

Causation, day 2

Probabilistic Theories of Causation

- According to the simplest probabilistic theories of causation, C (type) causes E iff C raises the probability of E .

$$C \text{ causes } E \leftrightarrow P(E | C) > P(E | \neg C)$$

- This account runs into problems immediately since probability raising is *symmetric* (you proved this on your second problem set).

$$P(E | C) > P(E | \neg C) \leftrightarrow P(C | E) > P(C | \neg E)$$

- Patrick Suppes fixed this problem by requiring that causes *precede* their effects.

Simpsons's Paradox

- Another problem: Just because C raises the probability of E , this doesn't mean that it is a *cause* of E .
- This is related to a statistical problem known as *Simpson's Paradox*. The paradox is that one factor C could raise the probability of another E *in general*, even though, once other relevant information is taken into account, C *lowers* the probability of E .
- Here's an example: suppose that we take an unbiased sample of 800 people, 450 of whom smoke (S), 350 of whom don't smoke ($\neg S$). Suppose that 350 get cancer (C) and 450 don't get cancer ($\neg C$). And suppose that the distribution of cancer amongst smokers and non-smokers was as shown below:

	S	$\neg S$
C	250	100
$\neg C$	100	250

From this, we should conclude that the probability that you get cancer, given that you smoke, is 5/7ths:

$$P(C | S) = \frac{250}{350} = \frac{5}{7} = 0.7142\dots$$

Whereas, the probability that you get cancer, given that you *don't* smoke, is 2/7ths:

$$P(C | \neg S) = \frac{100}{350} = \frac{2}{7} = 0.2857\dots$$

- So, the probabilistic theory tells us that smoking S causes cancer C .
- But—and here's the fascinating thing—it could very well turn out that there is a gene, X , such that having the gene makes you more likely to smoke, and that having the gene makes you more likely to get cancer. However, given that you have the gene, smoking actually *lowers* your probability of cancer; and, given that you don't have the gene, smoking *lowers* your probability of cancer.
- That is: it is consistent with the information given above that, if we separate people off into a group of people which have the gene X and a group of people which *don't* have the gene, $\neg X$, we would see the following distribution of smoking and cancer:

	X			$\neg X$	
	S	$\neg S$		S	$\neg S$
C	250	75		0	25
$\neg C$	50	10		50	240

- But *now*, given that you have the gene X , the probability that you get cancer if you smoke is 5/6ths,

$$P(C | S \wedge X) = \frac{250}{300} = \frac{5}{6} = 0.8333\dots$$

and the probability that you get cancer if you *don't* smoke is 15/17ths,

$$P(C | \neg S \wedge X) = \frac{75}{85} = \frac{15}{17} = .8823\dots$$

So, if you have the gene X , then smoking actually *lowers* the probability that you'll get cancer.

- And, given that you *don't* have the gene, $\neg X$, the probability that you get cancer if you smoke is 0,

$$P(C | S \wedge \neg X) = \frac{0}{50} = 0$$

and the probability that you get cancer if you *don't* smoke is 5/48ths,

$$P(C | \neg S \wedge \neg X) = \frac{25}{240} = \frac{5}{48} = 0.10416\dots$$

So, if you don't have the gene X , then smoking *lowers* the probability that you'll get cancer.

- So, either way—whether you have the gene or not—smoking *lowers* the probability of cancer.

- Notice that, in this case, if you have the gene, it's more likely that you'll be a smoker than if you don't have the gene:

$$P(S | X) = \frac{300}{385} = 0.7792\dots$$

$$P(S | \neg X) = \frac{50}{265} = 0.1886\dots$$

- And, if you have the gene, it's *also* more likely that you'll get cancer than if you don't have the gene.

$$P(C | X) = \frac{325}{385} = 0.8441\dots$$

$$P(C | \neg X) = \frac{25}{315} = 0.0793\dots$$

- This makes it seem plausible that it is the *gene* that is responsible for the correlation between smoking and cancer. It makes it seem as though the gene *promotes* both smoking and cancer, but that smoking actually *inhibits* cancer.
 - Notice that it is also perfectly compatible with this joint distribution between X , C , and S , that, while the gene promotes both smoking and cancer, the *cancer* inhibits *smoking*.
 - * We could, however, rule this causal order out on the basis that the smoking *precedes* the cancer.

Cartwright and Effective Strategies

- The thing to say about the relationship between probability and causation is that, if C causes E , then, holding fixed all the *other* causes of E , C raises the probability of E .

$$P(E | C \wedge K) > P(E | \neg C \wedge K)$$

where K is a specification of all of the *other* causes of E .

- But this just says something about the *relationship* between causation and probability. It doesn't *replace* the notion of causation with the notion of probability. It doesn't claim that causation *reduces* to probability-raising.
- So then why hang on to the notion of causation? Why not just make do with probabilities?
- Because, if we are to use our knowledge to *act* and to *control* the world, then we'll want more than mere *correlations*, we'll additionally want to know what our actions will *cause* to be the case.
 - Knowledge of the world's causal structure will tell us whether we ought to *smoke* or not. If smoking promotes cancer, then we ought *not*. Whereas, if smoking inhibits or is independent of cancer, then we *ought* to smoke.
 - Mere probabilistic information can't tell us about how we ought to *act*.
 - Even though the information that we're smoking ought to raise our credence that we're going to get cancer, that doesn't mean that smoking isn't a *good idea*, if we want to avoid cancer. It doesn't mean that smoking isn't an *effective strategy* for avoiding cancer.