Forthcoming in *Nous*

"The Thing To Do" Implies "Can"[1]

NICHOLAS SOUTHWOOD

Australian National University

I

Suppose that I have borrowed your prized first edition of Newton's *Principia Mathematica*. But I'm a nasty cove and rather than returning it to you I opt to consign it to the flames, thereby intentionally making it the case that I can't return it to you. The principle that "ought" implies "can" (OC) holds that

(OC) For any agent $a \in A$, a ought, at time t, to ϕ implies that, at t, a can ϕ.

From (OC) it seems that we are entitled to infer that it's not the case that I ought to return your prized first edition of Newton's *Principia*. This might seem absurd. Rather, it might seem that this is a case of my having intentionally made it the case that I cannot perform an act that I most certainly ought to perform (see Stocker 1971, pp. 314-5; Sinnott-Armstrong 1984, pp. 252-4; Young 1975, pp. 13-4).

Countless interpretations of (OC) have been proposed to respond to this familiar complaint. Some have focused on the "ought" (e.g. Streumer 2007; Vranas 2007), some on the "can" (e.g. McConnell 1989), some on the "implies" (e.g. Hare 1963, pp. 53-4; Sinnott-Armstrong 1984), some on the "ϕ" (e.g. Baltzly 2000;

Brennan and Southwood 2007), and some on the "*t*" (e.g. Howard-Snyder 2006). I suspect that all these interpretations either a) fail to block the complaint in a plausible and non-ad hoc way or b) throw out the baby with the bathwater,[2] though I won't make any attempt to defend this sweeping claim here.

My aim instead is to propose a *related* principle that I shall call the principle that "the thing to do" implies "can" (TDC). (TDC) holds that

(TDC) For any agent $a \in A$, at time t, ϕing is the thing for a to do only if, at t, a can ϕ.

I shall argue that (TDC) is entailed by important but underappreciated truths about practical reason, and that it is not vulnerable to the familiar complaint. This is interesting in its own right. Moreover, I shall suggest that it has interesting (though by no means straightforward) implications for (OC). Precisely what it implies for (OC) depends on the relation between claims about what we ought to do and claims about the thing to do -- in particular, whether the former imply the latter. If they do, then the argument for (TDC) that I shall offer also suggests a novel argument for (OC) and a novel response to the familiar complaint. If they don't, then (TDC) naturally suggests a novel error theory of (OC).

II

To explicate (TDC), we must say something about the core notion of a "claim about the thing to do." I am going to assume that there are such claims, and that they are

(or are made true by) claims about the correct answer to the question of *what to do*. There are two important ideas here that require some elaboration: a) the idea of the *question what to do*, and b) the idea of a *correct answer* to the question of what to do. Let us briefly consider each in turn.

The question of what to do is the question that confronts us in practical deliberation. When we engage in practical deliberation we ask ourselves and try to resolve the question of what to do: whether to go on holiday to Corsica or Sicily or Sardinia; whether to send one's children to a private or a public school; whether to have a sixth *pain au chocolat*; and so on.

Following Pamela Hieronymi (2009; 2011), I propose that the question of what to do is not the same as the question of *what one ought to do*.[3] One way to bring this out is to notice that we resolve the question of what we ought to do by forming a normative *belief*, whereas we resolve the question of what to do by forming an *intention* or *decision* (see Hieronymi 2009; Owens 2011). For example, I might resolve the question of whether I *ought* to have a sixth *pain au chocolat* by forming the belief that I ought not to. This doesn't yet mean that I have resolved the question of *whether* to have a sixth *pain au chocolat*. To resolve the question of *whether* to have a sixth *pain au chocolat* would be to decide to have a sixth *pain au chocolat* or to decide not to.[4]

Another way to bring out the contrast between the question of what to do and the question of what one ought to do is to note that it seems that we can ask ourselves and try to resolve the question of what we ought to do without asking ourselves and trying to resolve the question of what to do. This may happen, for

example, if we have *already* resolved to our satisfaction the question of what to do. Suppose that I have received a substantial Christmas bonus and intend to spend it on a new car. This does not rule out asking myself the question of what I *ought* to do with my bonus – say, whether I ought to give it to Oxfam. In asking myself the question of whether I ought to give the money to Oxfam, I needn't be reopening the question of what to do with my bonus.

Furthermore, it seems that we may also ask ourselves and try to resolve the question of what to do without asking ourselves and trying to resolve the question of what we ought to do. Again, this may happen if we have *already* resolved the question of what we ought to do. Perhaps I have *already* resolved the question of whether I ought to have a sixth *pain au chocolat*, yet the question of *whether* to have a sixth *pain au chocolat* remains open. Or it may happen if one is a consistent normative nihilist who holds that there are no truths about what one ought to do and who is consistent inasmuch as one refrains from having any beliefs about, or even interrogating the question of, what one ought to do. Being a consistent normative nihilist of this kind obviously doesn't mean that one is unable to interrogate the question of, say, *whether* to have a sixth *pain au chocolat* or *whether* to go on holiday to Corsica or Sicily or Sardinia.

III

Claims about the thing to do are (or are made true by) claims about the *correct answer* to the question of what to do. What does it take in order for an answer to the question of what to do to be *correct*?

One possibility is that

(1) For any agent $a \in A$, at time t, deciding to ϕ is the correct answer to the question of whether to ϕ for a only if, at t, a ought to ϕ.

On the assumption that I ought not to have a sixth *pain au chocolat*, (1) implies that if I answer the question of whether to have a sixth *pain au chocolat* by deciding to have one, then I have incorrectly resolved the question of what to do. Nothing that I shall say in what follows is inconsistent with (1).[5]

Another possibility is that the correctness of an answer to the question of what to do is tied more closely to the *attitudes* that agents happen to have. For example, it might be said that

(2) For any agent $a \in A$, at time t, deciding to ϕ is the correct answer to the question of whether to ϕ for a only if, at t, assuming that the attitudes upon which a's decision to ϕ are based are correct, then a ought to ϕ.[6]

Suppose that one resolves the question of whether to send one's children to a public or a private school by deciding to send them to a private school, and that one's decision is based on the belief that sending them to a private school will give them an advantage over others later in life. Perhaps this belief is false. Perhaps sending them to a public school would in fact confer upon them the desired advantage. Suppose, however, that on the false assumption that it is true, it follows that one

ought to send them to a private school -- because it is true that one ought to do what one will give one's children an advantage over others later in life. Under these circumstances, deciding to send them to a private school may count as the correct answer to the question of what to do, even if it is not true that one ought to send them to a private school. Alternatively, suppose that one ought *not* to do what one will give one's children an advantage over others later in life. (Perhaps doing so is contrary to some valid egalitarian principle of justice.) Under these circumstances, deciding to send them to a private school amounts to incorrectly resolving the question of what to do – since, assuming that one's (false) belief that sending them to a private school will give them an advantage over others later in life is true, it follows that one ought not to send them to a private school. This is so, irrespective of whether, in fact, one ought to send to send them to a public or a private school.

I won't try to resolve the question of which (if either) of these accounts of correctness is the right one. It suffices to note that what I shall say in what follows is consistent with a wide variety of views about the correctness of answers to the question of what to do.

IV

I now want to suggest two important correctness constraints.

Notice that an answer to the question of what to do may fail to be correct in two ways. First, it may be the case that some *rival* answer is the correct one. This is the kind of failure that happens, for example, if one resolves the question of whether

to have the sixth *pain au chocolat* by deciding to have it when, in fact, the correct answer to the question was to decide not to have it.

But, second, an answer may fail to be correct because there is something wrong with the *question*. Certain acts are such that it is not correct for an agent even to *ask* herself the question of whether to perform them. Where an act is such that it would be incorrect for an agent to ask herself the question of whether to perform it, it trivially follows that deciding to perform that act cannot be the correct answer to the question of what to do. This suggests the following constraint on the correctness of answers to the question of what to do:

(3) For any agent $a \in A$, at time t, deciding to ϕ is the correct answer to the question of whether to ϕ for a only if it is correct, at t, for a to ask herself the question of whether to ϕ.

Call this "the correctness of answer constraint."

What kinds of acts are such that it is not correct for an agent to ask herself the question of whether to perform them? Take *past acts*, such as my kissing my wife yesterday morning. I cannot correctly ask myself the question of whether to kiss my wife yesterday morning. That involves a kind of mistake. Past acts are beyond the jurisdiction of an agent's practical deliberation. (I can ask myself the question of what past acts I *ought* to have done, but not what past acts to do.)

Or take *others*' acts, such as your buying a new kayak. I cannot correctly ask myself the question of whether you are to buy a new kayak (absent some special

authorization). Again, that would involve a mistake. Others' acts lie beyond the jurisdiction of an agent's practical deliberation. (I can certainly ask myself the question of what others *ought* to do, or what they *will* do, but not *what to do for others*, as it were.)

Or take acts that we lack the *ability* to perform. If I don't speak Russian, then it would be incorrect for me to ask myself the question of whether to have a conversation in Russian with my next-door neighbor about the finer points of Dostoyevsky's pre-Siberian novels. If I don't play the piano, then it would be incorrect for me to ask myself the question of whether to play Ravel's "*Gaspard* de la nuit" this evening.

Or take acts that we have the ability to perform in general but where our ability is impaired or blocked by external obstacles. I may have the ability to keep promises in general. (I am not subject to any kind of promissory pathology.) Yet it may be that I cannot keep my promise to be at your wedding in Dunedin, because this would require me to take a flight from Australia, and all flights out of Australia have been grounded due to severe tropical cyclones. Under these circumstances, it would be incorrect for me to ask myself the question of whether to attend your wedding.

Can we say anything in general about the class of acts such that it is incorrect for an agent to ask herself the question of whether to perform them? It seems that we can. What all such acts have in common is that the agent *cannot* perform them. They are not *available* to the agent, as we might say. As such, they lie beyond the *jurisdiction* of the agent's practical deliberation. Past acts, others' acts, acts that we

lack the ability to perform, and acts where our abilities are impaired or blocked are not acts that we can perform in the relevant respect. This suggests the following constraint on the correctness of asking oneself the question of what to do:

(4) For any agent $a \in A$, at time t, it is correct for a to ask herself the question of whether to ϕ only if, at t, a can ϕ.

Call this "the correctness of question constraint."

V

We are now in a position to give the argument for (TDC).

First, take the correctness of answer constraint. This holds that

(3) For any agent $a \in A$, at time t, deciding to ϕ is the correct answer to the question of whether to ϕ for a only if it is correct, at t, for a to ask herself the question of whether to ϕ.

Next, take the correctness of question constraint. This holds that

(4) For any agent $a \in A$, at time t, it is correct for a to ask herself the question of whether to ϕ only if, at t, a can ϕ.

From the correctness of answer constraint and the correctness of question constraint we can infer that

(5) For any agent $a \in A$, at time t, deciding to ϕ is the correct answer to the question of whether to ϕ for a only if, at t, a can ϕ.

What this means is that if one cannot perform a particular act, then deciding to perform that act cannot be the correct answer to the question of what to do. Thus, for example, deciding (now) to kiss my wife yesterday, deciding that you are to buy a new kayak, deciding to have a literary chat in Russian, deciding to play "*Gaspard* de la nuit" this evening, and deciding to attend one's friend's wedding in New Zealand cannot be correct answers to the question of what to do.

Next, recall that claims about the thing for an agent to do just are (or are made true by) claims about the correct answer to the question of what to do. This implies that

(6) For any agent $a \in A$, at time t, ϕing is the thing for a to do if and only if, at t, *a's* deciding to ϕ is the correct answer to the question of whether to ϕ.

(5) and (6) imply (TDC).

Here, then, is the argument for (TDC):

(3) For any agent $a \in A$, at time t, deciding to ϕ is the correct answer to the question of whether to ϕ for a only if it is correct, at t, for a to ask herself the question of whether to ϕ.

(4) For any agent $a \in A$, at time t, it is correct for a to ask herself the question of whether to ϕ only if, at t, a can ϕ.

Therefore,

(5) For any agent $a \in A$, at time t, deciding to ϕ is the correct answer to the question of whether to ϕ for a only if, at t, a can ϕ (from (3) and (4)).

(6) For any agent $a \in A$, at time t, ϕing is the thing for a to do only if, at t, *a's* deciding to ϕ is the correct answer to the question of whether to ϕ.

Therefore

(TDC) For any agent $a \in A$, at time t, ϕing is the thing for a to do only if, at t, a can ϕ (from (5) and (6)).

VI

This argument establishes the truth of a non-trivial claim, namely, (TDC). It also helps to see why (TDC) is not vulnerable to the familiar complaint, namely, that it licenses impermissible inferences in situations where agents intentionally make it the case that they are unable to perform certain acts.

Suppose that by my act of consigning your prized first edition of Newton's *Principia* to the flames, I have intentionally made it the case that

(7) I cannot return your prized first edition of Newton's *Principia*.

(TDC) licenses us to infer from (7) that

(8) It is not the case that returning your prized first edition of Newton's *Principia* is the thing for me to do.

This inference can be seen to be perfectly unproblematic, once we spell out in full the chain of inferences by which we are entitled to infer (8) from (7).

First, given the correctness of question constraint, i.e. (4), we are entitled to infer from (7) that

(9) It is not the case that it is correct for me to ask myself the question of whether to return your prized first edition of Newton's *Principia*.

Notice that this inference is perfectly unproblematic. Returning your prized first edition of Newton's *Principia* simply isn't available to me. It lies beyond the jurisdiction of my practical deliberation. It doesn't matter why it isn't available to me. The fact that the explanation happens to involve intentional wrongdoing on my part is neither here nor there.

Second, given the correctness of answer constraint, i.e. (3), we are entitled to infer from (9) that

(10) It is not the case that deciding to return your prized first edition of Newton's *Principia* to you is the correct answer to the question of whether to return your prized first edition of Newton's *Principia*.

Again, this inference is perfectly unproblematic. Given that it would be incorrect even to *ask* myself the question of whether to return your first edition of Newton's *Principia*, clearly *answering* it by deciding to do so cannot be the correct answer to the question of what to do.

Third, given that claims about the thing to do entail claims about the correct answer to the question of what to do, i.e. (6), we are entitled to infer from (10) that

(8) It is not the case that returning your prized first edition of Newton's *Principia* is the thing for me to do.

This inference is also perfectly unproblematic. Claims about the thing to do just are (or at least are made true by) claims about the correct answer to the question of what to do.

So, the inference that (TDC) licenses from (7) to (8) is underwritten by a chain comprising three inferences: a) the inference licensed by the correctness of question constraint from (7) to (9); b) the inference licensed by the correctness of answer constraint from (9) to (10); and c) the inference licensed by the definition of claims about the thing to do from (10) to (8). Each of these three inferences is perfectly

unobjectionable. So too, therefore, is the inference from (7) to (8). This is why (TDC) is not vulnerable to the familiar complaint.

VII

I have argued that (TDC) follows from certain important truths about practical reason and that it is not vulnerable to the familiar complaint. This is interesting in its own right. But we might also wonder: Where does this leave us with respect to (OC)?

It depends. In particular, it depends on whether claims about what we ought to do imply claims about the thing to do, i.e. whether it is true that

(11) For any agent $a \in A$, if, at time t, a ought to ϕ then, at t, ϕing is the thing for a to do.

Some philosophers want to tie ought very closely to practical deliberation (see Kant 1998; Korsgaard 1996). For such philosophers, (11) may appear very natural, even truistic. Other philosophers insist upon sharply distinguishing ought and practical deliberation and are likely to be hostile to (11) (see Broome 2013). For my part, I am not going to take a stand on the issue of whether (11) is true or false.[7] This is partly because I am genuinely unsure about whether (11) is true or false.[8] More importantly, it doesn't matter for my purposes here. Either way, our argument for (TDC) has interesting implications for (OC).

Suppose, first, that (11) is true. This means that (OC) must also be true. Here is the argument:

(11) For any agent $a \in A$, if, at time t, a ought to ϕ then, at t, ϕing is the thing for a to do.

(TDC) For any agent $a \in A$, at time t, ϕing is the thing for a to do only if, at t, a can ϕ.

Therefore,

(OC) For any agent $a \in A$, at time t, a ought to ϕ only if, at t, a can ϕ.

So we have a new argument for (OC) – or at least a new argument for a version of (OC) that is restricted to actions.[9] This is interesting.

If (11) is true, it also means that (OC) is not vulnerable to the familiar complaint. The familiar complaint holds that (OC) licenses impermissible inferences in situations where we wrongfully make it the case that we cannot perform certain acts. We saw that (TDC) is not vulnerable to the familiar complaint, since the apparently objectionable inferences licensed by (TDC) can be seen to be perfectly unobjectionable once we spell out the complete chains of inference that underwrite the apparently objectionable inferences. If (11) is true, then the apparently objectionable inferences licensed by (OC) will be underwritten by the same chains with an extra link including those licensed by (11). So we also have a new and interesting response to the familiar complaint to (OC).[10]

Suppose instead that (11) is false. The falsity of (11) does not *entail* the falsity of (OC). Does this mean that our argument for (TDC) is *irrelevant* to (OC) if (11) is false? One way in which it might be thought to be relevant is by suggesting an *analogous* argument for (OC) – one that appeals to analogous of (3), (4) and (6), i.e.

(3*) For any agent $a \in A$, at time t, a's believing that she ought to ϕ is the correct answer to the question of whether she ought to ϕ only if it is correct, at t, for a to ask herself the question of whether she ought to ϕ.

(4*) For any agent $a \in A$, at time t, it is correct for a to ask herself the question of whether she ought to ϕ only if, at t, a can ϕ.

Therefore,

(5*) For any agent $a \in A$, at time t, a's believing that she ought to ϕ is the correct answer to the question of whether she ought to ϕ only if, at t, a can ϕ (from (3*) and (4*)).

(6*) For any agent $a \in A$, at time t, a ought to ϕ only if, at t, a's believing that she ought to ϕ is the correct answer to the question of whether she ought to ϕ.

Therefore,

(OC) For any agent $a \in A$, at time t, a ought to ϕ only if, at t, a can ϕ (from (5*) and (6*)).

I am skeptical, however, that this analogous argument will end up doing any serious dialectical work. Take (4*). (4*) implies that it is incorrect to ask oneself the question of whether one *ought* to become a better person than one is capable of

becoming. This is not obviously correct. I suspect, however, that insofar as (4*) strikes us as plausible, its plausibility is due to the evident truth of (4) proper (it is clearly incorrect to ask oneself the question of *whether* to become a better person than one is capable of becoming); and one's implicitly accepting the idea that there is a connection between the question of what we ought to do and the question of what to do, such that we ask and answer the former correctly only if we ask and answer the latter correctly. If we don't accept that there is this connection, then it is far from obvious why we should accept (4*). On the other hand, if we *do* accept that there is this connection, then it's hard to see how one could do so without also accepting (11). And if one accepts (11), then the analogous argument for (OC) is redundant, since, as we saw, (OC) follows directly from (11) and (TDC).

Still, there is a different way in which (TDC) might be relevant to (OC) even if (11) is false. Suppose that you are antecedently skeptical about (OC) – perhaps because of the familiar complaint, perhaps for some other reason. It would be nice to be able to explain why (OC) has nonetheless mistakenly struck so many people as compelling. In other words, it would be nice to be able to provide an *error theory* of (OC). Notice that, if (11) is false, (TDC) may help us to provide such an error theory. Even if (11) is false, it is presumably not obviously false. Indeed, it may seem true, even truistic. So here is how the error theory goes: We fall into error because we recognize implicitly the truth of (TDC) and we mistakenly conflate claims about what we ought to do and claims about the thing to do. The apparent truth of (OC) is parasitic on the truth of (TDC) coupled with the apparent truth of (11).

References

Alston, William. 1988. "The deontological conception of epistemic justification," *Philosophical Perspectives*, 2, 257-99.

Baltzly, Dirk. 2000. "Moral dilemmas are not a local issue," *Philosophy*, 75, 245-63.

Brennan, Geoffrey and Nicholas Southwood. 2007. "Feasibility in action and attitude." In *Homage à Wlodek: Philosophical Papers Dedicated to Wlodek Rabinowicz*, ed. Toni Rønnow-Rasmussen, Björn Petersson, Jonas Josefsson, and Dan Egonsson, URL = <http://www.fil.lu.se/hommageawlodek/site/papper/Brennan&Southwood.pdf>.

Broome, John. 2013. *Rationality Through Reasoning*. Oxford: Blackwell.

Chuard, Philippe and Nicholas Southwood. 2009. "Epistemic norms without voluntary control," *Nous*, 599-632.

Feldman, Richard. 2000. "The ethics of belief," *Philosophy and Phenomenological Research*, 60, 667–95.

Gibbard, Allan. 2003. *Thinking How to Live*. Cambridge, MA.: Harvard University Press.

Hare, R. M. 1963. *Freedom and Reason*. Oxford: Clarendon Press.

Hieronymi, Pamela. 2009. "The will as reason," *Philosophical Perspectives*, 23, 201-220.

Hieronymi, Pamela. 2011. "Reasons for action," *Proceedings of the Aristotelian Society*, 111: 407-27.

Howard-Snyder, Frances. 2006. "'Cannot' implies 'not ought,'" *Philosophical Studies*, 130: 233-46.

Kant, Immanuel. 1998. *Groundwork of the Metaphysics of Morals*, ed. Mary Gregor. Cambridge, UK; New York: Cambridge University Press.

Kolodny, Niko. 2005. "Why be rational?" *Mind,* 114, 509-63.

Korsgaard, Christine. 1996. *The Sources of Normativity*. Cambridge: Cambridge University Press.

McConnell, Terrance. 1989. "'Ought' implies 'can'" and the scope of moral requirements." *Philosophia*, 19: 437-54.

Moran, Richard. 2001. *Authority and Estrangement: An Essay on Self-Knowledge*. Princeton: Princeton University Press.

Owens, David. 2011. "Deliberation and the first person." In *Self Knowledge*, ed. Anthony Hatzimoysis. Oxford: Oxford University Press, pp. 261-77.

Ross, Jacob. 2012. "Rationality, normativity, and commitment," *Oxford Studies in Meta-Ethics*, 7, 138-81.

Schroeder, Mark. 2009. "Means-ends coherence, stringency, and subjective reasons," *Philosophical Studies*, 143, 223-48.

Sinnott-Armstrong, Walter. 1984. "'Ought' conversationally implies 'can,'" *Philosophical Review*, 93, 249-61.

Southwood, Nicholas. 2008. "Vindicating the normativity of rationality," *Ethics*, 119, 9-30.

Southwood, Nicholas. Ms. "The problem of rationality," unpublished paper.

Stocker, Michael. 1971. "'Ought' and 'can,'" *Australasian Journal of Philosophy*, 49: 303-16.

Streumer, Bart. 2003. "Does 'ought' conversationally implicate 'can'?" *European Journal of Philosophy*, 11, 219-28.

Streumer, Bart. 2007. "Reasons and impossibility," *Philosophical Studies*, 136, 351-84.

Vranas, Peter. 2007. "I ought, therefore I can," *Philosophical Studies*, 136, 167-216.

Young, Robert. 1975. *Freedom, Responsibility and God*. New York: Harper and Row.

Zimmerman, Michael M. 1996. *The Concept of Moral Obligation*. New York: Cambridge University Press.

---

[1] Earlier incarnations of this paper were presented at seminars and workshops at Aarhus, ANU, Cambridge, Monash and Oxford. On those and other occasions, I have benefitted immensely from the input of many (largely skeptical) friends and colleagues about the ideas in the paper. I would particularly like to thank Christian Barry, Geoff Brennan, John Broome, Tom Christiano, Ryan Cox, Garrett Cullity, Tom Dougherty, Ant Eagle, Dave Estlund, John Filling, Daniel Friedrich, Bob Goodin, Allison Hills, Karen Jones, Benjamin Kiesewetter, Holly Lawford-Smith, Leon Leontyev, Terry McDonald, Tristram McPherson, Mike Otsuka, Philip Pettit, David Plunkett, Jon Quong, Simon Rippon, Raffaele Rodogno, Mark Schroeder, François Schroeter, Laura Schroeter, Michael Smith, Annie Stilz, Patrick Tomlin, Laura Valentini, Stuart White, David Wiens, and a very astute anonymous referee for this journal. Research for the paper was supported by ARC Discovery Grant DP140102468.

[2] By which I mean that they succeed in blocking the complaint at the cost of surrendering the ability of (OC) to explain certain important, apparently legitimate inferences. For example, Walter Sinnott-Armstrong (1984) famously suggested the following interpretation of (OC): Claims about "ought" *conversationally implicate* claims about "can." Requirements of conversational implicature do not permit contraposition. So (OC), thus interpreted, does not license us to infer from the fact that I cannot return your first edition of Newton's *Principia Mathematica* that it's not the case that I ought. Good. But notice that not only does Sinnott-Armstrong's interpretation of (OC) rule out such inferences in the bad cases, it also rules out such inferences in the good cases. Suppose that I have promised the coach that I will open the bowling for the Coolangatta 1st XI, but then break both legs the day before the big game. It seems that in this case we are entitled to infer that it's not the case that I ought to open the bowling for the Coolangatta 1st XI. But Sinnott-Armstrong's interpretation of (OC) does not license this apparently legitimate inference.

[3] See also Southwood ms, Moran 2001 (and perhaps Gibbard 2003).

[4] The question of what to do must also be distinguished from the question of what one *will* do. We resolve the question of what we will do by forming a belief, a belief that we will or won't do this or that. Suppose that my more intrepid wife has persuaded me to go on a skydiving expedition with her on Saturday afternoon. As our plane begins its ascent, I might very well ask myself the question of what I will do when the door opens. Will I manfully throw myself out of the plane? Or will I chicken out? Notice that in asking myself this question, I am not asking myself the question of *whether* to throw myself out of the plane.

[5] While I have suggested that the *question* of what to do is not the same as the *question* of what one ought to do, nothing I say in what follows is supposed to be inconsistent with the claim that there is an important connection between the *correct answer* to the question of what to do and *what one ought to do*.

[6] This is based on the account of subjective rightness provided by Mark Schroeder (2009). Other kinds of attitude-relative criteria are offered by e.g. Kolodny 2005; Southwood 2008; and Ross 2012.

[7] I am very grateful to an anonymous referee for helping me to see the benefits of neutrality with respect to (11).

[8] One possible objection to (11) is that it might be thought to entail a kind of "negative bootstrapping." First, it might be argued that claims about the correct answer to the question of what to do (and hence claims about the thing to do) must be somehow tightly constrained by claims about our *actual mental states*, even when these mental states are objectionable. (For example, perhaps the correct answer to the question of what to do must not involve an akratic decision.) Second, if (11) is true, then it follows that claims about what we ought to do must *also* be equally tightly constrained by claims about our actual mental states. Third, from this it might seem to follow that that we can 'bootstrap' ourselves out of having certain obligations simply by having relevant objectionable actual mental states. (For example, if the correct answer to the question of what to do must not involve an akratic decision, then it would follow from (11) that I can "bootstrap" myself out of being obliged to look after my children simply by having the belief that I ought not to look after my children.) But, fourth, it does not seem that we may "bootstrap" ourselves out of having obligations in this way. Therefore, we must reject (11). Call this "the negative bootstrapping objection."

[9] Some philosophers have defended restricting (OC) to actions on other grounds (e.g. Alston 1988). Others have objected to this restriction (e.g. Feldman 2000 and Chuard and Southwood 2009). To be sure, we might wonder whether there is an analogous doxastic normative concept, "the thing to believe," and, if so, i) whether "the thing to believe" implies "can" and ii) whether claims about what we "ought" to believe imply claims about "the thing to believe." I won't try to address these interesting questions here.

[10] It may also mean that we are able to bolster certain existing response to the familiar complaint. For example, one appealing existing response holds that the familiar complaint conflates claims about what we *ought to do* and claims about what we *ought to have done* (Zimmerman 1996; Streumer 2003). Properly interpreted, (OC) only licenses inferences in the case of claims about what we ought to do. Thus, it licenses us to infer from the fact that I cannot return your prized first edition of Newton's *Principia* that it is not the case that I *ought to return it*. But it does not license us to infer that it is not the case that I *ought to have returned it*. If (11) is correct, then this provides us with a principled rationale for restricting (OC) to claims about what we ought to do and excluding claims about what we ought to have done, since clearly the latter do not entail claims about the thing to do. I am grateful to an anonymous referee for drawing my attention to the complementary nature of the relation between my response to the familiar complaint and the Zimmerman-Streumer response.