Days 2-3: Conditionals and probabilities

Conditionals, probability, and decision // *ESSLLI 2024*

Melissa Fusco and Matt Mandelkern

1 Decision and conditionals

Actual reasoning about decisions essentially involves *conditionals*:

(1) Should I do A? Well, what would happen if I do A?

Stalnaker's formulation of causal decision theory was put in terms of conditionals. In the Newcomb case, you should think about what *would* happen, *were you* to take both boxes — not what *will* happen, if you *do* take both boxes:

(2) a. What would happen if I were to take both boxes?
    b. What will happen if I take both boxes?

The claim is that you should think the answer to (2-a) is: I'll make more money; the answer to (2-b) is: I'll make less money.

Decision theorists have generally abjured or at least ignored this formulation of decision theory.

This is because many theories of the conditional predict that conditionals like (3-a) and (3-b) have zero probability, since they are some kind of universal quantifier over outcomes:

(3) a. If I were to take both boxes, I would get more money than if I were to take just one.
    b. If I take both boxes, I'll get less money than if I take just one.

However, this risks making decision theory look like a highly theoretical exercise with little connection to human decision. Conditionals play an essential role in human decision; we need a theory of the conditional, and of decision, that can make sense of this.

This should lead us to seek theories of the conditional on which (3-a) and (3-b) can play the role they appear to play in human reasoning; this, in turn, may have interesting upshots for decision theory.

2 Indicative vs. subjunctive conditionals

We can distinguish between *indicative* vs. *subjunctive* conditionals. The distinction is not grammatically precise or exhaustive, but roughly tracks the difference in morphology in pairs like (4):

Adams 1975. Recent work has called this *O-marking* vs. *X-marking*. A distinction is sometimes drawn between *indicative* and *counterfactual* conditionals. However, that's a confusing distinction since it is neither exhaustive nor exclusive: subjunctive conditionals can have true antecedents, and indicative conditionals can have false antecedents.

(4) a. If Oswald didn't shoot Kennedy, someone else did.
b. If Oswald hadn't shot Kennedy, someone else would have.

Though superficially similar, we should have different attitudes toward (4-a) and (4-b): given that we know Kennedy was shot, we should be sure of (4-a), but not of (4-b).

Intuitively, (4-a) is certain since the *probability* that someone shot Kennedy, *conditional on Oswald not doing it*, is ∼1. By contrast, (4-b) is not certain since that *chance* of someone shooting Kennedy, conditional on Oswald not having done it, was low.

3 Stalnaker's Thesis and Skyrms's Thesis

The following generalizations have seemed appealing to many, where Pr is any rational probability function, $>_i$ is the indicative conditional, $>_s$ the subjunctive conditional, and Pr any rational credence function:

- *Stalnaker's Thesis (StT)*: $Pr(p >_i q) = Pr(q|p)$ when $Pr(p) > 0$
- *Skyrms's Thesis (SkT)*: $Pr(p {>_s} q) = \mathbb{E}_{Pr}(Ch_t(q|p))$ where t is a salient time before the present such that $Ch_t(p)$ is sure to be > 0.

There's a lot of empirical evidence for Stalnaker's Thesis. The probability of (5) seems like it simply has to be the probability of landing n, conditional on the die being thrown:

See citations in Douven and Verbrugge 2013.

(5) If the die is thrown, it will land on side n.

Skyrms's Thesis is a little bit woollier, and thus harder to evaluate. But it seems to get at something important. Suppose the die is never thrown. The probability of (6) seems to be the *expectation* of the chance that (before it was determined it wouldn't be thrown) it lands on n, conditional on being thrown.

(6) If the die had been thrown, it would have landed on side n.

There are counterexamples to SkT in the form of counterlegals/temporals:

(7) If the initial conditions of the universe had been other than they are, then . . .

The problem is that there was no time where the antecedent has non-zero chance. Still, modulo these cases, SkT seems at least illuminating.

4 Some simple lessons

Suppose we take these theses seriously. There are some quick and striking upshots:

- $>_i$ is not the material conditional, since $Pr(p \supset q) = Pr(\overline{p} \vee q) \geq Pr(q|p)$ with equality iff $Pr(q|p) = 1$.

$Pr(\overline{p} \vee q) = Pr(\overline{p}) + Pr(pq) = Pr(\overline{p}) + Pr(q|p)Pr(p) = Pr(q|p)(\frac{Pr(\overline{p})}{Pr(q|p)} + Pr(p))$; and $\frac{Pr(\overline{p})}{Pr(q|p)} + Pr(p) \geq 1$, with equality iff $Pr(q|p) = 1$.

- $>_i$ / $>_s$ are not *modals* in the sense of being universal quantifiers over worlds.
  - Lewis (1973) thought $p >_s q$ means: *all* p-worlds most similar to actuality are q-worlds. But the probability that *all* die-throw worlds most similar to actuality are side n-worlds is obviously 0.
  - Similarly, many have thought $p >_i q$ means something like: *all* epistemically possible p-worlds most similar to actuality are q-worlds. But if you're unsure what happened with the die, again, then it would follow that (5) has 0 probability.

5 Lewisian Triviality

Lewis (1976) noticed a striking fact about StT: it is prima facie inconsistent with thinking that *rational updating goes by conditioning*.

For suppose StT held universally, as a fact about the connective $>_i$ and any probability measure. We have the following by the probability calculus:

$$Pr(p >_i q) = \underbrace{Pr(p >_i q|q)}_{1} Pr(q) + \underbrace{Pr(p >_i q|\neg q)}_{0} Pr(\neg q)$$

But note that $Pr(\cdot|q)$ and $Pr(\cdot|\neg q)$ are themselves probability measures, which we can write Pr_q and $Pr_{\neg q}$; by StT, $Pr_q(p >_i q) = Pr_q(q|p) = 1$ and $Pr_{\neg q}(p >_i q) = Pr_{\neg q}(q|p) = 0$, so the rhs reduces to $Pr(q)$!

But in general, $Pr(p >_i q)$ need not be $Pr(q)$; there is a .5 chance that I'll flip this fair coin, so the probability that the coin lands heads is .25, but the probability that the coin lands heads *if flipped* is .5.

A similar result: suppose you learn $p \vee (p >_i q)$. This can't change the conditional probability of q on p—since you kept all the p-worlds—but it will increase your credence in $p >_i q$ provided that $\neg p \wedge \neg(p >_i q)$ had non-zero prior probability.

Rothschild 2013

A similar result can be given for SkT.

6 Conditionals and context-sensitivity

This is an interesting result, but on reflection, it's not so obvious that it shows that StT is inconsistent with the idea that rational updating is conditioning.

For conditionals are, on the face of it, *context-sensitive*. Lewis's example: 'The kangaroo would fall over if it didn't have a tail'. Well, yes, in light of the laws of mechanics; no, in light of the laws of evolution.

And importantly, Lewisian triviality does not undermine the thesis that *probability talk* and *conditionals* are interpreted together so that StT is always true in a given context, that is, that we interpret 'probability' and 'if' in such a way that 'the probability that q if p is n' and 'the probability of q conditional on p is n' aways have the same truth-value in a given context.

7 Basic Tenability

Indeed, such a thesis is obviously *tenable*. Given a probability measure Pr over a Boolean algebra, we can just extend it iteratively as follows:

- Let $Pr^1 = Pr$
- when p, q are in the domain of Pr^n and $Pr^n(q|p)$ is defined, let $p >_i q$ be an arbitrary proposition s in the field of Pr^n such that $Pr^n(s) = Pr^n(q|p)$, if there is one; if not, simply extend the field of Pr^n with a new proposition s with measure $Pr^n(q|p)$.

The function obtained as the limit of this procedure obviously validates StT.

But the result isn't very interesting, because we want not just a StT-satisfying probability measure over just any algebra, but specifically over an algebra where we can plausibly identify $>_i$ as a *conditional* operator — hence satisfying certain basic logical principles.

8 Tenability in a basic conditional logic

Indeed, van Fraassen 1976 showed that StT is tenable in a weak conditional logic. But most have found that logic to be implausibly weak.

Though see Bacon 2015

9 Stalnaker's Conditional

A more plausible approach is compatible with a limited form of tenability. The approach is based on the following intuition:

Stalnaker 1968

$p > q$ is true iff q is true in the world *that would obtain if p were true*

This feels trivial, but it turns out to be enough to form the basis of a rich semantics/logic. We interpret conditionals with a *selection function* $f : (\wp(W) \setminus \varnothing) \times W \to W$ which, given a non-empty proposition φ and world w, tells us *how things would be*, at w, if φ were true.

We can identify inconsistent antecedents with the singleton of an absurd world λ which verifies everything.

$p > q$ is true at w iff q is true at $f([\![p]\!], w)$.

The selection function must satisfy certain intuitive constraints:

- *Strong centering*: $f(\varphi, w) = w$ when $w \in \varphi$
- *Success*: $f(\varphi, w) \in \varphi$
- *Reciprocity*: $f(\varphi, w) \in \psi$ & $f(\psi, w) \in \varphi \Rightarrow f(\varphi, w) = f(\psi, w)$

This semantics is sound and (weakly) complete for the logic C2, comprising the closure of the following set of axiom schemas:

Reciprocity is sometimes called CSO.

- PC: Every theorem of classical propositional logic
- Identity: $p > p$
- Reciprocity: $((p > q) \wedge (q > p) \wedge (p > r)) \rightarrow (q > r)$
- MP: $(p > q) \rightarrow (p \rightarrow q)$
- CEM: $(p > q) \vee (p > \neg q)$

under the following two inference rules, where $\rightarrow$ is the material conditional:

- Detachment $\vdash p \rightarrow q$ and $\vdash p$ together imply $\vdash q$
- Normality: $\vdash (p \wedge q) \rightarrow r$ implies $\vdash ((s > p) \wedge (s > q)) \rightarrow s > r$

10 Triviality in C2

StT isn't quite tenable in C2:

Proposition 1. *When there are A, B s.t. $Pr(A\overline{B}), Pr(\overline{A}), Pr(AB)$ are all > 0, then StT does not hold for Pr.*

(Stalnaker, 1974)

Proof. Consider any such A, B. For brevity let:

$$C := A \vee (A > \overline{B})$$

Then StT must fail, no matter the choice of selection function, for:

$$X := C > (\overline{A} \vee B)$$

Note that in C2, $\overline{C} \vdash X$. Hence $Pr(X|\overline{C}) = 1$ if defined. But $Pr(X|C)$ is, by Strong Centering, $Pr(\overline{A} \vee B|C)$, so if we also have $Pr(X) = Pr(\overline{A} \vee B|C)$ by StT, then by the law of total probability we have:

$$Pr(X) = Pr(\overline{A} \vee B|C) = \underbrace{Pr(X|C)}_{Pr(\overline{A} \vee B|C)} Pr(C) + \underbrace{Pr(X|\overline{C})}_{1 \text{ if defined}} Pr(\overline{C})$$

Hence

$$Pr(\overline{A} \vee B|C)(1 - Pr(C)) = Pr(\overline{C})$$

So

$$Pr(\overline{A} \vee B|C)(Pr(\overline{C})) = Pr(\overline{C})$$

We can moreover show that $Pr(\overline{C}) > 0$, since $Pr(\overline{A}) > 0$ and $Pr(A > B) > 0$ and these are independent by StT, so that $Pr(\overline{A} \wedge (A > B)) > 0$, which with the fact that $Pr(\neg(A > \bot)) = 1$ entails that $Pr(\overline{A} \wedge \neg(A > \overline{B})) > 0$.

Hence $Pr(\overline{A} \vee B|C) = 1$. But obviously C is true whenever $A\overline{B}$ is, while $\overline{A} \vee B$ is not. So $Pr(A\overline{B}) = 0$ after all. □

11 Limited StT in Stalnaker Semantics

With that in hand, it turns out that a limitation of StT *is* tenable:

van Fraassen 1976. See Khoo and Santorio 2018 for a helpful guide to a finitary version of the construction.

Proposition 2. *Any probability measure Pr over an algebra $\mathcal{B}$ on W can be extended to a probability measure Pr^* over a* C2*-conditional algebra extending $\mathcal{B}$, with $Pr^*(p > q) = Pr^*(q|p)$ provided $Pr^*(p) > 0$ and $p \in \mathcal{B}$.*

The proof idea is to model C2 with a set of sequences, where each sequence $\langle 1, 2, 3, \ldots \rangle$ represents a selection function: 2 is the way 1 would be if 1 weren't the case; 3 is the way 1 would be if 1 and 2 weren't the case; etc. So, when p is Boolean, it is true at a sequence iff true at the first world of that sequence; and $p > q$ is true at a sequence iff the first p-tail of that sequence is a q-tail (if there is one).

Then we can extend a measure over $\mathcal{B}$ to a measure over sets of ω-sequences over W, which serves as our conditional algebra: we set $Pr^*(W^\omega) = 1$ and then inductively define:

$$Pr^*(p \times \alpha) = Pr(p) \times Pr^*(\alpha) \text{ when } \alpha \subseteq W^\omega \text{ and } p \subseteq W$$

Sloppily writing $p \times \alpha$ for the sequences comprising an element of p concatenated to an element of α

We can show this is a probability measure on the conditional algebra. Then we can calculate $Pr^*(p > q)$ as follows, provided p is Boolean:

$$Pr^*(p > q) = Pr^*(p)Pr^*(p > q|p) + Pr^*(\overline{p})Pr^*(p > q|\overline{p}) =$$

$$Pr(p)Pr^*(q|p) + Pr(\overline{p})\frac{Pr(\overline{p} \wedge (p > q))}{Pr(\overline{p})} =$$

$$Pr(p)Pr^*(q|p) + Pr(\overline{p})\frac{\sum_{n>0} Pr(\overline{p})^n Pr^*(pq)}{Pr(\overline{p})} =$$

$$Pr(p)Pr^*(q|p) + Pr(\overline{p})Pr^*(pq)\sum_{n>0} Pr(\overline{p})^{n-1} =$$

$$Pr(p)Pr^*(q|p) + Pr(\overline{p})Pr^*(pq)\sum_{n\geq 0} Pr(\overline{p})^{n} =$$

$$Pr(p)Pr^*(q|p) + Pr(\overline{p})Pr^*(pq)\frac{1}{1 - Pr(\overline{p})} =$$

$$Pr(p)Pr^*(q|p) + Pr(\overline{p})\frac{Pr^*(pq)}{Pr(p)} = Pr^*(q|p)$$

So yes, a plausible conditional logic puts limits on StT. But those limits are relatively mild, since conditional-embedding antecedents are difficult to process, and it's not at all clear we have strong intuitions about them in favor of StT:

See Kaufmann 2023

(8) If the vase will break if dropped, then it is made of glass.

The tenability of SkT, modulo the same limits, is an immediate corollary, since we can simply model rational credence in subjunctives as the expectation of the probability of a conditional which obeys StT for a probability function representing chance, where you are subjectively uncertain about which chance function is actual.

12 Context-sensitivity

We've seen two responses to limitative results: (i) they aren't surprising, since conditionals, and hence probabilities of conditionals, are intrinsically context-sensitive, while conditional probabilities are not; (ii) they aren't worrisome, since StT is still tenable modulo minor caveats.

This brings us to an interesting question: even if we *can* have StT and SkT, or enough of them for our purposes, should we? They face some striking counterexamples. Consider this case:

Following Rothschild 2013, based on similar ones from McGee 2000; Kaufmann 2004.

(9) You are going to buy a car of a certain make with the following feature: almost all cars of this kind have software that both prevents crashes (with near certainty) and ensures that, in case of a crash, the airbag deploys (with near certainty). A small minority of those cars, however, have a bug which both makes it much more likely for the car to crash, and also much more likely that the airbag will never deploy.

   a. If the car you get crashes, the airbag will deploy.
   b. If the car you got were to crash, the airbag would deploy.

It seems easy to get a judgment that these are very probable: the car you get will almost certainly be such that the airbag is almost certain to deploy, conditional on crashing.

But the conditional probability/expectation of the conditional chance of the airbag deploying, conditional on crashing, is very low, since crashing is very strong evidence that the car was faulty.

If it is simply a matter of context-sensitivity whether StT/SkT holds, is it also a matter of corresponding context-sensitivity whether an act is rational?

References

Adams, E. (1975). *The Logic of Conditionals*. Dordrecht.
Bacon, A. (2015). Stalnaker's thesis in context. *Review of Symbolic Logic*, 8(1):131–163.
Douven, I. and Verbrugge, S. (2013). The probabilities of conditionals revisited. *Cognitive Science*, 37(4):711–730.
van Fraassen, B. (1976). Probabilities of conditionals. In Harper and Hooker, editors, *Foundations of Probability Theory, Statistical Inference, and Statistical Theories of Science*, volume I, pages 261–308. D. Reidel Publishing Company, Dordrecht-Holland.
Kaufmann, S. (2004). Conditioning against the grain. *Journal of Philosophical Logic*, 33:583–606.
Kaufmann, S. (2023). Bernoulli semantics and ordinal semantics for conditionals. *Journal of Philosophical Logic*, 52(1):199–220.
Khoo, J. and Santorio, P. (2018). Lecture notes: Probability of conditionals in modal semantics. Unpublished manuscript.
Lewis, D. (1973). *Counterfactuals*. Oxford: Blackwell.
Lewis, D. (1976). Probabilities of conditional and conditional probabilities. *Philosophical Review*, 8:297–315.
McGee, V. (2000). To tell the truth about conditionals. *Analysis*, 60(1):107–111.
Rothschild, D. (2013). Do indicative conditionals express propositions? *Noûs*, 47(1):49–68.
Stalnaker, R. (1968). A theory of conditionals. In Rescher, N., editor, *Studies in Logical Theory*, pages 98–112. Oxford: Blackwell.
Stalnaker, R. (1974). Letter to Bas Van Fraassen.