

Hypothetical Oughts

Jim Pryor, NYU
Christian Piller, York

1. First case: Alice has some evidence E that supports P, and P supports Q, and that's all. She should believe Q. Second case: Bob merely believes P without adequate justification. Natural to say things like: "Given his beliefs, he should also believe Q."

Typical examples: closure/inductive coherence, means/end, akrasia.

2. Advice to hardened criminals. (i) "At least say thank you." Same advice as one gives to saintly agents; versus advice distinctively for the imperfect, e.g. (ii) "Shoot to the head from concealment." (iii) Compromise between actual attitudes and the norms they're ignoring: "Use the gently poison." (iv) **Follow-through advice**: "Use the quick painful poison."

3. Does follow-through advisor just pretend to endorse agent's actual attitudes? (a) Can be explicit that he doesn't. (b) We think agent exhibits a further categorical defect by not following through. Why? Is follow-through advice committed to bootstrapping?

4. Suppose act utilitarianism, and that by it, Xing is good. Charlie believes Xing is bad, but Xs anyway. **Objectivist** says Charlie is morally lucky: in fact he's ended up doing good. Subjectivist says there is no moral luck: following your conscience (what you believe is good/bad) is always the good choice. Contra **subjectivist**, we do intuitively recognize moral luck. Contra the objectivist: even if Charlie did good in one sense, isn't there something wrong with him? Perhaps he did one good, and another bad.

5. Or reason here to **disambiguate**? In objective sense, Charlie ought to X. In subjective sense, he ought not.

This proposed disambiguation is not going to be adequate to the phenomenon.

6. Complex case #1. David believes he's in a kind of toxin puzzle. He believes that Xing will bring him money, and that Ming is necessary for Xing. He also believes that intending to X and Ming will lose him even more money. In real toxin puzzles, there are some things we want to say to commend Xing, and those things are inherited by known necessary means. But in fact, David is not in a toxin puzzle: Xing is just as bad as intending to X and Ming. So any commendation we'd give to Xing and so to Ming would be in a follow-through spirit: "Given you intend to X, you should M."

How should we understand that commendation? Objectively, David shouldn't X, shouldn't intend to X, shouldn't pursue means to Xing. By David's own lights, Ming while intending to X is a bad combination.

7. Complex case #2. Ella's credence in P is 0.5. She antecedently expects Frank to be as good a P-estimator as she is, and knows him to be working from the same evidence. Frank's credence is 0.3. In fact they're both off: their shared evidence, by itself, justifies credence 0.7 in P. The additional evidence that Frank thinks $\sim P$ may contribute some defeat, but suppose the net justified credence in P is still 0.6. How should Ella respond to her peer's disagreement? Common answers: (a) she should split the difference, and move from 0.5 to 0.4. (b) she should downgrade her credence somewhat, but not that far: say to 0.45; whereas Frank should move to 0.35. (c)

since she knows she's already taking into account all of Frank's evidence, his actual (by her lights) mis-assessment of it should have no further bearing; she should stay at 0.5.

In fact what her evidence supports is not moving towards Frank but away from him. Suppose a third peer Greta had come in with the correct $\text{cred}(P) = 0.7$ on their shared evidence. Now (a) (b) and (c) would agree Ella should stay at 0.5. That is not the objective verdict: (d) her evidence really supports $\text{cred}(P) = 0.7$ or 0.6. Perhaps it's subjective advice? But in original case with only Frank and Ella, isn't the subjective advice clearly (c): by her lights, Frank got this case wrong.

If there's a debate to which (a) and (b) are viable answers, it can't be about what the evidence objectively does call for, nor about what it subjectively seems the evidence calls for.

8. There's some normatively interesting relation between agent's actual attitudes and certain rational responses (beliefs, actions). Contra subjectivist, this doesn't yet amount to genuine justification. Contra the simple disambiguation proposal, it's not always captured by what genuine justification the agent has by her own subjective lights. What is it?

9. Three tasks: (a) semantics of the claims "If/given you believe/intend so-and-so, you ought to such-and-such." (b) why do we say such things? what are those speech-acts meant to be doing? **(c) refined model of what normative relations the described agents stand in.**

10. Wide-scope analysis of the follow-through advice: O(if you believe/intend so-and-so, then believe/intend/do such-and-such). Resist **factual detachment**: O(if A then B), $A \neq \Rightarrow O(B)$. **Deontic detachment** looks better: O(if A then B), $O(A) \Rightarrow O(B)$.

Narrow-scope analyses: if you believe/intend so-and-so, then O(believe/intend/do such-and-such). Wide-scopers object that this would commit us to bootstrapping.

Might we weaken the O from "decisively ought" to "have some reason"? Broome argues no, because one of the normative relations we're considering is an anti-akratic principle "If you believe you ought to X, you ought to (intend to) X." And if you violate this principle there definitely must be something wrong with you; it couldn't be that this was merely one prima minus which might be outweighed by other factors. So the "ought" in this principle has to be interpreted strictly, which on the narrow-scope analysis gives us bootstrapping (= what we called "subjectivism" in 4).

11. Problem for Broome's argument: Presumably we can have justified false beliefs about what we ought to do. That is, OBOX but not OX. However, if there's an anti-akratic principle that's both (a) *strict* and (b) *supports deontic detachment* (whether it's understood wide-scopely or otherwise), OBOX would \Rightarrow OX.

Perhaps (b) should yield; we'll see more trouble for it later. But without it, it's not clear what bearing the normative relations we're studying have on familiar justification. You justifiably intend X, and know M is a necessary means. Does the means/end principle we're considering have any connection to your having justification to M?

More natural to think (a) should yield. BOX normatively commends Xing, but may not strictly require it. Then Broome's argument against narrow-scopers fails.

12. Chisholm Puzzle: "If you're going to visit grandma, you ought to call ahead." Perhaps you ought to visit her, but in fact you're not going to. Deontic detachment says you ought to call ahead anyway.

Simon's Puzzle: "If you're going to visit grandma, you ought to start walking." Perhaps you ought to visit her, but in fact you're not going to. Deontic detachment says you ought to start walking anyway.

On reflection, original puzzle doesn't invite a wide-scope, deontic detachment analysis. The norm that tells you to call ahead is insensitive to whether you ought to be visiting. In this case, best to say: Will visit \implies O(call ahead). O(visit) but won't visit $\not\implies$ O(call ahead anyway).

In Simon's puzzle, on the other hand, the norm does seem sensitive to whether you ought to be visiting. In this case, factual detachment would incorrectly bootstrap: Will visit but O(\sim visit) $\not\implies$ O(start out). So narrow-scope analysis seems wrong. Deontic detachment not obviously right either: O(visit) but won't visit $\not\implies$ O(start out anyway)?

The difference between these cases shows more reliably than Broome's argument that a distinctive relation not captured by narrow-scope analysis is sometimes present.

13. That's not to say the wide-scope analysis is correct. Phenomenon may not merit a scopal analysis at all. Should identify what's essentially distinctive about it, separate from speculative implementations. So far: factual detachment no, deontic detachment maybe.

14. Kolodny defended some narrow-scope principles where Broome proposed wide-scope. Despite this, they agree about many further issues:

(a) What normative force does the normatively interesting relation we're investigating have? Do you have (decisive? any?) reason to conform to these requirements?

(b) Some ways of bringing yourself into conformance are rationally better than others.

(c) You don't reason *from propositions that you're in the antecedent states into* the conforming state. Instead you reason *from the contents of* the antecedent states.

15. What to say about Simon's puzzle?

(a) In one frame of mind, doesn't sound so bad to say: O(visit) \implies O(start out), even if you're not going to do either.

(b) On the other hand, someone who did respond by starting out, despite having no intention of visiting, would be exhibiting some kind of defect. Compare: Harry has evidence E that supports P, and P supports Q, and that's all. So he should believe Q. And he does believe Q, however he refuses to believe the intermediate premise P. What could he be thinking?

16. Perhaps the principle should be: O(visit) and intend to visit \implies O(start out). In your case, no means/end principle applies. (You'd only violate O(visit).)

However (a) is there *nothing* to be said on behalf of starting out for subjects who don't intend to visit? (More cautiously: O(visit) and don't intend to refrain from visiting \implies O(start out).)

(b) What about someone who O(visit), and does intend to visit (but not for the reasons they ought to---they'll be visiting in order to kill grandma). Ought they to start out? Tempted to say yes, but perhaps that's just 15(a). Tempted also to say they ought not be starting out, because they ought not be visiting for these reasons.

17. My proposal: (a) there's a network of facts that together contribute to what you should believe or do. The inputs here are what evidence/reasons you really have. Your actual beliefs/intentions are irrelevant except insofar as they're significant facts (as in toxin puzzle and Chisholm's "if you do intend, you should call ahead").

What (a) contributes to is what you're *ex ante* justified in believing/intending/doing. Or as it's (perhaps misleadingly) called in epistemology, your propositional justification. Here we do have straightforward deontic detachment. *Ex ante*, you *ought* to visit grandma, and therefore start out, though you lack any intention to visit her.

(b) When your actual attitudes diverge from the network, there's also another belief/intention/doing that you stand in a different normatively interesting relation to. This is determined by making minimal changes to the original network to make it conform to the actual attitudes we're holding fixed. In general, this will be very indeterminate; but in specific cases we'll often have focused enough answers. It may sometimes coincide with what the agent seems by her own lights to have reasons for, but this can't be relied on in general.

Call the belief/intention/doing that (b) commends **what you're justified in hypothetically on** the attitudes we held fixed. Leaving aside whether this is amenable to any kind of scopal analysis. Leaving aside also the semantics of the sentences we use to try to articulate this relation, and what we mean to be doing by so speaking. One claim: this is the relation we're trying to track in our follow-through advice.

18. How are the (b) relations normatively related to our familiar, non-hypothetical normative statuses? Not as the subjectivist says: Bob who merely believes P isn't thereby non-hypothetically justified in believing Q. Charlie who merely believes he ought not to X isn't thereby non-hypothetically justified in refraining from Xing.

Instead, (a) and (b) together are factors in what you can have *ex post* (in epistemology, "doxastic") justification to believe/intend/do. If you have *ex ante* justification to believe Q (because your evidence E supports P which supports Q, and that's all the evidence); but you also have hypothetical justification to refrain from believing Q, relative to your actual withholding of belief in P; then believing Q can't be *ex post* justified for you.

Neither can refraining from believing Q, for a different reason, namely that altogether, the only response your actual evidence commends is: believing Q.

19. So relative to your fixed attitude towards P, you've gotten yourself into a **epistemically tragic** relation to Q. No doxastic response will be right.

Sometimes we can say: "You shouldn't have gotten yourself into this pickle---and it's best to get yourself out if it if you can---but given that you're in it the best thing for you to do is..."

In tragic pickles, given that you're in the pickle, *nothing* you can do is normatively commended. (Not even refraining.)

20. So then what you should do is get yourself out of the pickle: here, get rid of your withholding belief in P. May not always be a rationally open option. Familiar worries: (a) voluntarism. Unfamiliar: (b) **you may have left yourself no attitudes from which that can be a reasonable step.**