Decision Theory

1.8 Causal vs. Evidential Decision Theory

George Mason University, Spring 2018

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 99% 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 [$\frac{99}{100}$]	0 [$\frac{1}{100}$]
two boxes	1M+1000 [$\frac{1}{100}$]	1000 [$\frac{99}{100}$]

$$EU(\text{one box}) = \frac{99}{100} \times 1\text{M} + \frac{1}{100} \times 0 = 990,000.$$

$$EU(\text{two boxes}) = \frac{1}{100} \times (1\text{M} + 1000) + \frac{99}{100} \times 1000 = 11,000.$$

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

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

	B contains \$1M	B is empty
one box	1M [$\frac{99}{100}$]	0 [$\frac{1}{100}$]

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.

Ex. Superbowl.

Who do you think will win the Super Bowl? If you correctly guess that the New England Patriots will win, then you receive \$100. If you correctly guess that the Philadelphia Eagles will win, then you receive \$200.

	You win the bet	You lose the bet
bet on Patriots	win \$100	win \$0
bet on Eagles	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 $[\frac{99}{100}]$	0 $[\frac{1}{100}]$
two boxes	1M+1000 $[\frac{1}{100}]$	1000 $[\frac{9}{100}]$

Using evidential probabilities, EU(one box) $>$ EU(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]

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

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

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

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 $Pr(\text{shoot} \mathrel{\Box\!\!\rightarrow} \text{Alfred dies})$ 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
$Pr(\text{press button} \,\Box\!\!\rightarrow \text{Paul lives and psychos die})$ 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.