PS681 - Intermediate Game Theory

Jason S. Davis

Winter 2015

Contents

1 Welcome/Introduction

- Jason Davis - E-mail: jasonsd@umich.edu - Office: Haven Hall 7730 - Telephone: (write in class)
- Most of you know me from all the math.
- If you don't, welcome!
- Plan is to have section Wednesdays from 3-5PM, but I'm open to suggestions.

1.1 Why learn formal theory/why take this course?

- If you want to use formal modeling in your work (obviously).
- If you want to be able to read all the work in your field (all of the subfields have subliteratures that use game theoretic methods).
- If you want to learn a particular "lens" for looking at the world (learning game theory has tranformed the way I think about the world more than learning any particular set of information).
- Disciplining one's thinking by adopting what Skip would/will call 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).
- All of this is true 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").

1.2 What formal theory is and is not

- It is NOT a theory predicated on the dogmatic insistence that actors are truely "rational".
- Rationality is not what many people think it is (most people don't think of transitivity, completeness, and IIA). Can incorporate many things that people don't think of as rational, e.g. regret, hatred, anger, etc.
- E.g. Ultimatum games; experiments don't follow predictions of the theory, but maybe there are "spite" payoffs. Isn't inconsistent with "rationality".
- Moreover, in some cases models are created with "true" irrationality, but the mathematics gets more complicated. Rationality as a well-behaved ordered relation is useful, and while it imposes some restrictions, these aren't as stringent as most think.
- Also: 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, as all models are literally false, so 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.
- Unlike in 598, they may require a bit more logical structure, and you won't be walked through them as much.
- 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 Quick review of formal logic

- Conditionals, biconditions, fallacies (e.g. affirming the consequent).
- Negations.
- 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 $A \rightarrow B$, 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.
- Alternatively, can prove by contrapositive: Assume $\neg B$ and show $\neg A$ as $\neg B \rightarrow \neg A \leftrightarrow A \rightarrow B$.
- 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 $xRy \wedge yRz \rightarrow xRz$. 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).
- 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.
- Skip's convex preferneces and uniqueness example. Strictly convex preferences imply that if there exists a maximum it is unique (you need other assumptions to get existence of a maximum, i.e. compactness and continuity).
- Convex preferences. Often an assumption of models for simplificity. Essentially equivalent to decreasing returns. Would I want 10 apples, 10 carrots, or some convex combination, if we assume I get equivalent utility from 10 apples and 10 oranges?
- Works in cases where one is indifferent between consumption bundles. Obviously might not want convex combination of bundle one like a lot with a bundle one doesn't like at all (I prefer 10 apples to a convex combination of 10 apples and 10 snakes).
- The proof is straightforward by contradiction: assume there are two elements in the maximal set, denoted x and y (note this implies that xIy) Strict convexity means that $(\lambda x + (1-\lambda)y)PxIy$. So this convex combination is strictly preferred to what were supposed to be the maximal elements, which means x and y can't be maximal elements. We've obtained a contradiction, and shown that strict convexity of preferences is inconsistent with multiple elements in the maximal set.

2.4.3 Induction

- Most complicated in terms of logical structure.
- Skip used one of these in class.
  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).

2.5 Preference Relations

- Transitivity and quasi-transitivity are relatively straightforward and can be defined in terms of triples.
- Acyclicity is defined on any finite subset of X (the choice space) i.e. R is acyclic on X if for any finite set $\{x_1, x_2, ..., x_n\} \in X,\ x_i P x_{i+1}\ \forall i < n \rightarrow x_1 R x_n$
- Why can't this be defined in terms of triples?
  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 $xPy, yPz \rightarrow xRz$ and then xRz, zPa, aPx which is not a cycle.
- Skip mentioned acyclicity only prevents "grand cycles", which is true for the logical relationship at hand. However, 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. $xPy \leftrightarrow xRy \wedge \neg yRx$)

2.5.1 Miscellaneous stuff

- Should we have preferences over "paths" instead of just outcomes? Why not just redefine the outcomes? If the "journey is just as important as the destination" then just define outcomes that take into account different journeys (flying first class to Paris is a different consumption bundle than swimming to Paris...)
- Ordinal versus cardinal? Depends on the application.
- Arg max versus max I think we covered fine, unless there are more questions.
- Reduction of compound lotteries: also pretty straightforward, maybe?
- Independence of alternatives w/ steak and cream example. We can redefine the outcomes again; steak & cream is a different good than just steak. If "consumption bundles" are the outcomes, then you can have complementarities without issue.
- Single peakedness versus single crossing. Single crossing is equivalent to projecting multidimensional space onto unidimentional space such that the projection is single-peaked (as I understand it). We can sometimes relax assumptions in a way that allows us to buy back a lot of what we wanted with the original assumptions.

2.5.2 Problem Set Thoughts

- I don't want to walk you through the steps as much as in 598 (particularly for this problem set, laters one may be different). There are often many ways to prove these things, and writing some of your own proofs is a useful exercise.
- 2.2 and 2.3 should be doable.

- 2.1 is harder. Positive Political Theory I by Austen-Smith and Banks is a good references. Don't hurt yourself doing 2.1. An earlier GSI for this course (who's super smart!) wrote a solution to it which is clearly wrong, in that it *assumes* that every subset $S \subseteq X$ has the antecedent properties of acyclicity (i.e. $x_i P x_{i+1}\ \forall i < n$). Keep in mind that acyclicity does not suggest the set HAS to have these properties, only derives implications for when a set does have those properties. Proving that $M(R,S) \neq \emptyset\ \forall S \subset X \rightarrow$ acyclicity is the easy part.

3 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)

3.1 Miscellaneous stuff from lecture

- When should we have discounting?

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

3.2 Problem Set Thoughts

- 3.5 is a little gross.

- Some thoughts: The most important point is that changing the question to be a two period game makes several things written in the original question incorrect. In particular, there are now more than two terminal states: not starting the war is a terminal state, surrendering after one battle leads to a terminal state, and any element of $\{W, L\} \times \{W, L\}$ is a terminal state. You need to compute the expected utility at every non-terminal state. Also, the expected utility at the initial node (i.e. deciding whether to start the war) will depend on whether you would choose to continue fighting after the first round, so you should figure this out first when computing the conditions on f. Lastly, the utility from surrender is zero (plus any costs from battles before surrendering), and the "first" cost f from the first possible battle will be discounted. The only nondiscounted payoff would be the payoff from not starting the war, which is zero.

4 Social Choice Theory

- Can we exclude preferences that are extremely unlikely?

- These kinds of things are discussed by people developing voting systems. Ontario Citizen's Assembly on Electoral Reform.

- Another important social choice theorem is Gibbard-Satterthwaite, which shows that voting rules with certain properties are subject to strategic misrepresentation of preferences (i.e. strategic voting). However, for evaluating systems, we might think that some rules are "more" subject to it than others, given what preference profiles would have to occur for strategic voting, the probability of these profiles occuring in reality, or given the information required in order to vote strategically.

- In general, social choice theorems often show you can't rule something out for all preference profiles, etc. but don't directly address the prevalence of different things.
- Different ways of addressing these kinds of question. Empirics. Ornstein 2013[1] uses a computational model to examine the prevalence of monotonicity failure in IRV elections.

4.1 Arrow's Theorem

- Unrestricted Domain.
  - Convex preferences, single-peakedness, etc. are all restrictions on this!
- Weakly Paretian.
- Independence of Irrelevant Alternatives
  - IIA in social choice theory versus in multinomial choice models.
  - The multinomial choice models were designed in some way to approximate utility function representations, but they're not exactly the same.
  - IIA in social choice theory means that z won't change $x > y$.
  - IIA in statistics is about relative probability of outcomes. The classic example is if the ratio of probability of a blue bus to a car is $1:1$. Now add in a red bus. Shouldn't the ratio of blue bus to car change? If people randomize with equal probability between blue and red buses, you would expect $0.5 : 1 \leftrightarrow 1 : 2$. This violates a statistical interpretation of IIA, but not a social choice theory interpretation of IIA. In fact, with some fairly weak assumptions, maintaining statistical IIA would entail a violation of social choice theory IIA.
- No dictator.
  - If we just use my preferences, and my preferences are rational, then of course the outcome will be complete, transitive, reflexive, IIA, etc. without any restrictions.
  - Moreover, it will be weakly Paretian, given that this is a statement about what happens if everyone prefers x to y, and I am part of everyone. So for the antecedent to be satisfied, everyone has to agree with me that $x > y$, and the outcome will be my preference, i.e. that $x > y$. If not everyone agrees with me, the antecedent will not be satisfied.
  - It's a preference "aggregation" rule in a similar way to how seven is an estimator; it can have a bunch of properties you want (unrestricted domain, transitivity, etc. for a dictator - unbiasedness, etc. for seven) even though it's throwing out a lot of information.

5 Game Theory

5.1 The Concept of the Solution Concept

- We talked about this a little in 598. Any strategy profile is a potential candidate for a equilibrium.
- Solution concepts are ways of restricting our attention to a set of strategy profiles in a systematic fashion. Skip called this a "filter".
- 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 stragies.

[1] Ornstein, Joseph T. Frequency of Monotonicity Failure under Instant Runoff Voting. Public Choice, 2013.

- "Rationalizability" is a particularly weak solution concept (even weaker than Nash. Every Nash is rationalizable but not vice versa). Involves iterated elimination of strictly dominated strategies.
- Follows from assumption of common knowledge of rationality. I know you're rational (eliminate your strictly dominated strategies). You know I know you're rational (so eliminate the strategies that are now dominated for me, given that I've eliminated your 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.
- Because all Nash are rationalizable, you can also use iterated elimination of strictly dominated strategies to simplify the process of looking for Nash equilibrium.
- Usually, Nash tends to be the starting point.
- 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 $S'' = \{(B,C)\}$.

5.2 Actions versus strategies

- Actions versus strategies. Does it ever matter? Ans: Rarely. Games of imperfect recall (i.e. when two sequential nodes are part of the same information set).

- Mixed strategy randomizes pure strategies, where pure strategies are actions for each information set. Doesn't work well for the imperfect recall game.

- "Behavioral strategies" actually involve randomizing actions, but these rarely are necessary.

5.3 Nash Equilibrium

- 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.

5.3.1 Mixed Strategies

- What is the strategy space? Convex hull of pure strategies.

- Strategy space is continuous because mixed strategies smooth things out. Convexifies the choice set.

- May hear of "feasible set" of payoffs. These are the technically feasible payoffs achievable from a combination of strategies.

- Do these exist in the real world?

- 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 $\sigma_A = (Pr(A) = 0.5, Pr(B) = 0.5, Pr(C) = 0)$. Keeping in mind that $EU_2(s_2 = C) = 1$ (i.e. payoff of 1 irrespective of player 1's strategy), $EU_2(s_2 = \sigma_A | s_1 = A) = (0.5)(5) + (0.5)(4) = 4.5 > 1$, $EU_2(s_2 = \sigma_a | s_1 = B) = (0.5)(4) + (0.5)(0) = 2 > 1$, $EU_2(s_2 = \sigma_a | s_1 = A) = (0.5)(0) + (0.5)(4) = 2 > 1$. So σ_A 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	$0, 0$	$2, 2$	$1, 1$

5.3.2 Nash Equilibria Outside Normal Form Games

- The same concept of Nash equilibrium can be applied to situations where the pure strategy choice space is continuous, or games with large numbers of players, i.e. cirucmstances.
- Example from the 2014 Midterm: two players divide a dollar s_1, S_2. What is the set of Nash equilibria? **Ans:** Anything where $s_1 + s_2 = 1$, as in this case, no player has an incentive to deviate.
- Threshold public goods game with five players: $k = 4$ players need to contribute in order for the public good to be provided, and the private benefit provided by the public good is $B > c$, where c is the cost of contributing. What are the Nash equilibria? **Ans:** No-one contributing is Nash, because then no player can get the public good by deviating, given the strategies adopted by the other players. Exactly 4 players contributing is also Nash, as in this case, any player deviating will obtain a strictly worse payoff given the other players' strategies, as they will either lose the public good, or contribute without increasing their benefit.
- Nash equilibria of two player median voter game (where candidates care only about winning)? What happens with three players? **Ans:** 3 player game does have Nash Equilibrium! Hard to see initially. Consider $s_1 = 0.25, s_2 = 0.4, s_3 = 0.75$. Player 3 wins with certainty, no other player can win by changing their position.

5.3.3 Miscellaneous

- What if you just "preferred" to get zero utility because you wanted to be lazy? You factor this into the payoffs. "0" doesn't necessarily mean getting a zero on the test, or whatever.
- It's like the whole ultimatum game experiments. On one level, the results seem to suggest that people do not act as the game theory would expect. On the other hand, the monetary outcomes probably don't directly correspond to their payoffs if, for instance, they get utility from "spite".

5.4 Extensive Form Games

5.4.1 Subgame Perfection

- 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.[2]
- 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?

[2] Morrow 1994

- Answer: Because the equilibrium *depends* on the strategies of the equilibrium path.
- 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:

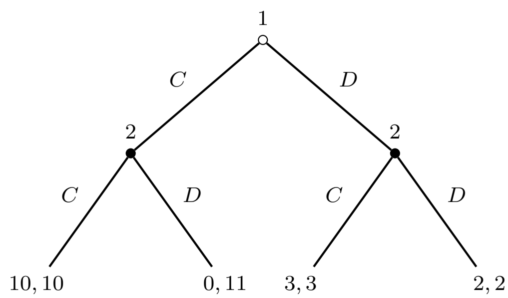

- In the above case, the only subgame perfect Nash equilibrium is $(\sigma_1, \sigma_2) = (D, DC)$.
- Note that although C is not played by player 1 in equilibrium, it is important to specify that Player 2 would play D if Player 1 played C in order for this equilibrium to hold. If we instead had the 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.

5.4.2 Information Sets

- Can only make one choice at any given information set, as you don't know which node you're at.
- Games of imperfect recall are an interesting implication of this (see actions versus strategies).
- How might we represent simultaneous games in extensive form?
- Example: Prisoner's dilemma in extensive form.

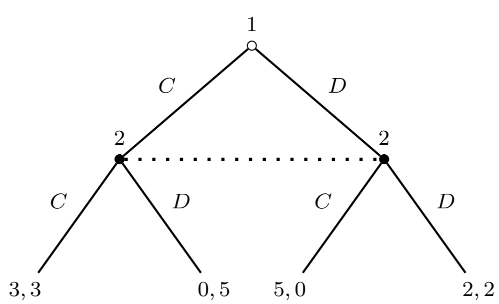

5.4.3 Example Questions

- k threshold public goods game in sequential form.
- Above is very similar to the bill voting game talked about in class.

- Here's a more 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.

5.5 Bayesian Nash equilibrium

- Here, we are dealing with uncertainty, but it should be noted we are not yet in a world where "signalling" occurs, as we are talking about simultaneous games instead of talking about sequential/dynamic games.
- Players may have different types, with common priors over the distribution of those types.
- A Bayesian game may include instances where some subset of players have information revealed to them, i.e. they get private information.
- In each case, 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 to different strategies.

- We'll do some simple examples. For instance, consider the following structure of game, that Skip likes:
  - Nature determines whether the payoffs are as in Game 1 or Game 2.
  - Player 1 (the row player) learns whether nature has drawn Game 1 or Game 2, but player 2 does not.
  - Player 1 chooses either T or B and Player 2 chooses either L or R.
  - 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:

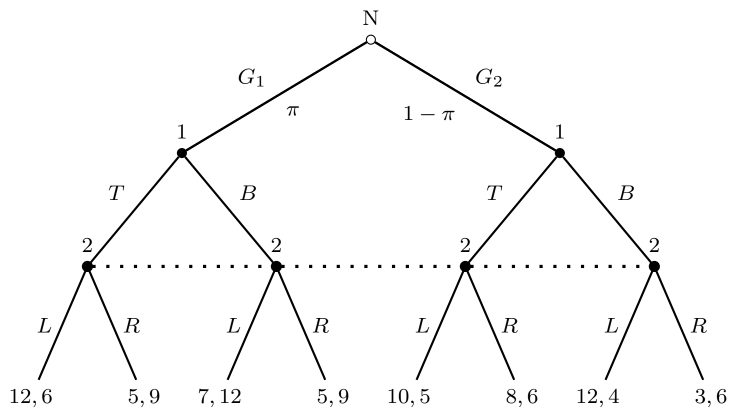

- To find equilibria, we need to check possible strategy profiles, where a strategy must specify an action at every information set.
- Player 1 has two information sets (see above) while player 2 has one information set.
- Let's assume $\pi = 1 - \pi = 0.5$ and solve the problem.
- I start by reducing 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.
- 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.
- 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)

5.5.1 Jury Voting

- In general, you solve the same way as you have with other applications of Bayesian Nash equilibria as a solution concept. Specify a strategy profile, and then compute expected utilities given that strategy profile to determine whether or not any player has an incentive to deviate.

- Individual strategies in this case include always vote guilty, always vote to acquit, or vote sincerely. Each player has one information set, i.e. does not observe other signals.

- Second question which applies to continuous signals is basically just an application of Bayes' rule.

- Note that it contains an error: $n-q-1$ should instead be $n-q$, otherwise it will not add up properly.

5.5.2 Palfrey Rosenthal k-contributions public goods game

- Let's look first at the game in the non-Bayesian (complete and perfect information) context.

- This is just a threshold public goods game. Each player's cost to contributing is less than their benefit, but they would prefer to get the benefit without contributing.

- Pure strategy Nash equilibria to this are straightforward. Q: What are they? **Ans**: k contribute or no-one contributes.

- Mixed strategy Nash equilibria (which MM say are "more compelling" without systematic justification of this claim...) will be at the point where each player is indifferent between contributing and not contributing. Note: there is an error in the book where it writes the payoff to contributing. In the first part of the following equation, it multiplies by zero instead of $-c$.
  Expected utility to contributing:

$$\begin{aligned} EU(c) &= Pr(x_{-i} < k-1) \cdot (-c) + Pr(x_{-i} \geq k-1) \cdot (1-c) \\ &= Pr(x_{-i} \geq k-1) - c(Pr(x_{-i} < k-1) + Pr(x_{-i} \geq k-1)) \\ &= Pr(x_{-i} \geq k-1) - c \end{aligned}$$

  Expected utility to not contributing:

$$EU(\neg c) = Pr(x_{-i} < k) \cdot 0 + Pr(x_{-i} \geq k) \cdot (1) = Pr(x_{-i} \geq k)$$

- Note what's different here. If you contribute, you decrease the likelihood of not getting the public good somewhat, but get $-c$ no matter what.

- To find point of indifference, set these equal to each other:

$$\begin{aligned} EU(c) &= EU(\neg c) \\ Pr(x_{-i} \geq k-1) - c &= Pr(x_{-i} \geq k) \\ \leftrightarrow Pr(x_{-i} \geq k-1) - Pr(x_{-i} \geq k) &= c \\ \leftrightarrow Pr(x_{-i} = k-1) &= c \end{aligned}$$

- Note the intuition here: any player only changes the outcome if they are pivotal, as this is the only case in which they get the public good when they would not have received it otherwise. Thus, the probability of this event, multiplied by the benefit (1), is the expected benefit which needs to be equal to the cost c.
- Now note that if the mixed strategy is symmetric (i.e. every player contributes with the same probability) we can rewrite the above by plugging in the formula for binomial probability: $\binom{n-1}{k-1}\sigma^{k-1}(1-\sigma)^{n-k} = c$
- Which implicity characterizes symmetric mixed strateg(ies) σ. May be up to two solutions.
- When we modify this for the Bayesian form with contributions drawn from $U\ [0,1]$ note that the contribution amount is *also* the probability, as, say, the probability that you draw a contribution cost less than 0.4 is just $Pr(c_i \in [0, 0.4]) = 0.4$.
- As noted in the book, it leads to the analogous condition: $\binom{n-1}{k-1}\hat{c}_n^{k-1}(1-\hat{c}_n)^{n-k} = \hat{c}_n$

5.6 Dynamic Games with Incomplete Information

- Incomplete versus imperfect information.
- Incomplete is where a player does not know the payoffs to the other player.
- Harsanyi transformation makes information complete but imperfect; "Nature" determines the type of the player with some probability.
- If this is immediately revealed to that player, one can imagine the common prior as common knowlege about what the other player believes.

5.6.1 Signalling games

- Below are two game trees that represent the same basic signalling 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.

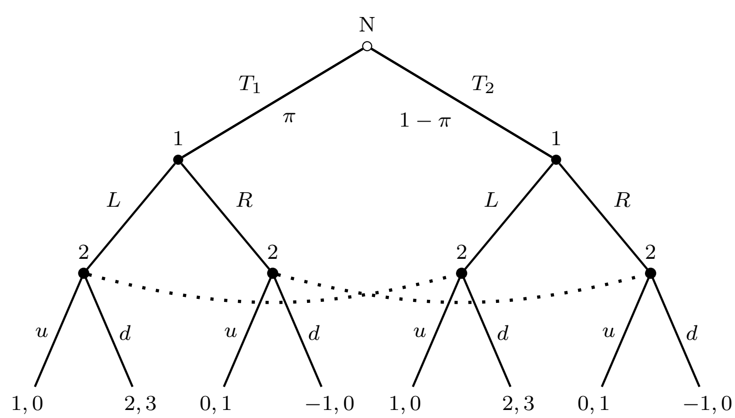

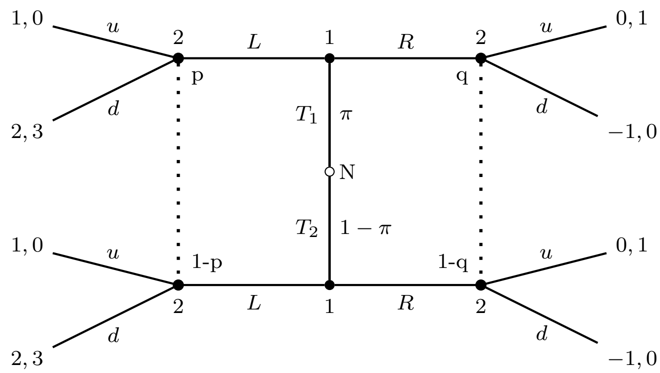

- What are Perfect Bayesian Nash Equilibria (PBNE) to these above game?
- If L is chosen, d dominates u for P2 irrespective of P1's type. Similarly, u dominates d if R is chosen. Thus, the equilibrium is is $\sigma = (L|T_1, L|T_2, u|L, d|R, p = \pi)$. Not super interesting. Let's do another, randomly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

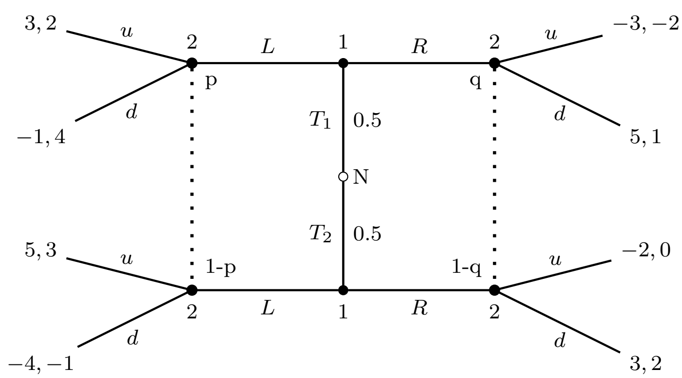

- If R is chosen, d dominates u. If L is chosen, there is no dominant strategy.
- To limit the set of strategy profiles we need to check, I first choose R and L (to determine separating or pooling) and then choose $P2$'s strategies such that they are a best response. Then I examine whether $P1$ would 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.

5.7 Repeated Games

- 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?

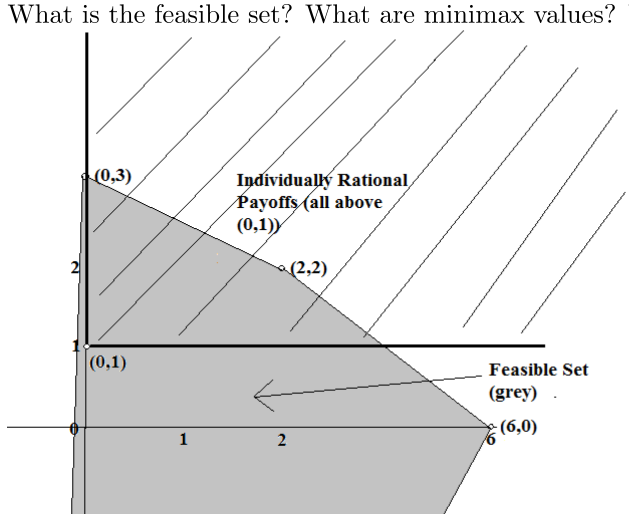

- Another folk theorem: any feasible and individually rational payoff vector can be obtained in a Nash equilibrium of a repeated game, given sufficiently "patient" players (i.e. high enough δ).

- 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?

5.8 Bargaining Theory

- Ultimatum games.
- Finite alternating offers game.
- With discounted payoffs? Give an example with three periods and discount rates $\delta_1 = 0.9, \delta_2 = 0.8$, i.e. different discount rates for each player.
- Nash bargaining solution. Maximize $(U_1(b) - U_1(\neg b))(U_2(b) - U_2(\neg b))$ for some bargain b. For any given surplus over reversion values $(U_i(\neg b))$, the amount that maximizes this expression will divide the surplus equally.
- This is a normative concept, and isn't rooted in players optimizing or equilibrium or whatnot.
- Tends to be used when there's a bargaining process in your game, but you're not particularly concerned about what the division of the surplus is. Can just appeal to the Nash bargaining solution to get some outcome and pick up from there, acknowledge that there are virtually infinity ways the bargaining process could actually be modeled.