

DO WE HAVE REASONS TO DO AS WE BELIEVE WE OUGHT TO DO?*

Julian Fink
Oxford University

Suppose you believe you ought to p. Does your failure to p imply that you are not entirely as you ought to be? Ought you to p if you believe you ought to p? This paper argues for a qualified version of this claim. It is qualified in two ways. First, I assume that this can be so only if ‘if you believe you ought to p’ appears within the scope of ‘you ought’. That is, you ought to [if you believe you ought to p, then p.] Second, I argue that you ought to do as you believe you ought to only *as far as it goes*; that is, unless there are exceptional reasons not to do so. In this sense, we have a *pro tanto* reason to do as we believe we ought to do.

My argument for this claim relies on two major premises: (i) you have a normative reason to discard any belief that you ought to p if it is not the case that you ought to p; (ii) if you have a reason to p or you have a reason to q, then you have a reason to (p or q). I defend (i) against a number of views in epistemology which exclude that the falsity of a normative proposition can be a reason not to believe this proposition. I argue agglomeration as in (ii) does not lead to the paradoxical results ascribed to it. After my defense of (i) and (ii) I show how these two premises entail the conclusion that we have a normative reason [to p if we believe we ought to p].

* * *

Suppose rationality requires us to do as we believe we ought to do. Are we only *irrational* in not satisfying this requirement? Or is it that we are also *not*

* The present paper is an adapted version of a longer paper, suitable for a 30 minutes conference presentation.

entirely as we ought to be? Do we have *normative reasons* to do what we believe we ought to do?

Most philosophers just take this for granted. You ought to do as rationality requires you to do. So, if rationality requires you to do as you believe you ought to do, then you ought to do what you believe you ought to do. For some this just follows *a priori*, assuming that ‘being rational’ is just *synonymous* with ‘being as one ought to be’.

Others hold that, though not having the same meaning, being rationally required to behave in a particular way implies necessarily that you ought to behave so. James Dreier even contends that ‘there is no sense at all to be made of the question of whether we have any reason to follow the rules of rationality’¹. He thinks that normative reasons, ought, and rationality are so intimately connected that asking this question would be like asking ‘if you ought to p, ought you to p?’, which admittedly makes no sense. It is simply unthinkable that ‘rationality requires you’ does not entail ‘you ought to’.

But it is not unthinkable. Two examples immediately reveal why. Suppose rationality requires you to p if you believe you ought to p. Also, suppose you believe that you ought to p. Is it true now that you ought to p? Clearly not. For this would imply that your ‘ought beliefs’ could never be false. They would make themselves true. Yet, this is surely implausible. There are false normative beliefs. Suppose you believe you ought to extinguish all mankind. In almost all conceivable situations, this will be false. Furthermore, suppose that an evil dictator withdraws from her attempt to extinguish all mankind only if you manage not to do what you believe you ought to do. In this situation, very plausibly, you ought *not* to do what you believe you ought to do; hence, you ought to be *irrational*.

Nonetheless, this paper argues for a qualified version of the claim that we ought to p if we believe we ought to p. My argument for this relies on two major premises: (i) that you have a reason not to believe that you ought to if it is not the case that you ought to; (ii) if you have a reason to p or you have a reason to q, then you have a reason to [p or q].

In the following I defend (i) and (ii) against putative objections. In particular, I defend (i) against ‘exclusive’ versions of evidentialism and teleology which both imply that the falsity of a normative proposition is no reason not to believe this proposition. Furthermore, I defend (ii) by showing that by agglomerating two disjunctively connected reasons to a disjunction governed

1 J. DREIER, ‘Humean doubts’, p. 29.

by a reason does not lead to paradoxical results. I show how (i) and (ii) entail that we have a normative reason [to p if we believe we ought to p].

Ought, reasons, and normative falsehoods

Suppose you believe that you ought to p. Do you now have a reason for bringing about that p? Or is it even that you ought to p? As pointed out in the introduction, if understood pre-theoretically, the answer to (at least) the last question must be ‘no’. Otherwise our normative beliefs would make themselves true. But this is clearly implausible. There are false normative beliefs, as I have shown above.

However, there is another way to understand the question of whether we ought to do what we believe we ought to do. Suppose it is not the case that [rationality requires you to p]² *whenever* you believe you ought to p. Instead rationality requires you to *avoid* situations in which you believe that you ought to p without bringing it about that you p. In other words, it requires you either to bring about that p or not to believe that you ought to p. This means that you can satisfy this requirement either by bringing about p or by ceasing to believe that you ought to p. Assuming that you ought to satisfy this ‘wide-scope’ requirement then no longer implies that your normative beliefs can never be false. All it would tell you is that you ought to avoid believing that you ought to p without bringing about that p. In other words, you ought to [if you believe you ought to p, then p]³.

I have defended the wide-scope reading of the requirements of rationality elsewhere, but due to limited space, I cannot repeat my argument here. I will just assume that the wide-scope reading represents the correct logical form of the requirement. This opens up the possibility that we have a reason to [if we believe we ought to p, then p].

I now come to my argument for this claim. First, consider two premises.

$$O(p) \vee \neg O(p),$$

2 The brackets are used to show that ‘rationality requires’ only governs ‘you to p’ and does not include ‘whenever you believe you ought to p’.

3 I owe the wide-scope interpretation of the requirements of rationality to John Broome. In particular, see his ‘Normative requirements’, his ‘Are intentions reasons?’, and his ‘Reasons’. For an earlier interpretation of the wide-scope form of requirements, see Jonathan Dancy’s ‘The logical conscience’.

where ‘O’ stands for ‘you ought to’, ‘p’ for a proposition, ‘v’ for the logical disjunction, ‘¬’ for the negation, ‘→’ for the material conditional, and ‘B’ for ‘you believe that’. I take (1) to be trivially true. It is the law of the excluded middle applied to normative propositions.

In contrast to (1), the second premise of my argument is not trivially true. Suppose it is not the case that you ought to p, but you believe that you ought to p. What should be your attitude towards this belief? Retain it or discard it? Mark Platts answers this question as follows: “(F)alsity is a decisive failing in a belief, and false beliefs *should* be discarded; beliefs should be changed to fit with the world”⁴.

Following Platts’s idea, the second premise of my argument says that you ought to discard your normative belief if it is false:

$$(2) \neg O(p) \rightarrow O\neg B(O(p))$$

(2) assigns an ‘ought’ to ‘not having a normative belief’ on the basis of a property of the content of your belief. However, other normative theories use other features to determine whether or not you ought to have it. (2) may very well conflict with these views. One view that clearly conflicts with (2) is what I call ‘conclusive teleology’. Conclusive teleology assigns an ‘ought’ to ‘not having a normative belief’ not on the basis of some features of the belief’s content. Instead, you ought not to have a normative belief if being in this *state* of believing is relevantly beneficial. For normative beliefs, conclusive teleology can roughly be defined as follows:

4 M. PLATTS, *Ways of Meaning*, p. 256; original emphasis.

5 Note that (2) should not be confused with two similar principles, (2a) and (2b):

(2a) $\neg O(p) \rightarrow O(B(\neg O(p)))$

(2b) $\neg O(p) \rightarrow \neg O(B(O(p)))$

(2a) says that if it is not the case that you ought to p, then you ought to believe that it is not the case that you ought to. Unless conjoined with a principle that you ought not [to believe that you ought to p and that it is not the case that you ought to p] (simultaneously), (2a), unlike (2), does not require you to rid yourself of any belief that you ought to p if it is not the case that you ought to p. It tells you to disbelieve that you ought to p if this is not the case. (2b) says that if it is not the case that you ought to p then it is not the case that you ought to believe that you ought to p. Unless conjoined with the principle that for every normative proposition N, you either ought to believe N or you ought not to believe N, (2b), unlike (2), does not tell you to discard your beliefs in normative falsehoods. Instead, it just says that it is not the case that you ought to believe it.

Conclusive teleology. Necessarily, if believing that you ought to p is sufficiently beneficial (in the relevant sense)⁶, then you ought to believe that you ought to p.

I call this view ‘teleology’ because it determines whether you ought to believe a normative proposition on the basis of the benefits of believing it. It is a ‘conclusive’ version of teleology because being relevantly beneficial suffices to ensure that there is a *conclusive* reason to believe something, as opposed to there being merely a *pro tanto* reason. By ‘conclusive reason’ I mean a consideration that explains why you ought to do something which, unlike a ‘*pro tanto* reason’, could not be overridden by other reasons. I add ‘necessarily’ to the formula to show that the conditional states a genuine entailment, and not only a coincidental connection⁷.

Here is a plausible example of when the benefits of a normative belief determine that you ought to have it. Suppose that an eccentric millionaire offers to donate £1 million to UNICEF if you manage to genuinely believe that you ought to fly to the moon. But it is not the case that you ought to fly to the moon. Suppose you have an obligation to pick up your son from school and that this is incompatible with you flying to the moon. Nevertheless, assuming that you can believe that you ought to fly to the moon, conclusive teleology then implies that you ought to believe that you ought to fly to the moon—even though this is not the case. This means that (2) must be wrong.

The second view (2) conflicts with is ‘conclusive evidentialism’. Roughly, this view defines roughly as follows:

Conclusive evidentialism. Necessarily, if there is conclusive evidence that you ought to p, then you ought to believe that you ought to p.

This view is obviously a version of evidentialism because it defines a strong link between evidence for a normative proposition and you being normatively required to believe it. It is a ‘conclusive’ version of evidentialism

6 I add ‘in the relevant sense’ to imply that your normative belief must be beneficial a particular way that contributes to determining what you ought to believe, all things considered. Further, I add ‘sufficiently’ to indicate some quantitative degree of being beneficial. For instance, being marginally beneficial may not suffice to ensure that you ought to have a normative belief. There may be other considerations that override the beneficial effects of your belief.

7 Unless specified otherwise, this is the reasons for why I will add ‘necessarily’ to my other formulas too.

because conclusive evidence for a normative belief provides a conclusive reason to believe a normative proposition.

Conclusive evidentialism entails that sometimes it will be the case that you ought to believe that you ought to *p*, even if it is not the case that you ought to *p*. This will be so if there is conclusive (yet misleading) evidence for the truth of ‘you ought to *p*’. If so, you ought to retain or form a belief that you ought to *p* even though it is false. Hence, (2) cannot be correct.

To save space, let me just assume conclusive teleology or conclusive evidentialism show that (2) to be false. I am sure this is so. Therefore, I need a modified version of (2) that is consistent with conclusive teleology and evidentialism. Here is my new version of (2):

$$(2^*) \neg O(p) \rightarrow R\neg(B(O(p))),$$

where ‘*R*’ stands for ‘you have a *pro tanto* reason’. (2*) says that if it is not the case that you ought to *p* then you have a reason not to believe that you ought to *p*. Unlike (2), (2*) is compatible with conclusive teleology and conclusive evidentialism. You may have a *pro tanto* reason to loose your belief that you ought to *p* even if, all benefits and evidence considered, you ought to retain it.

Before I come to discuss (2*), let me just state two general assumptions which, due to its length, I cannot argue for in this paper⁸. First, I simply assume that in (2*), ‘*R*’ is meant to represent a *normative* reason for not having a belief. However, one might argue that, though the falsity of *N* gives you ‘a reason’ not to believe *N*, the type of reason it gives you (epistemic, prudential, teleological, etc.) is not genuinely normative. Nevertheless, I will just assume if the falsity of *N* gives you a reason not to believe, this reason is genuinely normative.

However, assuming that ‘*R*’ in (2*) represents a normative reason forces me to make another assumption which I cannot justify here. I will just assume that in every possible world in which (i) it is not case that you ought to *p* and (ii) there is no (overriding) reason to believe you ought to *p*, you *can* not-believe that you ought to *p*. This assumption is necessary since the normativity of ‘*R*’ in (2*) entails that you *ought* not to believe that you ought to *p* in every possible world where both (i) and (ii) obtains. So, if ‘ought’ implies ‘can’, (i) and (ii) must also imply that you can not-believe that you

⁸ I discuss these two assumption in the long version of this paper.

ought to p. However, on a first view, (i) and (ii) do not seem to imply that; yet I will simply assume that they do.

I now turn to some more specific objections against (2*). These objections entail that the falsity of a proposition is simply no reason not to believe it. The following two views enforce this objection. I will call these views ‘exclusive teleology’ and ‘exclusive evidentialism’. Applied to normative beliefs, exclusive teleology is defined as follows:

Exclusive teleology. Necessarily, you have a reason not to believe that you ought to p only if not believing that you ought to p is beneficial (in the relevant sense).

Exclusive teleology says that *only* teleological considerations count in determining whether you ought to discard a normative belief. By assuming that it is not necessarily beneficial (in the relevant sense) to discard a false normative belief, exclusive teleology shows that (2*) cannot be correct. The falsehood of a normative belief is no reason not to believe it.

Applied to normative propositions, exclusive evidentialism defines as follows:

Exclusive evidentialism. Necessarily, you have a reason not to believe that you ought to p only if there is evidence that it is not the case that you ought to p.

Exclusive evidentialism says that only evidential considerations count as reasons for losing a normative belief. Unless you suppose that the falsehood of normative proposition itself counts as evidence for not believing it, exclusive evidentialism refutes (2*) straightaway. The falsity of a normative proposition does not give you a reason not to believe it. Only evidence for the falsity of a normative proposition does so.

I think both ‘exclusive’ versions of teleology and evidentialism are incorrect. They therefore fail to refute (2*). In fact, I have already given a counterexample to exclusive evidentialism. The fact that the eccentric millionaire’s offer to give £1 million to UNICEF implies that you ought to believe a normative falsehood shows that evidential considerations are not the

only reasons for believing something. Teleological considerations are also reasons.

However, why are teleological considerations not the *only* considerations counting in favour of believing something? I can only appeal to the intuition here that beliefs have a certain standard of correctness, and that this standard relates to the *content* of a belief and not to the state of believing. If this content-related standard is normative, as many philosophers argue⁹, it implies that you have a normative reason to form or modify your beliefs so that they fulfil the standard of correctness. Whatever the relevant standard of correct belief may be (truth, justification, or a combination of both, i.e. (arguably) knowledge, etc.) it will imply that teleological considerations are not the only considerations that count for or against believing something¹⁰.

So far, I have shown that (2*) is compatible with conclusive evidentialism and conclusive teleology. Furthermore, I argued that exclusive teleology and exclusive evidentialism do not threaten the truth of (2*)? But what is the positive evidence for (2*)? Why should we accept it? I think there are three potential ways of defending (2*).

The first defence comes from a weak version of evidentialism. Suppose that possessing evidence for the truth of a normative proposition N counts as a reason for believing N. Suppose also that one condition for this to be so is that we have a reason for believing N if N is true. Then it would follow that we have a reason to believe N if N —which arguable leads to a reason not to believe N if $\neg N$; hence (2*).

The key part of this argument can be formulated as follows.

- (3a) $E(O(p)) \rightarrow R(B(O(p)))$
 and (3b) $(E(O(p)) \rightarrow R(B(O(p)))) \rightarrow (O(p) \rightarrow R(B(O(p))))$.
 So, (3c) $O(p) \rightarrow R(B(O(p)))$.
 So, (2*) $\neg O(p) \rightarrow R(\neg(B(O(p))))$,

where ‘E’ stands for ‘you have evidence that’.

9 See, for example, Allan Gibbard’s ‘Thoughts and norms’; Ralph Wedgwood’s ‘The aim of belief’.

10 This is so because ‘being beneficial (in the relevant sense)’ is a *state-related* and not a *content-related* aspect of a belief.

In a nutshell, this argument derives (2*) from the fact that a particular version of evidentialism (3a) presupposes the truth of (2*). Should we accept this argument? One way to dismiss it would be to deny the truth of evidentialism. Evidence does not give us reasons to believe something. Yet in showing that exclusive teleology is incorrect I already accepted this version of evidentialism, and I will continue to do so. Also, one could doubt that the inference from (3c) to (2*) is incorrect. This is so because it presupposes that if you have a reason to believe that it is not the case that you ought to p, then you have a reason not to believe that you ought to p. These two reasons may sometimes come apart, however. Nevertheless, for explorative purposes, I will accept it too. Hence, the question becomes why (3b) is correct. The idea behind (3b) is the following. Asking why one has a reason to believe what evidence supports, the only tenable answer seems to be that doing so leads, rather reliably, to true beliefs. It may be even be the most reliable mechanism available to arrive at true beliefs. Contrarily, if believing in accordance with the evidence would not be a reliable means to come to true beliefs, there would be no reason to believe what the evidence supports. This may be so in a quirky world where evidence is often misleading and where there may be better means to acquire true beliefs (for instance, testimony, playing dice, etc.). It follows that we have a reason to align our beliefs according to our evidence only if we have a reason to believe what is true; hence (2*).

Unfortunately, I cannot rely on this argument in defending (2*). This is because I cannot exclude that the ultimate reason for why evidence provides us with reasons for believing something is not that this would reliably lead to true beliefs, but because it would, for instance, lead to beliefs that are justified. Justification may after all be the ultimate aim of a belief. Alternatively, evidence may guide us to beliefs that are (in the relevant sense) beneficial for us. Being beneficial could also be a belief's ultimate aim. Evidentialism would therefore no longer need to presuppose that we have a reason to believe N if N is true; it would therefore no longer support (2*).

Here is another attempt to defend (2*). One way commonly used to show that we have a reason to behave in a certain way is to show that a certain type of behaviour is a means to some desired end we have. Perhaps we all have some desire, given we ought to behave in a particular way, to believe that we ought to behave in this particular way. Perhaps we also believe that a necessary means to have true beliefs about how we ought to behave is not to believe a normative proposition if it is false. This would then give us a reason

not to believe that we ought to p if it is not the case that we ought to p , i.e. (2*); or so it is argued.

This argument can be formulated as follows:

- (4a) $D(O(p) \rightarrow B(O(p)))$
 and (4b) $B((O(p) \rightarrow B(O(p)))) \rightarrow (\neg O(p) \rightarrow \neg B(O(p)))$
 So, (4c) $R(\neg O(p) \rightarrow \neg B(O(p)))$.
 So, (2*) $\neg O(p) \rightarrow R\neg(B(O(p)))$,

where ‘D’ stands for ‘you desire that’. Unfortunately, this does also not provide us with a defence of (2*). First, it is not clear that we all share a desire, if we ought to p , to believe that we ought to p . Furthermore, not everyone believes that in order to believe that you ought to p whenever you ought to p , it is necessary not to believe that you ought to p if it is not the case that you ought to p . Someone who suffers from severe irrationality will perhaps deny this.

Nevertheless, for the sake of the argument, suppose everyone has such a desire and such a belief. Would we then have a reason not to believe a false normative proposition? I think the answer to this question is ‘no’ for two reasons. First, this argument falls prey to a version Michael Bratman’s ‘bootstrapping objection’¹¹. A desire and a means-end belief cannot ‘bootstrap’ a reason into existence. For example, suppose you desire to visit your friends, and you, strangely enough, happen to believe that a necessary means to do so is to kill your friend. It would be absurd to accept that you now have a reason to kill your friend. But even by supposing that you can pull a reason out of a contingent desire and a means-end belief, it would not give us the truth of (2*). This is because the inference from (4c) to (2*) is not valid. For the fact that you have a reason to make a conditional true (as in (4c)) does not mean that you have a reason to make true the consequent of this conditional. To get a better understanding of why this is so, let us write the conditional in (4c) as a disjunction and assume ‘ p ’ stands for ‘you help your neighbours’. Then, (4c) says that you have a reason to [not believe that you help your neighbours or you ought to help your neighbours]. This is plainly consistent with the conjunction that it is not the case that you ought to help your neighbours *and* it is not the case that you have a reason not to believe that

11 M. BRATMAN’S, *Intentions, Plans, and Practical Reason*, pp. 23-7. See also John Broome’s ‘Are intentions reasons?’.

you ought to help your neighbours. However, (2*) is *not* consistent with this conjunction. For if it is not the case that you ought to help your neighbours, (2*) implies that you have a reason not to believe that you ought to help your neighbours. Consequently, (4c) is consistent with a set of proposition (2*) is not consistent with. That is why (4c) cannot imply (2*).

I think (2*) ultimately needs to be premised on the following argument. One aspect of normative beliefs I have so far neglected is their practical character. By ‘practical character’ I mean their intimate connection with actions. A person free of weakness of the will and other forms of irrationality will end up doing what she believes she ought to do¹². Suppose you falsely believe that you ought to kill yourself. Whilst this belief alone may not be bad in itself, its consequences may be disastrous. In fact, given a sufficient degree of rationality, you will end up killing yourself, and your normative belief will be the cause of it. Without doubt, you have a reason to rid yourself of your belief that you ought to kill yourself. In general, you have a teleological reason to avoid believing normative falsehoods.

One may object that this example supports (2*) only because there is a strong reason not to do what you believe you ought to do (i.e. to kill yourself). This strong reason is what in fact explains why you have a reason not to believe that you ought to kill yourself, and not the fact that this belief is false. Hence, I need to look at cases where you falsely believe you ought to do something, yet there is *no* reason *not* to do what you believe you ought to do. Such cases are genuinely possible. For example, suppose it is not the case that you ought to bend your ring finger, yet you believe that you ought to do so. However, there is also no reason not to bend your finger. That is, there is no reason to do so but also no reason not to do so. Do you have a reason to rid yourself of this ‘normatively neutral’ false normative belief?

I think there is a reason not to believe ‘normatively neutral’ normative falsehoods. It would be uneconomical to clutter our minds with these beliefs; they pose, *ceteris paribus*, a waste of your mental capacity. They take away mental resources, and you have a reason not to waste your mental resources. Of course, this does not mean that there are not other reasons ‘to waste’ your mental resources. For example, it may have beneficial consequence for you or others; you may have strong evidence for it; etc. Nevertheless, there still is a reason for you not to have this belief.

12 On this point see, for example, John Broome’s ‘Normative practical reasoning’.

One may object to this by saying that this is trivially true for every belief, no matter whether true or false, epistemically, morally, or prudentially required. There is a reason to rid yourself of all of your beliefs because they all reduce your mental resources in some way. Clearly, this cannot be so. There is no context-independent (as it were) reason to rid ourselves of all our beliefs we have; or so it is argued.

I agree that this may sound incredible. However, I think it is just a plain consequence of taking the idea of ‘ought, all things considered’ seriously. ‘All things considered’ includes *every* detail that either ‘favours’ or ‘disfavours’ a particular line of action or belief. All I suggest is that the fact that believing something reduces your mental resources will play part in this weighing process. That it will often play only an insignificant part does not damage my argument here.

As a result, I will accept (2*) as correct. This allows me to proceed with my argument as follows.

Argument continued

Conjoining (1) and (2*) allows one to derive (5) by a valid inference:

$$(5) O(p) \vee R(\neg B(O(p))).$$

(5) says that either you ought to p or you have a reason not to believe you ought to p. I take ‘you ought to p’ to imply that ‘you have a reason to p’. This is a plain consequence of supposing that there is nothing you ought to do or to believe without there being an explanation of why this is so. Hence, (5) implies (6).

$$(6) R(p) \vee R(\neg B(O(p))).$$

(6) says that you have a reason to p or you have a reason not to believe that you ought to p. I content that (6) allows us to derive (7).

$$(7) R(p \vee \neg B(O(p))).$$

This inference presupposes the validity of the following agglomerations principle: if you have a reason to p or you have a reason to q, then you have a reason to (p \vee q). There are two ways to denounce this principle. First, one may argue that agglomeration is valid only if the reasons in the original disjunction are of the same kind. By ‘of the same kind’ I mean that they are both either epistemic, prudential, moral, practical, or etc. Second, one may reject this type of agglomeration as such.

Some may reject the inference from (6) to (7) for the first reason. (6) may be interpreted as having a ‘practical’ reason for an action on the left side of the disjunction and an ‘epistemic’ reasons for a belief on its right side. These two reasons are thus ‘inagglomerable’, or so it is argued. But I do not think this is correct. For, on a first view, both reasons govern propositions. In (6), the reason on the left side of the disjunction governs a proposition, namely ‘p’. Also, the reason on the right side of the disjunction governs a proposition, namely ‘ $\neg B(O(p))$ ’. Both reasons are thus reasons to make true a proposition. This suggests that they are both practical in the following sense. Both reasons are reasons for actions. The one on the left side of the disjunction is a reason for an action that leads to the truth of ‘p’; the one on the right side is a reason for a (mental) action that leads to the truth of ‘ $\neg B(O(p))$ ’. In this sense the two reasons may very well be ‘agglomerable’.

However, this will not answer those who principally doubt this type of agglomeration. One way to be doubtful is to show that this type of agglomeration leads to a cognate version of Ross’s paradox in deontic logic¹³. In Ross’s case the situation is the following. Suppose you ought to send a letter. Standard deontic logic allows us then to infer that you ought to [send the letter or burn the letter]¹⁴. Consequently, burning the letter seems to fulfil one of your obligations —even though this is (empirically) inconsistent with your original obligation.

A similar result seems obtainable with the agglomeration principle in question. Suppose you have a reason to send the letter. Via agglomeration it follows that you have a reason to [send it or burn it]. Suppose you burn it. It seems by doing so you indeed do something you have a reason for doing, and that is an incredible result.

Note, however, it does not entail that you fulfil any of your obligations, as in Ross’s example. You may have undefeated reasons to send the letter, and not to burn it. By then burning your letter you clearly fail to do one thing you ought to do even though you had a reason for failing to do so. In fact, I think we often encounter situations like this. A judge may conclude ‘surely, you had a reason to steal the money, but you should not have done it.’; A historian may argue that ‘Stalin had a reason to transport grain out of the Ukraine in 1932, but he ought not to have done so’; etc. There is nothing incredible in that we have a reason to behave in a way we ought not to behave in.

13 Alf Ross, “Imperatives and logic”.

14 See Risto Hilpinen’s “Deontic logic”.

However, there is another way to doubt the correctness of this type of agglomeration. Suppose you have a reason to p . Trivially, this implies ‘you have a reason to p or a reason not to p ’. Agglomeration then gives us ‘you have a reason to $[p \text{ or not-}p]$ ’. So, no matter what you do, you will do something you have a reason for doing. But this seems truly incredible. It is not the case that you have a reason to extinguish mankind, or the like.

I do not think that ‘you have a reason to $[p \text{ or not-}p]$ ’ implies that no matter what you do, you do have a reason for doing so. For this to be correct, ‘you have a reason to $[p \text{ or not-}p]$ ’ would need to entail the conjunction ‘you have a reason to p and you have a reason to not- p ’. Any semantics that would allow this inference would be clearly dubious. It would lead to flat contradictions. Suppose you have a reason to $[p \text{ or } q]$, but you have no reason to p . Suppose however that ‘a reason to $[p \text{ or } q]$ ’ implies ‘you have a reason to p and you have a reason to q ’, and, thus, ‘you have a reason for to p ’. Then it would be the case that you have a reason to p and you have no reason to p (at the same time)¹⁵. Consequently, the fact that agglomeration leads us from reason to p to a reason to $[p \text{ or not-}p]$ does not threaten the validity of the agglomeration principle. I therefore assume agglomeration to be correct.

I now come to the end of my argument. The last step of my argument leads me to transform the disjunction in (7) into a conditional. This results in

$$(8) R(B(O(p)) \rightarrow p).$$

(8) says that you have a reason to [if you believe you ought to p , then p]. In consequence, we indeed have a reason to do what we believe we ought to do. We have a reason to satisfy this major requirement of rationality.

Accordingly, my argument shows that a (significant) part of rationality is normative in a particular sense. We have a reason to satisfy one central requirement of rationality. But what about the other requirements of rationality? Does my argument show their normativity too? Perhaps it does. Theoretically, my argument may prove that we have a reason to satisfy *all* requirements of rationality. In fact, this depends on the following reduction. For all X and Y , if rationality requires you to [if you X , Y], X implies that you believe you ought to Y . If so, it seems that there is a reason for you to satisfy all requirements of rationality. Of course, for *some* requirements this reduction seems dubious. I suppose rationality requires you to believe p if you

¹⁵ I thank Geoff Ferrari for a very fruitful discussion of the agglomeration principle.

believe *not* not-p. This is just another way of saying that rationality requires you not to have contradictory beliefs. It seems incredible that believing *not* not-p implies that you believe that you ought to believe that p. For instance, I believe that it is not the case that my car is not parked in my street, but I do not believe that I ought to believe that my car is parked in my street. In fact, I have no normative belief about my ‘carparking belief’.

Nonetheless, for some requirements this reduction seems very plausible. Surely rationality requires you [if you intend to p and you believe that q is a necessary means to p, intend to q]. It seems plausible to suppose that an intention to p in conjunction with a belief that q is a necessary means for p entails that you believe that you ought to intend to q. For example, Donald Davidson’s account of intentions as unconditional or all-out evaluative judgements will more or less directly lead to this entailment¹⁶. Also, Joseph Raz’s denial of instrumental rationality as a separate category of rationality presupposed this entailment. Assuming then that this entailment is correct, my argument for having a reason to do what we believe we ought to do shows also that we have a reason to take the means necessary to our intend ends. It may therefore turn out that my argument establishes the normativity of an entire bundle of requirements of rationality.

¹⁶ See Donald Davidson’s ‘Intending’. Furthermore, in his ‘The myth of instrumental rationality’, Joseph Raz also presupposes this entailment.