

The rest of the course

Subtleties involved with maximizing expected utility:

- Finding the right state space:
 - The wrong state space leads to intuitively incorrect answers when conditioning
- Taking causality into account
 - If you don't, again you have problems
- Computational issues:
 - Computing probabilities efficiently using graphical representations
 - Computing utilities efficiently
 - Eliciting utilities efficiently
- Problems with maximizing expected utility
 - Effects of framing
 - Ellsberg paradox, Allais paradox
 - Dealing with large state/outcome spaces
- (Possibly) some discussion on alternatives to probability for representing belief (and how this impacts decision making)
- Current research by Blume, Easley, Halpern

Savage: Summary

Prof. Blume showed that if your preferences satisfy postulates A1-A5, then that determines a qualitative probability \triangleright on events:

- $A \triangleright B$ if A is more likely than B (where A and B are sets of states).

With two more (arguably less plausible) axioms, A6 and A7, Savage is able to prove his major theorem:

Theorem: If \succ is a preference order on acts satisfying A1–A7, then there exists a probability Pr on states and a utility u on acts such that

$$f \succ g \text{ iff } EU_{\text{Pr}}(f) > EU_{\text{Pr}}(g).$$

Moreover, Pr is unique and u is unique up to affine transformations.

Savage: Interpretation

- Savage's theorem says that, if a decision maker obeys Savage's postulates, she is acting *as if* she has a probability on states and a utility on outcomes, and is maximizing expected utility.
- Note surprisingly $\Pr(A) = 0$ iff A is null
- \Pr extends the qualitative probability determined by A1–A5.
- For each set B , \succ_B also satisfies Savage's postulates; $\Pr(\cdot \mid B)$ is the probability determined by this preference order.

Today's topic: Savage assumes you're handed a state space.

- How do we now we have the right state space?
- The wrong state space leads to funny results.
- The Blume-Easley-Halpern approach avoids states altogether.

Three-Prisoners Puzzle

Computing the value of information involves conditioning. Conditioning can be subtle ...

Consider the three-prisoner's puzzle:

- Two of three prisoners a , b , and c are chosen at random to be executed,
- a 's prior that he will be executed is $2/3$.
- a asks the jailer whether b or c will be executed
- The jailer says b .

It seems that the jailer gives a no useful information about his own chances of being executed.

- a already knew that one of b or c was going to be executed

But conditioning seems to indicate that a 's posterior probability of being executed should be $1/2$.

This is easily rephrased in terms of value of information ...

The Monty Hall Puzzle

- You're on a game show and given a choice of three doors.
 - Behind one is a car; behind the others are goats.
- You pick door 1.
- Monty Hall opens door 2, which has a goat.
- He then asks you if you still want to take what's behind door 1, or to take what's behind door 3 instead.

Should you switch?

- What's the value of Monty's information?

The Second-Ace Puzzle

Alice gets two cards from a deck with four cards: $A\spadesuit$, $2\spadesuit$, $A\heartsuit$, $2\heartsuit$.

$A\spadesuit A\heartsuit$	$A\spadesuit 2\spadesuit$	$A\spadesuit 2\heartsuit$
$A\heartsuit 2\spadesuit$	$A\heartsuit 2\heartsuit$	$2\spadesuit 2\heartsuit$

Alice then tells Bob “I have an ace”.

- Conditioning $\Rightarrow \Pr(\text{both aces} \mid \text{one ace}) = 1/5$.

She then says “I have the ace of spades”.

- $\Pr_B(\text{both aces} \mid A\spadesuit) = 1/3$.

The situation is similar if Alice says “I have the ace of hearts”.

Puzzle: Why should finding out which particular ace it is raise the conditional probability of Alice having two aces?

Protocols

Claim 1: conditioning is always appropriate here, but you have to condition in the right space.

Claim 2: The right space has to take the *protocol* (*algorithm, strategy*) into account:

- a protocol is a description of each agent's actions as a function of their information.
 - **if** receive message
then send acknowledgment

Protocols

What is the protocol in the second-ace puzzle?

- There are lots of possibilities!

Possibility 1:

1. Alice gets two cards
2. Alice tells Bob whether she has an ace
3. Alice tells Bob whether she has the ace of spades

There are six possible runs (one for each pair of cards that Alice could have gotten); the earlier analysis works:

- $\Pr_B(\text{two aces} \mid \text{one ace}) = 1/5$
- $\Pr_B(\text{two aces} \mid A_{\spadesuit}) = 1/3$

With this protocol, we can't say "Bob would also think that the probability was $1/3$ if Alice said she had the ace of hearts"

Possibility 2:


1. Alice gets two cards
2. Alice tells Bob she has an ace iff her leftmost card is an ace; otherwise she says nothing.
3. Alice tells Bob the kind of ace her leftmost card is, if it is an ace.

This protocol is not well specified. What does Alice do at step 3 if she has both aces?

Possibility 2(a):

- She chooses which ace to say at random:

Now there are seven possible runs.



- Each run has probability 1/6, except the two runs where Alice was dealt two aces, which each have probability 1/12.
- $\Pr_B(\text{two aces} \mid \text{one ace}) = 1/5$
- $\Pr_B(\text{two aces} \mid A♠) = \frac{1/12}{(\frac{1}{6} + \frac{1}{6} + \frac{1}{12})} = 1/5$
- $\Pr_B(\text{two aces} \mid A♥) = 1/5$

More generally: Possibility 2(b):

- She says “I have the ace of spades” with probability α
 - Possibility 2(a) is a special case with $\alpha = 1/2$

Again, there are seven possible runs.

- $\Pr_B(\text{two aces} \mid A_{\spadesuit}) = \alpha / (\alpha + 2)$
- if $\alpha = 1/2$, get $1/5$, as before
- if $\alpha = 0$, get 0
- if $\alpha = 1$, get $1/3$ (reduces to protocol 1)

Possibility 3:

1. Alice gets two cards
2. Alice tells Bob she has an ace iff her leftmost card is an ace; otherwise she says nothing.
3. Alice tells Bob the kind of ace her leftmost card is, if it is an ace.

What is the sample space in this case?

- has 12 points, not 6: the order matters
 - $(2\heartsuit, A\spadesuit)$ is not the same as $(A\spadesuit, 2\heartsuit)$

Now $\Pr(2 \text{ aces} \mid \text{Alice says she has an ace}) = 1/3$.

The Monty Hall puzzle

Again, what is the protocol?

1. Monty places a car behind one door and a goat behind the other two. (Assume Monty chooses at random.)
2. You choose a door.
3. Monty opens a door (with a goat behind it, other than the one you've chosen).

This protocol is not well specified.

- How does Monty choose which door to open if you choose the door with the car?
- Is this even the protocol? What if Monty does not have to open a door at Step 3?

Not too hard to show:

- If Monty necessarily opens a door at step 3, and chooses which one at random if Door 1 has the car, then switching wins with probability $2/3$.

But ...

- if Monty does not have to open a door at step 3, then all bets are off!

Naive vs. Sophisticated Spaces

Working in the sophisticated space gives the right answers, BUT ...

- the sophisticated space can be very large
- it is often not even clear what the sophisticated space is
 - What exactly is Alice's protocol?

When does conditioning in the naive space give the right answer?

- Hardly ever!

Formalization

Assume

- There is an underlying space W : the naive space
- The sophisticated space S consists of pairs (w, o) where
 - $w \in W$
 - o (the observation) is a subset of W
 - $w \in o$: the observation is always accurate.

Example: Three prisoners

- The naive space is $W = \{w_a, w_b, w_c\}$, where w_x is the world where x is not executed.
- There are two possible observations:
 - $\{w_a, w_b\}$: c is to be executed (i.e., one of a or b won't be executed)
 - $\{w_a, w_c\}$: b is to be executed

The sophisticated space consists of four elements of the form $(w_x, \{w_x, w_y\})$, where $x \neq y$ and $\{w_x, w_y\} \neq \{w_b, w_c\}$

- the jailer will not tell a that he won't be executed

Given a probability \Pr on S (the sophisticated space), let \Pr_W be the marginal on W :

$$\Pr_W(U) = \Pr(\{(w, o) : w \in U\}).$$

In the three-prisoners puzzle, $\Pr_W(w) = 1/3$ for all $w \in W$, but \Pr is not specified.

Some notation:

- Let X_O and X_W be random variables describing the agent's observation and the actual world:

$$X_O = U \text{ is the event } \{(w, o) : o = U\}.$$

$$X_W \in U \text{ is the event } \{(w, o) : w \in U\}.$$

Question of interest:

When is conditioning on U the same as conditioning on the observation of U ?

- When is $\Pr(\cdot \mid X_O = U) = \Pr(\cdot \mid X_W \in U)$?
- Equivalently, when is $\Pr(\cdot \mid X_O = U) = \Pr_W(\cdot \mid U)$?

This question has been studied before in the statistics community. The *CAR* (Conditioning at Random) condition characterizes when this happens.

The CAR Condition

Theorem: Fix a probability \Pr on \mathcal{R} and a set $U \subseteq W$. The following are equivalent:

(a) If $\Pr(X_O = U) > 0$, then for all $w \in U$

$$\Pr(X_W = w \mid X_O = U) = \Pr(X_W = w \mid X_W \in U).$$

(b) If $\Pr(X_W = w) > 0$ and $\Pr(X_W = w') > 0$, then

$$\Pr(X_O = U \mid X_W = w) = \Pr(X_O = U \mid X_W = w').$$

For the three-prisoners puzzle, this means that

- the probability of the jailer saying “ b will be executed” must be the same if a is pardoned and if c is pardoned.
- Similarly, for “ c will be executed”.

This is impossible no matter what protocol the jailer uses.

- Thus, conditioning *must* give the wrong answers.

CAR also doesn't hold for Monty Hall or any of the other puzzles.

Why CAR is important

Consider drug testing:

- In a medical study to test a new drug, several patients drop out before the end of the experiment
 - for *compliers* (who don't drop out) you observe their actual response; for dropouts, you observe nothing at all.

You may be interested in the fraction of people who have a bad side effect as a result of taking the drug three times:

- You can observe the fraction of compliers who have bad side effects
- Are dropouts “missing at random”?
 - If someone drops out, you observe W .
 - Is $\Pr(X_W = w \mid X_O = W) = \Pr(X_W = w \mid X_W \in W) = \Pr(X_W = w)$?

Similar issues arise in questionnaires and polling:

- Are shoplifters really as likely as non-shoplifters to answer a question like “Have you ever shoplifted?”
- concerns of homeless under-represented in polls

Newcomb's Paradox

A highly superior being presents you with two boxes, one open and one closed:

- The open box contains a \$1,000 bill
- Either \$0 or \$1,000,000 has just been placed in the closed box by the being.

You can take the closed box or both boxes.

- You get to keep what's in the boxes; no strings attached.

But there's a catch:

- The being can predict what humans will do
 - If he predicted you'll take both boxes, he put \$0 in the second box.
 - If he predicted you'll just take the closed box, he put \$1,000,000 in the second box.

The being has been right 999 of the the last 1000 times this was done.

What do you do?

The decision matrix:

- s_1 : the being put \$0 in the second box
- s_2 : the being put \$1,000,000 in the second box
- a_1 : choose both boxes
- a_2 : choose only the closed box

	s_1	s_2
a_1	\$1,000	\$1,001,000
a_2	\$0	\$1,000,000

Dominance suggests choosing a_1 .

- But we've already seen that dominance is inappropriate if states and acts are not independent.

What does expected utility maximization say:

- If acts and states aren't independent, we need to compute $\Pr(s_i | a_j)$.
 - Suppose $\Pr(s_1 | a_1) = .999$ and $\Pr(s_2 | a_2) = .999$.
- Then take act a that maximizes

$$\Pr(s_1 | a)u(s_1, a) + \Pr(s_2 | a)u(s_2, a).$$

- That's a_2 .

Is this really right?

- the money is either in the box, or it isn't ...

A More Concrete Version

The facts

- Smoking cigarettes is highly correlated with heart disease.
- Heart disease runs in families
- Heart disease more common in type A personalities

Suppose that type A personality is inherited and people with type A personalities are more likely to smoke.

- That's why smoking is correlated with heart disease.

Suppose you're a type A personality.

- Should you smoke?

Now you get a decision table similar to Newcomb's paradox.

- But the fact that $\Pr(\text{heart disease} \mid \text{smoke})$ is high shouldn't deter you from smoking.

More Details

Consider two causal models:

1. Smoking causes heart disease:

- $\Pr(\text{heart disease} \mid \text{smoke}) = .6$
- $\Pr(\text{heart disease} \mid \neg\text{smoke}) = .2$

2. There is a gene that causes a type A personality, heart disease, and a desire to smoke.

- $\Pr(\text{heart disease} \wedge \text{smoke} \mid \text{gene}) = .48$
- $\Pr(\text{heart disease} \wedge \neg\text{smoke} \mid \text{gene}) = .04$
- $\Pr(\text{smoke} \mid \text{gene}) = .8$
- $\Pr(\text{heart disease} \wedge \text{smoke} \mid \neg\text{gene}) = .12$
- $\Pr(\text{heart disease} \wedge \neg\text{smoke} \mid \neg\text{gene}) = .16$
- $\Pr(\text{smoke} \mid \neg\text{gene}) = .2$
- $\Pr(\text{gene}) = .3$

Conclusion:

- $\Pr(\text{heart disease} \mid \text{smoke}) = .6$
- $\Pr(\text{heart disease} \mid \neg\text{smoke}) = .2$

Both causal models lead to the same statistics.

- Should the difference affect decisions?

Recall:

- $\Pr(\text{heart disease} \mid \text{smoke}) = .6$
- $\Pr(\text{heart disease} \mid \neg\text{smoke}) = .2$

Suppose that

- $u(\text{heart disease}) = -1,000,000$
- $u(\text{smoke}) = 1,000$

A naive use of expected utility suggests:

$$\begin{aligned} & EU(\text{smoke}) \\ &= -999,000 \Pr(\text{heart-disease} \mid \text{smoke}) \\ &\quad + 1,000 \Pr(\neg\text{heart-disease} \mid \text{smoke}) \\ &= -999,000(.6) + 1,000(.4) \\ &= -593,600 \end{aligned}$$

$$\begin{aligned} & EU(\neg\text{smoke}) \\ &= -1,000,000 \Pr(\text{heart-disease} \mid \neg\text{smoke}) \\ &= -200,000 \end{aligned}$$

Conclusion: don't smoke.

- But if smoking doesn't cause heart disease (even though they're correlated) then you have nothing to lose by smoking!

Causal Decision Theory

In the previous example, we want to distinguish between the case where smoking causes heart disease and the case where they are correlated, but there is no causal relationship.

- the probabilities are the same in both cases

This is the goal of *causal decision theory*:

- Want to distinguish between $\Pr(s|a)$ and probability that *a causes s*.
 - What is the probability that smoking causes heart disease vs. probability that you get heart disease, given that you smoke.

Let $\Pr_C(s|a)$ denote the probability that *a causes s*.

- Causal decision theory recommends choosing the act *a* that maximizes

$$\sum_s \Pr_C(s | a)u(s, a)$$

as opposed to the act that maximizes

$$\sum_s \Pr(s | a)u(s, a)$$

So how do you compute $\Pr_C(s | a)$?

- You need a good model of causality . . .

Basic idea:

- include the causal model as part of the state, so state has form: (causal model, rest of state).
- put probability on causal models; the causal model tells you the probability of the rest of the state
- in the case of smoking, you need to know the probability that

In smoking example, need to know the probability that

- smoking is a cause of cancer: α
- the probability of heart disease given that you smoke, if smoking is a cause: .6
- the probability of no disease given that you don't smoke, if smoking is a cause: .2
- the probability that the gene is the cause: $1 - \alpha$
- the probability of heart disease if the gene is the cause (whether or not you smoke):
 $(.52 \times .3) + (.28 \times .7) = .352$.

$$EU(\text{smoke}) = \alpha(.6(-999,000) + .4(1,000)) + (1 - \alpha)(.352(-999,000) + .658(1,000))$$

$$EU(\neg\text{smoke}) + \alpha(.2(-990,000)) + (1 - \alpha)(.352(-990,000)).$$

- If $\alpha = 1$ (smoking causes heart disease), then gets the same answer as standard decision theory: you shouldn't smoke.
- If $\alpha = 0$ (there's a gene that's a common cause for smoking and heart disease), you have nothing to lose by smoking.

So what about Newcomb?

- Choose both boxes unless you believe that choosing both boxes *causes* the second box to be empty!