PHPE 400
Individual and Group Decision Making

Eric Pacuit
University of Maryland
pacuit.org

Strict Dominance

	s_1	s_2	s_3
A	2	3	1
B	1	2	0
C	1	4	0

Is there a way of assigning probabilities to the states s_1, s_2, and s_3 such that the decision maker strictly prefers B to A?

Is there a way of assigning probabilities to the states s_1, s_2, and s_3 such that the decision maker strictly prefers C to A?

Strict Dominance

	s_1	s_2	s_3
A	2	3	1
B	1	2	0
C	1	4	0

Is there a way of assigning probabilities to the states s_1, s_2, and s_3 such that the decision maker strictly prefers B to A? No!

Is there a way of assigning probabilities to the states s_1, s_2, and s_3 such that the decision maker strictly prefers C to A? Yes!

Strict Dominance

	s_1	s_2	s_3
A	2	3	1
B	1	2	0
C	1	4	0

Is there a way of assigning probabilities to the states s_1, s_2, and s_3 such that the decision maker strictly prefers B to A? No!

Is there a way of assigning probabilities to the states s_1, s_2, and s_3 such that the decision maker strictly prefers C to A? Yes!

X **strictly dominates** Y when for all states s, $u(X(s)) > u(Y(s))$.

- A strictly dominates B

Strict Dominance

	s_1	s_2	s_3
A	2	3	1
B	1	2	0
C	1	4	0

Is there a way of assigning probabilities to the states s_1, s_2, and s_3 such that the decision maker strictly prefers B to A? No!

Is there a way of assigning probabilities to the states s_1, s_2, and s_3 such that the decision maker strictly prefers C to A? Yes!

X **strictly dominates** Y when for all states s, $u(X(s)) > u(Y(s))$.

- A strictly dominates B
- A does not strictly dominate C

Weak Dominance

	s_1	s_2	s_3
A	2	3	1
B	1	2	1
C	2	3	1

Is there a way of assigning probabilities to the states s_1, s_2, and s_3 such that the decision maker strictly prefers B to A?

Is there a way of assigning probabilities to the states s_1, s_2, and s_3 such that the decision maker strictly prefers C to A?

Weak Dominance

	s_1	s_2	s_3
A	2	3	1
B	1	2	1
C	2	3	1

Does the decision maker strictly prefer A to B?

Does the decision maker strictly prefer A to C?

Weak Dominance

	s_1	s_2	s_3
A	2	3	1
B	1	2	1
C	2	3	1

Does the decision maker strictly prefer A to B? Depends...

Does the decision maker strictly prefer A to C? No!

X **weakly dominates** Y when for all states s, $u(X(s)) \geq u(Y(s))$ and there is some s' such that $u(X(s')) > u(Y(s'))$.

- A weakly dominates B

Weak Dominance

	s_1	s_2	s_3
A	2	3	1
B	1	2	1
C	2	3	1

Does the decision maker strictly prefer A to B? Depends...

Does the decision maker strictly prefer A to C? No!

X **weakly dominates** Y when for all states s, $u(X(s)) \geq u(Y(s))$ and there is some s' such that $u(X(s')) > u(Y(s'))$.

- A weakly dominates B
- A does not weakly dominate C

According to expected utility theory, preferences over lotteries should satisfy the Independence Axiom.

But, what about observed failures of the Independence Axiom, such as the Allais paradox or the Ellsberg paradox?

Evaluating Rational Choice Axioms

What should we make of the patterns found by psychologists and behavioral economists? Are these descriptive issues relevant for decision theory or rational choice theory?

Evaluating Rational Choice Axioms

What should we make of the patterns found by psychologists and behavioral economists? Are these descriptive issues relevant for decision theory or rational choice theory?

Any apparent violation of an axiom of the theory can always be interpreted in three different ways:

1. the subjects' preferences *genuinely* violate the axioms of the theory;

Evaluating Rational Choice Axioms

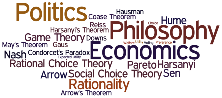

What should we make of the patterns found by psychologists and behavioral economists? Are these descriptive issues relevant for decision theory or rational choice theory?

Any apparent violation of an axiom of the theory can always be interpreted in three different ways:

1. the subjects' preferences *genuinely* violate the axioms of the theory;
2. the subjects' preferences have changed during the course of the experiment;

Evaluating Rational Choice Axioms

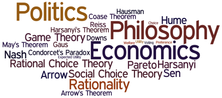

What should we make of the patterns found by psychologists and behavioral economists? Are these descriptive issues relevant for decision theory or rational choice theory?

Any apparent violation of an axiom of the theory can always be interpreted in three different ways:

1. the subjects' preferences *genuinely* violate the axioms of the theory;
2. the subjects' preferences have changed during the course of the experiment;
3. the experimenter has overlooked a relevant feature of the context that affects the subjects' preferences.

Recommending Behavior

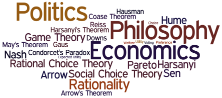

- One the one hand, that fact that many people have faulty reasoning about probabilities or deviate from EU theory does not mean that the theories are wrong (Hume's Law: *is* **does not** imply *can*). It could simply be that people are not naturally good at all kinds of reasoning, which is part of the reason why we study rational choice in the first place.

Recommending Behavior

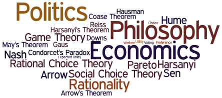

- One the one hand, that fact that many people have faulty reasoning about probabilities or deviate from EU theory does not mean that the theories are wrong (Hume's Law: *is* **does not** imply *can*). It could simply be that people are not naturally good at all kinds of reasoning, which is part of the reason why we study rational choice in the first place.

- On the other hand, *ought* **does** imply *can*, meaning that if we're going to say that people should follow EU theory, it needs to be possible that they actually do so.

Recommending Behavior

- One the one hand, that fact that many people have faulty reasoning about probabilities or deviate from EU theory does not mean that the theories are wrong (Hume's Law: *is* **does not** imply *can*). It could simply be that people are not naturally good at all kinds of reasoning, which is part of the reason why we study rational choice in the first place.

- On the other hand, *ought* **does** imply *can*, meaning that if we're going to say that people should follow EU theory, it needs to be possible that they actually do so.

- The question then becomes, 'Can people consistently follow EU theory? If not, when and why not?'.

Explaining/Predicting Behavior

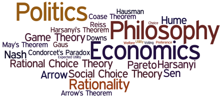

Stability: Individuals' preferences are stable over the period of the investigation.

Explaining/Predicting Behavior

Stability: Individuals' preferences are stable over the period of the investigation.

Invariance: Individuals' preferences are invariant to irrelevant changes in the context of making the decision.

A Dilemma

Either stick to the “formal axioms” of completeness, transitivity, Independence, etc. and refuse to assume the principles of stability and invariance.

A Dilemma

Either stick to the “formal axioms” of completeness, transitivity, Independence, etc. and refuse to assume the principles of stability and invariance. But then rational choice theory will be useless for all explanatory and predictive purposes because people could have fully rational preferences that constantly change or are immensely context-dependent.

A Dilemma

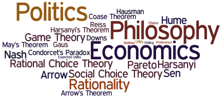

Either stick to the "formal axioms" of completeness, transitivity, Independence, etc. and refuse to assume the principles of stability and invariance. But then rational choice theory will be useless for all explanatory and predictive purposes because people could have fully rational preferences that constantly change or are immensely context-dependent. Alternatively, an economists can assume stability and invariance but only at the expense of making rational-choice theory a substantive theory, a theory laden not just with values but with *the economist's* values.

R. Nozick. *Newcomb's Problem and Two Principles of Choice*. 1969.

There are two boxes in front of us:

- box A, which contains \$1,000;
- box B, which contains either \$1,000,000 or nothing.

There are two boxes in front of us:

- box A, which contains \$1,000;
- box B, which contains either \$1,000,000 or nothing.

We have two choices:

- we open only box B.
- we open both box A and box B;

There are two boxes in front of us:

- box A, which contains \$1,000;
- box B, which contains either \$1,000,000 or nothing.

We have two choices:

- we open only box B.
- we open both box A and box B;

You can see inside box A, but not inside box B. We can keep whatever is inside any box we open, but we may not keep what is inside a box that we do not open.

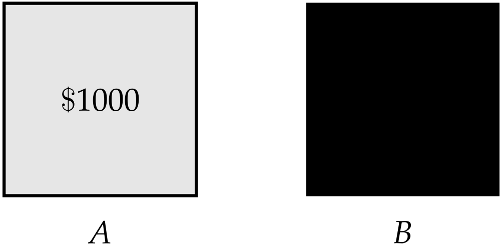

Choice:

one-box: choose box B

two-box: choose box A and B

A famous example: Newcomb's paradox

A very powerful being, who has been invariably accurate in his predictions about our behavior in the past, has already acted in the following way:

1. If he has predicted we will open just box B, he has put \$1,000,000 in box B.
2. If he has predicted we open both boxes, he has put nothing in box B.

What should we do?

Newcomb's paradox

you		*pred_B*	*pred_AB*
	B	$1M$	0
	AB	$1M+1T$	$1T$

Newcomb's paradox

you		$pred\_B$	$pred\_AB$
	B	$1M$	0
	AB	$1M$+$1T$	$1T$

Principle of dominance: take both boxes.

Newcomb's paradox

- $P(pred\_B \mid B)$: The probability that the wizard predicted you would choose box B *given that you decided to choose box* B.
- $P(pred\_AB \mid B)$: The probability that the wizard predicted you would choose both boxes *given that you decided to choose box* B.

Newcomb's paradox

- $P(pred\_B \mid B)$: The probability that the wizard predicted you would choose box B *given that you decided to choose box B*.
- $P(pred\_AB \mid B)$: The probability that the wizard predicted you would choose both boxes *given that you decided to choose box B*.

- $P(pred\_B \mid AB)$: The probability that the wizard predicted you would choose box B *given that you decided to choose both boxes*.
- $P(pred\_AB \mid AB)$: The probability that the wizard predicted you would choose both boxes *given that you decided to choose both boxes*.

Newcomb's paradox

- ✓ $P(pred\_B \mid B)$: The probability that the wizard predicted you would choose box B *given that you decided to choose box B*.
- ✗ $P(pred\_AB \mid B)$: The probability that the wizard predicted you would choose both boxes *given that you decided to choose box B*.

- ✗ $P(pred\_B \mid AB)$: The probability that the wizard predicted you would choose box B *given that you decided to choose both boxes*.
- ✓ $P(pred\_AB \mid AB)$: The probability that the wizard predicted you would choose both boxes *given that you decided to choose both boxes*.

Newcomb's paradox

you		*pred_B*	*pred_AB*
	B	$1M$	0
	AB	$1M{+}1T$	$1T$

Expected utility maximization: take box B.

$$P(pred\_B \mid B)1M + P(pred\_AB \mid B)0 > P(pred\_B \mid AB)(1M + 1T) + P(pred\_AB \mid AB)1T$$

Newcomb’s Paradox

What the Predictor did yesterday is *probabilistically dependent* on the choice today, but *causally independent* of today’s choice.

Act-state independence: For all states s and actions X, $P(s) = P(s \mid X)$

J. Collins. *Newcomb’s Problem*. International Encyclopedia of Social and Behavorial Sciences, 1999.