Decisions, Games, and Social Choice

Johns Hopkins University, Spring 2016

Housekeeping Details

Instructor	Justin Bledin Assistant Professor Philosophy Department Gilman 206 `jbledin@jhu.edu`
Office Hours	F 11:00am-12:00pm & by appt
Teaching Assistants	John Brandau Ph.D. Student Philosophy Department `john.brandau@gmail.com`
	David Jacobs Ph.D. Student Philosophy Department `djacobs83@gmail.com`
Office Hours	W 12:00am-1:00pm Gilman 269A
Class Code	AS.150.330
Lecture Time	TTh 3:00pm-4:15pm
Lecture Location	Gilman 50
Section Time	W 4:30pm-5:30pm
Section Location	Shaffer 100

Course Description

We investigate rational decision making at the individual and group level. In the first section of the course on *decision theory*, we consider how a single agent ought to act in a choice situation given her knowledge, or lack thereof, about the world and her particular risk profile. In the second section on *game theory*, we explore different kinds of competitive and cooperative strategic interactions between agents, and we define different kinds of solutions, or equilibria, of these games. We also apply game theory to the study of morality, convention, and the social contract. In the final section of the course on *social choice theory*, we turn to group decision making. Specifically, we discuss impossibility results by Arrow and Sen.

Prerequisites

While no prior knowledge of decision/game/social choice theory is required, students should be comfortable with mathematical formalism, probability, and basic methods of mathematical proof.

Readings

The following texts are required:

- Lara Buchak. *Risk and Rationality.* Oxford University Press, Oxford, 2013.
- Andy Egan. Some Counterexamples to Causal Decision Theory. *Philosophical Review*, 116(1):93–114, 2007.
- David Gauthier. Why Contractarianism? In Peter Vallentyne, editor, *Contractarianism and Rational Choice: Essays on David Gauthier's Morals by Agreement*, pages 15–30. Cambridge University Press, 1991.
- David Lewis. *Convention: A Philosophical Study.* Harvard University Press, Cambridge, 1969.
- David Lewis. Prisoners' Dilemma is a Newcomb Problem. *Philosophy and Public Affairs*, 8(3):235–240, 1979.
- Michael D. Resnik. *Choices: An Introduction to Decision Theory.* University of Minnesota Press, Minneapolis, 1987.
- Amartya Sen. The Impossibility of a Paretian Liberal. *Journal of Political Economy*, 78(1):152–157, 1970.
- Amartya Sen. Rationality and Social Choice. *American Economic Review*, 85(1):1–24, 1995.
- Brian Skyrms. Evolution and the Social Contract. In Grethe B. Peterson, editor, *The Tanner Lectures on Human Values 28*, pages 47–69. University of Utah Press, 2009.
- Holly Smith. Deriving Morality From Rationality. In Peter Vallentyne, editor, *Contractarianism and Rational Choice: Essays on David Gauthier's Morals by Agreement*, pages 229–253. Cambridge University Press, 1991.

Both *Choices* and *Convention* should be purchased from the campus bookstore. The other required texts will be posted on Blackboard.

The following textbooks are recommended for supplementary reading:

- Martin J. Osborne and Ariel Rubinstein. *A Course in Game Theory.* MIT Press, Cambridge, 1994.
- Martin Peterson. *An Introduction to Decision Theory.* Cambridge University Press, Cambridge, 2009.

The Osborne and Rubinstein textbook is pitched at the graduate level and should be consulted by students looking for a technically sophisticated treatment of game theory. The Peterson textbook covers much of the course material in a non-technical, accessible manner and should be consulted by students who are having difficulty grasping the underlying concepts introduced in the course.

Schedule

The following schedule projects the lectures over the course of the semester. It is subject to revision as the semester progresses.

Introduction		Jan 28
1. Decision Theory		
1.1 Setup	Resnik 1987, Ch. 1	Feb 2
1.2 Decisions Under Ignorance	Resnik 1987, Ch. 2	Feb 2, Feb 4
1.3 Decisions Under Risk	Resnik 1987, Ch. 3 and 4.1-4.3	Feb 9 & 11
1.4 Paradoxes and Problems	Resnik 1987, 4.4 and 4.5	Feb 16
1.5 Causal vs. Evidential Decision Theory	Resnik 1987, 4.6 / Egan 2007	Feb 18
1.6 Risk-Weighted Expected Utility	Buchak 2013, Ch. 1 and 2	Feb 23
2. Game Theory		
2.1 Zero Sum Games	Resnik 1987, 5.1-5.3	Feb 25, Mar 1
2.2 Nonzero Sum Games	Resnik 1987, 5.4 / Lewis 1979	Mar 3 & 8
2.3 Extensive Games	N/A	Mar 22 & 24
2.4 Cooperative Games	Resnik 1987, 5.5	Mar 29 & 31
2.5 Application: Morality	Gauthier 1991 / Smith 1991	Apr 5 & 7
2.6 Application: Convention	Lewis 1979	Apr 7, 12 & 14
2.7 Application: Social Contract	Skyrms 2009	Apr 19
3. Social Choice Theory		
3.1 Arrow's Theorem	Resnik 1987, 6.1 and 6.2	Apr 21 & 26
3.2 Sen on Liberalism	Sen 1970 / Sen 1995	Apr 28
Conclusion		Apr 28

Requirements

There are three requirements for taking this course.

The first requirement is to complete a series of exercise sets (worth 30% of your final grade). These exercise sets will be assigned roughly every two weeks and you will have a week to complete them.

The second requirement is to take a midterm exam on March 10 (worth 30% of your final grade).

The third requirement is to take a final exam on May 12 between 9:00am-12:00pm (worth 40% of your final grade).

Since we will cover a lot of ground this semester and much of the course material is cumulative, attendance is strongly encouraged.

Academic Integrity

Please do not cheat. This would be depressing. Cheating hurts the Johns Hopkins community by undermining academic and personal integrity, creating mistrust, and fostering unfair competition. Ethical violations include cribbing on exams, plagiarism, reuse of assignments, improper use of the internet and electronic devices, unauthorized collaboration, and alteration of graded assignments. Cheaters may receive a grade of F in the course and can face direr consequences in extreme cases.

Report any violations you witness to Justin, John, or David. You may consult the associate dean of student affairs and/or the chairman of the Ethics Board beforehand. For more information, see the guide *Academic Ethics for Undergraduates* and the Ethics Board website `www.jhu.edu/ethics`.

Disability Accommodations

If you are a student with a disability or believe that you might have a disability that requires special accommodations, please contact Student Disability Services to obtain a letter from a specialist: `studentdisabilityservices@jhu.edu`. Then pass on this letter to Justin, John, or David.

Enjoy the course!

Decisions, Games, and Social Choice

Johns Hopkins University, Spring 2016

Instructor	Justin Bledin Assistant Professor Philosophy Department Gilman 206 `jbledin@jhu.edu`
Office Hours	TBD

Teaching Assistants	John Brandau Ph.D. Student Philosophy Department TBD `john.brandau@gmail.com`
Office Hours	TBD
	David Jacobs Ph.D. Student Philosophy Department TBD `djacobs83@gmail.com`
Office Hours	TBD

Prerequisites

While no prior knowledge of decision/game/social choice theory is required, you should be comfortable with mathematical formalism, probability, and basic proof methods.

Readings

The following texts should be purchased from the campus bookstore:

- Michael D. Resnik. *Choices: An Introduction to Decision Theory*. University of Minnesota Press, Minneapolis, 1987.
- David Lewis. *Convention: A Philosophical Study*. Harvard University Press, Cambridge, 1969.

The other required texts will be posted on Blackboard.

Schedule

Introduction	Jan 28
1. Decision Theory	
1.1 Setup	Feb 2
1.2 Decisions Under Ignorance	Feb 2, Feb 4
1.3 Decisions Under Risk	Feb 9 & 11
1.4 Paradoxes and Problems	Feb 16
1.5 Causal vs. Evidential Decision Theory	Feb 18
1.6 Risk-Weighted Expected Utility	Feb 23

Schedule

2. Game Theory

2.1 Zero Sum Games	Feb 25, Mar 1
2.2 Nonzero Sum Games	Mar 3 & 8
2.3 Extensive Games	Mar 22 & 24
2.4 Cooperative Games	Mar 29 & 31
2.5 Application: Morality	Apr 5 & 7
2.6 Application: Convention	Apr 7, 12 & 14
2.7 Application: Social Contract	Apr 19

Schedule

3. Social Choice Theory

3.1 Arrow's Theorem	Apr 21 & 26
3.2 Sen on Liberalism	Apr 28

Conclusion

Apr 28

Requirements

There are three requirements for taking this course:

- Exercise sets (30% of final grade)
- Midterm exam on Mar 10 (30% of final grade)
- Final exam on May 12 (40% of final grade)

Academic Integrity

Please do not cheat. This would be depressing. Cheating hurts the Johns Hopkins community by undermining academic and personal integrity, creating mistrust, and fostering unfair competition. Ethical violations include cribbing on exams, plagiarism, reuse of assignments, improper use of the internet and electronic devices, unauthorized collaboration, and alteration of graded assignments. Cheaters may receive a grade of F in the course and can face direr consequences in extreme cases.

Report any violations you witness to Justin, John, or David. You may consult the associate dean of student affairs and/or the chairman of the Ethics Board beforehand. For more information, see the guide *Academic Ethics for Undergraduates* and the Ethics Board website `www.jhu.edu/ethics`.

Disability Accommodations

If you are a student with a disability or believe that you might have a disability that requires special accommodations, please contact Student Disability Services to obtain an accommodation letter from a specialist: `studentdisabilityservices@jhu.edu`. Then pass on this letter to Justin, John, or David.

Decision Theory

- An individual must choose between two or more actions.
- The outcome of each act depends on the state of the individual's external environment.
- The individual might be completely ignorant or quite knowledgable about the environment.
- The individual has preferences over the set of outcomes.
- The individual has a particular risk profile.

Decision Theory

- **Descriptive Decision Theory:** The study of how decisions are actually made by us ordinary mortals with our limited cognitive resources, biases, and other psychological deficiencies.
- **Normative/Evaluative Decision Theory:** The study of what choices *ought* to be made or how decisions *should* be made—that is, the study of decision making by ideal *rational* agents.

Decision Theory

Ex. Oysters.

Driving along a desolate stretch of the Pacific coast, you stop at a small wooden shack by the side of the road with an 'Oysters' sign. Inside, an old man is selling both raw and grilled oysters. You are very hungry, and you prefer eating a fresh raw oyster to a grilled one. But you'd rather eat nothing than eat a spoiled oyster and get sick. Do you eat raw oysters, grilled oysters, or do you just continue down the road?

Decision Theory

Ex. Pharmaceutical Company.

You are the CEO of a pharmaceutical company that has developed an insomnia drug which just received FDA approval. However, there is a 10% probability that the drug has bad side effects that were not detected in the FDA trials, and you are considering whether to run an additional \$1M test to find out if the drug has these effects (should the test reveal the bad side effects, you will not market the drug as planned). If you market the drug and there are no problems, then you stand to make \$5M in sales. But if you market the drug without further testing and it has problems, then you stand to lose \$15M. Do you run the test?

Decision Theory

Ex. Love vs. Work.

You are a graduate student at the University of Wisconsin and have been offered a tenure-track faculty position at Yale, but your romantic partner must remain in Madison for work and is not willing to do long-distance. Your post-graduate job prospects in Madison are not very promising but they might improve in a year or two. Do you take the job at Yale or stay with your partner in Madison?

Decision Theory

Ex. St. Petersburg Game (Bernoulli).

You are given a choice between the following two payouts:

A: \$100.

B: A fair coin is flipped until it lands tails. If the coin lands tails on the first toss, then you receive \$2. If the coin lands tails on the second toss, then you receive \$4. In general, if the coin lands tails on the nth toss, then you receive \$$2^n$.

Do you choose A or B?

Game Theory

- Multiple individuals, or *players*, are engaged in a competitive or cooperative strategic interaction, or *game*.
- Players must choose how to act in pursuit of their own individual objectives.
- These actions might be simultaneously or sequentially performed and agents might have full or only partial information about each others' moves.
- The outcome of the game depends on the actions of all the players (and possibly chance) and each player will take into account their expectations of other players' behavior in deciding how to act.
- Each player has preferences over the set of outcomes.

Game Theory

Ex. Stag Hunt (Rousseau).

You and an acquaintance have gone hunting and you must each decide whether to pursue a large stag or to pursue a small hare. Alone, you can each capture a hare. However, a stag requires two people to capture. If you cooperate and both hunt stag, then you will end up with much meat. But if you hunt stag while your acquaintance hunts hare, then you return home empty-handed. Do you hunt stag or hare?

Game Theory

Ex. Traveller's Dilemma (Basu).

An airline loses your suitcase and the suitcase of your doppelgänger that has the exact same contents. An airline manager separates you and your doppelgänger and asks you both to estimate the value of your lost luggage at no less than \$2 and no more than \$100 which is the maximum that the airline will reimburse you. If you both write down the same number, then the manager will treat this as the true value of your luggage and reimburse you both this amount. But if you write down different numbers, then the manager will treat the lower number as the true value. Moreover, whichever one of you wrote down the lower number will be awarded \$2 extra for your honesty, and whichever one of you wrote down the higher number will have \$2 deducted from your payout. What number do you write down?

Game Theory

Ex. Chain-Store Game (Selten).

You own a chain-store with branches in fifty different U.S. cities. In each city, you face a single potential competitor who might encroach on your turf. If a competitor moves in, you can choose to fight or cooperate. Since fighting is costly, you prefer cooperating to fighting in the face of competition, but your favored outcome is when a potential rival does not move in at all. A potential rival is best off when they move in and you cooperate, but they would rather not compete than fight with you. Now suppose that this all plays out sequentially, city after city, and each potential competitor knows what has happened previously. What is your strategy?

Social Choice Theory

- A heterogeneous group of individuals, or *citizens*, must choose between two or more group actions or policies.
- Each citizen has their own preferences over the set of alternatives and these individual preferences must be amalgamated into a single social preference relation.

Social Choice Theory

Ex. City Council.

The City of Toronto has some extra cash in its municipal budget and the mayor forms a three-person committee to help decide how to spend this money. The options on the table are a new park, a new homeless shelter, and the development of more bike lanes. The committee members have the following preference profiles:

Member 1	Member 2	Member 3
Park	Shelter	Bike Lanes
Shelter	Bike Lanes	Park
Bike Lanes	Park	Shelter

What group preference relation over the options should the committee submit to the mayor?

Decision Theory

1.1 Setup

Johns Hopkins University, Spring 2016

Faced with a real-world or hypothetical choice situation, the first task of a decision theorist is to set out precisely the relevant actions of the decision maker, the relevant states of the worlds, the outcome of each act in each state, and the probability of each outcome (when these probabilities are available).

Decision Under Certainty: A choice situation where you know the outcomes of your actions.

Decision Under Risk: A choice situation in which it is possible to assign probabilities to all of the potential outcomes of each action.

Decision Under Ignorance: A choice situation in which it is impossible to assign probabilities to the potential outcomes of one or more of your actions.

For the time being, we will set up the machinery for analyzing decisions under total ignorance. Later on, we will add probabilities to the picture.

In the Oysters example, the situation can be broken down like this:

- **Actions:** eat raw oysters, eat grilled oysters, eat nothing.
- **States:** oysters are fresh, oysters are unfresh.
- **Outcomes:**

  If you eat raw fresh oysters, then you have a delicious meal.

  If you eat raw unfresh oysters, then you get sick.

  If you eat grilled fresh oysters, then you have a decent meal.

  If you eat grilled unfresh oysters, then you have a decent meal.

  If you eat nothing and the oysters are fresh, then you stay hungry.

  If you eat nothing and the oysters are unfresh, then you stay hungry.

We can explicate this informal choice situation using a formal mathematical model.

Def 1.1.1. A *decision model* $\mathcal{M} = \langle \mathcal{A}, \Omega, \mathcal{O}, g \rangle$ consists of a set of actions $\mathcal{A}$, a set of states Ω, a set of outcomes $\mathcal{O}$, and a function $g : \mathcal{A} \times \Omega \to \mathcal{O}$ that maps each action $a \in \mathcal{A}$ and state $s \in \Omega$ to an outcome $o \in \mathcal{O}$.

For instance, a decision model $\mathcal{M}$ of Oysters consists of:

- $\mathcal{A} = \{\text{eat raw}, \text{eat grilled}, \text{eat nothing}\}$
- $\Omega = \{\text{fresh}, \text{unfresh}\}$
- $\mathcal{O} = \{\text{delicious meal}, \text{sick}, \text{decent meal}, \text{hungry}\}$
- $g(\text{eat raw}, \text{fresh}) = \text{delicious meal}$
  $g(\text{eat raw}, \text{unfresh}) = \text{sick}$
  $g(\text{eat grilled}, \text{fresh}) = \text{decent meal}$
  And so forth.

A decision model can be visualized in a table or graph. For example, the model $\mathcal{M}$ for Oysters corresponds to the following *decision matrix*:

	fresh	unfresh
eat raw	delicious meal	sick
eat grilled	decent meal	decent meal
eat nothing	hungry	hungry

This model also corresponds to the following *decision tree* (where boxes are *decision nodes* that call for action and circles are *chance nodes* where nature 'decides' what happens next):

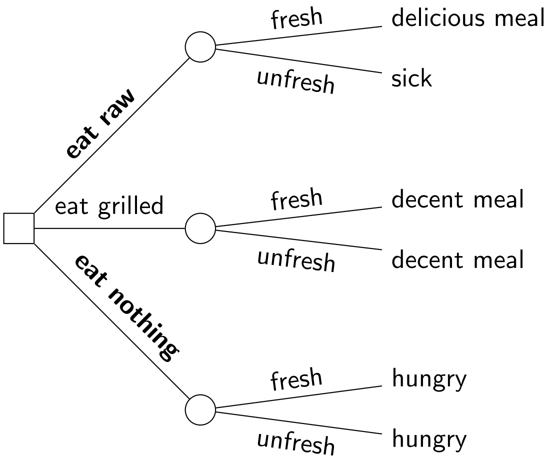

Aside: Decision trees are a particularly nice way to represent a *sequential* choice situation.

Ex. Socks.

In your bedroom, there are two drawers each of which contains exactly one red sock and one blue sock. Getting dressed in the morning, you must first select a sock from the first drawer and then select a sock from the second. You would very much like to wear matching socks. Which socks do you choose?

N.B. Assuming that you can see the color of the socks, this is a decision under certainty.

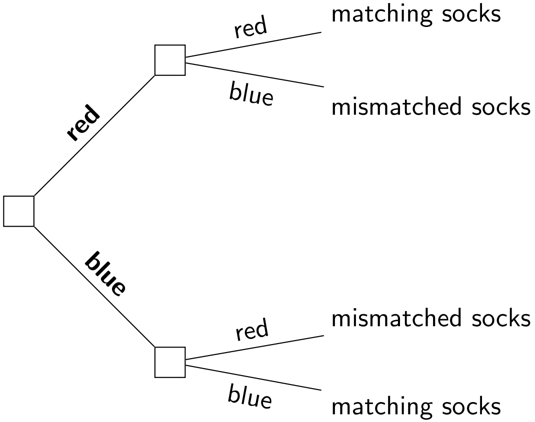
red
blue
red
blue
red
blue
matching socks
mismatched socks
mismatched socks
matching socks

Though the decision-theoretic information for Socks is more naturally expressed using a tree, this information can also be expressed using a matrix (with *strategies* on the left hand side):

	reality (@)
$\langle red, red \rangle$	matching socks
$\langle red, blue \rangle$	mismatched socks
$\langle blue, red \rangle$	mismatched socks
$\langle blue, blue \rangle$	matching socks

In general, we can go back and forth between matrices and trees. These are just two different means for visually organizing the same information in the underlying decision model.

Ex. Gray Day.

You walk out of your house without an umbrella on the way to the nearby bus stop. All of a sudden the sky darkens. If it starts to rain but your bus comes on time, then you will avoid getting too wet. But if your bus is late, then you will get drenched. On the other hand, if you go back inside and grab an umbrella, then you'll have to schlep it around all day. Do you get an umbrella?

Given the intricate nature of many choice situations, it is often a highly non-trivial matter to specify an appropriate formal decision model.

In Pharmaceutical Company, should the outcomes be specified in purely monetary terms? Should they incorporate human suffering and/or the company's goodwill?

In Love vs. Work, what possible states of the world are even relevant?

Unfortunately, there are often many plausible candidate decision models of a particular choice situation, and there is no hard and fast rule for selecting a single model in this set. Moreover, the cost of working with an inappropriate model can be enormous.

Applied decision theory is hard.

Our primary aim in this course is not to learn how to accurately model complex real-life choice situations, but rather to learn how to work with the abstract decision models themselves.

This being said, we will dabble in applied theory. So it is worth briefly mentioning some rough modeling principles.

1. The states in Ω should be mutually exclusive and exhaustive. Also, do not use a state space that is unnecessarily fine-grained.

	extremely fresh	fresh	unfresh
eat raw	delicious meal	delicious meal	sick
eat grilled	decent meal	decent meal	decent meal
eat nothing	hungry	hungry	hungry

You can always add a catch-all 'None of the previous states obtain' to Ω but it might be difficult to associate outcomes with acts in this state.

2. Try to avoid specifying states in Ω that are causally dependent on actions in $\mathcal{A}$, especially when probabilities are unavailable.

Ex. Superbowl.

Who do you think will win the Super Bowl? If you correctly guess that the Denver Broncos will win, then you receive \$100. If you correctly guess that the Carolina Panthers will win, then you receive \$200.

	Denver wins	Carolina wins
bet on Denver	win \$100	win \$0
bet on Carolina	win \$0	win \$200

2. Try to avoid specifying states in Ω that are causally dependent on actions in $\mathcal{A}$, especially when probabilities are unavailable.

Ex. Superbowl.

Who do you think will win the Super Bowl? If you correctly guess that the Denver Broncos will win, then you receive \$100. If you correctly guess that the Seattle Seahawks will win, then you receive \$200.

	You win the bet	You lose the bet
bet on Denver	win \$100	win \$0
bet on Seattle	win \$200	win \$0

In order to use a decision model $\mathcal{M}$ to prescribe or evaluate a particular choice, we must supplement $\mathcal{M}$ with information regarding the decision maker's *preferences*.

Many philosophers think that rationality does not mandate any single preference. Hume: "'Tis not contrary to reason to prefer the destruction of the whole world to the scratching of my finger."

But philosophers think that the preference profile of a rational agent will have a nice structure. In general, if a decision maker's preferences meet certain structural constraints (more on these later in the course), then aspects of her preferences can be usefully exhibited by *utility functions* that associate real numbers with outcomes in $\mathcal{O}$.

Time to introduce preference relations over $\mathcal{O}$.

$o_1 \succ o_2$ designates that o_1 is preferred to o_2.

$o_1 \sim o_2$ designates that o_1 and o_2 are preferred equally.

$o_1 \succcurlyeq o_2$ designates that o_1 is at least as preferred as o_2.

In Oysters, delicious meal $\succ$ decent meal $\succ$ hungry $\succ$ sick.

Note that these relations are interdefinable:

$o_1 \succ o_2$ if and only if $o_1 \succcurlyeq o_2$ and $o_1 \not\sim o_2$.

$o_1 \sim o_2$ if and only if $o_1 \succcurlyeq o_2$ and $o_2 \succcurlyeq o_1$.

$o_1 \succcurlyeq o_2$ if and only if $o_1 \succ o_2$ or $o_1 \sim o_2$.

Def 1.1.2. A function $u : \mathcal{O} \to \mathbb{R}$ is an *ordinal utility function* just in case it satisfies the following conditions:

(i) $u(o_1) > u(o_2)$ if and only if $o_1 \succ o_2$.

(ii) $u(o_1) = u(o_2)$ if and only if $o_1 \sim o_2$.

(iii) $u(o_1) \geq u(o_2)$ if and only if $o_1 \succcurlyeq o_2$.

(**N.B.** condition (i) implies (ii) and (iii).)

An ordinal utility function captures the *ordering* of the agent's preferences over the outcomes in $\mathcal{O}$.

It needn't capture the *strength* or *intensity* of the agent's preferences.

Ordinal utility functions are invariant up to *positive monotone transformations* (functions $t : \mathbb{R} \to \mathbb{R}$ where $t(x) \geq t(y)$ if and only if $x \geq y$). That is, if you apply a positive monotone transformation t to an ordinal utility function u, then the resulting composite function $t \circ u$ is still an ordinal utility function.

For example, the following function $u_1 : \mathcal{O} \to \mathbb{R}$ is an ordinal utility function:

$u_1(\text{delicious meal}) = 3$
$u_1(\text{decent meal}) = 2$
$u_1(\text{hungry}) = 0$
$u_1(\text{sick}) = -10$

Substituting these values for outcomes in the matrix for Oysters yields

	fresh	spoiled
eat raw	3	-10
eat grilled	2	2
eat nothing	0	0

The new function $u_2 : \mathcal{O} \mapsto \mathbb{R}$ obtained from u_1 by applying the positive monotone transformation $t(x) = x^3$ is also an ordinal utility function:

$u_2(\text{delicious meal}) = 27$
$u_2(\text{decent meal}) = 8$
$u_2(\text{hungry}) = 0$
$u_2(\text{sick}) = -1000$

Substituting these values for outcomes in the matrix for Oysters yields

	fresh	spoiled
eat raw	27	-1000
eat grilled	8	8
eat nothing	0	0

Again: ordinal utility functions needn't encode information about the strength of an agent's preferences.

The utility function u_1 suggests that the preference for a decent meal over being hungry is *stronger* than the preference for a delicious meal over a decent meal.

The utility function u_2 suggests that the preference for a decent meal over being hungry is *weaker* than the preference for a delicious meal over a decent meal.

Def 1.1.3. A function $u : \mathcal{O} \rightarrow \mathbb{R}$ is an *interval utility function* just in case $|u(o_1) - u(o_2)|$ reflects the relative strength of the agent's preference between $o_1, o_2 \in \mathcal{O}$.

Interval utility functions are invariant up to *positive linear transformations* (functions $t : \mathbb{R} \rightarrow \mathbb{R}$ where $t(x) = ax + b$ for $a > 0$). That is, if you apply a positive linear transformation t to an interval utility function u, then the resulting composite function $t \circ u$ is still an interval utility function.

Note that interval utility functions are ordinal utility functions but the converse is not true. For an agent's preferences to be representable with an interval utility function, her preference profile must satisfy even more constraints than are necessary for her preferences to be representable with an ordinal utility function.

We are all set up to start recommending and evaluating choices under ignorance. Keep in mind that we are interested in *rational* decisions that needn't lead to the best outcome when all is said and done.

Right Decision: A decision is *right* when there is no alternative action that leads to a better outcome in the actual world.

	fresh (@)	unfresh
eat raw	5	-10
eat grilled	2	2
eat nothing	0	0

Rational Decision: ???

Decision Theory

1.2 Decisions Under Ignorance

Johns Hopkins University, Spring 2016

What should you do in a choice situation where you do not know the probabilities of the outcomes of one or more of your potential actions?

	fresh	unfresh
eat raw	5	-10
eat grilled	2	2
eat nothing	0	0

Def 1.2.1. $a_1 \in \mathcal{A}$ *dominates* $a_2 \in \mathcal{A}$ if and only if $u(g(a_1, s)) \geq u(g(a_2, s))$ for every state $s \in \Omega$ and $u(g(a_1, s)) > u(g(a_2, s))$ for some state $s \in \Omega$.

Thesis 1.2.1 (Dominance Principle). Rationality forbids any dominated action.

The Dominance Principle presupposes an ordinal utility function.

	fresh	unfresh
eat raw	5	-10
eat grilled	2	2
eat nothing	0	0

	fresh	unfresh
eat raw	5	-10
eat grilled	2	2
eat nothing	0	0

	fresh	unfresh
eat raw	5	-10
eat grilled	2	2

	s_1	s_2	s_3	s_4
a_1	1	-2	3	5
a_2	9	-3	-2	15
a_3	10	7	4	5
a_4	8	-3	-3	12

	s_1	s_2	s_3	s_4
a_1	1	-2	3	5
a_2	9	-3	-2	15
a_3	10	7	4	5
a_4	8	-3	-3	12

	s_1	s_2	s_3	s_4
a_2	9	-3	-2	15
a_3	10	7	4	5
a_4	8	-3	-3	12

	s_1	s_2	s_3	s_4
a_2	9	-3	-2	15
a_3	10	7	4	5
a_4	8	-3	-3	12

	s_1	s_2	s_3	s_4
a_2	9	-3	-2	15
a_3	10	7	4	5

As we will see, the Dominance Principle is one of the few uncontroversial theses in decision theory (though it must be applied with caution when actions and states are causally dependent).

However, the Dominance Principle usually doesn't pin down a single action in $\mathcal{A}$. So it must be supplemented with additional principles.

Let $min(a)$ designate the minimum utility obtainable by performing act $a \in \mathcal{A}$. That is, $min(a) = min_{s \in \Omega}(u(g(a, s)))$.

Thesis 1.2.2 (Maximin Rule). Rationality forbids any action $a_i \in \mathcal{A}$ such that $min(a_i) \neq max_{a \in \mathcal{A}}(min(a))$.

In other words, maximize the minimum utility obtainable.

The Maximin Rule presupposes an ordinal utility function.

	s_1	s_2	s_3	s_4
a_2	9	-3	-2	15
a_3	10	7	4	5

	s_1	s_2	s_3	s_4
a_2	9	-3	-2	15
a_3	10	7	4	5

$$min(a_3) = 4 > min(a_2) = -3.$$

	s_1	s_2	s_3	s_4
a_3	10	7	4	5

	s_1	s_2	s_3	s_4
a_1	9	1	0	0
a_2	5	-2	3	5
a_3	0	1	2	20
a_4	-3	10	-5	0

	s_1	s_2	s_3	s_4
a_1	9	1	0	0
a_2	5	-2	3	5
a_3	0	1	2	20
a_4	-3	10	-5	0

$$min(a_1) = min(a_3) = max_{a \in \mathcal{A}}(min(a)) = 0.$$

	s_1	s_2	s_3	s_4
a_1	9	1	0	0
a_3	0	1	2	20

To break ties, we can consider second worst outcomes, third worst outcomes, and so forth.

Let $min^1(a)$ designate the lowest utility obtainable with $a \in \mathcal{A}$.
Let $min^2(a)$ designate the second lowest utility obtainable with $a \in \mathcal{A}$.
Let $min^n(a)$ designate the nth lowest utility obtainable with $a \in \mathcal{A}$.

Thesis 1.2.3 (Lexical Maximin Rule). Rationality forbids any action $a_1 \in \mathcal{A}$ such that there is some $n > 0$ and alternative action $a_2 \in \mathcal{A}$ where $min^n(a_2) > min^n(a_1)$ and $min^m(a_2) = max_{a \in \mathcal{A}}(min^m(a))$ for all $m < n$.

In other words, maximize the worst utility obtainable. But in the case of a tie, maximize the second worst utility obtainable. And so forth.

The Lexical Maximin Rule presupposes an ordinal utility function.

	s_1	s_2	s_3	s_4
a_1	9	1	0	0
a_3	0	1	2	20

	s_1	s_2	s_3	s_4
a_1	9	1	0	0
a_3	0	1	2	20

$$min^1(a_1) = min^1(a_3) = 0.$$

	s_1	s_2	s_3	s_4
a_1	9	1	0	0
a_3	0	1	2	20

$$min^2(a_1) = min^2(a_3) = 1.$$

	s_1	s_2	s_3	s_4
a_1	9	1	0	0
a_3	0	1	2	20

$$min^3(a_1) = 9 > min^3(a_3) = 2.$$

	s_1	s_2	s_3	s_4
a_1	9	1	0	0

Problem. The maximin rules are overly conservative and pessimistic. It seems eminently rational to accept a slightly lower minimum utility for the chance of massive gain.

	s_1	s_2
Gamble A	\$101	\$101
Gamble B	\$100	\$10000

This kind of decision matrix suggests that we should also consider the maximum utility obtainable.

Let $max(a)$ designate the maximum utility obtainable by performing act $a \in \mathcal{A}$. That is, $max(a) = max_{s \in \Omega}(u(g(a, s)))$.

Thesis 1.2.4 (Maximax Rule). Rationality forbids any action $a_1 \in \mathcal{A}$ such that $max(a_1) \neq max_{a \in \mathcal{A}}(max(a))$.

In other words, maximize the maximum utility obtainable.

The Maximax Rule presupposes an ordinal utility function.

	s_1	s_2	s_3	s_4
a_1	5	10	3	0
a_2	20	-2	-50	-10
a_3	0	1	4	7
a_4	-30	20	-5	0

	s_1	s_2	s_3	s_4
a_1	5	10	3	0
a_2	20	-2	-50	-10
a_3	0	1	4	7
a_4	-30	20	-5	0

$$max(a_2) = max(a_4) = max_{a \in \mathcal{A}}(max(a)) = 20.$$

	s_1	s_2	s_3	s_4
a_2	20	-2	-50	-10
a_4	-30	20	-5	0

To break ties, we could also introduce a Lexical Maximax Rule and consider second best outcomes, third best outcomes, and so forth.

The Maximax Rule is quite optimistic. But it has few adherents.

A more flexible rule allows for compromise between the maximum and minimum utility obtainable.

Def 1.2.2. Given an *optimism index* $\alpha \in \mathbb{R}[0, 1]$ that represents the degree of optimism of the decision maker, the α-*index* of $a \in \mathcal{A}$ is $\alpha \times max(a) + (1 - \alpha) \times min(a)$.

Thesis 1.2.5 (Optimism-Pessimism Rule). Rationality forbids any act $a_1 \in \mathcal{A}$ whose α-index is lower than that of some alternative act $a_2 \in \mathcal{A}$.

In other words, maximize the α-index.

When $\alpha = 1$, the OP Rule collapses to the Maximax Rule.

When $\alpha = 0$, the OP Rule collapses to the Maximin Rule.

The Optimism-Pessimism Rule presupposes an interval utility function.

	s_1	s_2	s_3
a_1	10	4	0
a_2	2	5	6

If $\alpha = 0.5$, then the α-index of a_1 is $0.5 \times 10 + 0.5 \times 0 = 5$ and the α-index of a_2 is $0.5 \times 6 + 0.5 \times 2 = 4$. So the OP Rule requires a_1.

If $\alpha = 0.2$, then the α-index of a_1 is $0.2 \times 10 + 0.8 \times 0 = 2$ and the α-index of a_2 is $0.2 \times 6 + 0.8 \times 2 = 2.8$. So the OP Rule requires a_2.

	s_1	s_2	s_3
a_1	8	4	0
a_2	3	5	7

If $\alpha = 0.5$, then the α-index of a_1 is $0.5 \times 8 + 0.5 \times 0 = 4$ and the α-index of a_2 is $0.5 \times 7 + 0.5 \times 3 = 5$. So the OP Rule requires a_2.

The fact that a rule presupposes an interval utility function is sometimes considered a disadvantage.

Problem. It seems rational to consider non-extreme utilities when making a decision.

	s_1	s_2	s_3
Gamble A	\$100	\$1	\$0
Gamble B	\$20	\$79	\$80

A natural response is to assign α-values to non-extreme outcomes.

Problem. Since α is subjective, the Optimism-Pessimism Rule is unstable.

An agent's optimism index can change over time, so the OP Rule imposes no consistency on an agent's decisions over time.

Moreover, we could rationalize impulsive behavior by declaring an appropriate degree of optimism.

The next rule focuses on *missed opportunities*.

Let $max(s)$ designate the maximum utility obtainable in state $s \in \Omega$. That is, $max(s) = max_{a \in \mathcal{A}}(u(g(a, s)))$.

Def 1.2.3. The *regret value* of $a \in \mathcal{A}$ in $s \in \Omega$ is $r(a, s) = max(s) - u(g(a, s))$.

Let $max\text{-}regret(a)$ designate the maximum regret value obtainable by performing act $a \in \mathcal{A}$. That is, $max\text{-}regret(a) = max_{s \in \Omega}(r(a, s))$.

Thesis 1.2.6 (Minimax Regret Rule). Rationality forbids any act $a_1 \in \mathcal{A}$ such that $max\text{-}regret(a_1) \neq min_{a \in \mathcal{A}}(max\text{-}regret(a))$.

In other words, minimize the maximum possible regret.

The Minimax Regret Rule presupposes an interval utility function.

	s_1	s_2	s_3
a_1	5	-2	10
a_2	-1	-1	20
a_3	-3	-1	5
a_4	0	-4	1

Step 1. Identify $\max(s)$ for each $s \in \Omega$.

	s_1	s_2	s_3
a_1	5	-2	10
a_2	-1	-1	20
a_3	-3	-1	5
a_4	0	-4	1

Step 1. Identify $\max(s)$ for each $s \in \Omega$.

	s_1	s_2	s_3
a_1	5	-2	10
a_2	-1	-1	20
a_3	-3	-1	5
a_4	0	-4	1

Step 2. Calculate $r(a, s)$ for each $a \in \mathcal{A}$ and $s \in \Omega$.

	s_1	s_2	s_3
a_1	0	1	10
a_2	6	0	0
a_3	8	0	15
a_4	5	3	19

The resulting matrix is called a *regret table*.

	s_1	s_2	s_3
a_1	0	1	10
a_2	6	0	0
a_3	8	0	15
a_4	5	3	19

Step 3. Identify *max-regret*(a) for each $a \in \mathcal{A}$.

	s_1	s_2	s_3
a_1	0	1	10
a_2	6	0	0
a_3	8	0	15
a_4	5	3	19

Step 3. Identify *max-regret*(a) for each $a \in \mathcal{A}$.

	s_1	s_2	s_3
a_1	0	1	10
a_2	6	0	0
a_3	8	0	15
a_4	5	3	19

Step 4. Perform act that minimizes *max-regret*.

	s_1	s_2	s_3
a_2	6	0	0

To break ties, we could introduce a Lexical Minimax Regret Rule.

Problem. The addition of a non-optimal alternative act can affect what the Minimax Regret Rule mandates.

	s_1	s_2	s_3
a_1	0	10	4
a_2	5	2	10

Decision Table

	s_1	s_2	s_3
a_1	5	0	6
a_2	0	8	0

Regret Table

Problem. The addition of a non-optimal alternative act can affect what the Minimax Regret Rule mandates.

	s_1	s_2	s_3
a_1	0	10	4
a_2	5	2	10

Decision Table

	s_1	s_2	s_3
a_1	5	0	6
a_2	0	8	0

Regret Table

Problem. The addition of a non-optimal alternative act can affect what the Minimax Regret Rule mandates.

	s_1	s_2	s_3
a_1	0	10	4
a_2	5	2	10

Decision Table

	s_1	s_2	s_3
a_1	5	0	6
a_2	0	8	0

Regret Table

Problem. The addition of a non-optimal alternative act can affect what the Minimax Regret Rule mandates.

	s_1	s_2	s_3
a_1	0	10	4
a_2	5	2	10
a_3	10	5	1

Decision Table

	s_1	s_2	s_3
a_1	10	0	6
a_2	5	8	0
a_3	0	5	9

Regret Table

Problem. The addition of a non-optimal alternative act can affect what the Minimax Regret Rule mandates.

	s_1	s_2	s_3
a_1	0	10	4
a_2	5	2	10
a_3	10	5	1

Decision Table

	s_1	s_2	s_3
a_1	10	0	6
a_2	5	8	0
a_3	0	5	9

Regret Table

Problem. The addition of a non-optimal alternative act can affect what the Minimax Regret Rule mandates.

	s_1	s_2	s_3
a_1	0	10	4
a_2	5	2	10
a_3	10	5	1

Decision Table

	s_1	s_2	s_3
a_1	10	0	6
a_2	5	8	0
a_3	0	5	9

Regret Table

The final decision-theoretic principle brings probabilities into the picture.

Guiding idea: If we are in a state of total ignorance, then we should treat all states in Ω as equally likely since there is *insufficient reason* for any other probability measure.

	$[\frac{1}{10}]$ s_1	$[\frac{1}{10}]$ s_2	$[\frac{1}{10}]$ s_3	$[\frac{1}{10}]$ s_4	$[\frac{1}{10}]$ s_5	$[\frac{1}{10}]$ s_6	$[\frac{1}{10}]$ s_7	$[\frac{1}{10}]$ s_8	$[\frac{1}{10}]$ s_9	$[\frac{1}{10}]$ s_{10}
a_1	10	9	8	7	6	5	4	3	2	1
a_2	10	0	0	0	0	0	0	0	0	0

Combining this guiding idea with the principle that we should maximize *expected utility* (more on this later) gives us:

Thesis 1.2.7 (Principle of Insufficient Reason). Rationality forbids any action $a_1 \in \mathcal{A}$ such that there is some alternative action $a_2 \in \mathcal{A}$ where $\sum_{s \in \Omega} \frac{1}{|\Omega|} \times u(g(a_2, s)) > \sum_{s \in \Omega} \frac{1}{|\Omega|} \times u(g(a_1, s))$.

$|\Omega|$ is the number of states in Ω.

The PIR presupposes an interval utility function.

	$[\frac{1}{3}]$	$[\frac{1}{3}]$	$[\frac{1}{3}]$
	s_1	s_2	s_3
a_1	6	3	0
a_2	12	12	9
a_3	0	9	0

The expected utility of a_1 is $\frac{1}{3} \times 6 + \frac{1}{3} \times 3 + \frac{1}{3} \times 0 = 3$.

	$[\frac{1}{3}]$	$[\frac{1}{3}]$	$[\frac{1}{3}]$
	s_1	s_2	s_3
a_1	6	3	0
a_2	12	12	9
a_3	0	9	0

The expected utility of a_2 is $\frac{1}{3} \times 12 + \frac{1}{3} \times 12 + \frac{1}{3} \times 9 = 11$.

	$[\frac{1}{3}]$	$[\frac{1}{3}]$	$[\frac{1}{3}]$
	s_1	s_2	s_3
a_1	6	3	0
a_2	12	12	9
a_3	0	9	0

The expected utility of a_3 is $\frac{1}{3} \times 0 + \frac{1}{3} \times 9 + \frac{1}{3} \times 0 = 3$.

	s_1	s_2	s_3
a_2	12	12	9

Problem. The Principle of Insufficient Reason is extremely sensitive to how states are individuated.

	$[\frac{1}{2}]$ fresh	$[\frac{1}{2}]$ unfresh
eat raw	5	-3
eat grilled	2	2
eat nothing	0	0

The expected utility of eating raw oysters is 1 while the expected utility of eating grilled oysters is 2. Thus, the PIR forbids eating raw oysters.

Problem. The Principle of Insufficient Reason is extremely sensitive to how states are individuated.

	$[\frac{1}{3}]$ extremely fresh	$[\frac{1}{3}]$ fresh	$[\frac{1}{3}]$ unfresh
eat raw	5	5	-3
eat grilled	2	2	2
eat nothing	0	0	0

The expected utility of eating raw oysters is $\frac{7}{3}$ while the expected utility of eating grilled oysters is 2. Thus, the PIR forbids eating grilled oysters.

Problem. There is arguably insufficient reason to adopt *any* probability measure over Ω including the uniform distribution.

Problem. The Principle of Insufficient Reason leads to counterintuitive results.

	s_1	s_2
Gamble A	\$300	-\$100
Gamble B	\$90	\$90

We have now seen 7 different decision-theoretic principles. Unfortunately, these often conflict.

	s_1	s_2	s_3
a_1	1	14	13
a_2	-1	17	11
a_3	0	20	6

The Maximin Rule allows only a_1.

We have now seen 7 different decision-theoretic principles. Unfortunately, these often conflict.

	s_1	s_2	s_3
a_1	1	14	13
a_2	-1	17	11
a_3	0	20	6

The Minimax Regret Rule allows only a_2.

We have now seen 7 different decision-theoretic principles. Unfortunately, these often conflict.

	s_1	s_2	s_3
a_1	1	14	13
a_2	-1	17	11
a_3	0	20	6

The Optimism-Pessimism Rule (with $\alpha = 0.5$) allows only a_3.

How to decide between the 7 principles?

Some decision theorists have tried to provide an *axiomatic* justification of their preferred decision rule. They argue that only their preferred rule satisfies certain reasonable conditions or axioms.

But this just pushes back the debate to the axioms themselves.

Resnik: "The debate could go back and forth over the conditions in this fashion with results no more conclusive than our previous discussions of the rules themselves. The situation is likely to remain this way, I think, until we have amassed a rich backlog of studies of genuine real-life examples of good decision making under ignorance. My hope is that these would help us sharpen our judgments concerning the various rules. I would also conjecture that we will ultimately conclude that no rule is always the rational one to use but rather that different rules are appropriate to different situations. If this is so, it would be more profitable to seek conditions delimiting the applicability of the various rules rather than to seek ones that will declare in favor of a single rule."

Decision Theory

1.3 Decisions Under Risk

Johns Hopkins University, Spring 2016

What should you do in a choice situation where you can assign probabilities to the outcomes of each of your potential actions?

Ex. Pharmaceutical Company.

You are the CEO of a pharmaceutical company that has developed an insomnia drug which just received FDA approval. However, there is a 10% probability that the drug has bad side effects that were not detected in the FDA trials, and you are considering whether to run an additional \$1M test to find out if the drug has these effects (should the test reveal the bad side effects, you will not market the drug as planned). If you market the drug and there are no problems, then you stand to make \$5M in sales. But if you market the drug without further testing and it has problems, then you stand to lose \$15M. Do you run the test?

For the time being, we assume that acts and states are independent, and probabilities can be assigned to each state.

Def 1.3.1. A *decision model* $\mathcal{M} = \langle \mathcal{A}, \Omega, \mathcal{P}, \mathcal{O}, g \rangle$ consists of a set of actions $\mathcal{A}$, a set of states Ω, a *probability measure* $Pr : 2^\Omega \to \mathbb{R}[0,1]$, a set of outcomes $\mathcal{O}$, and a function $g : \mathcal{A} \times \Omega \to \mathcal{O}$ that maps each action $a \in \mathcal{A}$ and state $s \in \Omega$ to an outcome $o \in \mathcal{O}$.

2^Ω is the *power set* of Ω (the set of all sets of states in Ω).

Assuming that your utility is linear in dollars, we can model Pharmaceutical Company as follows:

	$[\frac{1}{10}]$ bad effects	$[\frac{9}{10}]$ no bad effects
test	-1	4
no test	-15	5

Should you run the additional test or not?

Def 1.3.2. The *expected utility* of $a \in \mathcal{A}$ is $\sum_{s \in \Omega} Pr(s) \times u(g(a, s))$.

This will be designated by $EU(a)$.

Thesis 1.3.1 (Principle of Maximizing Expected Utility). Rationality forbids any act $a_1 \in \mathcal{A}$ such that $EU(a_1) < EU(a_2)$ for some alternative act $a_2 \in \mathcal{A}$.

In other words, maximize EU.

The Principle of Maximizing EU presupposes an interval utility function.

	$[\frac{1}{10}]$ bad effects	$[\frac{9}{10}]$ no bad effects
test	-1	4
no test	-15	5

$$EU(\text{test}) = \frac{1}{10} \times -1 + \frac{9}{10} \times 4 = 3.5.$$

	$[\frac{1}{10}]$ bad effects	$[\frac{9}{10}]$ no bad effects
test	-1	4
no test	-15	5

$$EU(\text{no test}) = \frac{1}{10} \times -15 + \frac{9}{10} \times 5 = 3.$$

	$[\frac{1}{10}]$ bad effects	$[\frac{9}{10}]$ no bad effects
test	-1	4
no test	-15	5

$$EU(\text{test}) = 3.5 > EU(\text{no test}) = 3.$$

	$[\frac{1}{10}]$ bad effects	$[\frac{9}{10}]$ no bad effects
test	-1	4

Ex. Monte Hall.

You are a contestant in a game show hosted by Monte Hall. There are three doors 1, 2, and 3. A car has been placed randomly behind one of these doors. Behind the other two doors are goats. Monte explains the rules: “First you will pick a door. Then I will open one of the doors that you did not pick. I know what is behind the doors so I will always reveal a goat. After I show you the goat, you will then have the option of switching your initial choice to the other door that I did not open.” You initially pick A and Monte then opens C to reveal a goat. Do you stick with A or switch to B?

	[?]	[?]
	car behind A	car behind B
stick with A	win car	win goat
switch to B	win goat	win car

	$[\frac{1}{3}]$ car behind A	$[\frac{2}{3}]$ car behind B
stick with A	win car	win goat
switch to B	win goat	win car

	$[\frac{1}{3}]$ car behind A	$[\frac{2}{3}]$ car behind B
stick with A	15	3
switch to B	3	15

$$EU(\text{stick with A}) = \frac{1}{3} \times 15 + \frac{2}{3} \times 3 = 7.$$

	$[\frac{1}{3}]$ car behind A	$[\frac{2}{3}]$ car behind B
stick with A	15	3
switch to B	3	15

$$EU(\text{switch to B}) = \tfrac{1}{3} \times 3 + \tfrac{2}{3} \times 15 = 11.$$

	$[\frac{1}{3}]$ car behind A	$[\frac{2}{3}]$ car behind B
stick with A	15	3
switch to B	3	15

$$EU(\text{switch to B}) = 11 > EU(\text{stick with A}) = 7.$$

	$[\frac{1}{3}]$ car behind A	$[\frac{2}{3}]$ car behind B
switch to B	win goat	win car

Note that the Principle of Maximizing Expected Utility brings together two different ingredients—*viz.*, probability and utility. For a proper treatment of decisions under risk, then, we need a theory of probability and a theory of utility.

Let us begin by theorizing a bit about probability.

The Kolmogorov Probability Axioms:

Ax1. $0 \leq Pr(X) \leq 1$.

Ax2. $Pr(\Omega) = 1$.

Ax3. If $X \cap Y = \emptyset$, then $Pr(X \cup Y) = Pr(X) + Pr(Y)$.

Thm 1.3.1. $Pr(X) + Pr(\overline{X}) = 1$.

Proof.

Since $X \cap \overline{X} = \emptyset$, $Pr(X \cup \overline{X}) = Pr(X) + Pr(\overline{X})$ by Ax3.

$Pr(X \cup \overline{X}) = Pr(\Omega) = 1$ by Ax2.

Thus, $Pr(X) + Pr(\overline{X}) = 1$. $\square$

Thm 1.3.2. $Pr(X \cup Y) = Pr(X) + Pr(Y) - Pr(X \cap Y)$.

Proof.

$Pr(X \cup Y) = Pr((X \cap Y) \cup (\overline{X} \cap Y) \cup (X \cap \overline{Y}))$.

$Pr(X \cup Y) = Pr((X \cap Y) \cup (\overline{X} \cap Y)) + Pr(X \cap \overline{Y})$ by Ax3.

$Pr(X \cup Y) = Pr(Y) + Pr(X \cap \overline{Y})$.

$Pr(X) = Pr((X \cap \overline{Y}) \cup (X \cap Y))$.

$Pr(X) = Pr(X \cap \overline{Y}) + Pr(X \cap Y)$ by Ax3.

$Pr(X \cap \overline{Y}) = Pr(X) - Pr(X \cap Y)$.

$Pr(X \cup Y) = Pr(X) + Pr(Y) - Pr(X \cap Y)$. $\square$

Def 1.3.3. The *conditional probability of X given Y* is $Pr(X|Y) = \frac{Pr(X \cap Y)}{Pr(Y)}$ given that $Pr(Y) \neq 0$.

Aside: Some philosophers argue that we should treat conditional probability as primitive and define unconditional probability in terms of it.

Def 1.3.4. X is *probabilistically independent* of Y if and only if $Pr(X) = Pr(X|Y)$.

Thm 1.3.3. If X is probabilistically independent of Y, then $Pr(X \cap Y) = Pr(X) \times Pr(Y)$.

Thm 1.3.4 (Inverse Probability Law). $Pr(X|Y) = \frac{Pr(X) \times Pr(Y|X)}{Pr(Y)}$ given that $Pr(Y) \neq 0$.

Thm 1.3.5 (Bayes' Theorem).
$Pr(X|Y) = \frac{Pr(X) \times Pr(Y|X)}{Pr(X) \times Pr(Y|X) + Pr(\overline{X}) \times Pr(Y|\overline{X})}$ given that $Pr(Y) \neq 0$.

Ex. Lung Spot.

You are a physician and have just observed a spot on an X-ray of your patient's lung. You think that the patient might have tuberculosis and you are considering a treatment for this disease with slightly harmful side effects. If the patient has tuberculosis and you treat it, then you will cure the disease. If you do not treat it, then the disease will get worse and the patient will end up suffering considerably. You know that the probability of observing a lung spot given that a patient has tuberculosis is 20%. You also know that the unconditional probability of observing a lung spot is 10% and the incidence of tuberculosis in the general population is 5%. Do you administer the treatment?

	[?]	[?]
	tuberculosis	no tuberculosis
treatment	mild harm	mild harm
no treatment	extreme harm	no harm

$$Pr(LS|TB) = 0.2$$

$$Pr(LS) = 0.1$$

$$Pr(TB) = 0.05$$

$$Pr(TB|LS) = \frac{Pr(TB) \times Pr(LS|TB)}{Pr(LS)} = \frac{0.05 \times 0.2}{0.1} = 0.1$$

	$[\frac{1}{10}]$ tuberculosis	$[\frac{9}{10}]$ no tuberculosis
treatment	mild harm	mild harm
no treatment	extreme harm	no harm

$Pr(TB)$ is the *prior probability* of tuberculosis.

$Pr(TB|LS)$ is the *posterior probability* of tuberculosis.

$Pr(TB|LS)$ can be used as a new prior in further applications of the Inverse Probability Law.

What if you know little about the probability of tuberculosis at the onset?

'Washing of the priors': Different prior probabilities will converge to the same value after repeated application of the Inverse Probability Law or Bayes' Theorem on new data.

Where do the probabilities come from that enter into expected utility calculations?

Classical Laplacean View: The probability of X is the ratio of the number of X-cases to the total number of relevant cases.

This is an *objective logical* view of probability.

Problem. It is assumed that each of the relevant cases associated with X is equally likely. What justifies this assumption?

Long Run Frequency View: The probability of X is the frequency of X-cases in repeated trials in the limit.

This is an *objective empirical* view of probability.

Problem. Observed frequencies can be far from long run frequencies.

Problem. This view does not apply nicely, if at all, to single events that are not amenable to trials.

Subjective View: The probability of X is an agent's degree of belief, or *credence*, in X's occurrence.

Problem. How to measure credences?

Problem. An agent's credences needn't satisfy the probability axioms.

Ramsey and De Finetti:

Credences are reflected in willingness to bet. An agent has credence $Cr(X)$ in X's occurrence just in case this agent is willing to take either side of a bet B where for any stake \$$S$, the loser pays the winner \$$(1 - Cr(X)) \times S$ if X and \$$Cr(X) \times S$ if $\overline{X}$:

Bet B	
X	\$$(1 - Cr(X)) \times S$
$\overline{X}$	\$$Cr(X) \times S$

Problem. Placing a bet on an event can sometimes affect whether this event occurs.

Problem. Betting can have collateral benefits or costs.

Problem. Utility needn't be linear in dollars.

Problem. A bet on an event only becomes winning when its occurrence becomes known. But this can be difficult or even impossible to verify.

A *rational* agent's credences obey the probability axioms.

Many arguments have been offered for this thesis. Here is the most famous of them:

Thm 1.3.6 (Dutch Book Theorem). An agent's credence function $Cr : 2^{\Omega} \to \mathbb{R}[0,1]$ is a probability measure just in case there doesn't exist a set of bets, a *Dutch Book*, each of which the agent is indifferent between purchasing and selling that collectively guarantee a monetary loss.

A rational agent is not Dutch Bookable. Thus, by the Dutch Book Theorem, a rational agent has probabilistically coherent credences.

Partial Proof. Assume that an agent's credence function Cr violates Ax3 of the probability calculus; specifically, $X \cap Y = \emptyset$ but $Cr(X) + Cr(Y) > Cr(X \cup Y)$. We show that this agent is Dutch Bookable.

A bookie can bet against X and Y but bet for $X \cup Y$:

Bet B_1	
X	\$$(1 - Cr(X))$
$\overline{X}$	\$$Cr(X)$

Bet B_2	
Y	\$$(1 - Cr(Y))$
$\overline{Y}$	\$$Cr(Y)$

Bet B_3	
$X \cup Y$	\$$(1 - Cr(X \cup Y))$
$\overline{X \cup Y}$	\$$Cr(X \cup Y)$

Then the agent's total payoffs are as follows:

$$\begin{array}{c|c} \overline{X} \cap Y & -\$Cr(X) + \$(1 - Cr(Y)) - \$(1 - Cr(X \cup Y)) \\ X \cap \overline{Y} & \$(1 - Cr(X)) - \$Cr(Y) - \$(1 - Cr(X \cup Y)) \\ \overline{X} \cap \overline{Y} & -\$Cr(X) - \$Cr(Y) + \$Cr(X \cup Y) \end{array}$$

Since $Cr(X \cup Y) - Cr(X) - Cr(Y) < 0$, the agent loses money in all possible situations.

There is plenty more to say about probability, but let us now turn to utility.

We can also ask: where do the utilities come from that enter into expected utility calculations?

Von Neumann and Morgenstern:

If an agent's preferences over outcomes in $\mathcal{O}$ and *lotteries* involving these outcomes satisfy certain structural axioms, then these preferences can be exhibited by an interval utility function that is invariant up to positive linear transformations and where the utility of a lottery is its expected utility.

Let $L(p, x, y)$ designate the *lottery* that gives x with probability p and gives y with probability $1 - p$.

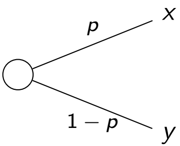

Def 1.3.5. The set of *lotteries* $\mathfrak{L}$ is built up as follows:

- $o \in \mathfrak{L}$ for each outcome $o \in \mathcal{O}$.
- If $L_1 \in \mathfrak{L}$ and $L_2 \in \mathfrak{L}$, then $L(p, L_1, L_2) \in \mathfrak{L}$ for any $p \in \mathbb{R}[0, 1]$.
- Nothing else is in $\mathfrak{L}$.

The Ordering Axioms:

Ax1 (Completeness). $L_1 \succ L_2 \vee L_1 \sim L_2 \vee L_2 \succ L_1$.

Ax2 (Asymmetry). If $L_1 \succ L_2$, then $L_2 \not\succ L_1$.

Ax3 (Negative Transitivity). If $L_1 \not\succ L_2$ and $L_2 \not\succ L_3$, then $L_1 \not\succ L_3$.

These axioms ensure that the agent's preferences can be represented by an ordinal utility function.

Ax4 (Continuity). If $L_1 \succ L_2$ and $L_2 \succ L_3$, then there is $p \in \mathbb{R}[0,1]$ such that $L_2 \sim L(p, L_1, L_3)$.

Ax5 (Better Prizes). $L_1 \succ L_2$ if and only if $L(p, L_1, L_3) \succ L(p, L_2, L_3)$, and $L_1 \succ L_2$ if and only if $L(p, L_3, L_1) \succ L(p, L_3, L_2)$.

Ax6 (Better Chances). If $L_1 \succ L_2$, then $p > q$ if and only if $L(p, L_1, L_2) \succ L(q, L_1, L_2)$.

Ax7 (Reduction of Compound Lotteries).
$L(p, L(q, L_1, L_2), L(r, L_1, L_2)) \sim L(pq + (1-p)r, L_1, L_2)$.

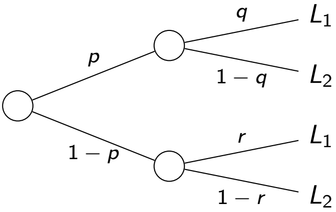

Ax4 (Continuity). If $L_1 \succ L_2$ and $L_2 \succ L_3$, then there is $p \in \mathbb{R}[0,1]$ such that $L_2 \sim L(p, L_1, L_3)$.

Ax5 (Better Prizes). $L_1 \succ L_2$ if and only if $L(p, L_1, L_3) \succ L(p, L_2, L_3)$, and $L_1 \succ L_2$ if and only if $L(p, L_3, L_1) \succ L(p, L_3, L_2)$.

Ax6 (Better Chances). If $L_1 \succ L_2$, then $p > q$ if and only if $L(p, L_1, L_2) \succ L(q, L_1, L_2)$.

Ax7 (Reduction of Compound Lotteries).
$L(p, L(q, L_1, L_2), L(r, L_1, L_2)) \sim L(pq + (1-p)r, L_1, L_2)$.

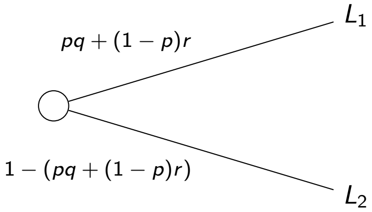

Thm 1.3.7 (Expected Utility Theorem). If an agent's preferences satisfy Ax1-Ax7, then there is a utility function $u : \mathfrak{L} \to \mathbb{R}[0,1]$ such that:

(i) $u(L_1) > u(L_2)$ if and only if $L_1 \succ L_2$.

(ii) $u(L(p, L_1, L_2)) = p \times u(L_1) + (1-p) \times u(L_2)$.

(iii) Any u' satisfying (i) and (ii) is a positive linear transformation of u.

Partial Proof of Thm 1.3.7. Assume that an agent's preferences satisfy Ax1-Ax7. We construct a utility function u satisfying (i) and (ii).

Let $B \in \mathcal{O}$ designate a *best* outcome—that is, $B \succcurlyeq o$ for all $o \in \mathcal{O}$.

Let $W \in \mathcal{O}$ designate a *worst* outcome—that is, $o \succcurlyeq W$ for all $o \in \mathcal{O}$.

Assume that $B \succ W$.

Let $u(B) = 1$ and $u(L) = 1$ for any lottery $L \in \mathfrak{L}$ such that $B \sim L$.

Let $u(W) = 0$ and $u(L) = 0$ for any lottery $L \in \mathfrak{L}$ such that $W \sim L$.

Now consider any non-extreme $L \in \mathfrak{L}$ such that $B \succ L$ and $L \succ W$.

By Ax4, $L \sim L(p, B, W)$ for some $p \in \mathbb{R}[0, 1]$ (moreover, p is unique).

Let $u(L) = p$.

The utility function $u : \mathfrak{L} \to \mathbb{R}[0, 1]$ constructed in this way satisfies (i).

By Ax6, $u(L_1) > u(L_2)$ if and only if $L(u(L_1), B, W) \succ L(u(L_2), B, W)$.

$L_1 \sim L(u(L_1), B, W)$.

$L_2 \sim L(u(L_2), B, W)$.

By Ordering Axioms, if $L_1 \sim L_3$ and $L_2 \sim L_4$, then $L_3 \succ L_4$ if and only if $L_1 \succ L_2$.

Thus, $u(L_1) > u(L_2)$ if and only if $L_1 \succ L_2$.

The utility function $u : \mathfrak{L} \to \mathbb{R}[0,1]$ also satisfies (ii).

Using the Ordering Axioms and Ax5, it is fairly straightforward to show the following:

Lem 1.3.1 (Substitution of Lotteries). If $L_1 \sim L(p, L_2, L_3)$, then both $L(q, L_1, L_4) \sim L(q, L(p, L_2, L_3), L_4)$ and $L(q, L_4, L_1) \sim L(q, L_4, L(p, L_2, L_3))$.

$L(p, L_1, L_2) \sim L(u(L(p, L_1, L_2)), B, W)$.

$L_1 \sim L(u(L_1), B, W)$

$L_2 \sim L(u(L_2), B, W)$

By Substitution of Lotteries, $L(p, L_1, L_2) \sim L(p, L(u(L_1), B, W), L(u(L_2), B, W))$.

By Ax7, $L(p, L(u(L_1), B, W), L(u(L_2), B, W)) \sim L(p \times u(L_1) + (1-p) \times u(L_2), B, W)$.

By Ordering Axioms, $L(u(L(p, L_1, L_2)), B, W) \sim L(p \times u(L_1) + (1-p) \times u(L_2), B, W)$.

By Better Chances, $u(L(p, L_1, L_2)) = p \times u(L_1) + (1-p) \times u(L_2)$.

Decision Theory

1.4 Paradoxes and Problems

Johns Hopkins University, Spring 2016

When making a decision under risk, *why* maximize expected utility?

Answer. In the long run, you will be better off by maximizing EU.

Reply. No real-life decision maker will ever face a decision an infinite number of times. Keynes: "In the long run we are all dead."

Reply. Many decisions are unique, such as the decision to marry a particular partner, the decision to start a particular war, and so on.

Answer. Recall the Expected Utility Theorem:

Thm 1.3.7. If an agent's preferences over lotteries satisfy Ax1-Ax7, then there is a utility function $u : \mathfrak{L} \to \mathbb{R}[0,1]$ such that:

(i) $u(L_1) > u(L_2)$ if and only if $L_1 \succ L_2$.

(ii) $u(L(p, L_1, L_2)) = p \times u(L_1) + (1-p) \times u(L_2)$.

(iii) Any u' satisfying (i) and (ii) is a positive linear transformation of u.

Notice that each act $a \in \mathcal{A}$ is itself a lottery whose expected utility is its utility. By maximizing EU, then, an agent is choosing an act that is at least as preferred as all others. Resnik: "In choosing an act whose expected utility is maximal an agent is simply doing what he wants to do!"

Reply. It now seems that an agent whose preferences satisfy Ax1-Ax7 doesn't even need decision theory.

Counter-reply. The maxim 'Maximize EU!' should be understood as 'Have preferences that satisfy the structural constraints Ax1-Ax7 (and then just do what you most prefer to do)!'

Counter-counter-reply. This just pushes back our initial question: Why have preferences that satisfy Ax1-Ax7?

Ex. Allais Paradox.

You are given a choice between the following two payouts:

A: \$1M.

B: You receive \$5M with 10% probability, \$1M with 89% probability, and nothing with 1% probability.

Do you choose A or B?

You are next given a choice between the following two payouts:

C: You receive \$5M with 10% probability and nothing with 90% probability.

D: You receive \$1M with 11% probability and nothing with 89% probability.

Do you choose C or D?

$EU(A) = 1 \times u(\$1\text{M})$.

$EU(B) = 0.1 \times u(\$5\text{M}) + 0.89 \times u(\$1\text{M}) + 0.01 \times u(\$0\text{M})$.

$EU(C) = 0.1 \times u(\$5\text{M}) + 0.9 \times u(\$0\text{M})$.

$EU(D) = 0.11 \times u(\$1\text{M}) + 0.89 \times u(\$0\text{M})$.

$EU(A) - EU(B) = 0.11 \times u(\$1\text{M}) - 0.1 \times u(\$5\text{M}) - 0.01 \times u(\$0\text{M})$.

$EU(D) - EU(C) = 0.11 \times u(\$1\text{M}) - 0.1 \times u(\$5\text{M}) - 0.01 \times u(\$0\text{M})$.

If you choose A over B, then presumably
$EU(A) - EU(B) = EU(D) - EU(C) > 0$.

So if you are an EU-maximizer, then you choose D over C.

The Allais Paradox plays on the common preference for a good for certain over a risky chance for a more valuable good.

Reply. The outcomes shouldn't be specified solely in terms of money. The outcome in offer B is \$0 *plus serious disappointment*.

Counter-reply. Every objection to the Principle of Maximizing EU might be thwarted by fiddling with the outcomes.

Reply. Bite the bullet. Savage: An agent who chooses A and C is irrational because they violate the *sure-thing principle* (note that they also violate Better Prizes).

	Ticket 1	Tickets 2-11	Tickets 12-100
gamble A	\$1M	\$1M	\$1M
gamble B	\$0M	\$5M	\$1M
gamble C	\$0M	\$5M	\$0M
gamble D	\$1M	\$1M	\$0M

The third column should be ignored when deciding between the gambles.

Counter-reply. Why satisfy the sure-thing principle?

Ex. Ellsberg Paradox.

An urn contains 90 balls. You know that 30 of these are yellow. You also know that the remaining 60 balls are either red or blue, but you do not know the proportion. I am about to draw a ball from the urn and I give you a choice between the following two payouts:

A: You receive \$100 if a yellow ball is drawn and \$0 otherwise.

B: You receive \$100 if a red ball is drawn and \$0 otherwise.

Do you choose A or B?

Suppose that I had instead offered a choice between the following two payouts:

C: You receive \$100 if either a red or blue ball is drawn and \$0 otherwise.

D: You receive \$100 if either a yellow or blue ball is drawn and \$0 otherwise.

Do you choose C or D?

Suppose that you assign a conditional probability of p to getting a red ball given that you get a red or blue ball.

$EU(A) = \frac{1}{3} \times u(\$100) + \frac{2}{3} \times u(\$0)$.

$EU(B) = \frac{1}{3} \times u(\$0) + \frac{2}{3} \times p \times u(\$100) + \frac{2}{3} \times (1-p) \times u(\$0)$.

$EU(C) = \frac{1}{3} \times u(\$0) + \frac{2}{3} \times u(\$100)$.

$EU(D) = \frac{1}{3} \times u(\$100) + \frac{2}{3} \times p \times u(\$0) + \frac{2}{3} \times (1-p) \times u(\$100)$.

$EU(A) - EU(B) = \frac{1-2p}{3} \times u(\$100) + \frac{2p-1}{3} \times u(\$0)$.

$EU(D) - EU(C) = \frac{1-2p}{3} \times u(\$100) + \frac{2p-1}{3} \times u(\$0)$.

If you choose A over B, then presumably
$EU(A) - EU(B) = EU(D) - EU(C) > 0$.

So if you are an EU-maximizer, then you choose D over C.

The Ellsberg Paradox plays on the common preference for known risks over unknown risks.

Reply. Bite the bullet. An agent who chooses A and C is irrational because they violate the sure-thing principle (and Better Prizes).

	$[\frac{1}{3}]$ Yellow	$[\frac{2}{3} \times p]$ Red	$[\frac{2}{3} \times (1-p)]$ Blue
gamble A	\$100	\$0	\$0
gamble B	\$0	\$100	\$0
gamble C	\$0	\$100	\$100
gamble D	\$100	\$0	\$100

Counter-reply. Why satisfy the sure-thing principle?

Reply. This is a decision under ignorance so the Principle of Maximizing EU does not apply.

Counter-reply. This reply is not open to subjectivists about probability who think that there are no real decisions under ignorance.

Ex. St. Petersburg Paradox.

You are given a choice between the following two payouts:

A: \$100.

B: A fair coin is flipped until it lands tails. If the coin lands tails on the first toss, then you receive \$2. If the coin lands tails on the second toss, then you receive \$4. In general, if the coin lands tails on the nth toss, then you receive $\$2^n$.

Do you choose A or B?

Let $EMV(a)$ designate the *expected monetary value* of $a \in \mathcal{A}$.

$EMV(A) = \$100$.

$EMV(B) = \frac{1}{2} \times \$2 + \frac{1}{4} \times \$4 + \frac{1}{8} \times \$8 + ... = \$1 + \$1 + \$1 + ... = \$\infty$.

Reply. There are diminishing returns to money.

Counter-reply. The St. Petersburg Paradox can be reframed in terms of utilities. If the coin lands tails on the first toss, then you receive a prize worth 2 utiles, and so forth.

Reply. There is, or should be, an upper bound on utility.

Counter-reply. This upper bound is ad hoc.

Reply. Jeffrey: "Anyone who offers to let the agent play the St. Petersburg game is a liar, for he is pretending to have an indefinitely large bank."

Counter-reply. All sorts of hypothetical prizes can be allowed.

Ex. St. Petersburg Paradox 2.

You are given a choice between the following two payouts:

A: A fair coin is flipped until it lands tails. If the coin lands tails on the first toss, then you receive \$2. If the coin lands tails on the second toss, then you receive \$4. In general, if the coin lands tails on the nth toss, then you receive $\$2^n$.

B: A biased coin that lands tails with probability 0.4 is flipped until it lands tails. If the coin lands tails on the first toss, then you receive \$2. If the coin lands tails on the second toss, then you receive \$4. In general, if the coin lands tails on the nth toss, then you receive $\$2^n$.

Do you choose A or B?

Intuitively, B is preferable.

However, $EMV(A) = EMV(B) = \infty$.

Ex. Two Envelope Paradox.

You are offered a choice between two envelopes A and B. You know that one of these envelopes contains twice as much money as the other, but you do not know how much money is in either envelope. You pick A. But right before you open this envelope, you are offered the opportunity to switch to B. Do you switch?

If you are an EU-Maximizer, then it seems that you should switch.

Suppose that there is $\$S$ in envelope A.

$EMV(\mathrm{A}) = \$S$.

$EMV(\mathrm{B}) = \frac{1}{2} \times \$\frac{1}{2}S + \frac{1}{2} \times \$2S = \$\frac{5}{4}S$.

$EMV(\mathrm{B}) > EMV(\mathrm{A})$.

But now suppose I offer you the opportunity to switch back...

Reply. The Two Envelope Paradox requires that there is an infinite amount of money in the world. If there is only $\$T$ available, then the envelope with more money can contain no more than $\$\frac{2}{3}T$. If this envelope contains $\$\frac{2}{3}T$, then the other envelope must contain $\$\frac{1}{3}T$.

Counter-reply. We can work with utilities instead.

Ex. Newcomb's Paradox.

You are standing in front of a table on top of which are two boxes A and B. Box A is transparent and you can see that it contains \$1000. Box B is opaque and contains either \$1M or nothing. You are offered the choice between taking only box B or both of these boxes. But before selecting, you are also given one extra piece of information. At some point in the past, a prophetic being called *The Predictor* predicted what you will do. If The Predictor predicted that you will take only box B, then \$1M was placed inside this box. If The Predictor predicted that you will take both boxes, then box B was left empty. The Predictor is almost always right. So what do you select?

Decision Theory

1.5 Causal vs. Evidential Decision Theory

Johns Hopkins University, Spring 2016

Ex. Newcomb's Paradox.

You are standing in front of a table on top of which are two boxes A and B. Box A is transparent and you can see that it contains \$1000. Box B is opaque and contains either \$1M or nothing. You are offered the choice between taking only box B or both of these boxes. But before selecting, you are also given one extra piece of information. At some point in the past, a prophetic being called *The Predictor* predicted what you will do. If The Predictor predicted that you will take only box B, then \$1M was placed inside this box. If The Predictor predicted that you will take both boxes, then box B was left empty. The Predictor is right 90% of the time. So what do you select?

Leaving out the information about The Predictor and assuming that your utility is linear in dollars, the decision table for Newcomb's Paradox is this:

	B contains \$1M	B is empty
one box	1M	0
two boxes	1M+1000	1000

Leaving out the information about The Predictor and assuming that your utility is linear in dollars, the decision table for Newcomb's Paradox is this:

	B contains \$1M	B is empty
one box	1M	0
two boxes	1M+1000	1000

Note that two boxes dominates one box.

Leaving out the information about The Predictor and assuming that your utility is linear in dollars, the decision table for Newcomb's Paradox is this:

	B contains \$1M	B is empty
two boxes	1M+1000	1000

So the Dominance Principle forbids taking only box B.

If we add the information about The Predictor, then the table becomes:
(N.B. This requires a more complex decision model)

	B contains \$1M	B is empty
one box	1M [$\frac{9}{10}$]	0 [$\frac{1}{10}$]
two boxes	1M+1000 [$\frac{1}{10}$]	1000 [$\frac{9}{10}$]

$$EU(\text{one box}) = \frac{9}{10} \times 1\text{M} + \frac{1}{10} \times 0 = 900000.$$

$$EU(\text{two boxes}) = \frac{1}{10} \times (1\text{M} + 1000) + \frac{9}{10} \times 1000 = 101000.$$

$$EU(\text{one box}) > EU(\text{two boxes}).$$

If we add the information about The Predictor, then the table becomes:

	B contains \$1M	B is empty
one box	1M [$\frac{9}{10}$]	0 [$\frac{1}{10}$]

So the Principle of Maximizing Expected Utility forbids taking both of the boxes.

How to resolve the conflict between the Dominance Principle and the Principle of Maximizing EU?

Recall that the Dominance Principle leads to trouble when actions and states are causally dependent.

Ex. Superbowl.

Who do you think will win the Super Bowl? If you correctly guess that the Carolina Panthers will win, then you receive \$100. If you correctly guess that the Seattle Seahawks will win, then you receive \$200.

	You win the bet	You lose the bet
bet on Carolina	win \$100	win \$0
bet on Seattle	win \$200	win \$0

But in Newcomb's Paradox, actions and states are *causally independent*.

You make your choice *after* The Predictor predicts what you will do. Standing in front of the table, \$1M has already been placed in box B or this box has been left empty.

Resnik: "The dilemma is this. If we use the dominance principle to decide, we must ignore all the empirical data pointing to the folly of choosing both boxes; but if we follow the data and maximize expected utility, we are at a loss to explain why the data are relevant."

Ex. Calvinist Theology.

Suppose that God has already determined who will go to heaven and who will go to hell. Nothing we do in our lifetimes can change this. However, while you do not know whether you are destined for heaven or hell, you do know that there is a high correlation between going to heaven and being devout. Should you sin?

	destined for heaven	destined for hell
sin	heaven + pleasure [low]	hell + pleasure [high]
do not sin	heaven [high]	hell [low]

The Dominance Principle forbids refraining from sinning.

The Principle of Maximizing EU (using the correlations) forbids sinning.

Ex. Smoking.

Suppose that smoking is strongly correlated with lung cancer because of a common cause—a genetic defect that tends to cause both smoking and lung cancer. In fact, smoking does not cause any of the diseases that it is strongly correlated with. These diseases are also caused by the genetic defect. Do you smoke?

	genetic defect	no genetic defect
smoke	disease + pleasure [high]	good health + pleasure [low]
do not smoke	disease [low]	good health [high]

The Dominance Principle forbids refraining from smoking.

The Principle of Maximizing EU (using the correlations) forbids smoking.

Decision theorists have typically not responded to Newcomb's Paradox by abandoning Expected Utility Theory. Instead, this paradox has generated vigorous debate about which probabilities should enter into expected utility calculations.

Causal Decision Theory: The probabilities used should be *causal* probabilities—that is, the probabilities should reflect the propensity for acts to produce or prevent outcomes.

Evidential Decision Theory: The probabilities used should be *evidential* probabilities—that is, the probabilities should reflect the likelihood, given the agent's total available evidence, that outcomes will occur given acts.

There are many variants of CDT and EDT. One popular approach uses *subjunctive conditionals*.

Let $a \mathrel{\Box\!\!\rightarrow} o$ abbreviate 'If the agent were to perform $a \in \mathcal{A}$, then outcome $o \in \mathcal{O}$ would occur.'

Let $Pr(a \mathrel{\Box\!\!\rightarrow} o)$ designate the probability of $a \mathrel{\Box\!\!\rightarrow} o$.
(N.B. This requires more complex decision models).

Since the truth conditions of subjunctive conditionals (arguably) track causal relations, $Pr(a \mathrel{\Box\!\!\rightarrow} o)$ is a causal probability.

According to the Causal Decision Theorist, the expected utility of $a \in \mathcal{A}$ is $\sum_{o \in \mathcal{O}_a} Pr(a \mathrel{\Box\!\!\rightarrow} o) \times u(o)$ where $\mathcal{O}_a$ is the set of outcomes achievable by performing a (formally: $\mathcal{O}_a = \{o \in \mathcal{O} : \exists s \in \Omega(g(a, s) = o)\}$).

According to the Evidential Decision Theorist, the expected utility of $a \in \mathcal{A}$ is $\sum_{o \in \mathcal{O}_a} Pr(a \mathrel{\Box\!\!\rightarrow} o|a) \times u(o)$.

Back to Newcomb's Paradox...

	B contains \$1M	B is empty
one box	1M	0
two boxes	1M+1000	1000

	B contains \$1M	B is empty
one box	1M $[\frac{9}{10}]$	0 $[\frac{1}{10}]$
two boxes	1M+1000 $[\frac{1}{10}]$	1000 $[\frac{9}{10}]$

Using evidential probabilities, $EU(\text{one box}) > EU(\text{two boxes})$.

	B contains $1M	B is empty
one box	1M [?]	0 [?]
two boxes	1M+1000 [?]	1000 [?]

What if we use causal probabilities?

	B contains \$1M	B is empty
one box	1M [p]	0 [1-p]
two boxes	1M+1000 [p]	1000 [1-p]

Since the acts and states are causally independent,

$Pr(\text{one box} \mathrel{\Box\!\!\rightarrow} \$1\text{M}) = Pr(\text{B contains } \$1\text{M}) = p.$
$Pr(\text{two boxes} \mathrel{\Box\!\!\rightarrow} \$1\text{M}+\$1000) = Pr(\text{B contains } \$1\text{M}) = p.$
$Pr(\text{one box} \mathrel{\Box\!\!\rightarrow} \$0) = Pr(\text{B is empty}) = 1 - p.$
$Pr(\text{two boxes} \mathrel{\Box\!\!\rightarrow} \$1000) = Pr(\text{B is empty}) = 1 - p.$

	B contains $1M	B is empty
one box	1M [p]	0 [1-p]
two boxes	1M+1000 [p]	1000 [1-p]

$$EU(\text{one box}) = p \times 1\text{M} + (1 - p) \times 0.$$

$$EU(\text{two boxes}) = p \times (1\text{M} + 1000) + (1 - p) \times 1000.$$

$$EU(\text{two boxes}) > EU(\text{one box}).$$

While the EDT-style Principle of Maximizing EU conflicts with the Dominance Principle, the CDT-style Principle of Maximizing EU does not.

In Calvinist Theology, both the Dominance Principle and the CDT-style Principle of Maximizing EU forbid refraining from sinning.

In Smoking, both the Dominance Principle and the CDT-style Principle of Maximizing EU forbid refraining from smoking.

CDTheorist. We should all be causal decision theorists like me.

EDTheorist. Remember Hume. There is no causality in the world, only statistical regularities. Causal Decision Theory is without foundation.

CDTheorist. If our actions do not cause outcomes to occur, then what is the point of choosing? What is the point of decision theory?

EDTheorist. Decision theory does not tell one to act so as to bring about the best outcome, but rather to choose the act that one would be happiest to learn that one had performed.

EDTheorist. Moreover, there are counterexamples to Causal Decision Theory.

Ex. Murder Lesion (Egan).

"Mary is debating whether to shoot her rival, Alfred. If she shoots and hits, things will be very good for her. If she shoots and misses, things will be very bad. (Alfred always finds out about unsuccessful assassination attempts, and he is sensitive about such things.) If she doesn't shoot, things will go on in the usual, okay-but-not-great kind of way. Though Mary is fairly confident that she will not actually shoot, she has, just to keep her options open, been preparing for this moment by honing her skills at the shooting range. Her rifle is accurate and well maintained. In view of this, she thinks that it is very likely that if she were to shoot, then she would hit. So far, so good. But Mary also knows that there is a certain sort of brain lesion that tends to cause both murder attempts and bad aim at the critical moment. If she has this lesion, all of her training will do her no good—her hand is almost certain to shake as she squeezes the trigger. Happily for most of us, but not so happily for Mary, most shooters have this lesion, and so most shooters miss. Should Mary shoot?"

	brain lesion	no brain lesion
shoot	failed assassination	Alfred dies
do not shoot	okay	okay

Since the evidential probability $Pr(\text{shoot} \mathrel{\Box\!\!\rightarrow} \text{failed assassination}|\text{shoot})$ is high, the EDT-style Principle of Maximizing EU forbids shooting.

	brain lesion	no brain lesion
shoot	failed assassination	Alfred dies
do not shoot	okay	okay

Since the causal probability $Pr(\text{shoot} \mathbin{\Box\!\!\rightarrow} \text{Alfred dies})$ is high, the CDT-style Principle of Maximizing EU forbids refraining from shooting.

Ex. The Psychopath Button (Egan).

"Paul is debating whether to press the 'kill all psychopaths' button. It would, he thinks, be much better to live in a world with no psychopaths. Unfortunately, Paul is quite confident that only a psychopath would press such a button. Paul very strongly prefers living in a world with psychopaths to dying. Should Paul press the button?"

	Paul is psycho	Paul is not psycho
press button	Paul dies	Paul lives and psychos die
do not press button	everyone lives	everyone lives

Since the evidential probability
$Pr(\text{press button} \mathrel{\Box\!\!\rightarrow} \text{Paul dies}|\text{press button})$ is high, the EDT-style Principle of Maximizing EU forbids pressing the button.

	Paul is psycho	Paul is not psycho
press button	Paul dies	Paul lives and psychos die
do not press button	everyone lives	everyone lives

Since the causal probability
Pr(press button $\Box\!\!\rightarrow$ Paul lives and psychos die) is high, the CDT-style Principle of Maximizing EU forbids refraining from pressing the button.

Egan: “Here is the moral that I think we should draw from all of this: Evidential decision theory told us to perform the action with the best expected outcome. Examples like [Smoking] show us that having the best expected outcome comes apart from having the best expected causal impact on how things are and that rationality tracks the latter rather than the former. So, they show us that evidential decision theory is mistaken. Causal decision theory told us to perform the action that, holding fixed our current views about the causal structure of the world, has the best expected outcome. Examples like The Murder Lesion and The Psychopath Button show us that this too comes apart from having the best expected causal impact on how things are. So, they show us that causal decision theory is mistaken.”

Peacemaker. Must we choose between CDT and EDT? Perhaps there are just two kinds of rationality—one captured by CDT and the other captured by EDT. Solidarity forever.

Decision Theory

1.6 Risk-Weighted Expected Utility

Johns Hopkins University, Spring 2016

How to capture an agent's attitudes towards *risk* in decision theory?

Ex. Elvis Stamp & Gloves.

You are a stamp collector who lives in a cold climate. Two fair coins are about to be flipped and you are given a choice between the following gambles:

A: If coin 1 lands heads, then you receive a rare Elvis stamp; otherwise, you receive nothing. If coin 2 lands tails, then you receive a nice pair of gloves; otherwise, you receive nothing.

B: If coin 1 lands heads, then you receive a rare Elvis stamp; otherwise, you receive a nice pair of gloves.

Do you choose A or B?

	$[\frac{1}{4}]$ HH	$[\frac{1}{4}]$ HT	$[\frac{1}{4}]$ TH	$[\frac{1}{4}]$ TT
gamble A	stamp	stamp + gloves	nothing	gloves
gamble B	stamp	stamp	gloves	gloves

Ex. Dice.

A dice is about to be rolled and you are given a choice between the following gambles:

A: If the dice lands on an even number, then you receive \$100; otherwise, you receive nothing.

B: \$49.

Do you choose A or B?

	$[\frac{1}{2}]$ lands even	$[\frac{1}{2}]$ lands odd
gamble A	\$100	\$0
gamble B	\$49	\$49

It seems *prima facie* reasonable to prefer B to A in Elvis Stamp & Gloves since the former gamble guarantees that you win a prize but the latter gamble does not. It also seems reasonable to prefer B to A in Dice should you wish to avoid risk.

How to capture these preferences in Expected Utility Theory?

We have utilities and probabilities to work with, but the probabilities are fixed in both examples. So risk-aversion must be encoded in your utility function.

In Elvis Stamp & Gloves, we might assume that your utility function displays *dependence* with respect to the two goods.

$u(\text{stamp} + \text{gloves}) - u(\text{stamp}) < u(\text{gloves}) - u(\text{nothing})$.

$$\begin{aligned} EU(\text{gamble A}) = {} & \tfrac{1}{4} \times u(\text{stamp} + \text{gloves}) + \tfrac{1}{4} \times u(\text{stamp}) \\ & + \tfrac{1}{4} \times u(\text{gloves}) + \tfrac{1}{4} \times u(\text{nothing}). \end{aligned}$$

$EU(\text{gamble B}) = \frac{1}{2} \times u(\text{stamp}) + \frac{1}{2} \times u(\text{gloves})$.

Since $u(\text{gloves}) + u(\text{stamp}) > u(\text{stamp} + \text{gloves}) + u(\text{nothing})$,

$EU(\text{gamble B}) > EU(\text{gamble A})$.

In Dice, we might assume that your utility function is *concave*—that is, your utility diminishes marginally in monetary value.

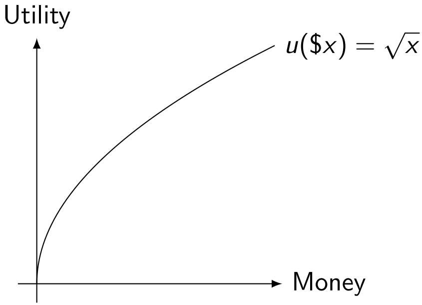

$EU(\text{gamble A}) = \frac{1}{2} \times u(\$100) + \frac{1}{2} \times u(\$0) = 5.$

$EU(\text{gamble B}) = u(\$49) {=} 7.$

$EU(\text{gamble B}) > EU(\text{gamble A}).$

Suppose that utilities correspond to the strength of an agent's desires for the various outcomes in $\mathcal{O}$. Then EU Theory suggests that risk-aversion amounts to the diminishing strength of desire as one accumulates goods or money.

But you might value the goods independently in Elvis Stamp & Gloves, and you might value small amounts of money linearly in Dice.

Intuitively, it seems that risk-avoidant agents (and risk-inclined agents, for that matter) are what Buchak calls "globally sensitive."

In Elvis Stamp & Gloves or Dice, you might be sensitive not just to the utility values of individual outcomes, but also to global features of the gambles such as the minimum utility they might yield, the maximum utility they might yield, the spread between the minimum and maximum utilities, and so forth.

A globally sensitive agent can still have a diminishing marginal utility function. Risk-averse preferences can result from either of these phenomena.

Given the Expected Utility Theorem, some decision theorists might think that global considerations are built into the utility function. EU Theory, they might claim, can account for global sensitivity.

But as we have already seen, EU Theory cannot handle the preferences of certain globally sensitive agents.

Ex. Allais Paradox.

You are given a choice between the following two payouts:

A: $1M.

B: You receive $5M with 10% probability, $1M with 89% probability, and nothing with 1% probability.

Do you choose A or B?

You are next given a choice between the following two payouts:

C: You receive $5M with 10% probability and nothing with 90% probability.

D: You receive $1M with 11% probability and nothing with 89% probability.

Do you choose C or D?

$EU(A) = 1 \times u(\$1\text{M})$.

$EU(B) = 0.1 \times u(\$5\text{M}) + 0.89 \times u(\$1\text{M}) + 0.01 \times u(\$0\text{M})$.

$EU(C) = 0.1 \times u(\$5\text{M}) + 0.9 \times u(\$0\text{M})$.

$EU(D) = 0.11 \times u(\$1\text{M}) + 0.89 \times u(\$0\text{M})$.

$EU(A) - EU(B) = 0.11 \times u(\$1\text{M}) - 0.1 \times u(\$5\text{M}) - 0.01 \times u(\$0\text{M})$.

$EU(D) - EU(C) = 0.11 \times u(\$1\text{M}) - 0.1 \times u(\$5\text{M}) - 0.01 \times u(\$0\text{M})$.

No utility function can rationalize strict preferences for A over B and C over D.

Buchak: "Taking the means to one's ends involves making *three* evaluations rather than two. First, an individual determines which ends he wants, and how much: these are the values captured by the utility function. Second, he determines how likely various actions are to lead to various ends: how effective each means would be of realizing each particular end. This evaluation is captured by the probability function. Finally, he determines which strategy to take when choosing among actions whose outcomes are uncertain: how effective each means is towards his general aim of realizing an end he wants (an aim which is satisfied by a particular end in proportion to the degree of his desire for it). Specifically, *he evaluates the extent to which he is generally willing to accept a chance of something worse in exchange for a chance of something better*."

Buchak's *Risk-Weighted Expected Utility Theory* includes three subjective parameters:

- utility function $u : \mathcal{O} \to \mathbb{R}$.
- probability function $Pr : 2^{\Omega} \to \mathbb{R}[0,1]$.
- *risk function* $r : \mathbb{R}[0,1] \to \mathbb{R}[0,1]$ where r is non-decreasing, $r(0) = 0$, and $r(1) = 1$.

The risk function r represents how the agent trades off securing good worst-case scenarios against allowing for good best-case scenarios.

Consider the following gamble where $u(o_2) \geq u(o_1)$:

	$[1-p]$	$[p]$
	s_1	s_2
a	o_1	o_2

$$EU(a) = (1-p) \times u(o_1) + p \times u(o_2).$$

$$EU(a) = u(o_1) + p(u(o_2) - u(o_1)).$$

Consider the following gamble where $u(o_2) \geq u(o_1)$:

$$\begin{array}{cc|c|c|} & [1-p] & [p] \\ & s_1 & s_2 \\ \cline{2-3} a & o_1 & o_2 \\ \cline{2-3} \end{array}$$

The *risk-weighted expected utility* of a is

$$REU(a) = u(o_1) + r(p)(u(o_2) - u(o_1)).$$

If $r(p) < p$, then $REU(a) < EU(a)$.

If $r(p) > p$, then $REU(a) > EU(a)$.

If $r(p) = p$, then $REU(a) = EU(a)$.

Consider the following gamble where $u(o_3) \geq u(o_2) \geq u(o_1)$:

	$[p_1]$	$[p_2]$	$[p_3]$
	s_1	s_2	s_3
a	o_1	o_2	o_3

The *risk-weighted expected utility* of a is

$$REU(a) = u(o_1) + r(p_2 + p_3)(u(o_2) - u(o_1)) + r(p_3)(u(o_3) - u(o_2)).$$

More generally, consider the following gamble where $u(o_n) \geq ... \geq u(o_1)$:

	$[p_1]$	...	$[p_n]$
	s_1	...	s_n
a	o_1	...	o_n

Def 1.6.1. The *risk-weighted expected utility* of a is

$$REU(a) = u(o_1) + r(\sum_{i=2}^{n} p_i)(u(o_2) - u(o_1)) + ... + r(p_n)(u(o_n) - u(o_{n-1})).$$

Thesis 1.6.1 (Principle of Maximizing Risk-Weighted EU). Rationality forbids any act $a_1 \in \mathcal{A}$ such that $REU(a_1) < REU(a_2)$ for some alternative act $a_2 \in \mathcal{A}$.

In other words, maximize REU.

Note that $min(a) \leq REU(a) \leq max(a)$.

Note that if $r(p) = p$ for all p, maximizing REU amounts to maximizing EU. So Risk-Weighted EU Theory can be regarded as a generalization of EU Theory.

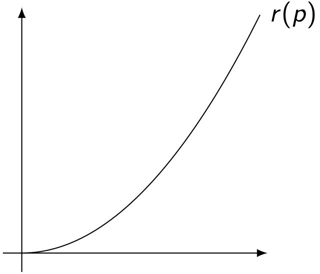

Risk-avoidant agents have *convex* risk functions.

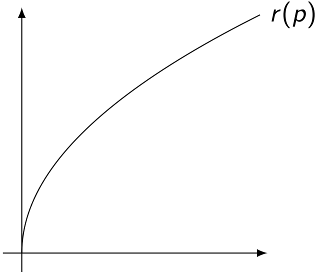

Risk-inclined agents have *concave* risk functions.

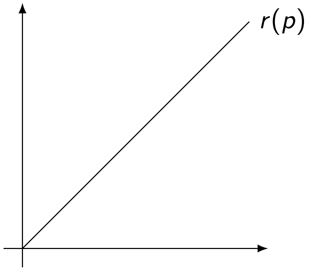

Globally neutral agents have *linear* risk functions where $r(p) = p$ for all p.

Buchak: "*On EU theory, to be generally risk-averse is to have a concave utility function, or, in the case of discrete goods, a utility function that displays non-independence. On rank-dependent theories like REU theory, to be generally risk-averse is to have a convex risk function.* The psychological intuition behind the diminishing marginal utility analysis of risk-aversion was that adding money to an outcome is of less value the more money that outcome already contains; or that getting an additional good is of less value if one already has some other good. The intuition behind the rank-dependent analysis of risk-aversion is that adding *probability* to a preferred outcome is of more value the more likely that outcome already is to obtain."

A maximinimizer has the risk function $r(p) = \begin{cases} 0 \text{ if } p \neq 1 \\ 1 \text{ if } p = 1 \end{cases}$

A maximaximizer has the risk function $r(p) = 1$ for all $p \in \mathbb{R}(0, 1]$.

An α-index-maximizer has the risk function $r(p) = \alpha$ for all $p \in \mathbb{R}(0, 1)$.

(Using any probability function as input to the REU calculation.)

Interestingly, experiments by Tversky and Kahneman suggest that most of us have S-shaped risk functions.

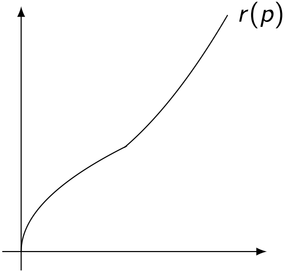

That is, most of us are risk-inclined in small probabilities and risk-avoidant in large probabilities. We care about the extremes.

Back to our examples...

Ex. Elvis Stamp & Gloves.

	$[\frac{1}{4}]$ HH	$[\frac{1}{4}]$ HT	$[\frac{1}{4}]$ TH	$[\frac{1}{4}]$ TT
gamble A	stamp	stamp + gloves	nothing	gloves
gamble B	stamp	stamp	gloves	gloves

Supposing that $u(\text{stamp+gloves}) > u(\text{stamp}) > u(\text{gloves}) > u(\text{nothing})$,

$REU(A) = u(\text{nothing}) + r(\frac{3}{4})(u(\text{gloves}) - u(\text{nothing}))$
$+r(\frac{1}{2})(u(\text{stamp}) - u(\text{gloves})) + r(\frac{1}{4})(u(\text{stamp+gloves}) - u(\text{stamp}))$.

$REU(B) = u(\text{gloves}) + r(\frac{1}{2})(u(\text{stamp}) - u(\text{gloves}))$.

Supposing that $u(\text{stamp+gloves}) - u(\text{stamp}) = u(\text{gloves}) - u(\text{nothing})$,

$REU(B) > REU(A)$ iff $1 > r(\frac{1}{4}) + r(\frac{3}{4})$.

This inequality holds when r is convex.

Ex. Dice.

	$[\frac{1}{2}]$ lands even	$[\frac{1}{2}]$ lands odd
gamble A	\$100	\$0
gamble B	\$49	\$49

$REU(A) = u(\$0) + r(\frac{1}{2})(u(\$100) - u(\$0))$.

$REU(B) = u(\$49)$.

Supposing that $u(\$x) = x$,

$REU(B) > REU(A)$ iff $\frac{49}{100} > r(\frac{1}{2})$.

This inequality can hold when r is convex.

Ex. Allais Paradox.

	$[\frac{1}{100}]$ Ticket 1	$[\frac{10}{100}]$ Tickets 2-11	$[\frac{89}{100}]$ Tickets 12-100
gamble A	\$1M	\$1M	\$1M
gamble B	\$0M	\$5M	\$1M
gamble C	\$0M	\$5M	\$0M
gamble D	\$1M	\$1M	\$0M

$REU(A) = u(\$1\text{M})$.

$REU(B) = u(\$0) + r(\frac{99}{100})(u(\$1\text{M}) - u(\$0)) + r(\frac{10}{100})(u(\$5\text{M}) - u(\$1\text{M}))$.

$REU(C) = u(\$0) + r(\frac{10}{100})(u(\$5\text{M}) - u(\$0))$.

$REU(D) = u(\$0) + r(\frac{11}{100})(u(\$1\text{M}) - u(\$0))$.

Supposing that $u(\$\text{xM}) = x$,

$REU(A) > REU(B)$ iff $1 > r(\frac{99}{100}) + 4r(\frac{10}{100})$.

$REU(C) > REU(D)$ iff $5r(\frac{10}{100}) > r(\frac{11}{100})$.

These inequalities can hold when r is convex.

Ex. St. Petersburg Paradox.

You are given a choice between the following two payouts:

A: \$100.

B: A fair coin is flipped until it lands tails. If the coin lands tails on the first toss, then you receive \$2. If the coin lands tails on the second toss, then you receive \$4. In general, if the coin lands tails on the nth toss, then you receive $\$2^n$.

Do you choose A or B?

$REU(A) = u(\$100)$.

$REU(B) = u(\$2) + r(\frac{1}{2})(u(\$4) - u(\$2)) + r(\frac{1}{4})(u(\$8) - u(\$4)) + ...$

If r is convex enough close to 0, $REU(B)$ can be finite.

Reply. But we can just offer prizes with higher utilities to compensate for the lower risk-adjusted weights.

If $r(p) = 0$ for $p \in [0, \epsilon]$ where $\epsilon > 0$, then $REU(B)$ is still finite.

Game Theory

2.1 Zero Sum Games

Johns Hopkins University, Spring 2016

Osborne and Rubinstein: “Game theory is a bag of analytical tools designed to help us understand the phenomena that we observe when decision-makers interact. The basic assumptions that underlie the theory are that decision-makers pursue well-defined exogenous objectives (they are *rational*) and take into account their knowledge or expectations of *other* decision-makers’ behavior (they *reason strategically*).”

Ex. Paper, Rock, & Scissors.

You are playing the game Paper, Rock, & Scissors. Paper beats rock. Rock beats scissors. Scissors beats paper. Two of the same is a draw. Which do you choose?

Ex. Stag Hunt.

You and an acquaintance have gone hunting and you must each decide whether to pursue a large stag or to pursue a small hare. Alone, you can each capture a hare. However, a stag requires two people to capture. If you cooperate and both hunt stag, then you will end up with much meat. But if you hunt stag while your acquaintance hunts hare, then you return home empty-handed. Do you hunt stag or hare?

We can explicate these informal interactive choice situations with formal mathematical models.

Def 2.1.1. A *strategic game* $\mathcal{G} = \langle \mathcal{N}, \{\mathcal{A}_i\}_{i\in\mathcal{N}}, \mathcal{O}, g, \{\succcurlyeq_i\}_{i\in\mathcal{N}} \rangle$ consists of the following ingredients:

- a finite set of players $\mathcal{N}$
- a set of actions $\mathcal{A}_i$ available to each player $i \in \mathcal{N}$
- a set of outcomes $\mathcal{O}$
- a function $g : \times_{i\in\mathcal{N}} \mathcal{A}_i \to \mathcal{O}$ that maps each action profile $a \in \times_{i\in\mathcal{N}} \mathcal{A}_i$ to an outcome $o \in \mathcal{O}$
- a preference relation $\succcurlyeq_i$ for each player over the set of outcomes $\mathcal{O}$ where $o_1 \succcurlyeq_i o_2$ just in case o_1 is at least as preferred as o_2 by player $i \in \mathcal{N}$

Def 2.1.2. $\mathcal{G}$ is *finite* iff $|\mathcal{A}_i|$ is finite for each $i \in \mathcal{N}$.

Going forward, it is assumed that the preferences of each player $i \in \mathcal{N}$ are well-behaved in the sense that they are representable with an interval utility function $u_i : \mathcal{O} \rightarrow \mathbb{R}$.

Keep in mind that $\succ_i$ and $\sim_i$ are definable in terms of $\succcurlyeq_i$.

Aside: We could model a situation where the outcomes of action profiles are affected by exogenous random variables by adding a set of states Ω.

The game $\mathcal{G}$ of Paper, Rock, & Scissors consists of:

- $\mathcal{N} = \{1, 2\}$
- $\mathcal{A}_1 = \{\text{paper}, \text{rock}, \text{scissors}\}$
  $\mathcal{A}_2 = \{\text{paper}, \text{rock}, \text{scissors}\}$
- $\mathcal{O} = \{1 \text{ wins}, 2 \text{ wins}, \text{draw}\}$
- $g(\langle \text{paper}, \text{paper} \rangle) = \text{draw}$
  $g(\langle \text{paper}, \text{rock} \rangle) = 1 \text{ wins}$
  $g(\langle \text{paper}, \text{scissors} \rangle) = 2 \text{ wins}$
  $g(\langle \text{rock}, \text{paper} \rangle) = 2 \text{ wins}$
  $g(\langle \text{rock}, \text{rock} \rangle) = \text{draw}$
  $g(\langle \text{rock}, \text{scissors} \rangle) = 1 \text{ wins}$
  $g(\langle \text{scissors}, \text{paper} \rangle) = 1 \text{ wins}$
  $g(\langle \text{scissors}, \text{rock} \rangle) = 2 \text{ wins}$
  $g(\langle \text{scissors}, \text{scissors} \rangle) = \text{draw}$
- $1 \text{ wins} \succ_1 \text{draw} \succ_1 2 \text{ wins}$
  $2 \text{ wins} \succ_2 \text{draw} \succ_2 1 \text{ wins}$

The game $\mathcal{G}$ of Stag Hunt consists of:

- $\mathcal{N} = \{1, 2\}$
- $\mathcal{A}_1 = \{\text{hunt stag}, \text{hunt hare}\}$
  $\mathcal{A}_2 = \{\text{hunt stag}, \text{hunt hare}\}$
- $\mathcal{O} = \{\text{both stag}, \text{both hare}, \text{1 hare only}, \text{2 hare only}\}$
- $g(\langle\text{hunt stag}, \text{hunt stag}\rangle) = \text{both stag}$
  $g(\langle\text{hunt stag}, \text{hunt hare}\rangle) = \text{2 hare only}$
  $g(\langle\text{hunt hare}, \text{hunt stag}\rangle) = \text{1 hare only}$
  $g(\langle\text{hunt hare}, \text{hunt hare}\rangle) = \text{both hare}$
- both stag $\succ_1$ both hare $\succcurlyeq_1$ 1 hare only $\succ_1$ 2 hare only
  both stag $\succ_2$ both hare $\succcurlyeq_2$ 2 hare only $\succ_2$ 1 hare only

A strategic game can be visualized in a table.

For example, the game $\mathcal{G}$ of Paper, Rock, & Scissors corresponds to the following *game matrix* (1 is the row player and 2 is the column player):

	paper	rock	scissors
paper	draw	1 wins	2 wins
rock	2 wins	draw	1 wins
scissors	1 wins	2 wins	draw

Substituting the pair of utilities $\langle u_1(o), u_2(o)\rangle$ for each outcome o in the matrix gives us:

	paper	rock	scissors
paper	0,0	1,-1	-1,1
rock	-1,1	0,0	1,-1
scissors	1,-1	-1,1	0,0

The game $\mathcal{G}$ of Stag Hunt corresponds to the following game matrix:

	hunt stag	hunt hare
hunt stag	both stag	2 hare only
hunt hare	1 hare only	both hare

Plugging in utilities:

	hunt stag	hunt hare
hunt stag	2,2	0,1
hunt hare	1,0	1,1

Def 2.1.3. A *strictly competitive* or *zero sum* two-player strategic game $\mathcal{G}$ is one where for any outcomes $o_1, o_2 \in \mathcal{O}$, $o_1 \succcurlyeq_1 o_2$ if and only if $o_2 \succcurlyeq_2 o_1$.

If player 1's preferences are represented with $u_1 : \mathcal{O} \to \mathbb{R}$, then player 2's preferences are representable with $u_2 : \mathcal{O} \to \mathbb{R}$ where $u_1(o) + u_2(o) = 0$ for all $o \in \mathcal{O}$.

Paper, Rock, & Scissors is zero sum.

Stag Hunt is not.

The most famous solution concept in game theory is *Nash equilibrium*.

Let $\langle a_j, a_{-j} \rangle$ designate the action profile where player j performs action $a_j \in \mathcal{A}_j$ and all of the other players perform $a_{-j} \in \times_{i \in \mathcal{N} \setminus \{j\}} \mathcal{A}_i$.

Def 2.1.4. A Nash equilibrium of $\mathcal{G}$ is an action profile $a^* \in \times_{i \in \mathcal{N}} \mathcal{A}_i$ such that for every player $j \in \mathcal{N}$, the following condition holds:

$g(\langle a_j^*, a_{-j}^* \rangle) \succcurlyeq_j g(\langle a_j, a_{-j}^* \rangle)$ for each $a_j \in \mathcal{A}_j$.

(Alternatively, $u_j(g(\langle a_j^*, a_{-j}^* \rangle)) \geq u_j(g(\langle a_j, a_{-j}^* \rangle))$ for each $a_j \in \mathcal{A}_j$.)

In other words, given the other players' equilibrium action profile a_{-j}^*, the equilibrium action a_j^* of player j is optimal.

Osborne and Rubinstein: "No player can profitably deviate, given the actions of the other players."

	paper	rock	scissors
paper	0,0	1,-1	-1,1
rock	-1,1	0,0	1,-1
scissors	1,-1	-1,1	0,0

Does Paper, Rock, & Scissors have any Nash equilibria?

	paper	rock	scissors
paper	0,0	1,-1	-1,1
rock	-1,1	0,0	1,-1
scissors	1,-1	-1,1	0,0

Does either player have incentive to deviate from $\langle \text{paper}, \text{paper} \rangle$?

	paper	rock	scissors
paper	0,0	1,-1	-1,1
rock	-1,1	0,0	1,-1
scissors	1,-1	-1,1	0,0

Yes. Player 1 does well to choose scissors instead of paper.

$$u_1(g(\langle \text{scissors}, \text{paper} \rangle)) > u_1(g(\langle \text{paper}, \text{paper} \rangle)).$$

	paper	rock	scissors
paper	0,0	1,-1	-1,1
rock	-1,1	0,0	1,-1
scissors	1,-1	-1,1	0,0

Yes. Player 2 also does well to choose scissors instead of paper.

$$u_2(g(\langle \text{paper}, \text{scissors} \rangle)) > u_2(g(\langle \text{paper}, \text{paper} \rangle)).$$

	paper	rock	scissors
paper	0,0	1,-1	-1,1
rock	-1,1	0,0	1,-1
scissors	1,-1	-1,1	0,0

In fact, Paper, Rock, & Scissors has *no* Nash equilibria.

	hunt stag	hunt hare
hunt stag	2,2	0,1
hunt hare	1,0	1,1

Does Stag Hunt have any Nash equilibria?

	hunt stag	hunt hare
hunt stag	2,2	0,1
hunt hare	1,0	1,1

Does either player have incentive to deviate from $\langle \text{hunt stag}, \text{hunt stag} \rangle$?

	hunt stag	hunt hare
hunt stag	2,2	0,1
hunt hare	1,0	1,1

No. So $\langle \text{hunt stag}, \text{hunt stag} \rangle$ is a Nash equilibrium.

	hunt stag	hunt hare
hunt stag	2,2	0,1
hunt hare	1,0	1,1

Does either player have incentive to deviate from $\langle \text{hunt hare}, \text{hunt hare} \rangle$?

	hunt stag	hunt hare
hunt stag	2,2	0,1
hunt hare	1,0	1,1

No. So $\langle$hunt hare, hunt hare$\rangle$ is a second Nash equilibrium.

In a zero sum game with players 1 and 2, $u_2(o) = -u_1(o)$ so we need include only the utility of player 1 in the game matrix.

	a_1	a_2	a_3
a_1	8,-8	8,-8	7,-7
a_2	0,0	-10,10	-4,4
a_3	9,-9	0,0	-1,1

	a_1	a_2	a_3
a_1	8	8	7
a_2	0	-10	-4
a_3	9	0	-1

Does this game have any Nash equilibria?

	a_1	a_2	a_3
a_1	8	8	7
a_2	0	-10	-4
a_3	9	0	-1

$\langle a_1, a_3 \rangle$ is the only Nash equilibrium.

	a_1	a_2	a_3
a_1	0	8	3
a_2	0	1	10
a_3	-2	6	5

Does this game have any Nash equilibria?

	a_1	a_2	a_3
a_1	0	8	3
a_2	0	1	10
a_3	-2	6	5

$\langle a_1, a_1 \rangle$ and $\langle a_2, a_1 \rangle$ are the Nash equilibria.

	a_1	a_2	a_3	a_4
a_1	1	2	3	1
a_2	0	5	0	0
a_3	1	6	4	1

Does this game have any Nash equilibria?

	a_1	a_2	a_3	a_4
a_1	1	2	3	1
a_2	0	5	0	0
a_3	1	6	4	1

$\langle a_1, a_1 \rangle$, $\langle a_1, a_4 \rangle$, $\langle a_3, a_1 \rangle$, and $\langle a_3, a_4 \rangle$ are the Nash equilbria.

The Nash equilibria of two-player zero sum games have various nice properties.

Let $min(a_j)$ designate the lowest utility obtainable by performing act $a_j \in \mathcal{A}_j$. That is, $min(a_j) = min_{a_{-j} \in \times_{i \in \mathcal{N} \setminus \{j\}} \mathcal{A}_i}(u_j(g(\langle a_j, a_{-j} \rangle)))$.

Def 2.1.5. A *maximinimizer* for player j is an action $a_j^* \in \mathcal{A}_j$ where $min(a_j^*) = max_{a_j \in \mathcal{A}_j}(min(a_j))$.

Thm 2.1.1. $\langle a_1^*, a_2^* \rangle$ is a Nash equilibrium of the two-player zero sum game $\mathcal{G}$ only if $a_1^* \in \mathcal{A}_1$ is a maximinimizer for player 1 and $a_2^* \in \mathcal{A}_2$ is a maximinimizer for player 2.

	a_1	a_2	a_3
a_1	8	8	7
a_2	0	-10	-4
a_3	9	0	-1

Recall that $\langle a_1, a_3 \rangle$ is the only Nash equilibrium of this game.

	a_1	a_2	a_3
a_1	8	8	7
a_2	0	-10	-4
a_3	9	0	-1

$a_1 \in \mathcal{A}_1$ is a maximinimizer for player 1.

	a_1	a_2	a_3
a_1	8	8	7
a_2	0	-10	-4
a_3	9	0	-1

$a_3 \in \mathcal{A}_2$ is a maximinimizer for player 2.

	a_1	a_2	a_3
a_1	0	8	3
a_2	0	1	10
a_3	-2	6	5

Recall that $\langle a_1, a_1 \rangle$ and $\langle a_2, a_1 \rangle$ are the Nash equilibria of this game.

	a_1	a_2	a_3
a_1	0	8	3
a_2	0	1	10
a_3	-2	6	5

$a_1 \in \mathcal{A}_1$ and $a_2 \in \mathcal{A}_1$ are maximinimizers for player 1.

	a_1	a_2	a_3
a_1	0	8	3
a_2	0	1	10
a_3	-2	6	5

$a_1 \in \mathcal{A}_2$ is a maximinimizer for player 2.

Proof of Thm 2.1.1.

If $\langle a_1^*, a_2^* \rangle$ is a Nash equilibrium, then $u_2(g(\langle a_1^*, a_2^* \rangle)) \geq u_2(g(\langle a_1^*, a_2 \rangle))$ for each $a_2 \in \mathcal{A}_2$, so $u_1(g(\langle a_1^*, a_2^* \rangle)) \leq u_1(g(\langle a_1^*, a_2 \rangle))$ for each $a_2 \in \mathcal{A}_2$.

Hence, $min(a_1^*) = u_1(g(\langle a_1^*, a_2^* \rangle)) \leq max_{a_1 \in \mathcal{A}_1}(min(a_1))$.

If $\langle a_1^*, a_2^* \rangle$ is a Nash equilibrium, then $u_1(g(\langle a_1^*, a_2^* \rangle)) \geq u_1(g(\langle a_1, a_2^* \rangle))$ for each $a_1 \in \mathcal{A}_1$, so $u_1(g(\langle a_1^*, a_2^* \rangle)) \geq min(a_1)$ for each $a_1 \in \mathcal{A}_1$.

Hence, $u_1(g(\langle a_1^*, a_2^* \rangle)) \geq max_{a_1 \in \mathcal{A}_1}(min(a_1))$.

Thus, $min(a_1^*) = u_1(g(\langle a_1^*, a_2^* \rangle)) = max_{a_1 \in \mathcal{A}_1}(min(a_1))$. That is, a_1^* is a maximinimizer for player 1.

Similar reasoning establishes that
$min(a_2^*) = u_2(g(\langle a_1^*, a_2^* \rangle)) = max_{a_2 \in \mathcal{A}_2}(min(a_2))$.

Thm 2.1.2. If the two-player zero sum game $\mathcal{G}$ has a Nash equilibrium, then a_1^* and a_2^* are maximinimizers for players 1 and 2 respectively only if $\langle a_1^*, a_2^* \rangle$ is a Nash equilibrium of $\mathcal{G}$.

Proof of Thm 2.1.2. Suppose that $\mathcal{G}$ has a Nash equilibrium.

Then from the Proof of Thm 2.1.1, we know that $max_{a_1 \in \mathcal{A}_1}(min(a_1)) = -max_{a_2 \in \mathcal{A}_2}(min(a_2))$.

Let $max_{a_1 \in \mathcal{A}_1}(min(a_1)) = v^*$. Then $max_{a_2 \in \mathcal{A}_2}(min(a_2)) = -v^*$.

Since a_1^* is a maximinimizer for 1, $u_1(g(\langle a_1^*, a_2 \rangle)) \geq v^*$ for all $a_2 \in \mathcal{A}_2$, so $u_1(g(\langle a_1^*, a_2^* \rangle)) \geq v^*$.

Since a_2^* is a maximinimizer for 2, $u_2(g(\langle a_1, a_2^* \rangle)) \geq -v^*$ for all $a_1 \in \mathcal{A}_1$, so $u_1(g(\langle a_1^*, a_2^* \rangle)) \leq v^*$.

Thus, $u_1(g(\langle a_1^*, a_2^* \rangle)) = v^*$ and $u_2(g(\langle a_1^*, a_2^* \rangle)) = -v^*$, so $\langle a_1^*, a_2^* \rangle$ is a Nash equilibrium of $\mathcal{G}$.

Thm 2.1.3. $\langle a_1^*, a_2^* \rangle$ and $\langle a_1^{**}, a_2^{**} \rangle$ are Nash equilibria of the two-player zero sum game $\mathcal{G}$ only if $u_1(g(\langle a_1^*, a_2^* \rangle)) = u_1(g(\langle a_1^{**}, a_2^{**} \rangle))$ (moreover, $u_2(g(\langle a_1^*, a_2^* \rangle)) = u_2(g(\langle a_1^{**}, a_2^{**} \rangle))$).

That is, all Nash equilibria of $\mathcal{G}$ have the same utilities. The equilibrium payoff v^* to player 1 is the *value* of the game.

Proof of Thm 2.1.3. Suppose that $\langle a_1^*, a_2^* \rangle$ and $\langle a_1^{**}, a_2^{**} \rangle$ are Nash equilibria of $\mathcal{G}$.

From the Proof of Thm 2.1.1, we know that
$u_1(g(\langle a_1^*, a_2^* \rangle)) = u_1(g(\langle a_1^{**}, a_2^{**} \rangle)) = max_{a_1 \in \mathcal{A}_1}(min(a_1))$.

We also know that
$u_2(g(\langle a_1^*, a_2^* \rangle)) = u_2(g(\langle a_1^{**}, a_2^{**} \rangle)) = max_{a_2 \in \mathcal{A}_2}(min(a_2))$.

	a_1	a_2	a_3
a_1	8	8	7
a_2	0	-10	-4
a_3	9	0	-1

The value of this game is 7.

	a_1	a_2	a_3
a_1	0	8	3
a_2	0	1	10
a_3	-2	6	5

The value of this game is 0.

Thm 2.1.4 (Coordination Theorem for Zero Sum Games). $\langle a_1^*, a_2^* \rangle$ and $\langle a_1^{**}, a_2^{**} \rangle$ are Nash equilibria of the two-player zero sum game $\mathcal{G}$ only if $\langle a_1^*, a_2^{**} \rangle$ and $\langle a_1^{**}, a_2^* \rangle$ are Nash equilibria of $\mathcal{G}$.

Proof of Thm 2.1.4. Suppose that $\langle a_1^*, a_2^* \rangle$ and $\langle a_1^{**}, a_2^{**} \rangle$ are Nash equilibria of $\mathcal{G}$.

By Thm 2.1.1, we know that $a_1^* \in \mathcal{A}_1$ and $a_1^{**} \in \mathcal{A}_1$ are maximinimizers for player 1, and $a_2^* \in \mathcal{A}_2$ and $a_2^{**} \in \mathcal{A}_2$ are maximinimizers for player 2.

Thus, by Thm 2.1.2, we know that both $\langle a_1^*, a_2^{**} \rangle$ and $\langle a_1^{**}, a_2^* \rangle$ are Nash equilibria of $\mathcal{G}$.

	a_1	a_2	a_3	a_4
a_1	1	2	3	1
a_2	0	5	0	0
a_3	1	6	4	1

$\langle a_1, a_1 \rangle$ and $\langle a_3, a_4 \rangle$ are Nash equilbria.

	a_1	a_2	a_3	a_4
a_1	1	2	3	1
a_2	0	5	0	0
a_3	1	6	4	1

So $\langle a_1, a_4 \rangle$ and $\langle a_3, a_1 \rangle$ are too.

Unfortunately, not every two-player zero sum game has an action profile in equilibrium.

	paper	rock	scissors
paper	0,0	1,-1	-1,1
rock	-1,1	0,0	1,-1
scissors	1,-1	-1,1	0,0

$$max_{a_1 \in \mathcal{A}_1}(min(a_1)) = -1$$

$$max_{a_2 \in \mathcal{A}_2}(min(a_2)) = -1$$

$$max_{a_1 \in \mathcal{A}_1}(min(a_1)) \neq -max_{a_2 \in \mathcal{A}_2}(min(a_2)).$$

Fortunately, once we allow for *mixed strategies*, every two-player zero sum game has a pair of pure or mixed strategies in Nash equilibirum.

A player's choices can now be *nondeterministic*.

		$[\frac{1}{4}]$	$[\frac{1}{2}]$	$[\frac{1}{4}]$
		paper	rock	scissors
$[\frac{1}{3}]$	paper	0	1	-1
$[\frac{1}{3}]$	rock	-1	0	1
$[\frac{1}{3}]$	scissors	1	-1	0

Player 1's mixed strategy is to play paper with probability $\frac{1}{3}$, rock with probability $\frac{1}{3}$, and scissors with probability $\frac{1}{3}$.

Player 2's mixed strategy is to play paper with probability $\frac{1}{4}$, rock with probability $\frac{1}{2}$, and scissors with probability $\frac{1}{4}$.

A player's choices can now be *nondeterministic*.

		$[\frac{1}{4}]$	$[\frac{1}{2}]$	$[\frac{1}{4}]$
		paper	rock	scissors
$[\frac{1}{3}]$	paper	0	1	-1
$[\frac{1}{3}]$	rock	-1	0	1
$[\frac{1}{3}]$	scissors	1	-1	0

Viewing mixed strategies naïvely, we can think of player 1 and player 2 as committing themselves to randomized chance mechanisms that select actions with various probabilities.

Each $a_i \in \mathcal{A}_i$ is a *pure strategy* of player i.

Let $\triangle(\mathcal{A}_i)$ designate the set of probability measures over $\mathcal{A}_i$.

Each $\alpha_i \in \triangle(\mathcal{A}_i)$ is a *mixed strategy* of player i.

$\alpha_i(a_i) = p_i$ iff player i chooses $a_i \in \mathcal{A}_i$ with probability p_i.

The mixed strategy α_i for player $i \in \mathcal{N}$ with $|\mathcal{A}_i| = n$ where $a_1 \in \mathcal{A}_i$ is chosen with probability p_1, $a_2 \in \mathcal{A}_i$ is chosen with probability p_2, ..., and $a_n \in \mathcal{A}_i$ is chosen with probability p_n can be written as $\langle a_1[p_1], ..., a_n[p_n]\rangle$.

$\langle \text{paper}[\frac{1}{3}], \text{rock}[\frac{1}{3}], \text{scissors}[\frac{1}{3}]\rangle \in \triangle(\mathcal{A}_1)$.

$\langle \text{paper}[\frac{1}{4}], \text{rock}[\frac{1}{2}], \text{scissors}[\frac{1}{4}]\rangle \in \triangle(\mathcal{A}_2)$.

Note that pure strategies can be regarded as special cases of mixed strategies. For instance, $\text{rock} = \langle \text{paper}[0], \text{rock}[1], \text{scissors}[0]\rangle$.

Before we worked with action profiles in $\times_{i \in \mathcal{N}} \mathcal{A}_i$.

We will now work with mixed strategy profiles in $\times_{i \in \mathcal{N}} \triangle(\mathcal{A}_i)$.

For instance,
$\langle\langle \text{paper}[\frac{1}{3}], \text{rock}[\frac{1}{3}], \text{scissors}[\frac{1}{3}]\rangle, \langle \text{paper}[\frac{1}{4}], \text{rock}[\frac{1}{2}], \text{scissors}[\frac{1}{4}]\rangle\rangle$ is a mixed strategy profile in $\triangle(\mathcal{A}_1) \times \triangle(\mathcal{A}_2)$.

What is the expected utility of a mixed strategy profile $\langle \alpha_1, ..., \alpha_{|\mathcal{N}|} \rangle$ for each player $i \in \mathcal{N}$?

Note that each mixed strategy profile determines a probability distribution over the set $\times_{i \in \mathcal{N}} \mathcal{A}_i$ of action profiles.

		$[\frac{1}{4}]$ paper	$[\frac{1}{2}]$ rock	$[\frac{1}{4}]$ scissors
$[\frac{1}{3}]$	paper	0	1	-1
$[\frac{1}{3}]$	rock	-1	0	1
$[\frac{1}{3}]$	scissors	1	-1	0

What is the expected utility of a mixed strategy profile $\langle \alpha_1, ..., \alpha_{|\mathcal{N}|} \rangle$ for each player $i \in \mathcal{N}$?

Note that each mixed strategy profile induces a probability distribution over the set $\times_{i \in \mathcal{N}} \mathcal{A}_i$ of action profiles.

	paper	rock	scissors
paper	0 $[\frac{1}{12}]$	1 $[\frac{1}{6}]$	-1 $[\frac{1}{12}]$
rock	-1 $[\frac{1}{12}]$	0 $[\frac{1}{6}]$	1 $[\frac{1}{12}]$
scissors	1 $[\frac{1}{12}]$	-1 $[\frac{1}{6}]$	0 $[\frac{1}{12}]$

In general, the probability of pure strategy profile $\langle a_1, ..., a_{|\mathcal{N}|} \rangle \in \times_{i \in \mathcal{N}} \mathcal{A}_i$ is $\Pi_{i \in \mathcal{N}} \alpha_i(a_i)$.

The expected utility of mixed strategy profile $\langle \alpha_1, ..., \alpha_{|\mathcal{N}|} \rangle \in \times_{i \in \mathcal{N}} \alpha_i$ for player $i \in \mathcal{N}$ is $\Sigma_{\langle a_1, ..., a_{|\mathcal{N}|} \rangle \in \times_{i \in \mathcal{N}} \mathcal{A}_i} \Pi_{i \in \mathcal{N}} \alpha_i(a_i) u_i(g(\langle a_1, ..., a_{|\mathcal{N}|} \rangle))$.

What is the expected utility of a mixed strategy profile $\langle \alpha_1, ..., \alpha_{|\mathcal{N}|} \rangle$ for each player $i \in \mathcal{N}$?

Note that each mixed strategy profile induces a probability distribution over the set $\times_{i \in \mathcal{N}} \mathcal{A}_i$ of action profiles.

	paper	rock	scissors
paper	0 $[\frac{1}{12}]$	1 $[\frac{1}{6}]$	-1 $[\frac{1}{12}]$
rock	-1 $[\frac{1}{12}]$	0 $[\frac{1}{6}]$	1 $[\frac{1}{12}]$
scissors	1 $[\frac{1}{12}]$	-1 $[\frac{1}{6}]$	0 $[\frac{1}{12}]$

For example,
$EU_1(\langle\langle \text{paper}[\frac{1}{3}], \text{rock}[\frac{1}{3}], \text{scissors}[\frac{1}{3}]\rangle, \langle \text{paper}[\frac{1}{4}], \text{rock}[\frac{1}{2}], \text{scissors}[\frac{1}{4}]\rangle\rangle) =$
$0 \times \frac{1}{12} + 1 \times \frac{1}{6} - 1 \times \frac{1}{12} - 1 \times \frac{1}{12} + 0 \times \frac{1}{6} + 1 \times \frac{1}{12} + 1 \times \frac{1}{12} - 1 \times \frac{1}{6} + 0 \times \frac{1}{12} = 0.$

What is the expected utility of a mixed strategy profile $\langle \alpha_1, \ldots, \alpha_{|\mathcal{N}|} \rangle$ for each player $i \in \mathcal{N}$?

Note that each mixed strategy profile induces a probability distribution over the set $\times_{i \in \mathcal{N}} \mathcal{A}_i$ of action profiles.

	paper	rock	scissors
paper	0 [$\frac{1}{12}$]	1 [$\frac{1}{6}$]	-1 [$\frac{1}{12}$]
rock	-1 [$\frac{1}{12}$]	0 [$\frac{1}{6}$]	1 [$\frac{1}{12}$]
scissors	1 [$\frac{1}{12}$]	-1 [$\frac{1}{6}$]	0 [$\frac{1}{12}$]

For example,
$EU_2(\langle\langle \text{paper}[\frac{1}{3}], \text{rock}[\frac{1}{3}], \text{scissors}[\frac{1}{3}]\rangle, \langle \text{paper}[\frac{1}{4}], \text{rock}[\frac{1}{2}], \text{scissors}[\frac{1}{4}]\rangle\rangle) =$
$0 \times \frac{1}{12} - 1 \times \frac{1}{6} + 1 \times \frac{1}{12} + 1 \times \frac{1}{12} + 0 \times \frac{1}{6} - 1 \times \frac{1}{12} - 1 \times \frac{1}{12} + 1 \times \frac{1}{6} + 0 \times \frac{1}{12} = 0.$

We can now extend the concept of Nash equilibrium to cover mixed strategies.

Def 2.1.6. A *mixed strategy Nash equilibrium* of $\mathcal{G}$ is a mixed strategy profile $\alpha^* \in \times_{i \in \mathcal{N}} \alpha_i$ such that for every player $j \in \mathcal{N}$, the following condition holds:

$EU_j(\langle \alpha_j^*, \alpha_{-j}^* \rangle) \geq EU_j(\langle \alpha_j, \alpha_{-j}^* \rangle)$ for each $\alpha_j \in \triangle(\mathcal{A}_j)$.

In other words, given the other players' equilibrium mixed strategy profile α_{-j}^*, the equilibrium mixed strategy α_j^* of player j is optimal.

Osborne and Rubinstein: "No player can profitably deviate, given the actions of the other players."

Pure strategy Nash equilibria can be regarded as special cases of mixed strategy Nash equilibria.

		$[\frac{1}{4}]$ paper	$[\frac{1}{2}]$ rock	$[\frac{1}{4}]$ scissors
$[\frac{1}{3}]$	paper	0	1	-1
$[\frac{1}{3}]$	rock	-1	0	1
$[\frac{1}{3}]$	scissors	1	-1	0

Is $\langle\langle \text{paper}[\frac{1}{3}], \text{rock}[\frac{1}{3}], \text{scissors}[\frac{1}{3}]\rangle, \langle \text{paper}[\frac{1}{4}], \text{rock}[\frac{1}{2}], \text{scissors}[\frac{1}{4}]\rangle\rangle$ a mixed strategy Nash equilibrium of Paper, Rock, & Scissors?

No. Player 1 does well to play $\langle \text{paper}[1], \text{rock}[0], \text{scissors}[0]\rangle$ instead.

$EU_1(\langle\langle \text{paper}[\frac{1}{3}], \text{rock}[\frac{1}{3}], \text{scissors}[\frac{1}{3}]\rangle, \langle \text{paper}[\frac{1}{4}], \text{rock}[\frac{1}{2}], \text{scissors}[\frac{1}{4}]\rangle\rangle) = 0.$

$EU_1(\langle\langle \text{paper}[1], \text{rock}[0], \text{scissors}[0]\rangle, \langle \text{paper}[\frac{1}{4}], \text{rock}[\frac{1}{2}], \text{scissors}[\frac{1}{4}]\rangle\rangle) = \frac{1}{4}.$

		$[\frac{1}{3}]$	$[\frac{1}{3}]$	$[\frac{1}{3}]$
		paper	rock	scissors
$[\frac{1}{3}]$	paper	0	1	-1
$[\frac{1}{3}]$	rock	-1	0	1
$[\frac{1}{3}]$	scissors	1	-1	0

Is $\langle\langle \text{paper}[\frac{1}{3}], \text{rock}[\frac{1}{3}], \text{scissors}[\frac{1}{3}]\rangle, \langle \text{paper}[\frac{1}{3}], \text{rock}[\frac{1}{3}], \text{scissors}[\frac{1}{3}]\rangle\rangle$ a mixed strategy Nash equilibrium of Paper, Rock, & Scissors?

Yes. If either player plays the equiprobable mixed strategy, the expected utilities of both players are 0.

$$\begin{array}{cc|c|c|}
 & & [\frac{1}{2}] & [\frac{1}{2}] \\
 & & a_1 & a_2 \\
\hline
[\frac{1}{2}] & a_1 & 6 & 3 \\
\hline
[\frac{1}{2}] & a_2 & 2 & 4 \\
\hline
\end{array}$$

Is $\langle\langle a_1[\frac{1}{2}], a_2[\frac{1}{2}]\rangle, \langle a_1[\frac{1}{2}], a_2[\frac{1}{2}]\rangle\rangle$ a mixed strategy Nash equilibrium?

No. $EU_1(\langle\langle a_1[\frac{1}{2}], a_2[\frac{1}{2}]\rangle, \langle a_1[\frac{1}{2}], a_2[\frac{1}{2}]\rangle\rangle) = \frac{15}{4}$.

But $EU_1(\langle\langle a_1[1], a_2[0]\rangle, \langle a_1[\frac{1}{2}], a_2[\frac{1}{2}]\rangle\rangle) = \frac{18}{4}$.

Does this game even have any mixed strategy Nash equilibria?

Thm 2.1.5 (Maximin Theorem for Two-player Zero Sum Games). Every two-person zero sum game has at least one mixed strategy Nash equilibrium. Moreover, the expected utilities of each Nash equilibrium mixed strategy profile are the same.

Partial Proof of Thm 2.1.5. We show that any 2x2 zero sum game in *standard form* has a mixed strategy Nash equilibrium.

	a_1	a_2
a_1	a	b
a_2	c	d

$a, b, c, d \in \mathbb{R}[0, \infty)$ where $a > b$, $a > c$, $d > b$, and $d > c$.

Note that this game does not have any pure strategy Nash equilibria.

$$\begin{array}{cc|c|c|}
 & & [q] & [1-q] \\
 & & a_1 & a_2 \\
\hline
[p] & a_1 & a & b \\
\hline
[1-p] & a_2 & c & d \\
\hline
\end{array}$$

$EU_1(\langle\langle a_1[p], a_2[1-p]\rangle, \langle a_1[q], a_2[1-q]\rangle\rangle) =$
$apq + bp(1-q) + c(1-p)q + d(1-p)(1-q) =$
$apq + bp - bpq + cq - cpq + d - dp - dq + dpq =$
$(a - c + d - b)pq - (d - b)p - (d - c)q + d =$
$Apq - Bp - Cq + D =$
$A((p - \frac{C}{A})(q - \frac{B}{A})) + \frac{DA-BC}{A}$

where

$A = a - b + d - c > 0,\ B = d - b > 0,\ C = d - c > 0,\ D = d > 0.$

$EU_1(\langle\langle a_1[p], a_2[1-p]\rangle, \langle a_1[q], a_2[1-q]\rangle\rangle) = A((p-\frac{C}{A})(q-\frac{B}{A})) + \frac{DA-BC}{A}$.

$EU_2(\langle\langle a_1[p], a_2[1-p]\rangle, \langle a_1[q], a_2[1-q]\rangle\rangle) = -A((p-\frac{C}{A})(q-\frac{B}{A})) - \frac{DA-BC}{A}$.

Player 1 can prevent EU_1 from falling below $\frac{DA-BC}{A}$ by setting $p = \frac{C}{A}$.

Player 2 can prevent EU_2 from falling below $-\frac{DA-BC}{A}$ by setting $q = \frac{B}{A}$.

$\langle\langle a_1[\frac{C}{A}], a_2[1-\frac{C}{A}]\rangle, \langle a_1[\frac{B}{A}], a_2[1-\frac{B}{A}]\rangle\rangle$ is a mixed strategy equilibrium.

The value of the game is $\frac{DA-BC}{A}$.

	a_1	a_2
a_1	6	3
a_2	2	4

What is a mixed strategy Nash equilibrium of this game?

	a_1	a_2
a_1	6	3
a_2	2	4

$$A = a - b + d - c = 5$$
$$B = d - b = 1$$
$$C = d - c = 2$$
$$D = d = 4$$

$$p = \frac{C}{A} = \frac{2}{5}$$
$$q = \frac{B}{A} = \frac{1}{5}$$
$$EU_1(\langle\langle a_1[\tfrac{2}{5}], a_2[\tfrac{3}{5}]\rangle, \langle a_1[\tfrac{1}{5}], a_2[\tfrac{4}{5}]\rangle\rangle) = \frac{2}{25} \times 6 + \frac{8}{25} \times 3 + \frac{3}{25} \times 2 + \frac{12}{25} \times 4 = \frac{90}{25}$$

$$EU_1(\langle\langle a_1[\tfrac{2}{5}], a_2[\tfrac{3}{5}]\rangle, \langle a_1[\tfrac{1}{5}], a_2[\tfrac{4}{5}]\rangle\rangle) = \frac{DA - BC}{A} = \frac{18}{5}$$

	a_1	a_2
a_1	-3	1
a_2	2	0

What is a mixed strategy Nash equilibrium of this game?

	a_1	a_2
a_1	0	4
a_2	5	3

First translation: add the constant $+3$.

	a_1	a_2
a_2	5	3
a_1	0	4

Second translation: switch rows.

	a_1	a_2
a_2	5	3
a_1	0	4

$$A = a - b + d - c = 6$$
$$B = d - b = 1$$
$$C = d - c = 4$$
$$D = d = 4$$

$$p = \frac{C}{A} = \frac{2}{3}$$
$$q = \frac{B}{A} = \frac{1}{6}$$

$$EU_1(\langle\langle a_1[\tfrac{1}{3}], a_2[\tfrac{2}{3}]\rangle, \langle a_1[\tfrac{1}{6}], a_2[\tfrac{5}{6}]\rangle\rangle) = \frac{DA - BC}{A} = \frac{20}{6}$$

	a_1	a_2
a_1	-3	1
a_2	2	0

$$EU_1(\langle\langle a_1[\tfrac{1}{3}], a_2[\tfrac{2}{3}]\rangle, \langle a_1[\tfrac{1}{6}], a_2[\tfrac{5}{6}]\rangle\rangle) = \tfrac{20}{6} - 3 = \tfrac{1}{3}$$

Game Theory

2.2 Nonzero Sum Games

Johns Hopkins University, Spring 2016

Things are no longer so clean once we move on to nonzero sum games. The equilibria of these games often lack the nice properties of equilibria of strictly competitive games.

Ex. Stag Hunt.

You and an acquaintance have gone hunting and you must each decide whether to pursue a large stag or to pursue a small hare. Alone, you can each capture a hare. However, a stag requires two people to capture. If you cooperate and both hunt stag, then you will end up with much meat. But if you hunt stag while your acquaintance hunts hare, then you return home empty-handed. Do you hunt stag or hare?

	hunt stag	hunt hare
hunt stag	2,2	0,1
hunt hare	1,0	1,1

Nash equilibria?

	hunt stag	hunt hare
hunt stag	2,2	0,1
hunt hare	1,0	1,1

$max_{a_1 \in \mathcal{A}_1}(min(a_1)) = 1$

$max_{a_2 \in \mathcal{A}_2}(min(a_2)) = 1$

Moral. Playing maximinimizers will result in the suboptimal Nash equilibrium $\langle \text{hunt hare}, \text{hunt hare} \rangle$. Individual rationality alone cannot lead to the optimal Nash equilibrium $\langle \text{hunt stag}, \text{hunt stag} \rangle$. It is useful to develop incentive mechanisms for cooperation.

Moral. Not all Nash equilibria of a nonzero sum game have the same utilities. We cannot always speak of the *value* of a game.

Ex. Bach or Stravinsky? (Osborne and Rubinstein)

"Two people wish to go out together to a concert of music by either Bach or Stravinsky. Their main concern is to go out together, but one person prefers Bach and the other person prefers Stravinsky."

This is a variant on the classic Battle of the Sexes (Luce and Raiffa).

	Bach	Stravinsky
Bach	3,2	1,1
Stravinsky	0,0	2,3

Nash equilibria?

	Bach	Stravinsky
Bach	3,2	1,1
Stravinsky	0,0	2,3

$max_{a_1 \in \mathcal{A}_1}(min(a_1)) = 1$

$max_{a_2 \in \mathcal{A}_2}(min(a_2)) = 1$

Moral. Playing maximinimizers will result in the non-equilibrium $\langle \text{Bach}, \text{Stravinsky} \rangle$. Individual rationality leads to an unstable solution.

Ex. Hawk & Dove (Osborne and Rubinstein).

"Two animals are fighting over some prey. Each can behave like a dove or like a hawk. The best outcome for each animal is that in which it acts like a hawk while the other acts like a dove; the worst outcome is that in which both animals act like hawks. Each animal prefers to be hawkish if its opponent is dovish and dovish if its opponent is hawkish."

This is a variant on the classic Chicken.

	dove	hawk
dove	2,2	1,3
hawk	3,1	0,0

Nash equilibria?

	dove	hawk
dove	2,2	1,3
hawk	3,1	0,0

$max_{a_1 \in \mathcal{A}_1}(min(a_1)) = 1$

$max_{a_2 \in \mathcal{A}_2}(min(a_2)) = 1$

Moral. Playing maximinimizers will result in the non-equilibrium $\langle \text{Dove}, \text{Dove} \rangle$. Again, individual rationality leads to an unstable solution.

Ex. Prisoner's Dilemma (Raiffa, Flood and Dresher, Tucker).

Two individuals are know to have committed minor offenses but they are also suspected of robbing a bank together. They are arrested and placed in separate interrogation rooms. If they both confess to the robbery, then they will be sentenced to 5 years in prison. If only one confesses, then he will be set free but used as a character witness against the other individual who will receive a sentence of 10 years. If neither confesses, then they will both be convicted of the minor offenses and sentenced to only 1 year in prison.

	confess	do not confess
confess	-5,-5	0,-10
do not confess	-10,0	-1,-1

Nash equilibria?

	confess	do not confess
confess	-5,-5	0,-10
do not confess	-10,0	-1,-1

Moral. The only Nash equilibrium $\langle \text{confess}, \text{confess} \rangle$ is suboptimal. Individual rationality can lead to outcomes that are not best for the group.

Real-life PDs:

- Arms races
- Vampire bats
- Doping in sports
- Drug addiction
- Carbon dioxide emissions

Lewis (and others): Prisoner's Dilemma is a Newcomb Problem.

To confess, you take a transparent box containing \$1000 (like taking the Queen's shilling).

If you do not confess, then the other prisoner receives \$1M.

	confess	do not confess
confess	\$1000,\$1000	\$1M + \$1000,\$0
do not confess	\$0,\$1M+\$1000	\$1M,\$1M

Suppose that whether a prisoner receives \$1M is causally independent of what they do, but that each prisoner believes for good reason that the other prisoner will act like them. Should they confess?

Ex. Traveller's Dilemma (Basu).

An airline loses your suitcase and the suitcase of your doppelgänger that has the exact same contents. An airline manager separates you and your doppelgänger and asks you both to estimate the value of your lost luggage at no less than \$2 and no more than \$100 which is the maximum that the airline will reimburse you. If you both write down the same number, then the manager will treat this as the true value of your luggage and reimburse you both this amount. But if you write down different numbers, then the manager will treat the lower number as the true value. Moreover, whichever one of you wrote down the lower number will be awarded \$2 extra for your honesty, and whichever one of you wrote down the higher number will have \$2 deducted from your payout. What number do you write down?

Partial game matrix:

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

Nash equilibria?

Partial game matrix:

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

The only Nash equilibrium is $\langle \$2, \$2 \rangle$

Partial game matrix:

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

Note that the \$2-\$3 region is a Prisoner's Dilemma.

Ex. Guess $\frac{2}{3}$ of the Average (Ledoux).

n individuals pick a real number between 0 and 1. A prize of \$1 is split between the individuals who pick the number closest to $\frac{2}{3}$ of the average of all the picks.

Nash equilibrium?

Everyone picks 0.

Game Theory

2.3 Extensive Games

Johns Hopkins University, Spring 2016

Interactive choice situations often have a sequential structure. This structure is made explicit in *extensive games*.

An extensive game has *perfect information* if each player, when making his or her moves, knows everything that has previously occurred in the game.

An extensive game has *imperfect information* if each player, when making his or her moves, may have only partial information about what has previously occurred in the game.

We consider only games with perfect information here.

Ex. Sharing Apples.

There are two apples. I will propose an allocation of these apples to you and me. You will decide whether to accept or reject it. If you reject my proposed allocation, then neither of us gets any apples.

Ex. Chain-Store Game (Selten).

You own a chain-store with branches in fifty different U.S. cities. In each city, you face a single potential competitor who might encroach on your turf. If a competitor moves in, you can choose to fight or cooperate. Since fighting is costly, you prefer cooperating to fighting in the face of competition, but your favored outcome is when a potential rival does not move in at all. A potential rival is best off when they move in and you cooperate, but they would rather not compete than fight with you. Now suppose that this all plays out sequentially, city after city, and each potential competitor knows what has happened previously. What is your strategy?

Def 2.3.1. An *extensive game with perfect information* $\mathcal{G} = \langle \mathcal{N}, \mathcal{H}, \mathcal{P}, \mathcal{O}, g, \{\succcurlyeq_i\}_{i \in \mathcal{N}} \rangle$ consists of the following ingredients:

- a finite set of players $\mathcal{N}$
- a set of *histories* $\mathcal{H}$ (sequences of actions) where some subset of these $\mathcal{Z} \subseteq \mathcal{H}$ are the *terminal histories*
- a *player function* $\mathcal{P} : \mathcal{H} \setminus \mathcal{Z} \to \mathcal{N}$ that sends each non-terminal history $h \in \mathcal{H} \setminus \mathcal{Z}$ to the player $i \in \mathcal{N}$ who acts at this juncture
- a set of outcomes $\mathcal{O}$
- a function $g : \mathcal{Z} \to \mathcal{O}$ that sends each terminal history $h \in \mathcal{Z}$ to an outcome $o \in \mathcal{O}$
- a preference relation $\succcurlyeq_i$ for each player over the set of outcomes $\mathcal{O}$

A history $h = \langle a_1, a_2, ... \rangle$ is a sequence of actions by the players of the game (this can be infinite).

$\mathcal{A}(h) = \{a : \langle h, a \rangle \in \mathcal{H}\}$ is the set of actions available after h.

The set of histories $\mathcal{H}$ must satisfy the following three properties:

- $\emptyset \in \mathcal{H}$.
- Every initial sequence of a history is a history: if $\langle a_1, ..., a_n, ... \rangle \in \mathcal{H}$, then $\langle a_1, ..., a_n \rangle \in \mathcal{H}$.
- If every finite initial sequence of an infinite sequence is a history, then the infinite sequence is a history.

Def 2.3.2. A history $h \in \mathcal{H}$ is *terminal* just in case either h is infinite or $h = \langle a_1, ..., a_n \rangle$ and there is no a_{n+1} such that $\langle a_1, ..., a_n, a_{n+1} \rangle \in \mathcal{H}$.

A terminal history is a kind of sequential action profile for all players.

Def 2.3.3. $\mathcal{G}$ is *finite* iff $|\mathcal{H}|$ is finite.

Def 2.3.4. $\mathcal{G}$ has a *finite horizon* iff the longest history in $\mathcal{H}$ is finite.

The extensive game starts at $\emptyset$.

$\mathcal{P}(\emptyset)$ is the player who acts first.

$\mathcal{P}(\emptyset)$ chooses an action in $\mathcal{A}(\emptyset) = \{a : \langle \emptyset, a \rangle \in \mathcal{H}\}$.

Suppose that $\mathcal{P}(\emptyset)$ chooses $a_1 \in \mathcal{A}(\emptyset)$.

$\mathcal{P}(\langle \emptyset, a_1 \rangle)$ is the player who acts next.

$\mathcal{P}(\langle \emptyset, a_1 \rangle)$ chooses an action in $\mathcal{A}(\langle \emptyset, a_1 \rangle) = \{a : \langle \emptyset, a_1, a \rangle \in \mathcal{H}\}$.

Suppose that $\mathcal{P}(\langle \emptyset, a_1 \rangle)$ chooses $a_2 \in \mathcal{A}(\langle \emptyset, a_1 \rangle)$.

$\mathcal{P}(\langle \emptyset, a_1, a_2 \rangle)$ is the player who acts next.

And so on.

If $\mathcal{G}$ has a finite horizon, then this process eventually culminates in a terminal history $\langle \emptyset, a_1, ..., a_n \rangle$.

$g(\langle \emptyset, a_1, ..., a_n \rangle) \in \mathcal{O}$ is the outcome of the game.

The game $\mathcal{G}$ of Sharing Apples consists of:

- $\mathcal{N} = \{1, 2\}$
- $\mathcal{H} = \{\emptyset, \langle\emptyset, 20\rangle, \langle\emptyset, 11\rangle, \langle\emptyset, 02\rangle, \langle\emptyset, 20, a\rangle, \langle\emptyset, 20, r\rangle, \langle\emptyset, 11, a\rangle,$ $\langle\emptyset, 11, r\rangle, \langle\emptyset, 02, a\rangle, \langle\emptyset, 02, r\rangle\}$
  where
  $\mathcal{H} \setminus \mathcal{Z} = \{\emptyset, \langle\emptyset, 20\rangle, \langle\emptyset, 11\rangle, \langle\emptyset, 02\rangle\}$
  $\mathcal{Z} = \{\langle\emptyset, 20, a\rangle, \langle\emptyset, 20, r\rangle, \langle\emptyset, 11, a\rangle, \langle\emptyset, 11, r\rangle, \langle\emptyset, 02, a\rangle, \langle\emptyset, 02, r\rangle\}$
- $\mathcal{P}(\emptyset) = 1$, $\mathcal{P}(\langle\emptyset, 20\rangle) = \mathcal{P}(\langle\emptyset, 11\rangle) = \mathcal{P}(\langle\emptyset, 02\rangle) = 2$
- $\mathcal{O} = \{20, 02, 11, 00\}$
- $g(\langle\emptyset, 20, a\rangle) = 20$
  $g(\langle\emptyset, 20, r\rangle) = 00$
  $g(\langle\emptyset, 11, a\rangle) = 11$
  And so forth.
- $20 \succ_1 11 \succ_1 02 \sim_1 00$
  $02 \succ_2 11 \succ_2 20 \sim_2 00$

An extensive game can be visualized in a tree.

For example, the game $\mathcal{G}$ of Sharing Apples corresponds to the following *game tree*.

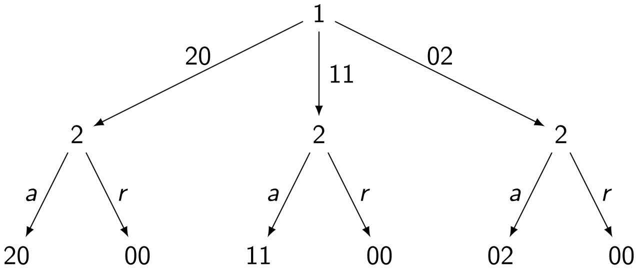

In an extensive game, players can have plans for how to act as the game unfolds.

Def 2.3.5. A *strategy* s_i of player $i \in \mathcal{N}$ is a function that maps each history $h \in \mathcal{H} \setminus \mathcal{Z}$ such that $\mathcal{P}(h) = i$ to an action $a \in \mathcal{A}(h)$.

That is, a strategy dictates how a player will act at *any* potential point in the game where this player is called to action.

In Sharing Apples, a strategy for player 1 will dictate how the apples are allocated. For example, $s_1(\langle \emptyset \rangle) = 11$.

A strategy for player 2 will dictate whether to accept or reject any proposed allocation. For example, $s_2(\langle \emptyset, 20 \rangle) = r$, $s_2(\langle \emptyset, 11 \rangle) = a$, and $s_2(\langle \emptyset, 02 \rangle) = a$.

A strategy must specify an action at every choice point, even those that are unreachable if this very strategy is followed.

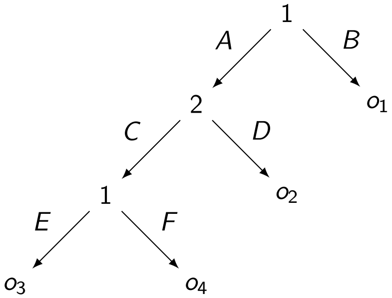

Player 1 has four strategies: AE, AF, BE, and BF (N.B. I will often use abbreviations for strategies).

BE and BF specify an action after history $\langle \emptyset, A, C \rangle$ even though this history will not be actualized if player 1 initially performs action B.

Let $\mathcal{S}_i$ designate the set of strategies for player $i \in \mathcal{N}$.

We will now work with strategy profiles in $\times_{i \in \mathcal{N}} \mathcal{S}_i$.

The outcome $o(s)$ of a strategy profile $s = \langle s_i, ..., s_{|\mathcal{N}|} \rangle$ is the outcome $o \in \mathcal{O}$ that results when each player $i \in \mathcal{N}$ follows their strategy s_i in this profile.

Formally, $o(s) = g(\langle a_1, a_2, .. \rangle)$ where $a_{n+1} = s_{\mathcal{P}(\langle a_1, ..., a_n \rangle)}(\langle a_1, ..., a_n \rangle)$.

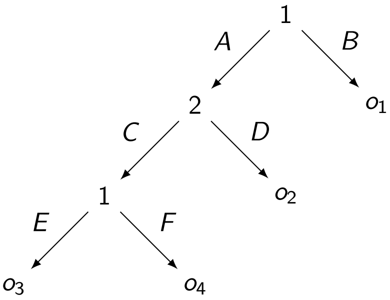

$$o(\langle AE, C\rangle) = ?$$

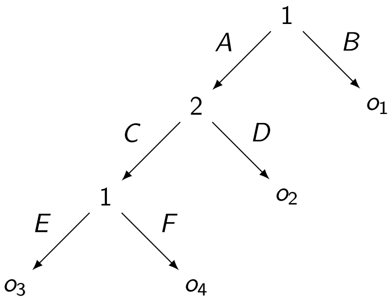

$$o(\langle AE, C \rangle) = o_3$$

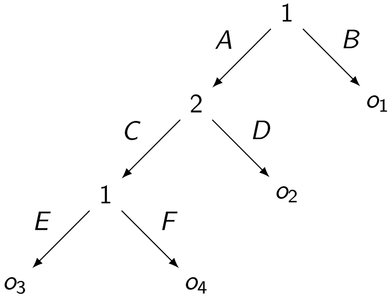

$$o(\langle AF, D\rangle) = ?$$

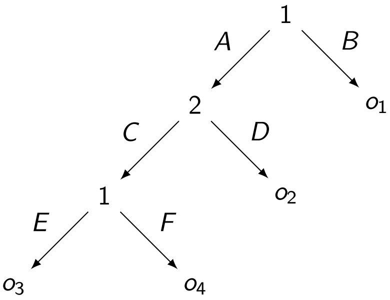

$$o(\langle AF, D\rangle) = o_2$$

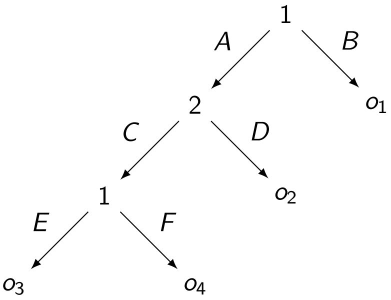

$$o(\langle BE, C\rangle) = ?$$

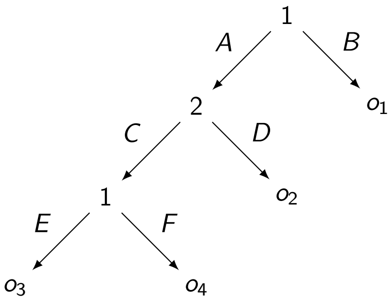

$$o(\langle BE, C \rangle) = o_1$$

Solution concepts such as *Nash equilibrium* for extensive games can be formulated in terms of strategy profiles.

Def 2.3.6. A Nash equilibrium of an extensive game with perfect information $\mathcal{G}$ is a strategy profile $s^* \in \times_{i \in \mathcal{N}} \mathcal{S}_i$ such that for every player $j \in \mathcal{N}$, the following condition holds:

$o(\langle s_j^*, s_{-j}^* \rangle) \succsim_j o(\langle s_j, s_{-j}^* \rangle)$ for each $s_j \in \mathcal{S}_j$.

In other words, given the other players' equilibrium strategy profile s_{-j}^*, the equilibrium strategy s_j^* of player j is optimal.

As before, we can also work with utilities.

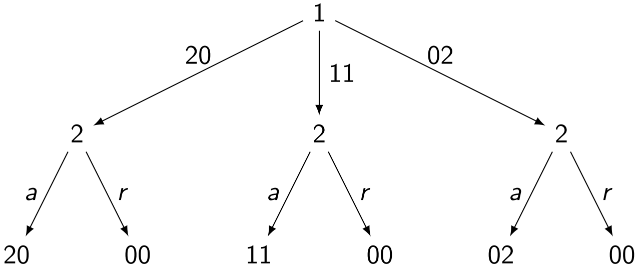

Assume that an apple is valued at a utile.

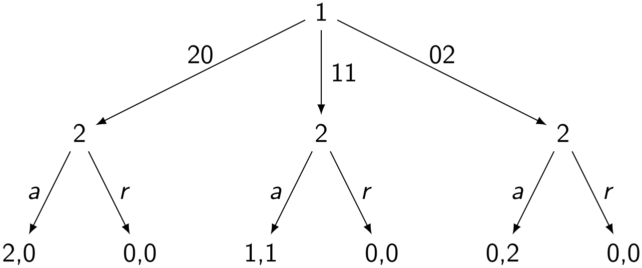

What are the Nash equilibria of Sharing Apples?

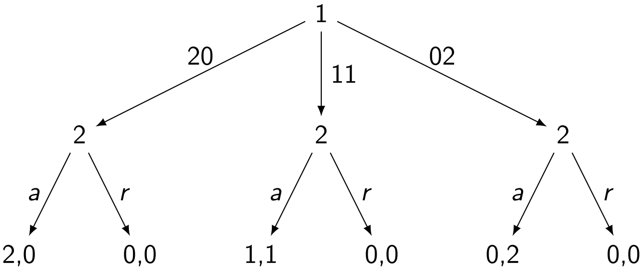

$$\langle 20, aaa\rangle, \langle 20, aar\rangle, \langle 20, ara\rangle, \langle 20, arr\rangle, \langle 20, rra\rangle, \langle 20, rrr\rangle,$$
$$\langle 11, raa\rangle, \langle 11, rar\rangle$$
$$\langle 02, rra\rangle$$

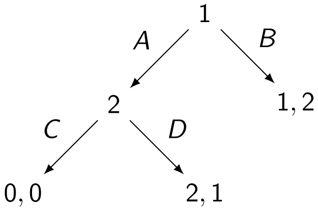

What are the Nash equilibria of this game?

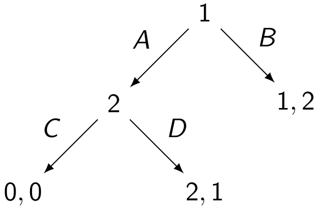

$$\langle A, D \rangle,\ \langle B, C \rangle$$

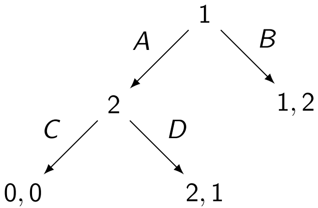

But $\langle B, C \rangle$ is not a stable solution to this game. It is sustained by the threat of player 2 to perform action C. But this action is irrational once player 1 has performed action A.

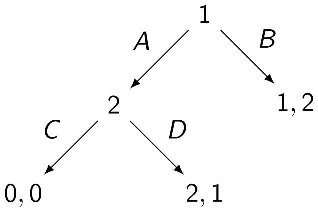

Moral. The solution concept of Nash equilibrium is ill-suited for extensive games since it ignores the sequential structure of these games.

Def 2.3.7. The *subgame* of the extensive game with perfect information $\mathcal{G} = \langle \mathcal{N}, \mathcal{H}, \mathcal{P}, \mathcal{O}, g, \{\succcurlyeq_i\}_{i \in \mathcal{N}} \rangle$ *that follows history h* is the game $\mathcal{G}(h) = \langle \mathcal{N}, \mathcal{H}|_h, \mathcal{P}|_h, \mathcal{O}, g|_h, \{\succcurlyeq_i\}_{i \in \mathcal{N}} \rangle$ consisting of the following ingredients:

- a finite set of players $\mathcal{N}$
- a set of histories $\mathcal{H}|_h$ where $h' \in \mathcal{H}|_h$ iff $\langle h, h' \rangle \in \mathcal{H}$
- a player function $\mathcal{P}|_h$ where $\mathcal{P}|_h(h') = \mathcal{P}(\langle h, h' \rangle)$
- a set of outcomes $\mathcal{O}$
- a function $g|_h : \mathcal{Z}|_h \to \mathcal{O}$ where $g|_h(h') = g(\langle h, h' \rangle)$
- a preference relation $\succcurlyeq_i$ for each player over the set of outcomes $\mathcal{O}$

The game $\mathcal{G}$ of Sharing Apples consists of:

- $\mathcal{N} = \{1, 2\}$
- $\mathcal{H} = \{\emptyset, \langle\emptyset, 20\rangle, \langle\emptyset, 11\rangle, \langle\emptyset, 02\rangle, \langle\emptyset, 20, a\rangle, \langle\emptyset, 20, r\rangle, \langle\emptyset, 11, a\rangle,$ $\langle\emptyset, 11, r\rangle, \langle\emptyset, 02, a\rangle, \langle\emptyset, 02, r\rangle\}$
  where
  $\mathcal{H} \setminus \mathcal{Z} = \{\emptyset, \langle\emptyset, 20\rangle, \langle\emptyset, 11\rangle, \langle\emptyset, 02\rangle\}$
  $\mathcal{Z} = \{\langle\emptyset, 20, a\rangle, \langle\emptyset, 20, r\rangle, \langle\emptyset, 11, a\rangle, \langle\emptyset, 11, r\rangle, \langle\emptyset, 02, a\rangle, \langle\emptyset, 02, r\rangle\}$
- $\mathcal{P}(\emptyset) = 1$, $\mathcal{P}(\langle\emptyset, 20\rangle) = \mathcal{P}(\langle\emptyset, 11\rangle) = \mathcal{P}(\langle\emptyset, 02\rangle) = 2$
- $\mathcal{O} = \{20, 02, 11, 00\}$
- $g(\langle\emptyset, 20, a\rangle) = 20$
  $g(\langle\emptyset, 20, r\rangle) = 00$
  $g(\langle\emptyset, 11, a\rangle) = 11$
  And so forth.
- $20 \succ_1 11 \succ_1 02 \sim_1 00$
  $02 \succ_2 11 \succ_2 20 \sim_1 00$

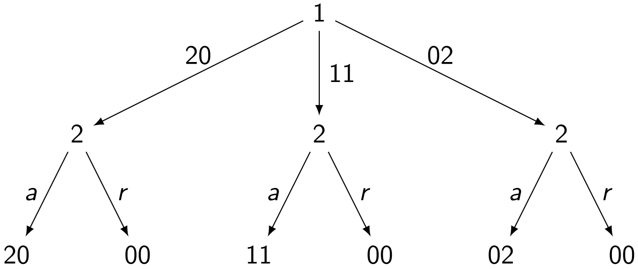
1
20
11
02
2
2
2
a
r
a
r
a
r
20
00
11
00
02
00

The subgame $\mathcal{G}(\langle\emptyset, 11\rangle)$ of Sharing Apples that follows history $\langle\emptyset, 11\rangle$ consists of:

- $\mathcal{N} = \{1, 2\}$
- $\mathcal{H} = \{\emptyset, \langle\emptyset, a\rangle, \langle\emptyset, r\rangle\}$
  where
  $\mathcal{H} \setminus \mathcal{Z} = \{\emptyset\}$
  $\mathcal{Z} = \{\langle\emptyset, a\rangle, \langle\emptyset, r\rangle\}$
- $\mathcal{P}(\emptyset) = 2$
- $\mathcal{O} = \{20, 02, 11, 00\}$
- $g(\langle\emptyset, a\rangle) = 11$
  $g(\langle\emptyset, r\rangle) = 00$
- $20 \succ_1 11 \succ_1 02 \sim_1 00$
  $02 \succ_2 11 \succ_2 20 \sim_2 00$

$$
\begin{tikzcd}
 & 2 \arrow[dl, "a"'] \arrow[dr, "r"] & \\
11 & & 00
\end{tikzcd}
$$

Given a strategy s_i for player $i \in \mathcal{N}$ and a history $h \in \mathcal{H}$ in the extensive game $\mathcal{G}$, let $s_i|_h$ designate the strategy induced by s_i in the subgame $\mathcal{G}(h)$. That is, $s_i|_h(h') = s_i(\langle h, h' \rangle)$ for each $h' \in \mathcal{H}|_h$.

Def 2.3.8. The *subgame perfect equilibrium* of an extensive game with perfect information $\mathcal{G} = \langle \mathcal{N}, \mathcal{H}, \mathcal{P}, \mathcal{O}, g, \{\succcurlyeq_i\}_{i \in \mathcal{N}} \rangle$ is a strategy profile $s^* \in \times_{i \in \mathcal{N}} \mathcal{S}_i$ such that for every player $j \in \mathcal{N}$ and non-terminal history $h \in \mathcal{H} \setminus \mathcal{Z}$ for which $\mathcal{P}(h) = j$, the following condition holds:

$o(\langle s_j^*|_h, s_{-j}^*|_h \rangle) \succcurlyeq_j o(\langle s_j|_h, s_{-j}^*|_h \rangle)$ for each $s_j \in \mathcal{S}_j$.

Equivalently, s^* is a subgame perfect equilibrium of $\mathcal{G}$ just in case $s^*|_h$ is a Nash equilibrium of $\mathcal{G}(h)$ for each $h \in \mathcal{H} \setminus \mathcal{Z}$.

A subgame perfect equilibrium is a Nash equilibrium but the converse needn't be true.

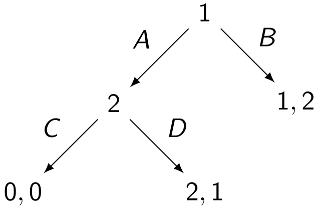

$\langle A, D\rangle$ and $\langle B, C\rangle$ are Nash equilibria.

Are either of these subgame perfect equilibria?

Only $\langle A, D\rangle$ since D but not C is a Nash equilibrium of $\mathcal{G}(\langle\emptyset, A\rangle)$.

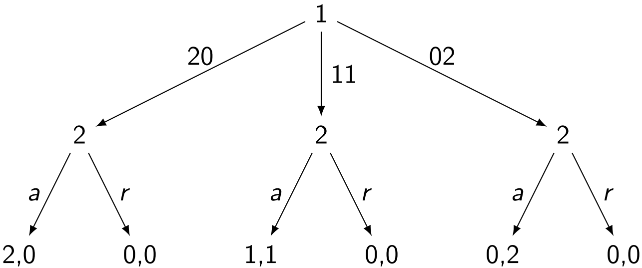

$\langle 20, aaa\rangle$,$\langle 20, aar\rangle$,$\langle 20, ara\rangle$,$\langle 20, arr\rangle$,$\langle 20, rra\rangle$,$\langle 20, rrr\rangle$,
$\langle 11, raa\rangle$,$\langle 11, rar\rangle$,$\langle 02, rra\rangle$ are Nash equilibria.

Are any of these subgame perfect equilibria?

Only $\langle 20, aaa\rangle$ and $\langle 11, raa\rangle$.

Thm 2.3.1. Every finite extensive game with perfect information has a subgame perfect equilibrium.

Since subgame perfect equilibria are Nash equilibria, every finite extensive game with perfect information has a Nash equilibrium.

Example: chess with the rule that the game is a draw when a board position is repeated three times.

Ex. Chain-Store Game (Selten).

You own a chain-store with branches in fifty different U.S. cities. In each city, you face a single potential competitor who might encroach on your turf. If a competitor moves in, you can choose to fight or cooperate. Since fighting is costly, you prefer cooperating to fighting in the face of competition, but your favored outcome is when a potential rival does not move in at all. A potential rival is best off when they move in and you cooperate, but they would rather not compete than fight with you. Now suppose that this all plays out sequentially, city after city, and each potential competitor knows what has happened previously.

What are the subgame perfect equilibria of this game?

The subgame tree for any particular city (where you are player 1 and your rival is player 2):

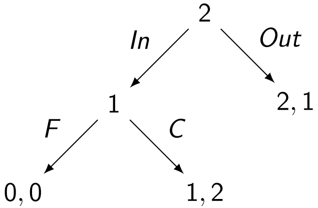

The only subgame perfect equilibrium is $\langle In, C \rangle$.

In the full Chain-Store Game, the unique subgame perfect equilibrium is where the rival enters each city and you cooperate.

But isn't it better to fight at first and develop a reputation for being aggressive?

Ex. Centipede Game (Rosenthal).

$$\begin{tikzcd}
1 \arrow[r] \arrow[d] & 2 \arrow[r] \arrow[d] & 1 \arrow[r] \arrow[d] & 2 \arrow[r] \arrow[d] & 1 \arrow[r] \arrow[d] & 2 \arrow[r] \arrow[d] & 6,5 \\
1,0 & 0,2 & 3,1 & 2,4 & 5,3 & 4,6 &
\end{tikzcd}$$

What are the subgame perfect equilibria of this game?

Each player stops the game and cashes out in each period.

Moral. It is useful to develop incentive mechanisms for cooperation.

Game Theory

2.4 Cooperative Games

Johns Hopkins University, Spring 2016

In a noncooperative game, each player acts alone and the outcome depends on their individual actions.

In a cooperative or *coalition* game, the outcome depends on the joint action of a coalition.

Coalition games ignore the details of how groups of players function internally. Contracts, promises, threats, etc., are not modeled explicitly in these games.

Def 2.4.1. A *coalition game with transferable payoff* $\mathcal{G} = \langle \mathcal{N}, v \rangle$ consists of the following ingredients:

- a finite set of players $\mathcal{N}$
- a function $v : 2^{\mathcal{N}} \setminus \emptyset \mapsto \mathbb{R}$ that assigns every nonempty subset $S \subseteq \mathcal{N}$ of players a real number $v(S)$

$v(S)$ is the *worth* of the *coalition* S. It is the amount available for distribution to the members of S should they act together as a unit.

In other words, each joint action that S can take results in some distribution of $v(S)$ to its members.

It is assumed that the actions of players in $\mathcal{N} \setminus \mathcal{S}$ do not influence $v(S)$.

We consider games with the following property:

Def 2.4.2. A coalition game with transferable payoff $\mathcal{G} = \langle \mathcal{N}, v \rangle$ is *cohesive* iff $v(\mathcal{N}) \geq \sum_{k=1}^{K} v(S_k)$ for every partition $\{S_1, ..., S_K\}$ of $\mathcal{N}$.

This ensures that it is optimal that a coalition of all players forms.

The most basic solution concept for coalition games is *the core*.

Let $x = \langle x_i \rangle_{i \in \mathcal{N}}$ designate a payoff vector across all players.

$x(S) = \sum_{i \in S} x_i$ is the combined payoff to members of the coalition S.

Def 2.4.3. The payoff vector x is S-*feasible* iff $x(S) = v(S)$.

Def 2.4.4. The payoff vector x is *feasible* iff x is $\mathcal{N}$-feasible.

Def 2.4.5. The *core* of the coalition game with transferable payoff $\mathcal{G} = \langle \mathcal{N}, v \rangle$ is the set of payoff vectors $\{x : x$ is feasible and there is no coalition S and S-feasible payoff vector y such that $y_i > x_i$ for each $i \in S\}$.

Equivalently, the core is $\{x : x$ is feasible and $v(S) \leq x(S)$ for each $S\}$.

The core consists of those distributions to all players such that no proper subset of the players would do better to break away and form a coalition.

As with the concept of Nash equilibrium, the basic idea is that each vector in the core is stable, in the sense that no deviation is profitable.

Ex. Three Player Majority.

Working together, three players can obtain 1 util. Any pair of them can obtain $\alpha \in \mathbb{R}[0, 1]$. Alone, each player can obtain nothing.

The game $\mathcal{G}$ of Three Player Majority consists of:

- $\mathcal{N} = \{1, 2, 3\}$
- $v(\{1, 2, 3\}) = 1$
  $v(\{1, 2\}) = v(\{1, 3\}) = v(\{2, 3\}) = \alpha$
  $v(\{1\}) = v(\{2\}) = v(\{3\}) = 0$

The core is $\{x : x(\mathcal{N}) = 1 \text{ and } x(S) \geq \alpha \text{ when } |S| = 2\}$.

If $\alpha > \frac{2}{3}$, then the core is $\emptyset$.

Ex. Treasure Expedition.

A group of n explorers has found treasure in the mountains. Each pair can carry out one piece of treasure.

The game $\mathcal{G}$ of Treasure Expedition consists of:

- $\mathcal{N} = \{1, ..., n\}$
- $v(S) = \begin{cases} \frac{|S|}{2} \text{ if } |S| \text{ is even} \\ \frac{|S|-1}{2} \text{ if } |S| \text{ is odd} \end{cases}$

If $|\mathcal{N}|$ is ≥ 3 and odd, then the core is $\emptyset$.

If $|\mathcal{N}|$ is ≥ 4 and even, then the core is $\{\langle \frac{1}{2}, ..., \frac{1}{2} \rangle\}$.

Def 2.4.6. A coalition game with transferable payoff $\mathcal{G} = \langle \mathcal{N}, v \rangle$ is *simple* iff $v(S) \in \{0, 1\}$ for each S and $v(\mathcal{N}) = 1$.

Def 2.4.7. In a simple game, a coalition S where $v(S) = 1$ is a *winning coalition*.

Def 2.4.8. In a simple game, a *veto player* in $\mathcal{N}$ is part of every winning coalition.

Thm 2.4.1. In every payoff vector in the core of a simple game, each non-veto player must receive a payoff of 0.

Thus, if there are no veto players, then the core is $\emptyset$.

Def 2.4.9. A coalition game with transferable payoff $\mathcal{G} = \langle \mathcal{N}, v \rangle$ is *zero sum* iff $v(S) + v(\mathcal{N} \setminus S) = v(\mathcal{N})$ for each S.

Def 2.4.10. A coalition game with transferable payoff $\mathcal{G} = \langle \mathcal{N}, v \rangle$ is *additive* iff $v(S) + v(T) = v(S \cup T)$ for all S, T such that $S \cap T = \emptyset$.

An additive game is cohesive, but the converse needn't hold.

Thm 2.4.2. The core of a zero sum game that is not additive is $\emptyset$.

Partial Proof of Thm 2.4.2. Assume that $\mathcal{G}$ is zero sum and $v(S) + v(T) < v(S \cup T)$ for some S, T such that $S \cap T = \emptyset$.

Consider some payoff vector x.

If x is in the core, then
$x(S) + x(T) = x(S \cup T) \geq v(S \cup T) > v(S) + v(T)$.

So $x(S) > v(S)$ or $x(T) > v(T)$.

Without loss of generality, assume that $x(S) > v(S)$.

$x(\mathcal{N} \setminus S) = x(\mathcal{N}) - x(S) < v(\mathcal{N}) - v(S)$ since x is feasible.

$v(\mathcal{N}) - v(S) = v(\mathcal{N} \setminus S)$ since $\mathcal{G}$ is zero sum.

Thus, $x(\mathcal{N} \setminus S) < v(\mathcal{N} \setminus S)$ and x cannot be in the core after all.

Def 2.4.11. A *coalition game* $\mathcal{G} = \langle \mathcal{N}, \mathcal{O}, g, \{\succsim_i\}_{i \in \mathcal{N}} \rangle$ consists of the following ingredients:

- a finite set of players $\mathcal{N}$
- a set of outcomes $\mathcal{O}$
- a function $g : 2^{\mathcal{N}} \setminus \emptyset \mapsto 2^{\mathcal{O}}$ that assigns every nonempty subset $S \subseteq \mathcal{N}$ a set of outcomes $g(S) \subseteq \mathcal{O}$
- a preference relation $\succsim_i$ for each player $i \in \mathcal{N}$

Coalition games with transferable payoff are a special case of coalition games:

- $\mathcal{O} = \mathbb{R}^{|\mathcal{N}|}$
- $g(S) = \{x \in \mathbb{R}^{|\mathcal{N}|} : x(S) = v(S) \text{ and } x_j = 0 \text{ for } j \in \mathcal{N} \setminus S\}$
- $x \succsim_i y$ iff $x_i \geq y_i$

Def 2.4.12. The *core* of the coalition game $\mathcal{G} = \langle \mathcal{N}, \mathcal{O}, g, \{\succsim_i\}_{i \in \mathcal{N}} \rangle$ is the set $\{o \in \mathcal{O} : \text{there is no coalition } S \text{ and } o' \in g(S) \text{ such that } o' \succsim_i o \text{ for each } i \in S\}$.

Game Theory

2.5 Application: Morality

Johns Hopkins University, Spring 2016

Gauthier [1991]

"Morality faces a foundational crisis. Contractarianism offers the only plausible resolution of this crisis. These two propositions state my theme. What follows is elaboration."

MacIntyre: "In the actual world which we inhabit the language of morality is in...[a] state of grave disorder...we have—very largely, if not entirely—lost our comprehension, both theoretical and practical, of morality."

Harman: "Moral hypotheses do not help explain why people observe what they observe. So ethics is problematic and nihilism must be taken seriously...An extreme version of nihilism holds that morality is simply an illusion...In this version, we should abandon morality, just as an atheist abandons religion after he has decided that religious facts cannot help explain observations."

Gauthier [1991]

Morality is supposed to justify our choices and actions.

However, it is not clear what grounds moral principles.

Moreover, Expected Utility Theory is also supposed to justify our choices and actions.

Deliberative justification based on EU Theory is more basic than *moral justification* since the former mode "relates to our deep sense of self."

"Deliberative justification does not refute morality. Indeed, it does not offer morality the courtesy of a refutation. It ignores morality, and seemly replaces it. It preempts the arena of justification, apparently leaving morality no room to gain purchase."

If you can't beat them, join them: morality must find its place within EU Theory.

Gauthier [1991]

Recall that individual rationality can lead to suboptimal outcomes in many interactive choice situations.

	confess	do not confess
confess	-5,-5	0,-10
do not confess	-10,0	-1,-1

Gauthier [1991]

The role of morality is to constrain individuals so that we can end up in the (Pareto) optimal outcomes.

	confess	do not confess
confess	-5,-5	0,-10
do not confess	-10,0	-1,-1

Gauthier [1991]

"Each person can see the benefit, to herself, of participating with her fellows in practices requiring each to refrain from the direct endeavor to maximize her own utility, when such mutual restraint is mutually advantageous. No one, of course, can have reason to accept any unilateral constraint on her maximizing behavior; each benefits from, and only from, the constraint accepted by her fellows. But if one benefits more from a constraint on others than one loses by being constrained oneself, one may have reason to accept a practice requiring everyone, including oneself, to exhibit such a constraint. We may represent such a practice as capable of gaining unanimous agreement among rational persons who were choosing the terms on which they would interact with each other. And this agreement is the basis of morality."

Gauthier [1991]

"Morality is not to be understood as a constraint arising from reason alone on the fulfillment of nonrational preferences. Rather, a rational agent is one who acts to achieve the maximal fulfillment of her preferences, and morality is a constraint on the manner in which she acts, arising from the effects of interaction with other agents."

Example: Assisting one's fellows.

Gauthier [1991]

Gauthier is not the first to attempt to derive morality from rationality. The novelty in his approach is really in how he addresses Hobbes' Foole who asks: Granting that it is rational to agree to certain constraints on individual choice, why is it rational to adhere to these constraints in any particular choice situation where you would be better off by breaking them?

Hobbes' answer: a supreme sovereign is needed to detect and punish cheaters.

Gauthier's main ideas are these:

- Accepting a moral practice is a matter of having a certain kind of disposition.
- *Constrained maximizers* who are disposed to comply with moral practices can expect to do better than *straightforward maximizers* who are not disposed to be moral.

Gauthier [1991]

"In plausible circumstances, persons who are genuinely disposed to a more rigorous compliance with moral practices than would follow from their interests at the time of performance can expect to do better than those who are not so disposed. For the former, constrained maximizers as I call them, will be welcome partners in mutually advantageous cooperation, in which each relies on the voluntary adherence of others, from which the latter, straightforward maximizers, will be excluded. Constrained maximizers may thus expect more favorable opportunities than their fellows. Although in assisting their fellows, keeping their promises, and complying with other moral practices, they forgo preference fulfillment that they might obtain, yet they do better overall than those who always maximize expected utility, because of their superior opportunities."

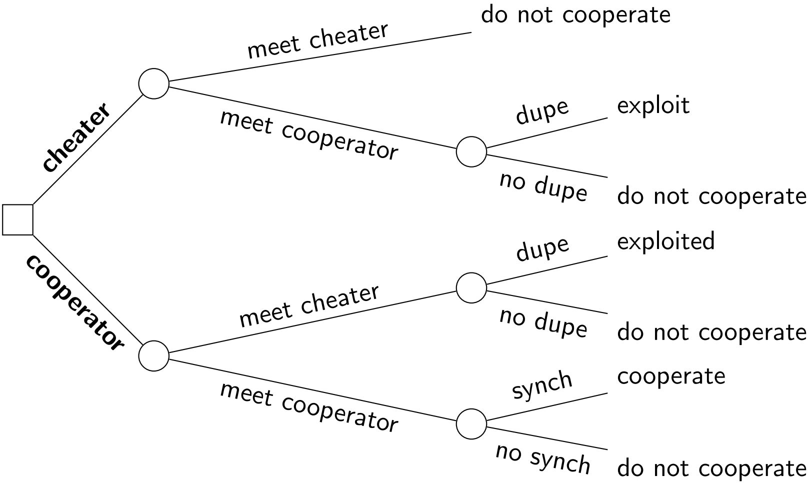
cheater
cooperator
meet cheater
meet cooperator
meet cheater
meet cooperator
do not cooperate
dupe
no dupe
exploit
do not cooperate
dupe
no dupe
exploited
do not cooperate
synch
no synch
cooperate
do not cooperate

$u(\text{exploit}) = 1$
$u(\text{cooperate}) = u'$
$u(\text{do not cooperate}) = u$
$u(\text{exploited}) = 0$

where $0 < u < u' < 1$

$Pr(\text{meet cooperator}) = p$
$Pr(\text{dupe}) = q$
$Pr(\text{synch}) = r$

$EU(\text{cheater}) = u + pq(1 - u)$
$EU(\text{cooperator}) = u + pr(u' - u) - q(1 - p)u$

In the special case where $p = 1$,

$EU(\text{cheater}) = u + q(1 - u)$
$EU(\text{cooperator}) = u + r(u' - u)$

So $EU(\text{cooperator}) > EU(\text{cheater})$ iff $\frac{r}{q} > \frac{1-u}{u'-u}$

Gauthier [1991]

Gauthier's theory is a social contract theory in the spirit of Rawls.

We have not explicitly agreed to our existing moral practice. Rather, moral principles "are those that would secure our agreement ex ante, in an appropriate premoral situation. They are those to which we should have agreed as constituting the terms of our future interaction, had we been, per impossible, in a position to decide those terms. Hypothetical agreement thus provides a test of the justifiability of our existent moral practices."

Gauthier [1991]

Objection. Why is it rational to dispose oneself to accept the constraints that would be agreed upon in a premoral original position? Why isn't it rational to dispose oneself to accept the constraints that are *actually* in play in our society? After all, the latter disposition seems more pertinent to mutually advantageous interaction with one's fellows.

Reply. Constraints that would not secure agreement ex ante are unstable upon reflection on the existing moral order. Individuals whose prospects would be improved by renegotiation can make a strong appeal.

Objection. If there are situations where $EU(\text{cheater}) > EU(\text{cooperator})$, why not agree in the original position to cooperate conditional on society being such that r is high, q is low, and so on? Why agree to cooperate across the board? In the original position, you do not know whether society will be structured in a way that rewards constrained maximization.

Objection (Smith [1999]). Constrained and straightforward maximization are not the only options. Once other kinds of dispositions are taken into account, it is not clear that we should be constrained maximizers.

Smith [1991]

Ex. Fishermen.

"You and I are two fishermen inhabiting adjoining properties on a dangerous coastline. Hidden sandbars often cause our boats to run aground and our catch to be lost. Each of us can expect two such accidents in the coming year, one on our own sandbar, and one on our neighbor's sandbar. If either of us erected a lighthouse, it would prevent any accidents on the adjacent sandbar. The cost to each of us of a single accident is $500, whereas the cost per year of erecting and maintaining a lighthouse is $600."

Smith [1991]

	build	do not build
build	-\$600,-\$600	-\$1100,-\$500
do not build	-\$500,-\$1100	-\$1000,-\$1000

Smith [1991]

	build	do not build
build	-\$600,-\$600	-\$1100,-\$500
do not build	-\$500,-\$1100	-\$1000,-\$1000

Smith [1991]

In Smith's reconstruction of a simple version of Gauthier's argument where dispositions are transparent, each fisherman decides whether to be a constrained maximizer or straightforward maximizer.

CM: Forming the intention to build if you build, or not build if you do not; and then actually building if I expect you to build, or not building if I expect you not to build.

SM: Forming the intention not to build whatever you do; and then actually not building whatever I expect you to do.

Smith [1991]

	CM	SM
CM	-$600,-$600	-$1000,-$1000
SM	-$1000,-$1000	-$1000,-$1000

Smith [1991]

	CM	SM
CM	-$600,-$600	-$1000,-$1000
SM	-$1000,-$1000	-$1000,-$1000

Though $\langle \mathrm{CM}, \mathrm{CM} \rangle$ and $\langle \mathrm{SM}, \mathrm{SM} \rangle$ are both Nash equilibria of this game, CM dominates for each player.

Smith [1991]

However, there are other options besides CM and SM.

A fisherman might decide to be an *unconditional* or *radical cooperator*.

UC: Building one's lighthouse whatever one's partner does.

RC: Building one's lighthouse if and only if one's partner has chosen unconditional cooperation UC.

CM is not the best choice if your partner chooses UC or RC.

Smith [1991]

	CM	SM	UC	RC
CM	-\$600,-\$600	-\$1000,-\$1000	-\$600,-\$600	-\$1000,-\$1000
SM	-\$1000,-\$1000	-\$1000,-\$1000	-\$500,-\$1100	-\$1000,-\$1000
UC	-\$600,-\$600	-\$1100,-\$500	-\$600,-\$600	-\$600,-\$600
RC	-\$1000,-\$1000	-\$1000,-\$1000	-\$600,-\$600	-\$1000,-\$1000

Though $\langle \text{CM}, \text{CM} \rangle$ and $\langle \text{SM}, \text{SM} \rangle$ are still the only Nash equilibria of this game, CM no longer dominates for either player.

Game Theory

2.6 Application II: Convention

Johns Hopkins University, Spring 2016

Lewis' *Convention* [1969]

It is often said that language is conventional. But what does this mean exactly? It is not the case that all of our linguistic conventions could have been agreed upon by a board of syndics. After all, they would have had to speak a rudimentary language to hash out the terms of their agreement.

Lewis' aim is to analyze convention in its full generality, including tacit convention not created by explicit agreement.

I. Coordination and Convention

- Coordination Problems
- Solving Coordination Problems
- Convention (First Pass)

Ex. Meeting (cf. Bach or Stravinsky?).

“Suppose you and I both want to meet each other. We will meet if and only if we go to the same place. It matters little to either of us where (within limits) he goes if he meets the other there; and it matters little to either of us where he goes if he fails to meet the other there. We must each choose where to go. The best place for me to go is the place where you will go, so I try to figure out where you will go and to go there myself. You do the same. Each chooses according to his expectation of the others choice. If either succeeds, so does the other; the outcome is one we both desired.”

Ex. Telephone.

"Suppose you and I are talking on the telephone and we are unexpectedly cut off after three minutes. We both want the connection restored immediately, which it will be if and only if one of us calls back while the other waits. It matters little to either of us whether he is the one to call back or the one to wait. We must each choose whether to call back, each according to his expectation of the other's choice, in order to call back if and only if the other waits."

Ex. Rowing (Hume).

"Suppose you and I are rowing a boat together. If we row in rhythm, the boat goes smoothly forward; otherwise the boat goes slowly and erratically, we waste effort, and we risk hitting things. We are always choosing whether to row faster or slower; it matters little to either of us at what rate we row, provided we row in rhythm. So each is constantly adjusting his rate to match the rate he expects the other to maintain."

Ex. Driving.

"Suppose several of us are driving on the same winding two-lane roads. It matters little to anyone whether he drives in the left or the right lane, provided the others do likewise. But if some drive in the left lane and some in the right, everyone is in danger of collision. So each must choose whether to drive in the left lane or in the right, according to his expectations about the others: to drive in the left lane if most or all of the others do, to drive in the right lane if most or all of the others do (and to drive where he pleases if the others are more or less equally divided)."

Ex. Stag Hunt (Rousseau).

"Suppose we are in a wilderness without food. Separately we can catch rabbits and eat badly. Together we can catch stags and eat well. But if even one of us deserts the stag hunt to catch a rabbit, the stag will get away; so the other stag hunters will not eat unless they desert too. Each must choose whether to stay with the stag hunt or desert according to his expectations about the others, staying if and only if no one else will desert."

Ex. Language.

"Suppose that with practice we could adopt any language in some wide range. It matters comparatively little to anyone (in the long run) what language he adopts, so long as he and those around him adopt the same language and can communicate easily. Each must choose what language to adopt according to his expectations about his neighbors' language: English among English speakers, Welsh among Welsh speakers, Esperanto among Esperanto speakers, and so on."

What do these *coordination problems* have in common? What are their important features?

Def 2.6.1. An n-player game of *pure conflict* $\mathcal{G}$, or *zero sum* game, is one where for any outcome $o \in \mathcal{O}$, $u_1(o) + ... + u_n(o) = 0$ (under some linear rescaling).

Def 2.6.2. An n-player game of *pure coordination* $\mathcal{G}$ is one where for any outcome $o \in \mathcal{O}$, $u_1(o) = ... = u_n(o)$ (under some linear rescaling).

There is a spectrum between games of pure conflict and games of pure coordination. Coordination games are closer to the latter end of this spectrum.

The players' interests rise and fall together.

Ex. Meeting.

	London	Tokyo	Anchorage
London	3,2	1,0	1,1
Tokyo	1,0	2,2	0,1
Anchorage	2,1	0,1	1,1

The maximum difference in a single cell is 1.
The maximum difference across cells is 3.

Ex. Telephone.

	call back	wait
call back	-1,-1	1,1
wait	1,1	-1,-1

The maximum difference in a single cell is 0.
The maximum difference across cells is 2.

This is a game of pure coordination.

Ex. Rowing.

	row fast	row slow
row fast	2,2	-2,-1
row slow	-1,-2	1,1

The maximum difference in a single cell is 1.
The maximum difference across cells is 4.

Ex. Stag Hunt.

	hunt stag	hunt hare
hunt stag	2,2	0,1
hunt hare	1,0	1,1

The maximum difference in a single cell is 1.
The maximum difference across cells is 2.

Def 2.6.3. A *coordination equilibrium* of $\mathcal{G}$ is an action profile $a^* \in \times_{i \in \mathcal{N}} \mathcal{A}_i$ such that for players $j, k \in \mathcal{N}$, the following condition holds:

$u_k(g(\langle a_j^*, a_{-j}^* \rangle)) \geq u_k(g(\langle a_j, a_{-j}^* \rangle))$ for each $a_j \in \mathcal{A}_j$.

No player can profit if any single player deviates from the Nash.

Coordination equilibria are clearly Nash equilibria because of the cases where $j = k$.

However, the converse does not hold.

Ex. Meeting.

	London	Tokyo	Anchorage
London	3,2	1,0	1,1
Tokyo	1,0	2,2	0,1
Anchorage	2,1	0,1	1,2

Ex. Meeting.

	London	Tokyo	Anchorage
London	3,2	1,0	1,1
Tokyo	1,0	2,2	0,1
Anchorage	2,1	0,1	1,2

$\langle$London, London$\rangle$ is a coordinated equilibrium.

Ex. Meeting.

	London	Tokyo	Anchorage
London	3,2	1,0	1,1
Tokyo	1,0	2,2	0,1
Anchorage	2,1	0,1	1,2

$\langle$Tokyo, Tokyo$\rangle$ is a coordinated equilibrium.

Ex. Meeting.

	London	Tokyo	Anchorage
London	3,2	1,0	1,1
Tokyo	1,0	2,2	0,1
Anchorage	2,1	0,1	1,2

$\langle$Anchorage, Anchorage$\rangle$ is a Nash equilibrium but not a coordinated equilibrium.

Ex. Telephone.

	call back	wait
call back	-1,-1	1,1
wait	1,1	-1,-1

Ex. Telephone.

	call back	wait
call back	-1,-1	1,1
wait	1,1	-1,-1

$\langle$call back, wait$\rangle$ and $\langle$wait, call back$\rangle$ are coordinated equilibria.

Ex. Rowing.

	row fast	row slow
row fast	2,2	-2,-1
row slow	-1,-2	1,1

Ex. Rowing.

	row fast	row slow
row fast	2,2	-2,-1
row slow	-1,-2	1,1

$\langle$row fast, row fast$\rangle$ and $\langle$row slow, row slow$\rangle$ are coordinated equilibria.

Ex. Stag Hunt.

	hunt stag	hunt hare
hunt stag	2,2	0,1
hunt hare	1,0	1,1

Ex. Stag Hunt.

	hunt stag	hunt hare
hunt stag	2,2	0,1
hunt hare	1,0	1,1

$\langle$hunt stag, hunt stag$\rangle$ and $\langle$hunt hare, hunt hare$\rangle$ are coordinated equilibria.

Any game of pure coordination will have at least one coordinated equilibrium.

For a problem to be a *coordination* problem, it must have two or more different coordinated equilibria.

But this requirement is not quite strong enough.

	a_1	a_2
a_1	1,1	1,1
a_2	0,0	0,0

$\langle a_1, a_1 \rangle$ and $\langle a_1, a_2 \rangle$ are coordinated equilibria but there is no real problem here: player 1 can ensure coordination by choosing act a_1.

Def 2.6.4. A *proper equilibrium* of $\mathcal{G}$ is an action profile $a^* \in \times_{i \in \mathcal{N}} \mathcal{A}_i$ such that for $j \in \mathcal{N}$, the following condition holds:

$u_j(g(\langle a_j^*, a_{-j}^* \rangle)) > u_j(g(\langle a_j, a_{-j}^* \rangle))$ for each $a_j \in \mathcal{A}_j$.

	a_1	a_2
a_1	1,1	1,1
a_2	0,0	0,0

Neither $\langle a_1, a_1 \rangle$ nor $\langle a_1, a_2 \rangle$ is a proper equilibrium.

A coordination problem will have two or more proper coordinated equilibria.

Isn't this a genuine coordination problem?

	a_1	a_2	a_3
a_1	1,1	1,1	0,0
a_2	1,1	1,1	0,0
a_3	0,0	0,0	1,1

Intuitively, yes. But this problem can be reformulated as follows:

	a_1 or a_2	a_3
a_1 or a_2	1,1	0,0
a_3	0,0	1,1

Def 2.6.5. A *coordination problem* is a situation of "interdependent decision by two or more agents in which coincidence of interest predominates and in which there are two or more proper coordinated equilibria."

Agents might succeed in coordinating by a stroke of good luck.

However, they are more likely to reach a coordinated equilibrium "through the agency of a system of suitably concordant mutual expectations."

"In general, each may do his part of one of the possible coordination equilibria because he expects the others to do theirs, thereby reaching that equilibrium."

"He has a decisive reason to do his own part if he is *sufficiently* confident in his expectation that the others will do theirs. The degree of confidence which is sufficient depends on all his payoffs and sometimes on the comparative probabilities he assigns to the different *ways* the others might not all do their parts, in case not all of them do."

	a_1	a_2	a_3
a_1	1,1	0,0	-8,-8
a_2	0,0	1,1	9,9

Suppose that player 1 knows that if player 2 does not do her part in reaching the coordinated equilibrium $\langle a_1, a_1 \rangle$ by choosing a_1, then she will choose a_2.

$EU_1(a_1) = 1 \times Cr_1(\text{player 2 performs } a_1) + 0 \times Cr_1(\text{player 2 performs } a_2)$.

$EU_1(a_2) = 0 \times Cr_1(\text{player 2 performs } a_1) + 1 \times Cr_1(\text{player 2 performs } a_2)$.

Assuming $Cr_1(\text{player 2 performs } a_1) + Cr_1(\text{player 2 performs } a_2) = 1$,

$EU_1(a_1) > EU_1(a_2)$ iff $Cr_1(\text{player 2 performs } a_1) > 0.5$.

Thus, player 1 will do her part in reaching the coordinated equilibrium $\langle a_1, a_1 \rangle$ iff $Cr_1(\text{player 2 performs } a_1) > 0.5$.

	a_1	a_2	a_3
a_1	1,1	0,0	-8,-8
a_2	0,0	1,1	9,9

Suppose that player 1 knows that if player 2 does not do her part in reaching the coordinated equilibrium $\langle a_1, a_1\rangle$ by choosing a_1, then she is just as likely to choose a_2 as a_3.

$EU_1(a_1) = 1 \times Cr_1(\text{player 2 performs } a_1) - 8 \times Cr_1(\text{player 2 performs } a_3)$.

$EU_1(a_2) = 1 \times Cr_1(\text{player 2 performs } a_2) + 9 \times Cr_1(\text{player 2 performs } a_3)$.

Since $Cr_1(\text{player 2 performs } a_2) = \frac{1 - Cr_1(\text{player 2 performs } a_1)}{2}$ and $Cr_1(\text{player 2 performs } a_3) = \frac{1 - Cr_1(\text{player 2 performs } a_1)}{2}$,

$EU_1(a_1) > EU_1(a_2)$ iff $Cr_1(\text{player 2 performs } a_1) > 0.9$.

Thus, player 1 will do her part in reaching the coordinated equilibrium $\langle a_1, a_1\rangle$ iff $Cr_1(\text{player 2 performs } a_1) > 0.9$.

In some cases, it might be near impossible to coordinate given the high level of mutual confidence required. Example: The millionaire giving away his fortune to 1000 people.

But we can often acquire the mutually concordant expectations required for coordination.

One way to do this is to put ourselves in each other's shoes.

Ex. Meeting.

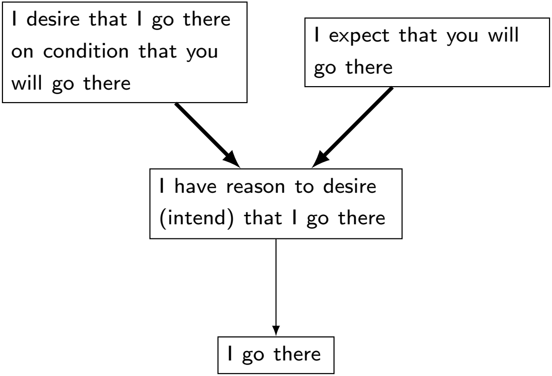
I desire that I go there on condition that you will go there
I expect that you will go there
I have reason to desire (intend) that I go there
I go there

Ex. Meeting.

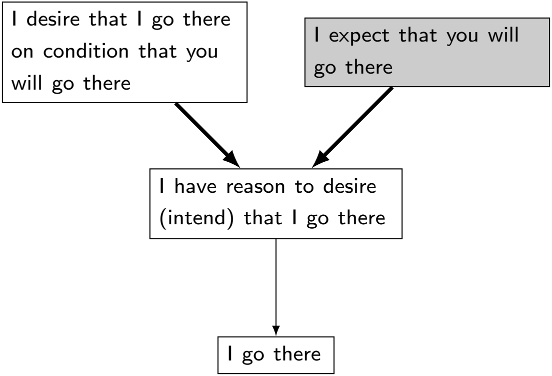

I desire that I go there on condition that you will go there
I expect that you will go there
I have reason to desire (intend) that I go there
I go there

Ex. Meeting.

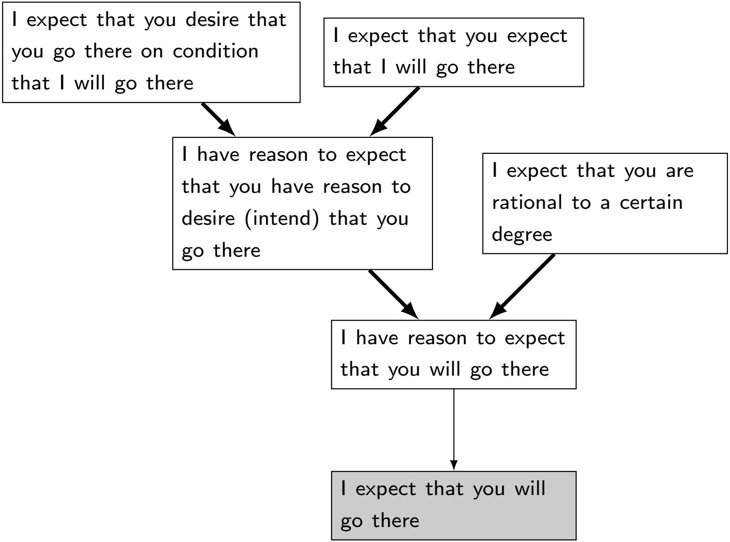
I expect that you desire that you go there on condition that I will go there
I expect that you expect that I will go there
I have reason to expect that you have reason to desire (intend) that you go there
I expect that you are rational to a certain degree
I have reason to expect that you will go there
I expect that you will go there

“Note that replication is *not* an interaction back and forth between people. It is a process in which *one* person works out the consequences of his beliefs about the world—a world he believes to include other people who are working out the consequences of their beliefs, including their belief in other people who...By our interaction in the world we acquire various high-order expectations that can serve us as premises. In our subsequent reasoning we are windowless monads doing our best to mirror each other, mirror each other mirroring each other, and so on.”

We will not generally solve a coordination problem by reasoning from 100th-order expectations. Nevertheless, higher-order expectations provide reasons to do one’s part.

“The more orders, the better.”

One way to engender the concordant mutual expectations required to solve a coordination problem is through explicit agreement—a promise, declaration of intention, and so forth.

If one promises to coordinate, one has a *second* independent incentive to coordinate. Indeed, strong promises might change the payoffs to such a degree that we no longer have a coordination problem.

Ex. Stag Hunt.

	hunt stag	hunt hare
hunt stag	2,2	0,1
hunt hare	1,0	1,1

If breaking a promise lowers one's utility by 2, then the only Nash equilibrium is $\langle$hunt stag, hunt stag$\rangle$.

One way to engender the concordant mutual expectations required to solve a coordination problem is through explicit agreement—a promise, declaration of intention, and so forth.

If one promises to coordinate, one has a *second* independent incentive to coordinate. Indeed, strong promises might change the payoffs to such a degree that we no longer have a coordination problem.

Ex. Stag Hunt.

	hunt stag	hunt hare
hunt stag	2,2	0,-1
hunt hare	-1,0	-1,-1

If breaking a promise lowers one's utility by 2, then the only Nash equilibrium is $\langle$hunt stag, hunt stag$\rangle$.

However, explicit agreement is not the only source of concordant expectations.

"Explicit agreement is an especially good and common means to coordination—so much so that we are tempted to speak of coordination otherwise produced as *tacit* agreement. But agreement (literally understood) is not the only source of concordant expectations to help us solve our coordination problems. We do without agreement by choice if we find ourselves already satisfied with the content and strength of our mutual expectations. We do without it by necessity if we have no way to communicate, or if we can communicate only at a cost that outweighs our improved chance of coordination (say, if we are conspirators being shadowed)."

Experiments by Schelling reveal that we often do well at solving novel coordination problems without communicating. Subjects try to reach a coordination equilibrium that is *salient* in some respect (salience, rather than goodness, is what matters).

One source of salience is *precedent*. If players have already faced the coordination problem before, or a similar coordination problem before, then one equilibrium might be unique in a preeminently conspicuous respect because it, or an analogous equilibrium, was reached in the previous problem.

A fictive precedent can also do the trick. Example: Fabricated story about meeting on Charles Street.

There might be many precedents to follow. Players' actions might conform to a noticeable *regularity*.

We can reach coordination equilibria in new coordination problems by all continuing to conform to this same regularity.

It does not matter why coordination was achieved at particular equilibria in the past.

If there is a regularity in action, each of us needn't be acquainted with the exact same past coordination problems in order to coordinate in future problems. Also, our acquaintance with a precedent needn't be detailed.

"Coordination by precedent, at its simplest, is this: achievement of coordination by means of shared acquaintance with the achievement of coordination in a single past case exactly like our present coordination problem. By removing inessential restrictions, we have come to this: achievement of coordination by means of shared acquaintance with a *regularity* governing the achievement of coordination in a class of past cases which bear some conspicuous analogy to one another and to our present coordination problem. Our acquaintance with this regularity comes from our experience with some of its instances, not necessarily the same ones for everybody."

“Each new action in conformity to the regularity adds to our experience of general conformity. Our experience of general conformity in the past leads us, by force of precedent, to expect a like conformity in the future. And our expectation of future conformity is a reason to go on conforming, since to conform if others do is to achieve a coordination equilibrium and to satisfy one’s own preferences. And so it goes—we’re here because we’re here because we’re here because we’re here. Once the process gets started, we have a metastable self-perpetuating system of preferences, expectations, and actions capable of persisting indefinitely. As long as uniform conformity is a coordination equilibrium, so that each wants to conform conditionally upon conformity by the others, conforming action produces expectation of conforming action and expectation of conforming action produces conforming action.”

Def 2.6.6. "A regularity R in the behavior of members of a population P when they are agents in a recurrent situation S is a *convention* if and only if, in any instance of S among members of P,

- everyone conforms to R
- everyone expects everyone else to conform to R
- everyone prefers to conform to R on condition that the others do, since S is a coordination problem and uniform conformity to R is a proper coordination equilibrium in S."

A convention consists of a regularity in behavior, a system of mutual expectations, and a system of preferences.

Examples:

- Meeting at same place each year.
- Oberlin phone convention: original caller calls back.
- Rowing at same speed.
- Driving in the right lane.
- Hunters always doing the same thing.
- Speaking English.

II. Convention Refined

- Common Knowledge
- Alternatives to Conventions
- Degrees of Convention

How does the system of concordant higher-order mutual expectations that fosters coordination arise?

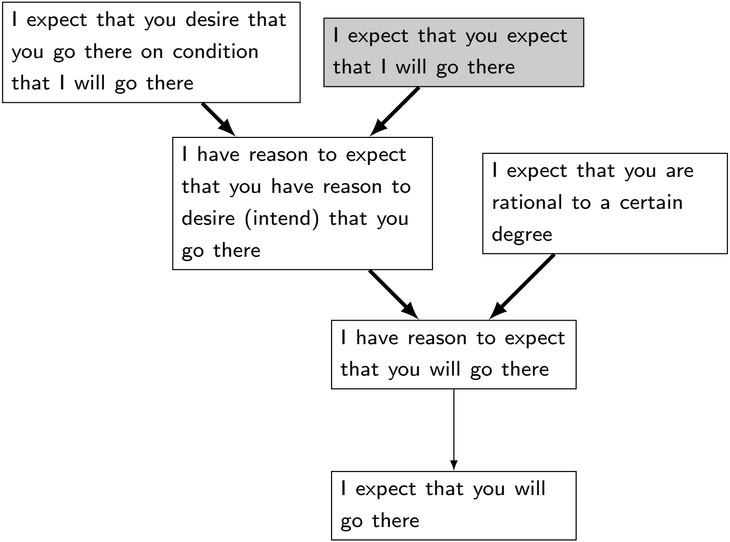

Suppose that we explicitly agree to coordinate. In particular, suppose that the following state of affairs A holds: we have just concluded our meeting in a cafe, we must meet again, and I tell you that I will return to the same cafe tomorrow.

A meets the following conditions:

(i) You and I have reason to believe that A holds
(ii) A indicates to both of us that you and I have reason to believe that A holds
(iii) A indicates to both of us that I will return

where A *indicates* to S that such and such just in case if S had reason to believe that A held then S would thereby have reason to believe that such and such.

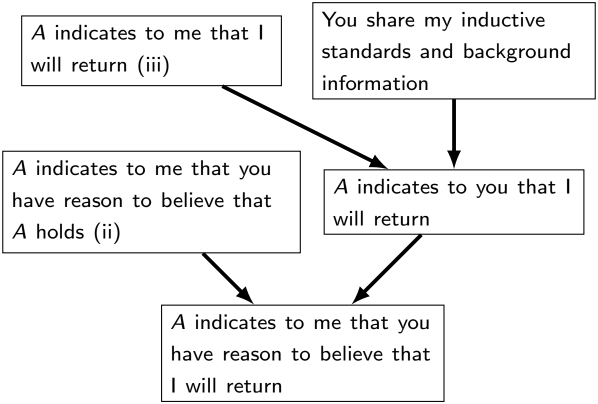
A indicates to me that I will return (iii)
You share my inductive standards and background information
A indicates to me that you have reason to believe that *A* holds (ii)
A indicates to you that I will return
A indicates to me that you have reason to believe that I will return

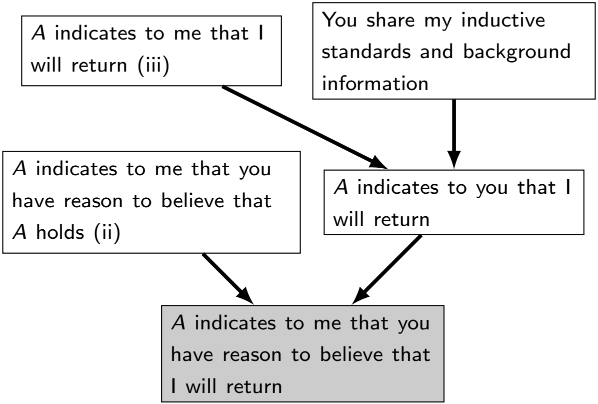

A indicates to me that I will return (iii)
You share my inductive standards and background information
A indicates to me that you have reason to believe that *A* holds (ii)
A indicates to you that I will return
A indicates to me that you have reason to believe that I will return

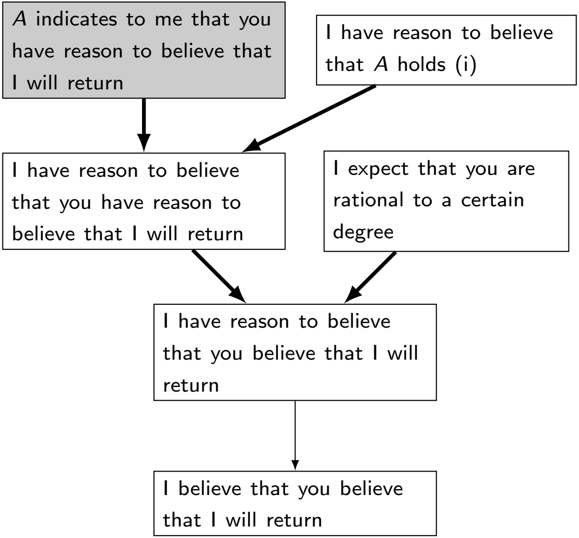
A indicates to me that you have reason to believe that I will return
I have reason to believe that A holds (i)
I have reason to believe that you have reason to believe that I will return
I expect that you are rational to a certain degree
I have reason to believe that you believe that I will return
I believe that you believe that I will return

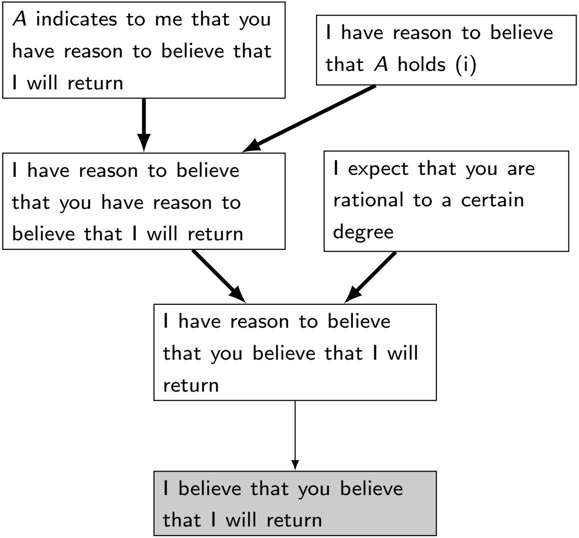
A indicates to me that you have reason to believe that I will return
I have reason to believe that A holds (i)
I have reason to believe that you have reason to believe that I will return
I expect that you are rational to a certain degree
I have reason to believe that you believe that I will return
I believe that you believe that I will return

"I take this example to be typical; all the higher-order expectations involved in sustaining conventions, and more or less all we ever have, seem to be produced in this way."

Def 2.6.7. "It is *common knowledge* in a population P that [such and such] if and only if some state of affairs A holds such that:

- everyone in P has reason to believe that A holds
- A indicates to everyone in P that everyone in P has reason to believe that A holds
- A indicates to everyone in P that [such and such]."

A is a *basis* for common knowledge in P that such and such. Along with mutual ascriptions of rationality, common inductive standards, and background information, A engenders concordant higher-order mutual expectations that such and such is the case.

Past conformity to a convention is a basis for common knowledge about future conformity.

Consider a conventional regularity R in population P and let A be the state of affairs that members of P have conformed to R in the past.

- everyone in P has reason to believe that A holds
- A indicates to everyone in P that everyone in P has reason to believe that A holds
- A indicates to everyone in P that members of P will conform to R going forward

Def 2.6.8. "A regularity R in the behavior of members of a population P when they are agents in a recurrent situation S is a *convention* if and only if, *it is true that, and it is common knowledge in P that*, in any instance of S among members of P,

- everyone conforms to R
- everyone expects everyone else to conform to R
- everyone prefers to conform to R on condition that the others do, since S is a coordination problem and uniform conformity to R is a proper coordination equilibrium in S." (emphasis added)

There is some state of affairs A that, *inter alia*, indicates that these three conditions hold.

Common knowledge is an important feature of the conventions already discussed.

Given the common knowledge requirement, certain regularities that do not seem to be conventions will not count as conventions. Example: The condescending driver case.

Common knowledge is not the only source of higher-order expectations.

“Suppose I am a resident of Ableton and I believe everything printed in the *Ableton Argus*. Today's *Argus* prints this story:

The *Bakerville Bugle* is totally unreliable; what it prints is as likely to be false as to be true. Yet the residents of Bakerville believe everything in it. Today's *Bugle* printed this story:

The *Charlie City Crier* is totally unreliable; what it prints is as likely to be false as to be true. Yet the residents of Charlie City believe everything in it. Today's *Crier* printed this story:

The *Dogpatch Daily* is totally unreliable; what it prints is as likely to be false as to be true. Yet the residents of Dogpatch believe everything in it. Today's *Daily* printed this story:

Tomorrow it will rain cats and dogs.”

I expect the residents of Bakerville to expect the residents of Charlie City to expect the residents of Dogpatch to expect that it will rain cats and dogs. But I do not have lower-order expectations.

The second refinement of the analysis of convention is motivated by regularities in action that cannot be broken down into self-contained coordination problems.

Ex. Oligopoly.

"Suppose that we are contented oligopolists. As the price of our raw material varies, we must each set new prices. It is to no one's advantage to set his prices higher than the others set theirs, since if he does he tends to lose his share of the market. Nor is it to anyone's advantage to set his prices lower than the others set theirs, since if he does he menaces his competitors and incurs their retaliation. So each must set his prices within the range of prices he expects the others to set."

Now, consider a convention where we follow a price leader—that is, one of us initiates price changes in a way that suits all of us. There is a regularity in action. But to think of us as repeatedly facing self-contained coordination problems is precious. Each of us can change our prices whenever we want.

Def 2.6.9. "A regularity R in the behavior of members of a population P when they are agents in a recurrent situation S is a *convention* if and only if, it is true that, and it is common knowledge in P that, in any instance of S among members of P,

- everyone conforms to R
- everyone expects everyone else to conform to R
- everyone has approximately the same preferences regarding all possible combinations of actions
- everyone prefers that everyone conform to R, on condition that at least all but one conform to R
- everyone would prefer that everyone conform to R', on condition that at least all but one conform to R'

where R' is some possible regularity in the behavior of members of P in S, such that no one in any instance of S among members of P could conform both to R' and to R."

- everyone has approximately the same preferences regarding all possible combinations of actions
- everyone prefers that everyone conform to R, on condition that at least all but one conform to R
- everyone would prefer that everyone conform to R', on condition that at least all but one conform to R'

If S is a self-contained interactive choice situation, then the first condition ensures that it is a game of coordination rather than a game of conflict.

If S is a self-contained interactive choice situation, then the second condition ensures that uniform conformity to R is a proper coordination equilibrium.

If S is a self-contained interactive choice situation, then the third condition ensures that uniform conformity to R' is another proper coordination equilibrium.

Thus, the new definition generalizes the old definition by removing the game-theoretic scaffolding.

If R is a convention, then R' might have been our convention instead. Conventions are inherently *arbitrary*.

Since convention requires common knowledge of this arbitrariness, whether a regularity counts as a convention can be sensitive to exposure and open-mindedness to alternatives:

"What is not conventional among narrow-minded and inflexible people, who would not know what to do if others began to behave differently, may be conventional among more adaptable people. What is not conventional may become conventional when news arrives of aliens who behave differently; or when somebody invents a new way of behaving, even a new way no one adopts. When children and the feeble-minded conform to our conventions, they may not take part in them *as* conventions, for they may lack any conditional preference for conformity to an alternative; or they may have the proper preferences, but not as an item of common knowledge. I find these corollaries of our analysis of convention neither welcome nor unwelcome. The analysis is settling questions hitherto left open."

The third and final refinement of the analysis of convention makes things less strict.

Def 2.6.10. "A regularity R in the behavior of members of a population P when they are agents in a recurrent situation S is a *convention* if and only if, it is true that, and it is common knowledge in P that, in *almost* any instance of S among members of P,

- *almost* everyone conforms to R
- *almost* everyone expects *almost* everyone else to conform to R
- *almost* everyone has approximately the same preferences regarding all possible combinations of actions
- *almost* everyone prefers that *any one more* conform to R, on condition that *almost everyone* conform to R
- *almost* everyone would prefer that *any one more* conform to R', on condition that *almost everyone* conform to R'

where R' is some possible regularity in the behavior of members of P in S, such that *almost* no one in *almost* any instance of S among members of P could conform both to R' and to R." (my emphasis)

IV. Convention and Communication

- Sample Signals
- Analysis of Signaling
- Meaning of Signals

Ex. Paul Revere.

The sexton of the Old North Church wants to communicate information about the British army to Paul Revere.

Both the sexton and Paul Revere must choose a contingency plan that will guide their behavior going forward. Each will choose his plan with regard to his expectation of the other's choice.

Some candidate plans for the sexton:

Plan R1:
If the redcoats are observed staying home, hang no lantern.
If the redcoats are observed setting out by land, hang one lantern.
If the redcoats are observed setting out by sea, hang two lanterns.

Plan R2:
If the redcoats are observed staying home, hang one lantern.
If the redcoats are observed setting out by land, hang two lanterns.
If the redcoats are observed setting out by sea, hang no lantern.

Plan R3:
If the redcoats are observed staying home, hang one lantern.
If the redcoats are observed setting out by land, hang no lantern.
If the redcoats are observed setting out by sea, hang two lanterns.

Some candidate plans for Revere:

Plan C1:
If no lantern is observed in the belfry, go home.
If one lantern is observed in the belfry, warn the countryside that the redcoats are coming by land.
If two lanterns are observed in the belfry, warn the countryside that the redcoats are coming by sea.

Plan C2:
If no lantern is observed in the belfry, warn the countryside that the redcoats are coming by sea.
If one lantern is observed in the belfry, go home.
If two lanterns are observed in the belfry, warn the countryside that the redcoats are coming by land.

Plan C3:
If no lantern is observed in the belfry, warn the countryside that the redcoats are coming by land.
If one lantern is observed in the belfry, go home.
If two lanterns are observed in the belfry, warn the countryside that the redcoats are coming by sea.

The choice of contingency plans is a coordination problem.

	C1	C2	C3
R1	2,2	0,0	1,1
R2	0,0	2,2	1,1
R3	1,1	1,1	2,2

In reality, the proper coordination equilibrium $\langle R1, C1 \rangle$ was reached through explicit agreement. The sexton and Paul Revere agreed upon signals for a single occasion.

“I have now described the character of a case of signaling without mentioning the meaning of the signals: that two lanterns meant that the redcoats were coming by sea, or whatever. But nothing important seems to have been left unsaid, so what has been said must somehow imply that the signals have their meanings.”

Examples of signaling conventions:

- International Code of Signals (for ships).
- Helping a truck park.
- Blazing a trail.

Def 2.6.11. A (two-sided) *signaling problem* is a situation S involving a *communicator* and *audience* such that it is common knowledge among them that:

- Exactly one of the *states* $s_1, ..., s_m$ holds and the communicator is well-positioned to know which state holds
- Each member of the audience can perform one of the *responses* $r_1, ..., r_m$
- There is a 1-1 function $F : \{s_i\} \mapsto \{r_i\}$ from states to responses such that everyone prefers that each member of the audience do $F(s_i)$ on condition that s_i holds

  (a function $f(x)$ is 1-1 just in case $f(x) \neq f(y)$ whenever $x \neq y$)
- The communicator can send one of the *signals* $\sigma_1, ..., \sigma_n$ $(n \geq m)$ and the audience is well-positioned to know which signal was sent

Def 2.6.12. A *communicator's contingency plan* $F_c : \{s_i\} \mapsto \{\sigma_i\}$ is a function from states to signals. If F_c is a 1-1 function, then this plan is *admissible*.

Def 2.6.13. An *audience's contingency plan* $F_a : \{\sigma_i\} \mapsto \{r_i\}$ is a 1-1 function from signals to responses. If the ranges of F_a and F coincide, then F_a is *admissible*.

The composition $f \circ g$ is the function $f(g(x))$.

Def 2.6.14. A *signaling system* $\langle F_c, F_a \rangle$ is a system of communicator's and audience's contingency plans where $F_a \circ F_c = F$.

s_1	r_1
s_2	r_2
s_3	r_3
s_4	r_4
s_5	r_5
$\vdots$	$\vdots$
s_m	r_m
states	**responses**

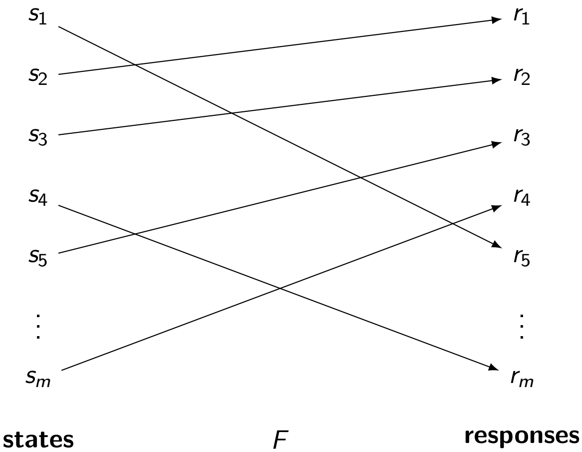
s_1
s_2
s_3
s_4
s_5
s_m
r_1
r_2
r_3
r_4
r_5
r_m
states
F
responses

s_1	σ_1	r_1
s_2	σ_2	r_2
s_3	σ_3	r_3
s_4	σ_4	r_4
s_5	σ_5	r_5
$\vdots$	$\vdots$	$\vdots$
s_m	σ_n	r_m
states	**signals**	**responses**

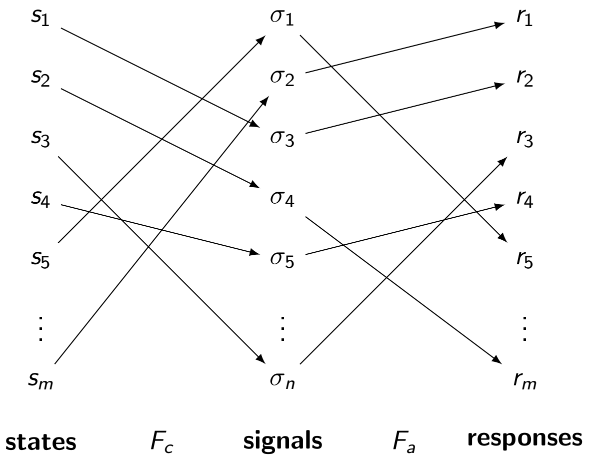
s_1
s_2
s_3
s_4
s_5
s_m
σ_1
σ_2
σ_3
σ_4
σ_5
σ_n
r_1
r_2
r_3
r_4
r_5
r_m
states
F_c
signals
F_a
responses

Ex. Paul Revere.

$$\begin{tikzcd}[column sep=large]
\text{redcoats at home} \arrow[rr] & & \text{go home} \\
\text{redcoats by land} \arrow[rr] & & \text{warn by land} \\
\text{redcoats by sea} \arrow[rr] & & \text{warn by sea} \\
\textbf{states} & F & \textbf{responses}
\end{tikzcd}$$

Ex. Paul Revere.

redcoats at home	⟶	hang no lantern	⟶	go home
redcoats by land	⟶	hang one lantern	⟶	warn by land
redcoats by sea	⟶	hang two lanterns	⟶	warn by sea
states	R1	**signals**	C1	**responses**

Ex. Paul Revere.

$$\begin{tikzcd}
\text{redcoats at home} \arrow[r] & \text{hang no lantern} \arrow[dr] & \text{go home} \\
\text{redcoats by land} \arrow[r] & \text{hang one lantern} \arrow[ur] & \text{warn by land} \\
\text{redcoats by sea} \arrow[r] & \text{hang two lanterns} \arrow[r] & \text{warn by sea} \\
\textbf{states} \arrow[r, phantom, "\text{R1}"] & \textbf{signals} \arrow[r, phantom, "\text{C3}"] & \textbf{responses}
\end{tikzcd}$$

Ex. Paul Revere.

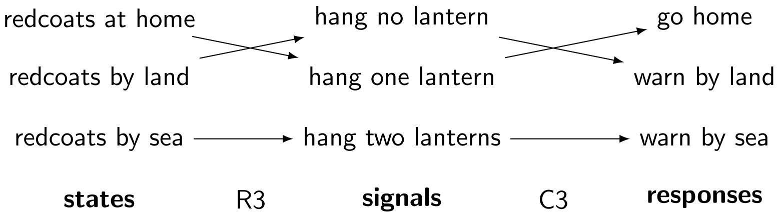

Fact. All and only admissible contingency plans belong to signaling systems.

Fact. In a signaling problem with m states and n signals, there are $\frac{n!}{(n-m)!}$ signaling systems.

In Paul Revere, there are $\frac{3!}{0!} = 6$ signaling systems.

The choice between contingency plans is a coordination problem and signaling systems are proper coordination equilibria (there may be other coordination equilibria besides).

Def 2.6.15. A *signaling convention* is any convention whereby communicators and audiences faced with a signaling problem S do their part of a signaling system $\langle F_c, F_a \rangle$. This system is a *conventional signaling system*.

Def 2.6.16. A *verbal expression* is a finite sequence of types of vocal sounds or marks.

Def 2.6.17. A *verbal signal* is an act of uttering or inscribing a verbal expression.

Much of our language falls under verbal signaling.

"If we endow a hypothetical community with a great many verbal signaling conventions for use in various activities, with verbal expressions suitably chosen *ad hoc*, we shall be able to simulate a community of language users—say, ourselves—rather well. An observer who stayed in the background watching these people use conventional verbal signals as they went about their business might take a long time to realize that they were not ordinary language users. But an observer who tried to converse with them would notice some deficiencies. He would find that every verbal expression they used was conventionally associated with some readily observable state of affairs, or with some definite responsive action, or both. And he would find that they could use only finitely many verbal expressions, so that the conventions governing their verbal signaling could be described by mentioning each expression used."

"Yet it remains true that our hypothetical verbal signalers do not do anything we do not do. We just do more. Their use of language duplicates a fragment of ours."

At this point, we can introduce *meaning* into the picture.

Consider a signaling system $\langle F_c, F_a \rangle$ where $F_c(s) = \sigma$ and $F_a(\sigma) = r$.

On the one hand, we might say that σ *means* in $\langle F_c, F_a \rangle$ that s holds.

On the other hand, we might say that σ *means* in $\langle F_c, F_a \rangle$ to do r.

In the former case, σ is an *indicative* signal.

In the latter case, σ is an *imperative* signal.

A signal is *neutral* if it is equally properly called an indicative and imperative signal.

Whether a signal counts as indicative, imperative, or neutral depends on whether the communicator or audience has to engage in significant deliberation in the signaling system.

We can also introduce *truth* into the picture.

If σ is an indicative signal that s holds in system $\langle F_c, F_a \rangle$, then we might say that σ is *true in* $\langle F_c, F_a \rangle$ when s holds and σ is *false in* $\langle F_c, F_a \rangle$ when s does not hold.

Since signals are actions, we are ascribing truth/falsity to actions. But in the case of verbal signals, we might also ascribe truth/falsity to verbal expressions or their tokenings.

To engage in a conventional signaling system $\langle F_c, F_a \rangle$ is to follow a convention of *truthfulness in* $\langle F_c, F_a \rangle$.

"In any instance of S among members of P, the communicator tries to give whichever signal is true *under the prevailing convention* in that instance, and the audience responds by doing whatever seems best on the assumption that he has succeeded in so doing...We can say that a signal σ is *true* in P in any instance of S if and only if there is some suitable signaling system that is conventionally adopted in P and σ is true in that signaling system in that instance of S."

If $\langle F_c, F_a \rangle$ is a verbal signaling system, then we can think of this system as a *language* $\mathcal{L}$.

Speakers of this language follow a convention of *truthfulness in* $\mathcal{L}$.

Social Choice Theory

3.1 Arrow's Theorem

Johns Hopkins University, Spring 2016

Suppose that a heterogeneous group of individuals, or *citizens*, must choose between two or more group actions or policies. Each citizen has their own preferences over the set of alternatives. The problem of social choice is to amalgamate their *group preference profile* into a single social preference ordering.

Ex. City Council.

The City of Toronto has some extra cash in its municipal budget and the mayor forms a three-person committee to help decide how to spend this money. The options on the table are a new park, a new homeless shelter, and the development of more bike lanes. The committee members have the following preferences:

Member 1	Member 2	Member 3
Park	Shelter	Bike Lanes
Shelter	Bike Lanes	Park
Bike Lanes	Park	Shelter

What group preference relation over the options should the committee submit to the mayor?

Some background assumptions:

- The set of citizens $\mathcal{N}$ is finite.
- The set of alternatives $\mathcal{O}$ is finite.
- Each citizen's preferences $\succcurlyeq_i$ over $\mathcal{O}$ can be exhibited by an ordinal utility function.

We will focus on a particular class of collective choice rules:

Def 3.1.1. A *social welfare function* (SWF) for $\langle \mathcal{N}, \mathcal{O} \rangle$ operates on a group preference profile $\{\succcurlyeq_i\}_{i \in \mathcal{N}}$ and yields a social ordering $\succcurlyeq_s$ over $\mathcal{O}$ that can be exhibited by an ordinal utility function.

Majority Rule:
If more than half of the citizens in $\mathcal{N}$ prefer o_1 to o_2, then $o_1 \succ_s o_2$.
If more than half of the citizens in $\mathcal{N}$ prefer o_2 to o_1, then $o_2 \succ_s o_1$.
Otherwise $o_2 \sim_s o_1$.

Majority Rule for $\langle\{\text{Member 1}, \text{Member 2}, \text{Member 3}\}, \{\text{Bike Lanes}, \text{Park}, \text{Shelter}\}\rangle$ does not always deliver an ordering.

Member 1	Member 2	Member 3
Park	Shelter	Bike Lanes
Shelter	Bike Lanes	Park
Bike Lanes	Park	Shelter

Park $\succ_s$ Shelter $\succ_s$ Bike Lanes $\succ_s$ Park.

This is known as the *voting paradox*.

We will want our SWFs to satisfy certain reasonable conditions. Some of these are discussed by the economist Kenneth Arrow.

Unrestricted Domain (U): The SWF produces a social ordering $\succcurlyeq_s$ for *every* possible preference profile for the citizenry.

This rules out Majority Rule.

Non-Dictatorship (D): It is not the case that there is some citizen $i \in \mathcal{N}$ such that for any $o_1, o_2 \in \mathcal{O}$, the SWF yields $o_1 \succ_s o_2$ whenever $o_1 \succ_i o_2$.

Curly	**Moe**	Larry	Society
Fish	Fish	Lamb	Fish
Burger,Lamb	Chicken	Fish	Chicken
Chicken	Burger	Burger,Chicken	Burger
	Lamb		Lamb

A *constant* SWF that always returns the same social ordering $\succcurlyeq_s$ satisfies U and D.

However, this kind of SWF is problematic because it *imposes* an ordering from the outside. The output of a constant SWF can fail to reflect the preferences of the citizens in society.

The following condition is designed to exclude constant SWFs:

Citizens' Sovereignty (CS): For every $o_1, o_2 \in \mathcal{O}$, there is at least one group preference profile for which the SWF yields a social ordering where $o_1 \succ_s o_2$.

However, some SWFs that satisfy U, D, and CS are still problematic.

For example, consider a SWF that yields $o_1 \succ_s ... \succ_s o_n$ when all of the citizens have the same preferences and yields $o_n \succ_s ... \succ_s o_1$ otherwise.

This SWF is still insensitive to the preferences of individual citizens.

The next condition is meant to restore this sensitivity:

Positive Association (PA): For every $o_1, o_2 \in \mathcal{O}$, if the SWF yields $o_1 \succ_s o_2$ for a given group preference profile, then the SWF yields $o_1 \succ_s o_2$ for any group preference profile that is just like the original profile except that one or more of the citizens has moved o_1 up in their ordering.

Positive Association (PA): For every $o_1, o_2 \in \mathcal{O}$, if the SWF yields $o_1 \succ_s o_2$ for a given group preference profile, then the SWF yields $o_1 \succ_s o_2$ for any group preference profile that is just like the original profile except that one or more of the citizens has moved o_1 up in their ordering.

Curly	Moe	Larry	Society
Fish	Fish	Lamb	Fish
Burger,Lamb	Chicken	Fish	**Chicken**
Chicken	Burger	Burger,**Chicken**	**Burger**,Lamb
	Lamb		

Positive Association (PA): For every $o_1, o_2 \in \mathcal{O}$, if the SWF yields $o_1 \succ_s o_2$ for a given group preference profile, then the SWF yields $o_1 \succ_s o_2$ for any group preference profile that is just like the original profile except that one or more of the citizens has moved o_1 up in their ordering.

Curly	Moe	Larry	Society
Fish,**Chicken**	Fish	Lamb	Fish,**Chicken**
Burger,Lamb	Chicken	Fish,**Chicken**	**Burger**,Lamb
	Burger	Burger	
	Lamb		

The final condition is this:

Independence of Irrelevant Alternatives (I): For every $o_1, o_2 \in \mathcal{O}$, if each citizen ranks $o_1, o_2 \in \mathcal{O}$ in the same order in the group preference profiles $\{\succsim_i\}_{i \in \mathcal{N}}$ and $\{\succsim'_i\}_{i \in \mathcal{N}}$, then $o_1, o_2 \in \mathcal{O}$ must be in the same order with respect to each other in the social orderings $\succsim_s$ and $\succsim'_s$ yielded by the SWF from the two group profiles.

Roughly, condition I requires that the SWF obtain $\succsim_s$ by comparing alternatives two at a time in isolation from the other alternatives.

Curly	Moe	Larry	Society
Fish	Fish	Lamb	Fish
Burger,Lamb	Chicken	Fish	Chicken
Chicken	Burger	Burger,Chicken	Burger,Lamb
	Lamb		

In determining the social ordering of Lamb and Fish, it does not matter that Lamb is Moe's worst option and Larry's best option. It matters only that Moe prefers Fish to Lamb and Larry prefers Lamb to Fish. Beef and Chicken are 'irrelevant'.

Condition I excludes rank-ordering methods.

Suppose that a best option gets 1 point, a second-best option gets 2 points, and so on, and that society prefers options with fewer total points.

Voter 1	Voter 2	Voter 3	Society
Reagan	**Carter**	Anderson	**Reagan**, **Carter**, Anderson
Carter	Anderson	**Reagan**	
Anderson	**Reagan**	**Carter**	

Condition I excludes rank-ordering methods.

Suppose that a best option gets 1 point, a second-best option gets 2 points, and so on, and that society prefers options with fewer total points.

Voter 1	Voter 2	Voter 3	Society
Reagan	**Carter**	**Reagan**	**Reagan**
Carter	**Reagan**	**Carter**	**Carter**
Anderson	Anderson	Anderson	Anderson

Though each individual ordering of Reagan and Carter remains unchanged, the social ordering of these candidates is different.

We are now in a position to state our main (negative) result:

Thm 3.1.1 (Arrow's Impossibility Theorem). Where $|\mathcal{O}| \geq 3$ and $|\mathcal{N}| \geq 2$, there is no SWF for $\langle \mathcal{N}, \mathcal{O} \rangle$ that meets conditions U, D, CS, PA, and I.

If $|\mathcal{O}| = 1$ or $|\mathcal{N}| = 1$, there is no social choice to be made.

If $|\mathcal{O}| = 2$ and $|\mathcal{N}| \geq 2$, then Majority Rule for $\langle \mathcal{N}, \mathcal{O} \rangle$ meets all five of Arrow's conditions.

Going forward, I assume that $|\mathcal{O}| \geq 3$ and $|\mathcal{N}| \geq 2$.

In order to prove Arrow's Theorem, we can replace conditions CS and PA with the following single condition:

Pareto (P): Given a group preference profile where $o_1 \succ_i o_2$ for each $i \in \mathcal{N}$, the SWF yields $o_1 \succ_s o_2$.

The Pareto condition is sometimes called the condition of *unanimity rule*.

Lem 3.1.1. P implies CS.

Lem 3.1.2. P does not imply PA.

Lem 3.1.3 (Pareto Lemma). CS, PA, and I together imply P.

Proof of Lem 3.1.2. Consider the following social choice situation:

$\mathcal{N} = \{\text{Huey}, \text{Dewey}, \text{Louie}\}$

$\mathcal{O} = \{\text{apples}, \text{oranges}\}$

Suppose that a SWF yields Huey's preferences if Dewey and Louie have the same preferences, and the SWF yields Dewey's preferences otherwise.

This SWF clearly satisfies condition P.

However, this SWF does not meet condition PA:

Huey	Dewey	Louie	Society
apples	apples,oranges	apples,oranges	apples
oranges			oranges

Huey	Dewey	Louie	Society
apples	apples,oranges	apples	apples,oranges
oranges		oranges	

Proof of Lem 3.1.3. Suppose that a SWF satisfies conditions CS, PA, and I, and consider a group preference profile $\{\succcurlyeq_i\}_{i\in\mathcal{N}}$ where $o_1 \succ_i o_2$ for each $i \in \mathcal{N}$. We must show that the SWF yields $o_1 \succ_s o_2$ for this profile.

Since condition I is in effect, we can restrict our attention to o_1 and o_2. Let us assume that these are the only alternatives: $\mathcal{O} = \{o_1, o_2\}$.

By condition CS, there is some group preference profile $\{\succcurlyeq_i'\}_{i\in\mathcal{N}}$ for which the SWF yields $o_1 \succ_s o_2$.

If $\{\succcurlyeq_i'\}_{i\in\mathcal{N}}=\{\succcurlyeq_i\}_{i\in\mathcal{N}}$, we are done.

If $\{\succcurlyeq_i'\}_{i\in\mathcal{N}} \neq \{\succcurlyeq_i\}_{i\in\mathcal{N}}$, then note that $\{\succcurlyeq_i\}_{i\in\mathcal{N}}$ is just like $\{\succcurlyeq_i'\}_{i\in\mathcal{N}}$ except that one or more of the citizens has moved o_1 up in their ordering.

By condition PA, the SWF yields $o_1 \succ_s o_2$ for $\{\succcurlyeq_i\}_{i\in\mathcal{N}}$ as desired.

Thm 3.1.1 (Arrow's Impossibility Theorem). There is no SWF for $\langle \mathcal{N}, \mathcal{O} \rangle$ that meets conditions U, D, CS, PA, and I.

To prove Thm 3.1.1, it suffices to prove the following theorem:

Thm 3.1.2 (Arrow's Impossibility Theorem, Pareto Version). There is no SWF for $\langle \mathcal{N}, \mathcal{O} \rangle$ that meets conditions U, D, P, and I.

Suppose that there is some SWF for $\langle \mathcal{N}, \mathcal{O} \rangle$ that meets conditions U, D, CS, PA, and I. By the Pareto Lemma, this SWF meets conditions U, D, P, and I. Thus, if no SWF meets U, D, P, and I, then no SWF meets U, D, CS, PA, and I.

Thm 3.1.2 can be restated thus:

Any SWF for $\langle \mathcal{N}, \mathcal{O} \rangle$ that meets conditions U, P, and I is dictatorial.

Proof Sketch of Thm 3.1.2.

Sen [1995] provides an elegant proof.

We need some definitions:

Def 3.1.2. A group of citizens $S \subseteq \mathcal{N}$ is *decisive for* $\langle o_1, o_2 \rangle$ just in case the SWF yields $o_1 \succ_s o_2$ whenever $o_1 \succ_i o_2$ for each $i \in S$.

Def 3.1.3. A group of citizens $S \subseteq \mathcal{N}$ is *decisive* just in case S is decisive for every ordered pair of alternatives in $\mathcal{O} \times \mathcal{O}$.

Def 3.1.4. A citizen $i \in \mathcal{N}$ is a *dictator* just in case $\{i\}$ is decisive.

Def 3.1.5. The SWF is *dictatorial* just in case some $i \in \mathcal{N}$ is a dictator.

Suppose that a SWF satisfies U, P, and I.

Lem 3.1.4 (Field-Expansion Lemma). If $S \subseteq \mathcal{N}$ is decisive for any ordered pair $\langle o_1, o_2 \rangle \in \mathcal{O} \times \mathcal{O}$, then S is decisive.

This can be used to prove:

Lem 3.1.5 (Group-Contraction Lemma). If $S \subseteq \mathcal{N}$ is decisive, then either $|S| = 1$ or some smaller group $S' \subset S$ is decisive.

The argument now proceeds as follows.

By condition P, $\mathcal{N}$ is decisive.

By the Group-Contraction Lemma, there is some group $S \subset \mathcal{N}$ that is decisive.

If $S = \{i\}$ for some $i \in \mathcal{N}$, then i is a dictator and the SWF is dictatorial.

Otherwise, by the Group-Contraction Lemma, there is some group $S' \subset S$ that is decisive.

And so forth. Since $\mathcal{N}$ is finite, this process halts.

By Arrow's Impossibility Theorem, we know that no social welfare function meets conditions U, D, CS, PA, and I (alternatively, no SWF meets conditions U, D, P, and I).

In the aftermath of Arrow's result, social choice theorists have studied which SWFs satisfy weaker sets of conditions.

Note that Majority Rule satisfies D, CS, PA, and I.

In fact, Majority Rule (defined only over the group preference profiles for which it delivers orderings) is the only SWF that satisfies certain strengthenings of Arrow's conditions.

Anonymity (A): If the SWF yields a social ordering $\succcurlyeq_s$ for $\{\succcurlyeq_i\}_{i\in\mathcal{N}}$, and $\{\succcurlyeq'_i\}_{i\in\mathcal{N}}$ is obtainable from $\{\succcurlyeq_i\}_{i\in\mathcal{N}}$ by exchanging some preference relations among the citizens, then the SWF yields $\succcurlyeq_s$ for $\{\succcurlyeq'_i\}_{i\in\mathcal{N}}$.

Tweedledee	Tweedledum	Society
Checkers	Chess	Checkers
Cards	Checkers	Cards
Croquet	Croquet,Cards	Croquet
Chess		Chess

Anonymity (A): If the SWF yields a social ordering $\succsim_s$ for $\{\succsim_i\}_{i\in\mathcal{N}}$, and $\{\succsim'_i\}_{i\in\mathcal{N}}$ is obtainable from $\{\succsim_i\}_{i\in\mathcal{N}}$ by exchanging some preferences among the citizens, then the SWF yields $\succsim_s$ for $\{\succsim'_i\}_{i\in\mathcal{N}}$.

Tweedledee	Tweedledum	Society
Chess	Checkers	Checkers
Checkers	Cards	Cards
Croquet,Cards	Croquet	Croquet
	Chess	Chess

A SWF that meets condition A responds only to the preferences of the citizens and not to their identities.

A implies D.

Neutrality (N): For every $o_1, o_2, o_3, o_4 \in \mathcal{O}$, if o_1 and o_2 occupy the same relative positions in $\{\succsim_i\}_{i \in \mathcal{N}}$ as o_3 and o_4 occupy in $\{\succsim'_i\}_{i \in \mathcal{N}}$, then o_1 and o_2 must occupy the same relative position in the social ordering $\succsim_s$ yielded by the SWF for $\{\succsim_i\}_{i \in \mathcal{N}}$ as o_3 and o_4 occupy in the social ordering $\succsim'_s$ yielded for $\{\succsim'_i\}_{i \in \mathcal{N}}$.

Tweedledee	Tweedledum	Society
Chess	**Croquet**	**Chess**,**Croquet**
Checkers,**Croquet**	**Chess**	Cards
Cards	Cards	Checkers
	Checkers	

Neutrality (N): For every $o_1, o_2, o_3, o_4 \in \mathcal{O}$, if o_1 and o_2 occupy the same relative positions in $\{\succcurlyeq_i\}_{i \in \mathcal{N}}$ as o_3 and o_4 occupy in $\{\succcurlyeq'_i\}_{i \in \mathcal{N}}$, then o_1 and o_2 must occupy the same relative position in the social ordering $\succcurlyeq_s$ yielded by the SWF for $\{\succcurlyeq_i\}_{i \in \mathcal{N}}$ as o_3 and o_4 occupy in the social ordering $\succcurlyeq'_s$ yielded for $\{\succcurlyeq'_i\}_{i \in \mathcal{N}}$.

Tweedledee	Tweedledum	Society
Cards	**Checkers**	**Cards**,**Checkers**
Croquet	**Cards**	Chess
Chess	Chess,Croquet	Croquet
Checkers		

A SWF that meets condition N disregards the nature of the alternatives under consideration.

N implies I (let $o_1 = o_3$ and $o_2 = o_4$).

Positive Response (PR): For every $o_1, o_2 \in \mathcal{O}$, if the SWF yields $o_1 \succsim_s o_2$ for a given group preference profile, then the SWF yields $o_1 \succ_s o_2$ for any group preference profile that is just like the original profile except that one or more of the citizens has moved o_1 up in their ordering.

PR implies PA.

Thm 3.1.3 (May's Theorem). The only SWF that meets conditions A, N, and PR is Majority Rule.

However, Majority Rule doesn't satisfy condition U. Practical difficulties still lurk in the voting paradox.

Social Choice Theory

3.2 Sen on Liberalism

Johns Hopkins University, Spring 2016

Majority Rule leads to the voting paradox. But there is another difficulty with this collective choice rule: it is *illiberal*.

Sen: "Given other things in the society, if you prefer to have pink walls rather than white, then society should permit you to have this, even if a majority of the community would like to see your walls white. Similarly, whether you should sleep on your back or on your belly is a matter in which the society should permit you absolute freedom, even if a majority of the community is nosey enough to feel that you must sleep on your back."

Liberalism (L): For each citizen $i \in \mathcal{N}$, there are some $o_1, o_2 \in \mathcal{O}$ such that the SWF yields $o_1 \succ_s o_2$ whenever $o_1 \succ_i o_2$ and the SWF yields $o_2 \succ_s o_1$ whenever $o_2 \succ_i o_1$.

That is, citizen i is decisive for both $\langle o_1, o_2 \rangle$ and $\langle o_2, o_1 \rangle$.

"Condition L represents a value involving individual liberty that many people would subscribe to."

Thm 3.2.1 (Sen's Impossibility Theorem). There is no SWF for $\langle \mathcal{N}, \mathcal{O} \rangle$ that meets conditions U, P, and L.

Note that the controversial condition I does not come into play.

In fact, even a weaker condition than L is unsatisfiable in conjunction with U and P.

Minimal Liberalism (L*): There are at least two citizens in $\mathcal{N}$ such that for each of them, there are some $o_1, o_2 \in \mathcal{O}$ such that the citizen is decisive for both $\langle o_1, o_2 \rangle$ and $\langle o_2, o_1 \rangle$.

L* implies D.

Thm 3.2.2 (Sen's Impossibility Theorem, Minimal Version). There is no SWF for $\langle \mathcal{N}, \mathcal{O} \rangle$ that meets conditions U, P, and L*.

“It turns out that a principle reflecting liberal values even in a very mild form cannot possibly be combined with the weak Pareto principle, given an unrestricted domain. If we do believe in these other conditions, then the society cannot permit even minimal liberalism. Society cannot then let more than one individual be free to read what they like, sleep the way they prefer, dress as they care to, etc., *irrespective* of the preferences of others in the community.”

Ex. *Lady Chatterly's Lover*

"There is one copy of a certain book, say *Lady Chatterly's Lover*, which is viewed differently by 1 and 2. The three alternatives are: that individual 1 reads it (x), that individual 2 reads it (y), and that no one reads it (z). Person 1, who is a prude, prefers most that no one reads it, but given the choice between either of the two reading it, he would prefer that he read it himself rather than exposing gullible Mr. 2 to the influences of Lawrence. (Prudes, I am told, tend to prefer to be censors rather than being censored.) In decreasing order of preference, his ranking is z, x, y. Person 2, however, prefers that either of them should read it rather than neither. Furthermore, he takes delight in the thought that prudish Mr. 1 may have to read Lawrence, and his first preference is that person 1 should read it, next best that he himself should read it, and worst that neither should. His ranking is, therefore, x, y, z."

Ex. *Lady Chatterly's Lover*

Mr. 1	Mr. 2
No one reads	Mr. 1 reads
Mr. 1 reads	Mr. 2 reads
Mr. 2 reads	No one reads

Liberal values seemingly require that the choice between No one reads and Mr. 1 reads is up to Mr. 1: No one reads $\succ_s$ Mr. 1 reads.

Liberal values seemingly require that the choice between No one reads and Mr. 2 reads is up to Mr. 2: Mr. 2 reads $\succ_s$ No one reads.

But then Mr. 2 reads $\succ_s$ Mr. 1 reads, contradicting condition P.

There is tension between Liberalism and the Pareto principle.

“What is the moral? It is that in a very basic sense liberal values conflict with the Pareto principle. If someone takes the Pareto principle seriously, as economists seem to do, then he has to face problems of consistency in cherishing liberal values, even very mild ones. Or, to look at it in another way, if someone does have certain liberal values, then he may have to eschew his adherence to Pareto optimality. While the Pareto criterion has been thought to be an expression of individual liberty, it appears that in choices involving more than two alternatives it can have consequences that are, in fact, deeply illiberal.”

“The dilemma posed here may appear to be somewhat disturbing. It is, of course, not necessarily disturbing for every conceivable society, since the conflict arises with only particular configurations of individual preferences. The ultimate guarantee for individual liberty may rest not on rules for social choice but on developing individual values that respect each other’s personal choices.”