$$(Beer, Beer, N|Beer, F|Quiche, p = 0.9, q \leq 0.5)$$

- Which perhaps seems more reasonable! It in fact makes sense that observing Quiche off the equilibrium path would lead Player 2 to think it's more likely that the player is wimpy, given the payoff structure chosen.
- This is just one example of a **belief refinement**.
- In general, recall that uniqueness is one potential way of evaluating a solution concept: we might hope that we could generate more precise predictions about behavior from a situation.
- Belief refinements are a way of trying to systematically get us to more precise predictions, by establishing rules that can extract the more sensible equilibria.

- Using a belief refinement involves a trade-off: you're imposing more structure on the model by ruling out certain kinds of beliefs, so you better hope that you're right that your restriction on beliefs is correct.
- However, they do allow for more precise predictions, so it may be a trade-off worth making depending on the circumstances. Or, as happens in many cases, an author may simply report both the set of PBNE, and those that survive certain belief refinements, allowing readers to make their own decisions about which equilibria are the most reasonable.

4.10.4 Signaling Games With Non-Finite Spaces

- The examples of PBNE so far involve finite strategy spaces and finite type spaces.
- But the same logic can be applied to infinite and/or continuous spaces for both.
- Continuous type spaces tend to require integrals, since that's what you use to compute expected utilities in those cases. So given that we have avoided these for most of this course, we won't focus on those now.[6]
- Continuous strategy spaces with finite type spaces are often a bit more straightforward, though thinking through them can be a little tricky.
- Consider, for instance, the classic job market signaling framework from Spence 1978.
- In it, there are two types of workers, High Type and Low Type, so $\Theta_1 = \{\theta_L, \theta_h\}$. These types correspond to the value of hiring such a worker to the employer.
- $Pr(\theta_H) = \lambda$, $Pr(\theta_L) = 1 - \lambda$
- Education (e) costs for workers are:
$$c(e,\theta) = \begin{cases} e & \text{if } \theta_L \\ \frac{e}{2} & \text{if } \theta_H \end{cases}$$
- Where the utility to a worker if offered a job at wage w is:
$$u_w(e,\theta) = w - c(e,\theta)$$
- Utility for the employer is:
$$u_r(w,\theta) = \theta - w$$
- Let's say the employer has to offer a market wage of $\bar{w} = \lambda\theta_H + (1-\lambda)\theta_L$ if they hire someone, but can choose whether to hire someone or not.
- There are *many* PBNE of this game, given that we can set off equilibrium path beliefs however we want. But it's worth keeping two things in mind that make things simpler:
  1. Player 1 (the worker) has only two info sets: they condition on their type, and choose an education level e.
  2. Player 2 (the employer) observes e and decides whether to hire. So their strategy, and their beliefs, are both functions of e.
- So let's think about how we'd approach finding different kinds of equilibria:
- Pooling equilibria:

[6] Though if you're itching for an example, you can check out one of my papers:
`http://www-personal.umich.edu/~jasonsd/DavisJason.ScreeningforLosers.pdf`

- Because of how we've set the wage, the employer is exactly indifferent between hiring and not hiring in this case.
- So we can consider looking for an equilibrium in which hiring occurs. We can sustain *any* equilibrium in which e is sufficiently low that both types would rather be educated than not be hired.
- We do this by setting this $e' \leq \bar{w}$ and then setting off-eq-path beliefs such that $\mu(e) < \lambda$ for $e < e'$. This means that if an employer observes an e lower than the equilibrium, they believe they are facing a low type with sufficiently high probability that they prefer not to hire the worker.
- This keeps the workers from deviating to a lower e. Not wanting to pay the cost of higher education keeps them from deviating to higher e.
- We can't sustain a pooling equlibrium where $e > \bar{w}$, because then the low types would prefer to deviate to no education, even if it means they won't be hired.
- What about a "no hire" equilibrium? Because we set $\bar{w}$ to make the employer exactly indifferent, this is possible at $e = 0$, if we set off-eq-path beliefs such that $\mu(e) < \lambda$, $\forall e > 0$. Thus, the potential employee has no way of convincing an employer they are worth hiring.
- However, if w were even slightly lower that $\bar{w}$, this would no longer be the case, since now the employer would strictly prefer to hire in any pooling equilibrium. And if w were even slightly higher that $\bar{w}$, you conversely couldn't have a pooling equilibrium that includes hiring. But setting $w = \bar{w}$ allows us to explore both kinds of pooling equilibria.

- Separating equilibria:
  - In order for a separating equilibrium to hold, we need to choose an e^* such that high types want to choose it in order to be hired, but low types don't have an incentive to imitate them, because it would be too expensive.
  - This means we need an e such that $w \geq \frac{e}{2}$ but $w \leq e$. Or $e \in [w, 2w]$.
  - Then, once again, we set off-eq-path beliefs such that there's no incentive to deviate. So we make sure than $\mu(e) < \lambda$ for $e < e^*$.
  - So once again, there are many many equilibria, though they all share certain properties.
- Somewhat generically, the job market signaling game produces a lot of equilibria.
- Do all of these equilibria survive the intuitive criterion? No!
- Consider that a high type can make the claim, for any $e > w$, that they are the only type that could conceivably benefit from this choice, and thus they must be a high type.
- Intuitive criterion rules out the off-eq-path beliefs that keep high types from deviating to a lower e.
- So there's really only *one* separating equilibrium that survives intuitive criterion, in which $e = w$ exactly.

4.11 Repeated Games

- Our analysis of repeated games mostly relies on the same solution concepts used in other areas, e.g. NE and SGPE.
- For finite games, there's nothing especially distinct about the analysis.
- Infinitely repeated games can get weirder. Often get "anything goes" type results.
- Folk Theorems. One (Friedman 1971) is about subgame perfection. Any payoff vector with greater payoffs to each player than some NE can be supported as a subgame perfect Nash equilibrium

- Minimax strategy minimizes the payoff that another player obtains, given that this player is playing a best response.
- Minimax values are the payoffs to each player that they obtain if every other player minimaxes them.
- *Individually* rational payoffs are those "above" the minimax payoff vector. This makes sense; the lowest payoff attainable by a player who is playing a best response is their minimax payoff, so they will never accept a payoff lower than that if they are rational..
- Feasible set is the convex hull of payoffs to pure strategy profiles. Can obtain any of these via mixed strategies.
- Recall that convex hull of a set A is the smallest set B that contains all convex combinations of A. So obtaining any point in the convex hull just involves specifying the weights of the convex combination.
- Example game:

	L	M	R
U	(6,0)	(-1,-100)	(0,1)
D	(2,2)	(0,3)	(1,1)

- What is the feasible set? What are minimax values? What are individually rational payoffs?

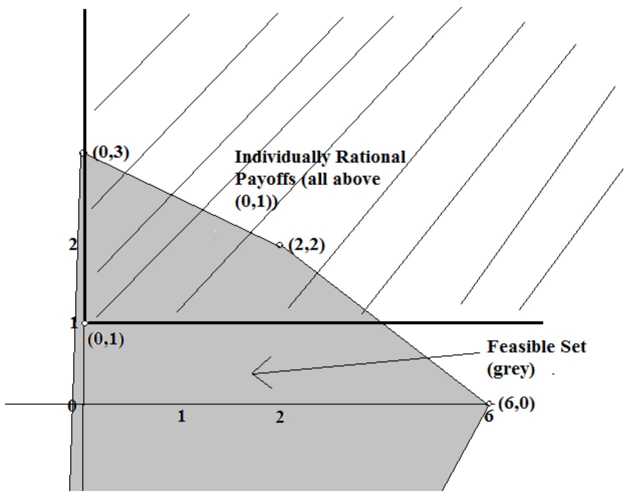

- Another folk theorem: any feasible and individually rational payoff vector can be obtained in a Nash equilibrium of a repeated game, given sufficiently "patient" players (i.e. high enough δ).
- How? Play mixed strategies with weightings required to obtain payoff vector, and switch to minimax strategies if anyone deviates from this.
- Will these be subgame perfect? Certainly if the minimax strategy is Nash (see Friedman 1971).
- If not, we can provide another folk theorem that allows for any individually rational and feasible payoff vector to be obtained in a subgame perfect equilibrium! Proof is more difficult, relies on "full dimensionality" condition. See Fudenberg and Maskin 1986.
- Involves punishing those who fail to punish.
- One shot deviation principle is important for subgame perfection, and also in Markov perfect equilibria (which basically entails subgame perfection, but with Markov strategies).
- Can, for instance, defection be a profitable one-shot deviation from repeated prisoner's dilemma when the other player is playing Grim trigger?

5 Appendix: Math Review

5.1 Sets and numbers

- Sets are collections of elements. Note: this is different than vectors.
- Different kinds of numbers are sets of all numbers. "Subsets" in particular.
- Examples: $\mathbb{N}$ = natural numbers = $\{1, 2, 3, ...\}$. These include either all positive integers, or all positive integers plus zero.
- $\mathbb{Z}$ = integers = $\{..., -2, -1, 0, 1, 2, ...\}$.
- Can include any kind of elements, and can be characterized through rosters, sets, or intervals.
- Roster: $X = \{A, B, cat, phonebook, 7\}$
- Set-builder: $X = \{X|x = x^2\}$. What is this set?
- Interval notation. $X = (0, 1)$ or $X = (0, 2]$
- Sum, difference, and products of two integers are also integers. Taking quotients allows us to develop set of "rational numbers": $\mathbb{Q}$ = rational numbers = $\{\frac{a}{b}|a, b \in Z, b \neq 0\}$
- Is every number a rational number? No: some like $\sqrt{2}, e, \pi$ are not. These are called "irational numbers".
- Total set of rational and irrational numbers if the set of "real numbers", denoted $\mathbb{R}$.
- Imaginary or complex numbers are also a thing, but we don't tend to deal with them very much
- Another question; what is this set? $X = \{X|x \in \mathbb{R}, x^2 < x\}$.
- Notation: say x is an element of a set Y with $x \in Y$.
- $\forall$ means for all or for each element in something. E.g. $\forall x \in \{1, 2, 3, 4\}, x > 0$.
- $\exists$ means there is at least one element that satisfies a property, e.g. $\exists x \in \{1, 2, 3, 4\} s.t. x < 3$.
- Subsets: $(\forall x \in S, x \in S') \to S \subseteq S'$
- Unions combine sets, e.g. if $X = \{A, B, tree, saxophone\}$ and $Y = \{B, saxophone, computer, 9\}$ means $X \cup Y = \{A, B, tree, saxophone, computer, 9\}$.
- Intersections include all elements that are in both sets, i.e. $X \cap Y = \{B, saxophone\}$.
- Draw Venn diagrams.
- How would we write this using the notation before? $(x \in X) \wedge (x \in Y) \to x \in X \cap Y$
- $A \backslash B$ means elements in A but not B, e.g. $X \backslash Y = \{A, tree\}$.
- What is $\mathbb{Z} \backslash \mathbb{N}$? How about $\mathbb{N} \backslash \mathbb{Z}$? This is null set, or $\varnothing$.
- Other notation: $|$ means such that (I also use "s.t.")
- $\equiv$ means "defined as" or "equivalent to".
- Cartesian products of sets: all ordered pairs of sets.
- Equivalent sets.
- Summation operator: $\sum 1^4 x = x + x + x + x$. Can generalize with x_i.
- Product operator: $\Pi_i^4 x = x^4$. Generalize with x_i.

5.2 Exponents

- $x^3 = x \cdot x \cdot x$. E.g. $2^3 = 8$.
- $x^m x^n = x^{m+n}$
- $\frac{x^m}{x^n} = x^{m-n}$.
- $x^{-y} = \frac{1}{x^y}$.
- $x^0 = 1$.
- $(xy)^n = x^n y^n$.
- $(x^n)^m = x^{mn}$
- Practice problem: $((xy)^4)^{0.5} yz^{-3}$
- Square roots are just fractional exponents, e.g. $\sqrt{2} = 2^{1/2}$. Often easier to deal with them in this fashion (in my opinion).

5.3 Introduction to Logarithms and e

- Log is inverse of exponent. Can be written for any base.
- I will focus on logarithms of base e, as these are the most common.
- $e = lim_{n\to\infty}(1 + \frac{1}{n})^n$
- $ln(e) = 1$, $ln(e^2) = 2$, $ln(e^3 23.383) = 323.383$
- $ln(x_1 \cdot x_2) = ln(x_1) + ln(x_2)$. E.g. $ln(e^2 \cdot e^3) = ln(e^3) + ln(e^2) = 5$
- $ln(\frac{x_1}{x_2}) = ln(x_1) - ln(x_2)$
- $ln(x_1^b) = bln(x_1)$. E.g. $ln(e^{12}) = 12ln(e) = 12$
- $ln(1) = 0$
- Logarithms thus can be very useful for dealing with products. E.g. if $y = \Pi_i^n x_i$, $ln(y) = \Sigma_i^n ln(x_i)$.
- First difference of log is also often used to approximate percentage changes. Note that for some variable x_t, $ln(x_t) - ln(x_{t-1}) = ln(\frac{x_t}{x_{t-1}})$

5.4 Inequalities and Absolute Values

- As we've already done a little, we often can represent sets by inequalities.
- Inequalities satisfy many of the same properties we get from equations, but with a few rule changes to consider.
- Keep in mind though that our end outcomes are not single values but sets of values.
- Strict versus weak inequalites. $x > y$ versus $x \geq y$.
- Transitivity: $x > y$ and $y > z$ implies $x > z$.
- Flip signs when multiplying or adding by a negative number.
- Absolute values: distance from zero on number line.

- Absolute value of -6 equals absolute value of 6 for that reason.
- $|x-3|$. Graph it.
- Including both inequalities and absolute values can sometimes be confusing.
- Q: $|x-1|>2$? $|x-1|<2$?
- Ans: $(-\infty,-1)\wedge(3,\infty)$. Second is $(-1,3)$ (interval centered on 1 with 2 on each side)

5.5 Elements of Real Analysis

- Displaying coordinates in 2-space, 3-space, etc.
- R^1, R^2, R^3, R^n. R^1 is the set of real numbers, and n generalizes this to n dimensional space.
- $n=1$ is the number line, $n=2$ is plane, $n=3$ is space.
- Points in R^n are represented by n-tuples, which are vectors. E.g. in R^3 you have 3-tuple $(0,7,3)$
- These are "Euclidean" vectors, i.e. vectors in n-dimensional Euclidean space.
- Can have generalized vectors in general vector spaces, which we may touch on in the linear algebra section, but Euclidean vectors tend to be our domain of interest.
- Addition and subtraction with vectors is simple: subtract or add each element to its corresponding element in other vector.
- Scalar multiplication is when you have a number in $\mathbb{R}$ and you multiply it by a vector. It entails multiplying each element of the vector by the scalar, so for instance for $a\in\mathbb{R}$:

$$a\begin{pmatrix}x_1\\x_2\\x_3\end{pmatrix}=\begin{pmatrix}ax_1\\ax_2\\ax_3\end{pmatrix}$$

- Dot products multiply each element of a vector with the corresponding element of another vector, and then add them all together. So, for instance:

$$\boldsymbol{v_1}=\begin{pmatrix}x\\y\\z\end{pmatrix}\ \boldsymbol{v_2}=\begin{pmatrix}a\\b\\c\end{pmatrix}$$

$$\boldsymbol{v_1}\cdot\boldsymbol{v_2}=xa+yb+zc$$

- We might talk about points *within a neighborhood* of another point. This is about describing an interval, disk, or (generalizing to R^n) "ball" around the point.
- If in R^1, we have epsilon interval around c defined as $\{x:|x-c|<\epsilon\}$. Recall that this creates an interval centered on c with ϵ on each side, from inequalities.
- Can generalize easily into 2-space. Anyone remember Pythagorean theorem?
- Distance between two 2-vectors involves subtracting them and then applying Pythagorean theorem.
- *i.e.* between $(1,2)$ and $(4,5)$ we would have $\sqrt{(4-1)^2+(5-2)^2}$
- Helpfully, this generalizes into n-space.
- i.e. for R^3 with $(1,2,3)$ $(2,4,6)$ we would have $\sqrt{(2-1)^2+(4-2)^2+(6-3)^2}$

- An ϵ-ball is the set of points around some point in n-space. i.e. all y such that $d(x, y) < \epsilon$ is an ϵ-ball (or "open ball") around x, where x and y are vectors in n-space.
- Interior point: point is interior to set X if there exists an ϵ-ball around the point such that all the points in the ball are also in the set.
- Boundary point: Any ϵ-ball includes points both inside and outside the set.
- Open set: Every point in S is an interior point. The set still has boundary points, i.e. $x = 2$ for $(0, 2)$, but these points are not included in the set.
- Closed set: Includes all boundary points.
- Bounded: set can be contained in an ϵ-ball.
- Compact (in Euclidean geometry): is both closed and bounded.

5.6 Functions and correspondences

- Functions map elements of some input set to an element of another set.
- Simplest functions map R^1 to R^1, are sometimes represented as $f : R \to R$.
- More complicated are two variables functions $f : R^2 \to R$, n-variable functions $f : R^n \to R$, or vector values functions $f : R^n \to R^m$.
- For now, we will start with one-variable functions, move on to multivariable functions, and leave vector-valued functions for some point in the future.
- Domain: input set.
- Co-domain: set the input set is mapped to. **NB: This is often referred to as the range, but this is ambiguous, as range may mean the specific values obtained by the function across the domain**.
- Image: the unambiguous word for which values the function actually obtains over its domain.
- For $f : A \to B$, graph is ordered pairs $G = \{(a, b) \in A \times B | f(a) = b\}$. This is generalization of graphing that we are familiar with.
- Single variable function examples: $f(x) = 2x + 3$. Implies $f(2) = 2(2) + 3 = 7$. We take input of 2 and get output of 7.
- What is the domain, co-domain, and image of this function? How about $f(x) = x^2$? (image is R_+ not R in this second case).
- How about $f(x) = 1/x$? $f(x)\sqrt{x}$?
- Can also define two distinct functions by changing domain and range, e.g. if $f : \mathbb{Z} \to \mathbb{Z}$ is $f(x) = x$.
- Oftentimes infer domain and co-domain from context.
- Linear functions: $y = mx + b$. We love fitting lines to things!
- Polynomial functions, e.g. $y = ax^2 + bx + c$. Degree of polynomial is highest exponent.
- Increasing function: $x > y \to f(x) > f(y)$.
- Decreasing function: opposite.

- Can have strictly increasing and decreasing, or nonstrict.
- Monotonicity means the function is either nondecreasing (i.e. is monotonically increasing) or nonincreasing (i.e. is montonically decreasing).
- Monotone transformations preserve order.
- Local minima and maxima: construct an ϵ-ball such that the value is highest or lowest in that ball.
- Global minima and maxima are if the value is the highest in the image.
- Continuity. Smoothness of function.
- Talk about factoring and graphing functions. Intercepts, etc.
- $g(x)$ is an inverse function of $f(x)$ if and only if $f(g(x)) = x$.
- Inverse functions will sometimes be denoted $f^{-1}(x)$.
- Exists only if function is one-to-one.
- Functions that are one-to-one and "onto" (where onto means that every element of y is paired with an element of x) are called bijections.
- For instance, $\sqrt{x}$ is not onto.
- Composite functions, e.g.: $f(x) = g(x) + h(x)$
- Asymptotes and convergence: e.g. $\frac{1}{1+x}$.
- Maxima and minima at boundaries.

5.7 Introduction to Differential Calculus (single variable)

- Quick review of limits.
- As before with asymptote, e.g. $\frac{1}{1+x}$. The limit as $x \to \infty$ is defined, but we obviously cannot evalute "at" infinity.
- Not all limits are defined. May increase without bound for example, e.g. $f(x) = x$.
- Slopes: $m = \frac{y_2-y_1}{x_2-x_1}$, $y = mx + b$. Standard grade school stuff.
- This helps us determine the rate of change over some interval.
- What about when the function is nonlinear?
- Secant lines give us average rates of change over some interval. Example of distance traveled and speed.
- What if we want to know the rate of change at a particular point (e.g. instantaneous speed)?
- Find the value as the secant interval tends to zero. This is the tangent line. Slightly more formally...
- $\lim_{h\to 0} \frac{f(x+h)-f(x)}{h}$
- Show example with x^2. Class works through x^3.
- Power rule: $(x^n)' = nx^{n-1}$.
- Sum rule: $(f(x) + g(x))' = f'(x) + g'(x)$.

- Product rule: derivative of $f(x)g(x)$ is $f'(x)g(x) + f(x)g'(x)$.
- Constant rule: If a is a constant, $(a)' = 0$
- Derivative of $ln(x)$, $(ln(x))' = \frac{1}{x}$. Derivative of e^x, $(e^x)' = e^x$.
- Chain Rule: Derivative for $f(g(x)) = f'(g(x))g'(x)$.
- Lots of other rules; review on your own.
- Practice: $f(x) = ln(x^{234\pi})$. What is $f'(x)$?
- Different ways of writing derivatives: $f'(x) = f_x(x) = \frac{dy}{dx}$
- Higher order derivatives, e.g. $f''(x)$. Rate of change of rate of change.
- Some functions are not differentiable, or are not differentiable at particular points. These functions are not "well-behaved".

5.8 Using Calculus to Analyze Functions

- Strictly increasing: $\forall x,\ f'(x) > 0$
- Concavity and convexity.
- Concave: any secant line is below f. Formally, for $\lambda \in [0, 1]$, $f(\lambda x' + (1-\lambda)x) \geq \lambda f(x') + (1-\lambda)f(x)$
- Convexity: any secant line between two points on function is above f.
- For $\lambda \in [0, 1]$, $f(\lambda x' + (1-\lambda)x) \leq \lambda f(x') + (1-\lambda)f(x)$
- Strict concavity and convexity changes these inequalities to strict.
- Can use to determine whether function "curves up" or "curves down".
- More generally: a convex combination is a linear combination of points or numbers such that the "coefficients" (in this case, the λs) sum to 1.
- Convex sets contain all convex combinations of points within the set.
- Convex hull is the smallest convex set that contains a set of points.

5.9 Superficial introduction to optimization

- What happens when we want to find where a function reaches a maximum or minimum?
- Finding critical points (often maxima and minima).
- In general, this is one of the biggest things one does in economics, political science, statistics, engineering, industrial farming... take derivatives and set them equal to zero.
- Look at function $-x^2 + 4x$. Derivative equals zero when tangent line is horizontal, which in this case corresponds to the maximum.
- What about $x^2 - 4x$?
- How do we determine if maximum or minimum? Intuition versus math. Show instances where it is neither a minimum or maximum.

- If function is concave at critical point (i.e. $f''(x) < 0$) then we have found a maximum. If it is convex $f''(x) > 0$ we have found minimum.
- Furthermore, if the function is concave everywhere (i.e. $f''(x) < 0 \; \forall x \in D$, where D is the domain), then any maxima are also *global* maxima. If this relationship is strict, then the global maximum is unique.
- If $f''(x)$ is undefined or zero, we do not know whether it is a maximum or a minimum or neither. These are called saddle points.
- A simple least squares problem:

$$
\begin{aligned}
f(x) &= (x_1 - c)^2 + (x_2 - c)^2 + (x_3 - c)^2 \\
&= x_1^2 + x_2^2 + x_3^2 + 3c^2 - 2cx_1 - 2cx_2 - 2cx_3 \\
\leftrightarrow \frac{df}{dc} &= -2x_1 - 2x_2 - 2x_3 + 6c = 0 \\
\leftrightarrow 3c &= x_1 + x_2 + x_3 \\
\leftrightarrow c &= \frac{x_1 + x_2 + x_3}{3}
\end{aligned}
$$

5.10 Brief introduction to multivariable calculus

- Take derivative with respect to a particular variable in a function. Treat other variables as constants.
- E.g. $f(x, y) = x^2 + y^2 - xy$. $\frac{\partial f}{\partial x} = 2x - y$, $\frac{\partial f}{\partial y} = 2y - x$
- Can also take cross-partial derivatives. This shows how a rate of change changes with another variable.
- Application to interaction effects. Basic partial differentiation is used incorrectly enough that Brambor, Clark, and Golder wrote the most highly cited paper in Political Analysis (2346 citations and counting!) by explaining these.
- In their analysis of APSR, AJPS, and JOP articles from 1998 to 2002, they found that 62% of paper that included interactions did not interpret them correctly, i.e. interpreted coefficients as unconditional marginal effects (which again, is only true with linear model).
- Think of a model $y = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \beta_3 x_1 x_2$. What is the effect of x_1 on y?
- It's $\frac{\partial y}{\partial x_1}$. In linear model, we have special case where it's equal to β_1.
- What about optimization when we have a function of many variables?
- Consider $-x^2 + 4x - y^2 + 6x$. Take partial derivatives and set each to zero. What happens if we're below or above the solution values?
- Consider function $x^2 + y^2 - xy - 3y$. Take partial derivatives and set each to zero. Is this a minimum or maximum?
- **Hot take:** Most of the social sciences, statistics, engineering, and some parts of agriculture is taking derivatives and setting them equal to zero.
- This is because setting derivatives equal to zero is the core basis for most of **optimization theory**, which underlies everything from optimizing behavior, to minimizing least squares, to maximizing likelihoods, to choosing the right amount of concrete or irrigation.

- To speak a little more formally about the multivariate case: the partial derivatives of a function $f : \mathbb{R}^n \to \mathbb{R}$ can be organized into what's called a gradient vector, often represented with ∇f, i.e.:

$$\nabla f = \begin{pmatrix} \frac{\partial f}{\partial x_1} \\ \frac{\partial f}{\partial x_2} \\ \vdots \\ \frac{\partial f}{\partial x_n} \end{pmatrix}$$

- So the relevant condition for an (interior) maximum is now $\nabla f = \mathbf{0}$. This is a "critical point".
- We have analogous conditions to the "second derivative" test to check for a maximum that involve the **Hessian** of this function, i.e. the matrix of second-derivatives.
- We will not talk about Hessians, both because it would take more time than I want to allocate to it, *and* because we can sometimes simply impose conditions on preferences that allow us to assume the critical point is a global maximum (e.g. strict concavity in each variable and additive separability).

5.11 Brief introduction to integral calculus

- Indefinite integrals are antiderivatives. Go from $f'(x)$ to $f(x)$
- New power rule: $\int x^n = \frac{x^{n+1}}{n+1}$
- Definite integral $\int_a^b f'(x)dx = f(b) - f(a)$. Often written as $\int_a^b f(x)dx = F(b) - F(a)$
- This is a Riemann integral, it is the limit as the rectangles' widths go to zero.
- Formally $lim_{\Delta \to 0} \Sigma_i^N f(x_i)\Delta = \int_a^b f'(x)dx = f(b) - f(a)$
- Useful for lots of things: in statistics, particularly continous distributions.
- Practice problem: Find area under x^2 for $x \in [0, 3]$.
- Practice problem 2: Find area under $x^3 + x$ for $x \in [1, 2]$
- Multiple integrals! Can just evaluate these independently. Also, order of integration is not important for definite integrals (Fubini's Theorem).
- Practice problem: $\int_0^1 \int_0^1 x^2 y \; dxdy$

5.12 Sets, Probability, and Statistics

- Sample space is the set of all possible outcomes. Denote this S.
- An "event" is some subset of outcomes. Denote this $A \subseteq S$
- What is the sample space of rolling a die? How about rolling two dice? Write the ordered pairs $S = \{(1,1), (1,2), ..., (6,6)\}$.
- Say we wanted to define the event of rolling a six. What subset $A \subseteq S$ represents this event?

5.13 Algebra of Sets: Intersections and Unions

- Definition: The complement of a set A (denoted A^c) is the set of all elements of S that do not belong to A.
- In terms of events, this is when event A did not happen. Return to one die - what is the complement of rolling a six?
- Definition: The intersection of A and B, denoted $A \cap B$, is the set of all elements that belong to *both* A and B.
- Definition: The union of A and B, denoted $A \cup B$, is the set of all elements that belong to *either* A or B.
- Return to two dice. Let A be the event that the two dice add to 5. Let B be the event that both die are even numbers. What is $A \cup B$? What is $A \cap B$?
- Disjoint/mutually exclusive iff $A \cap B = \emptyset$
- Unions of multiple sets $\bigcup_{i=1}^n A_i$, intersections of multiple sets $\bigcap_{i=1}^n A_i$.
- Examples with dice.
- Countable versus uncountable sample spaces?
- **Partition:** A collection of nonempty subsets such that each element in the set is contained in one and only one subset.
- Note: this implies (1) the subsets are disjoint; (2) they "cover" the set, in some sense.
- Example: For number line $[0, 1]$, a partition would be $\{[0, 0.5), [0.5, 1]\}$.
- $\{[0, 0.5], (0.5, 1]\}$ would not be a partition, because the sets are not disjoint.
- $\{[0, 0.5), (0.5, 1]\}$ would not be a partition, because 0.5 is not in any subset.
- Some other properties of sets are as follows:
  - $A \cup B = B \cup A$
  - $A \cap B = B \cap A$
  - $A \cap (B \cup C) = (A \cap B) \cup (A \cap C)$
  - $A \cup (B \cap C) = (A \cup B) \cap (A \cup C)$
  - $(A \cup B)^c = A^c \cap B^c$
  - $(A \cap B)^c = A^c \cup B^c$
  - For any set B, $A = (A \cap B) \cup (A \cap B^c)$
  - $A \cup B = A \cup (B \cap A^c)$

5.14 Introduction to Probability

- Look to define a probability function that assigns probabilities to events.
- Probability axioms:
  1. For any event A, $P(A) \geq 0$
  2. If Ω is the sample space, $P(\Omega) = 1$ (this is the one that makes it different from other measures)

3. For a series of disjoint events A_i with $i \in \{1, \ldots, n\}$, $P(\bigcup_{i=1}^{n} A_i) = \sum_{i=1}^{n} P(A_i)$

- Every property of probabilities can be proven from these axioms, e.g.:

  1. $P(A^c) = 1 - P(A)$
     Ans: By definition, $\Omega = A \cup A^c$. Since A and A^c are disjoint, $P(A \cup A^c) = P(A) + P(A^c) = P(\Omega) = 1$. Rearrange to get above.
  2. $P(\varnothing) = 0$.
     Ans: Using (1), we have that $P(\varnothing) + P(\varnothing^c) = 1$. Rearrange to get $P(\varnothing) = 1 - P(\Omega) = 0$.
  3. If $A \subset B$, then $P(A) \leq P(B)$
     Ans: First, note that $A = (A \cap B) \cup (A \cap B^c)$ and $B = (B \cap A) \cup (B \cap A^c)$. Now, consider that because $A \subset B$, $A \cap B^c = \varnothing$. Thus $P(A) = P(A \cap B)$ and $P(B) = P((B \cap A) \cup (B \cap A^c))$ which by the union property of probabilities (given that these are disjoint), gives us $P(B) = P(A \cap B) + P(B \cap A^c)$. Since $P(B \cap A^c) \geq 0$ by non-negativity of probabilities, it is the case that $P(B) \geq P(A)$.

- Other properties are as follows.
- $P(A \cup B) = P(A) + P(B) - P(A \cap B)$
- $P(A \cap B) = P(A) + P(B) - P(A \cup B)$
- What does it mean to say that two variables are independent?
- Answer: $P(A \cap B) = P(A)P(B)$
- If they are not independent, then the probability of one depends on whether the other occurs or not. Simplest example involves mutually exclusive events, e.g. $P(B|A) = 0$.
- Bayes Rule used to find conditional probabilities.

5.15 Conditional probability

- Conditional probability is the probability of some event A given that some other event B has already occurred.
- This has the effect of shrinking the sample space.
- Consider a simple example with two die. What's the unconditional probability of rolling a twelve? What is the conditional probability of rolling a twelve, given that your first roll returned a six?
- Written as $P(A|B) = \frac{P(A \cap B)}{P(B)}$
- This can also be written $P(A|B) = \frac{P(A \cap B)}{P(B \cap A) + P(B \cap A^C)}$
- Conditional probability can sometimes be counterintuitive. Consider the Monty Hall problem (from the show Let's Make a Deal).
- Prize is behind randomly selected door of three doors. Monty Hall (who, incidentally, was born in Winnipeg, Manitoba, Canada) would then open one of the other doors to show that there was no prize behind it. The contestant would then be offered the opportunity to switch doors.
- Consider an example case where you choose Door 2. At the point of choosing, you have a one-third chance of it being the correct door, and each other door has a one-third chance
- Then Monty (that knave) opens a door. Consider that the sample space is now (Prize Location, Door Opened)= $\{(1, 3), (2, 1), (2, 3), (3, 1)\}$.

- These respectively have the probabilities $\{1/3, 1/6, 1/6, 1/3\}$. Why?
- Find $P(door2win|door3open)$ and $P(door1win|door3open)$.
- Use Bayes' Theorem for calculating conditional probabilities.
- $P(A|B) = \frac{P(B|A)P(A)}{P(B)} = \frac{P(B|A)P(A)}{P(B|A)P(A)+P(B|\neg A)P(\neg A)} = \frac{P(B|A_j)P(A_j)}{\Sigma_i P(B|A_i)P(A_i)}$
- Law of total probability: $P(A) = \Sigma_i P(A|B_i)P(B_i)$
- Practice example: Say that 1% of the population has cancer. Now say you have a test for cancer that correctly states you have cancer 99% of the time if you have cancer, but returns a false positive 1% of the time even if you don't. You take the test and it says you have cancer. What is the probability that you have cancer?
- $P(cancer|+) = \frac{P(+|cancer)P(cancer)}{P(+|cancer)P(cancer)+P(+|\neg cancer)P(\neg cancer)} = \frac{(1)(0.01)}{(1)(0.01)+(0.01)(0.99)} = 0.5$
- Intuition: Say you randomly selection 100 people from the population. On average, approximately one will have cancer and test positive, and one will not have cancer and test positive. So if you are tested positive, you have a 50/50 chance of being either.
- Independence: $P(A|B) = P(B)$.

5.16 Quick introduction to distributions

- In general, random variables map events in the sample space to real numbers.
- In assigning numerical values to objects in the sample space, can simplify the sample space substantially.
- Quick review: What is a sample space?
- We may want to know what the probability is that this variable will take on certain values, or certain intervals of values. This is known as the variable's distribution.
- We can consider this for discrete or continuous random variables.
- For instance, let's consider coin tosses, where the probability of heads is $1/2$. Let's define the random variable h as the number of heads. Now let's consider a case where we flip five coins.
- Distribution of this random variable h can be written $P(h = K) = p_h(K) = \binom{5}{k}\ p^k(1-p)^{5-k}$ for $k = 0, 1, \ldots, n$.
- This is a special case of what's known as the binomial distribution. Discuss the arbitrary form.
- This is known as the probability distribution function, and for the discrete case, it assigns a probability to each of a finite number of realizations of the random variables.
- In class: find the probability of each realization of the variable.
- The cumulative distribution function is the probability that we achieve any value less than or equal to a particular value. E.g. for the case defined above: $F(h) = \Sigma_{i=0}^{k} \binom{5}{k}\ p^k(1-p)^{5-k}$ where $p = 1/2$
- To find the probability that the random variable falls between two values, say $2 \leq h \leq 4$ use $F(4) - F(2-1)$
- Find cumulative distribution of random variable h.
- Analogous case for continuous random variables, but uses integrals given that the probability of any one value is zero.

- Instead, define probabilities over intervals, e.g. $P([a, b]) = \int_a^b f(t)dt$.
- Given this form, what would the cumulative distribution look like?
- Closing thought; what does it mean to say that variables are independently and identically distributed?

5.17 More on distributions

- What does it mean to say that variables are independently and identically distributed?
- What are some ways that we can connect a linear model to a distribution?
- Say we have a linear model and some cumulative distribution $\phi(x)$ (this could be the standard normal distribution, for instance). We can wrap this distribution around the model by writing it $\phi(X\beta)$
- Models like logit and probit allow us to do this in a way that ensures estimated $\hat{y}$s are between zero and one.
- Integrals of probability distributions from negative infinity to a number? What does this mean? Consider example of uniform distribution: $p(x) = \frac{1}{b-a}$; $\forall x \in [a, b]$, $p(x) = 0$ elsewhere

5.18 Expected Values

- In discrete space: $E(x) = \mu = \Sigma_k k \cdot p_x(k)$
- Example of single die, where we want the expected value of rolling a die.
- What about when we have unequal probabilities? Say we flip a coin and add five to the dice total if it lands on heads?
- Analogous continuous case: $E(Y) = \mu = \int_{-\infty}^{\infty} y \cdot f(y)dy$
- This, as you know, is the "average", which is a measure of central tendency.
- Rules of expectation operator:
  1. $E(a) = a$
  2. $E(bX) = bE(X)$
  3. $E(a + bX) = a + bE(X)$
  4. $\Sigma E(g(X)) = E(\Sigma g(x))$
  5. $E(E(X)) = E(X)$
- Conditional expectation: $E(Y|X)$
- Example: Dice when six has already been rolled. What is conditional expectation of the value?
- Regression function: $E(y|x)$

5.19 Variance and Other Moments

- m^{th} moment of X is $E(X^m)$. m^{th} central moment is $E(X - E(X))^m$
- Variance is second moment.
- Covariance: $E(x - E(x))(y - E(Y))$
- Also can be expressed as $E(X^2) - (E(X))^2$. See proof below.