Normative Practical Reasoning

John Broome (john.broome@philosophy.ox.ac.uk)

This is a corrected version of the paper that appeared in the *Proceedings of the Aristotelian Society*, Supplementary Volume 75 (2001), pp. 175–93.

1. Instrumental reasoning.

Aristotle identified practical reasoning as reasoning that concludes in an act, and he was thinking of a non-mental act such as tasting food. But reasoning is a mental process, which takes place in the reasoner's mind. Its conclusion must be a mental state or a mental event; it cannot be a non-mental act. I therefore identify practical reasoning as reasoning that concludes in an intention, which is a mental state. This is as practical as reasoning can get. When reasoning concludes in an intention to act, the intention is in turn likely to cause the intended act. But that last bit of causation is not part of the reasoning process.

One can intend many things besides acts. For example, one can intend to be alert, or to be asked to join the club. I shall call reasoning practical if it concludes in an intention of any sort, not just an intention to act. So in one way my notion of practical reasoning is wider than Aristotle's.

This paper is about one particular sort of practical reasoning, which I shall call 'normative'. Normative practical reasoning sets out from a normative belief – a belief in a normative proposition – and concludes in an intention.

However, I need to do some preliminary work. Many philosophers think that reasoning must conclude in a belief; it cannot conclude in an intention. They doubt the existence of practical reasoning as I understand it. In view of this widespread scepticism, I shall start by offering an example of practical reasoning that seems to me incontrovertible, and explain how it works. Since this is only a preliminary to my discussion of normative practical reasoning, my argument here will be sketchy and leave gaps. I have closed some of the gaps in another article.[1]

Here is the example:

I am going to leave the next buoy to starboard
and In order to leave the next buoy to starboard, I must tack,
so I shall tack.

In this piece of reasoning, the first and third sentences express intentions, and the second a belief. You might go through this reasoning while approaching the next buoy. You already intend to leave it to starboard. When the wind shifts, you acquire the belief that you must tack in order to do so. Reasoning then leads you to form the intention of tacking.

Your reasoning takes you from two mental states, an intention and a belief, to a new mental state, an intention. I shall call them the 'premise-states' and the 'conclusion-state' respectively. The states have contents, which I shall assume to be propositions. Until section 2, I shall also assume that propositions about you can be expressed in the first, second or third person. Then, if your name is 'Leslie', your reasoning can be described like this:

I(Leslie will leave the next buoy to starboard)
and B(In order to leave the next buoy to starboard, Leslie must tack).
As a result, I(Leslie will tack). (α)

'I' and 'B' are predicates of propositions, standing for 'you intend that' and 'you believe that'. (α) is a description of your reasoning. It is not an inference.

Your reasoning is a special type of practical reasoning: it is *instrumental* reasoning, which means it is reasoning from an end to a means. Moreover, it is a special type of instrumental reasoning: it is reasoning from an end to a means that you believe is necessary. So this is a

very special example. I mention it only to show there is at least one sort of genuine practical reasoning, concluding in an intention. Intuitively, it does indeed seem to be genuine, correct, reasoning; it seems you are right to reason this way. Furthermore, that is really so. Here is an explanation of why.

Your reasoning is a sequence of mental states that have propositions as their contents. The content of the reasoning as a whole is this sequence of propositions:

Leslie will leave the next buoy to starboard.
In order to leave the next buoy to starboard, Leslie must tack.
Leslie will tack.

This is a valid syllogism. It would have been valid even if the major premise had been a simple material conditional. As it is, the major premise has a stronger modal form, but the syllogism is valid all the same. If the first two propositions are true, the third must be true.

The validity of this syllogism means that the following would be correct reasoning:

B(Leslie will leave the next buoy to starboard)
and B(In order to leave the next buoy to starboard, Leslie must tack).
As a result, B(Leslie will tack). (β)

This is reasoning that concludes in a belief, which is generally called *theoretical* reasoning. A spectator on shore might run through reasoning of this form. If she has the first two beliefs, she would thereby acquire the third. She would be right to reason this way. The propositions that constitute the content of her reasoning stand in the relation of validity: if the first two are true, then so must the third be. The spectator believes the first two. That is to say, she takes them to be true. Since they cannot be true without the third's also being true (and since this is obvious), she cannot rationally take them to be true without taking the third to be true. This is why her theoretical reasoning is correct.

Your position on the boat is different. Like the spectator, you take the major premise of the syllogism to be true. But you do not simply take the minor premise to be true. You intend it, which is to say you are set to make it true. Since the syllogism is valid, the two premises cannot be true without the conclusion's being true. So you cannot rationally be set to make one premise true and take the other as true, without being set to make the conclusion true. That is to say, you cannot rationally intend one premise and believe the other without intending the conclusion. This is why your reasoning is correct practical reasoning. It is correct because of the very same validity as makes the spectator's theoretical reasoning correct.

David Hume argued that there is no such thing as practical reasoning. His argument has a flaw. It is based on the assumption that 'reason is the discovery of truth or falsehood'.[2] From this it follows immediately that reasoning cannot conclude in an intention. But the assumption is overstated. We might grant Hume that reason is concerned with truth and falsehood, but we should not grant him that it is only the discovery of truth and falsehood. Both theoretical and practical reasoning are concerned with the truth or falsehood of the propositions that constitute their contents. The difference between them is that they regulate different attitudes to truth. One regulates the attitude of taking as true; the other the attitude of being set to make true. In my example, both sorts of reasoning follow the transmission of truth through the valid syllogism. The theoretical reasoning (β) follows it in a truth-taking way; the practical reasoning (α) in a truth-making way.

I have not been entirely accurate. In some circumstances you can rationally intend one premise of a valid syllogism, and believe the other, and yet not intend the conclusion. For example, you may believe the conclusion is true anyway, without your making it true. If so,

you need not intend it. This is a complication, which must be sorted out in a full account of what makes practical reasoning correct.[3] But here I am offering only an example of correct practical reasoning, so I shall ignore the complication. What I have said is right so long as you do not believe that the conclusion of the syllogism is true anyway, without your intending it.

I said that you cannot rationally believe one premise and intend the other without intending the conclusion. More precisely, I meant to say that you ought, if you intend one premise and believe the other, to intend the conclusion. This sentence, too, needs to be made more precise. I mean the conditional clause, 'if you intend one premise and believe the other', to be within the scope of 'you ought'. This will be my usage throughout this paper, even though it is not standard English. As I mean it, the sentence is not equivalent to: 'If you intend one premise and believe the other, you ought to intend the conclusion.' I mean to say, in symbols:

$$O_p((I_p(\text{Leslie will leave the next buoy to starboard}) \;\&\; B_p(\text{In order to leave the next buoy to starboard, Leslie must tack})) \supset I_p(\text{Leslie will tack})). \qquad (\gamma)$$

'O_p' stands for 'you ought'.

Till now, I have been sloppy about the owner of the belief and intentions; I simply took it to be you. But in formulating (γ), I needed to show that the subject of the 'ought' is the person who has the belief and intentions. I added the suffix '*p*' for that purpose. Informally, I shall continue to assume *p* is you.

In my particular example, I am also assuming you are Leslie, but I have not represented that assumption in the notation. This is because I believe you may intend a proposition whose subject is not you, and you may reason instrumentally with that intention. For instance, you may intend that your daughter gets a good education, and as a result you may intend that she goes to a good school. I have left that possibility open in my notation.

A useful piece of terminology allows me to rephrase (γ). When $O_p(X \supset Y)$, I say that X *normatively requires* Y of *p*. So your intention to leave the next buoy to starboard and your belief that to do so you must tack together normatively require you to intend to tack. This sentence means the same as (γ).

(γ) does not say you ought to be in any particular mental state. For example, it does not say you ought to intend to tack. Indeed, that might be false. Perhaps you ought to leave the next buoy to port, and consequently perhaps it is not the case that you ought to tack.

(γ) specifies a relation that ought to hold between your mental states. It does not directly concern your reasoning, which is a process of causal interaction amongst your mental states. However, through a process of reasoning, you can ensure that the relation between your mental states is as it ought to be; reasoning can bring you to satisfy a normative requirement such as (γ). I say a reasoning process is *correct* if it achieves this result. Correctness is a merit of reasoning. It is a derivative merit, derived from normative requirements that hold between mental states.

Reasoning is correct if it makes your mental states conform to normative requirements you are under. This means that my descriptions of correct reasoning processes, such as (α) and (β), are really too simple. You can conform to a normative requirement such as (γ) in two ways. You may enter the conclusion-state or alternatively you may leave one of the premise-states; you may give up one of your existing beliefs or intentions. So a correct reasoning process is not truly as linear as my descriptions imply. It sets out from the premise-states, but it may not end in your entering a new conclusion-state; it may end in your leaving a premise-state. A process of reasoning that has either of these results will be correct. Still, when the

premise-states are secure, a simple linear description such as (α) or (β) is a good approximation, and I shall continue to use descriptions of this sort.

2. Normative practical reasoning.

That concludes my preliminary example of practical reasoning. I have already said it exemplifies a very special type of practical reasoning: instrumental reasoning to a means you believe is necessary. One feature of this type is that its content – the content of the belief and intentions that participate in the reasoning – includes no normative propositions.

Nevertheless, it seems plausible that there is such as thing as correct practical reasoning that does have a normative content. We often have normative beliefs; we may believe we ought to tread carefully, for example. Beliefs like this must surely participate in our practical reasoning. We often deliberate in order to arrive at a normative belief about what to do, and the point of our deliberation is ultimately to bring us to a decision – the forming of an intention. If a normative belief cannot lead to an intention, this sort of deliberation cannot achieve its purpose, and it is plausible that the way a normative belief leads to an intention is by reasoning.

By 'normative practical reasoning' I mean reasoning that has a normative belief ineliminably amongst its premise-states, and concludes in an intention. The purpose of this paper is to begin an investigation into normative practical reasoning. It is plausible there is such a thing. But is there really?

I shall tentatively propose the answer 'yes'. I suggest that

I ought to tack
so I shall tack

is an example of correct normative practical reasoning. Here the first sentence expresses a normative belief and the second an intention. Commonly when you reason like this, you have arrived at your belief that you ought to tack by a previous process of reasoning. This previous process concludes in a belief: the normative belief that you ought to tack. So it is not practical reasoning by my definition, and it is not my concern in this paper. I am concerned only with the reasoning that occurs once you have the normative belief.

A description of your reasoning is:

B_pO_I(Leslie will tack))
As a result, I_p(Leslie will tack).

The subscripts on the 'B' and the 'I' in this description show that the belief and the intention belong to the same person *p*. As in (γ), my notation does not imply that the subject of the embedded proposition 'Leslie will tack' is also the owner of the belief and intention. This is because I believe you may intend a proposition whose subject is not you; I gave an example in section 1. Moreover, I believe 'you ought' can govern a proposition whose subject is not you. Grammatically, it cannot; grammatically, 'you ought' must govern an infinitive, and the implicit subject of the infinitive must be you. But I believe this is only a quirk of grammar, and my notation leaves other possibilities open. In the next few paragraphs, I shall be forced to use the ungrammatical expression 'ought that', since otherwise I could not make my point accurately. I feel justified in doing so, since only a quirk of grammar disallows it.

A significant feature of the description is that'O' has the subscript '*I*', which stands for the pronoun 'I'. Normative practical reasoning has a normative belief amongst its premise-states. If it is to be correct reasoning, this belief must be self-ascribing.[4] You cannot correctly derive an intention from a normative belief unless you believe that you yourself are the subject of the belief's normative content. Suppose you are not Leslie but you believe Leslie ought that

Leslie will tack. (You believe that Leslie ought to tack, that is to say.) From your belief, you cannot correctly derive an intention that Leslie will tack. It may have nothing to do with you. Perhaps you are a rival in the race and will be delighted if Leslie does not tack.

Now suppose you are Leslie but do not know it. (Some fantastic story could be told about how you find yourself in this state.) Suppose you believe Leslie ought that Leslie will tack. Your belief has the content that you would express as 'Leslie ought that Leslie will tack'. This is exactly the same as the content of the belief I attributed to you in the case where you are not Leslie. So you still cannot correctly derive from it an intention that Leslie will tack. Only if your belief has the content that you would express as 'I ought that Leslie will tack', can you correctly derive an intention that Leslie will tack. In this case, we can report your belief using the reflexive pronoun 'you*': you believe that you* ought that Leslie will tack.[5]

In general, this type of normative practical reasoning takes the form

$$B_pO_IA$$

As a result, I_pA. (δ)

where A is some proposition. There may be other types of normative practical reasoning, but I shall concentrate on this simple type. I claim tentatively that (δ) describes correct reasoning. I shall call this *the reasoning claim*.

Please understand that I make this claim in a spirit of exploration. My positive case for it is slender. In section 1, I was able to ground my example of instrumental reasoning on logical validity, which strongly supports the case for its being genuine reasoning. I can do nothing so effective for (δ); I can show only that it has some of the characteristics of reasoning. This paper mainly has the negative purpose of responding to some of the objections that can be raised against the reasoning claim, and I do not even have a definitive response to one of those.

My reasoning claim has two components. First, I claim that, if you believe you ought to tack, your belief normatively requires you to intend to tack. You ought, if you believe you ought to tack, to intend to tack. In general:

$$O_p(B_pO_IA \supset I_pA)$$

I call this *the requirement claim*.

The requirement claim is not that, if you believe you ought to tack, you ought to intend to tack. That may be false. Suppose your belief is false; it is not case that you ought to tack. Then it might well not be the case that you ought to intend to tack. In general, the requirement claim is not

$$B_pO_IA \supset O_pI_pA.$$

The reasoning claim is the conjunction of the requirement claim and the further claim that the process set out in (δ) is reasoning. Some processes that bring you to satisfy a normative requirement are not reasoning. Hearing a loud noise normatively requires you to believe there is a loud noise, but I do not think the process that takes you from hearing a loud noise to believing there is a loud noise is reasoning. However, I claim that (δ) is reasoning.

What extra does the reasoning claim add to the requirement claim? Granted that a process brings you to satisfy a normative requirement, what other features must it have if it is to be correct reasoning?

First, reasoning is a process that occurs amongst mental states. It is sometimes thought these states can only be beliefs, but that is not so; the example in section 1 showed they may also be intentions. So (δ) satisfies this condition.

Second, the mental states have propositional contents. That is why the process that goes from hearing a loud noise to believing there is a loud noise is not reasoning; hearing a loud

noise does not have a propositional content.[6] However, (δ) does satisfy this condition.

Third, a correct process of reasoning (which is one that brings you to satisfy a normative requirement) is made correct – I shall say 'validated' – by the formal relations that hold between the mental states' contents. 'Formal' includes more than just 'logical'. My earlier examples (α) and (β) of reasoning are validated by the relation of logical validity, but not all reasoning is validated by logic. For example, the reasoning 'It's red, so it's not blue' is correct, but not validated by logic. If (δ) is indeed correct reasoning, it is not validated by logic.

I do not need to specify the whole meaning of 'formal'. All that matters for my purposes is that, if (δ) is correct reasoning, it is correct whatever proposition is substituted for '*A*'. Furthermore, its correctness cannot be affected by anything that is not within the contents of the mental states; for instance, it cannot be affected by the abilities of the reasoner. A doubt can be raised about whether (δ) satisfies this condition; this is part of the inability objection below.

Fourth, a correct reasoning process is in some way guided by its own correctness. The nature of this guidance is difficult to specify, and I shall not try to specify it in this paper. However, I see no reason why the process set out in (δ) should not satisfy this condition as well as any process of reasoning does. Once you believe you ought to tack, you might decide to do so, guided by the appropriateness of this decision.

This leads me to mention a complication. (δ) describes a process in which your normative belief leads you to a decision. This process might be guided by its own correctness. On the other hand, it might not be; your normative belief might cause you to make the appropriate decision, but in some deviant way that is not guided by the correctness of the process.[7] In that case, you are not reasoning correctly even if you go through the process (δ). My reasoning claim, more precisely stated, is that (δ) may be correct reasoning, not that it must be correct reasoning.

In the rest of this paper, I shall concentrate on simple cases where the beliefs and intentions in question are yours, and the proposition in question is the proposition that you ϕ, where ϕ holds the place for some verb. For example: you tack. This will allow me to avoid too much strain on grammar. Your reasoning in this special case is:

I ought to ϕ
so I shall ϕ (ϵ)

I claim this is correct reasoning.

3. The misdirection objection.

Your premise-state in (ϵ) is a belief that you ought to ϕ. It is not a belief that you ought to intend to ϕ. So how come it normatively requires you to intend to ϕ? Perhaps it might normatively require you to ϕ, but surely not to intend to ϕ. That seems like a misdirection of the normative requirement. This is a objection to the requirement claim.

Some unfortunate possibilities can reinforce the objection. Even if your belief that you ought to ϕ is true, it may not be the case that you ought to intend to ϕ, and it may even be the case that you ought not to intend to ϕ. Indeed, in some circumstances, intending to ϕ may actually prevent you from ϕing.

For example,[8] suppose that at noon tomorrow you will be offered a toxic drink. If you drink it, you will suffer severe pain for a while, but no permanent damage will ensue. Suppose that tonight you will be given a million pounds by an eccentric millionaire, if and only if tonight you intend to drink the toxin tomorrow. The prize will be awarded just for having the

intention, even if in the end you do not fulfil it. It is plausible that you ought not to drink the toxin when it is offered, since at that time nothing will be gained by doing so. If that is so and, tonight, you correctly believe you ought not to drink the toxin tomorrow, then according to the requirement claim, this belief normatively requires you to intend, tonight, not to drink the toxin tomorrow. Yet it is also plausible that you ought not to have this intention. Having it will prevent you from having the opposite intention of drinking the toxin, which will win you a million pounds.

For another example, suppose you ought to sleep, and you believe this. I claim you are normatively required to intend to sleep. But suppose that intending to sleep will actually prevent you from sleeping. (Assume you do not know this. Otherwise, your knowledge may prevent you from intending to sleep; this possibility gives rise to the inability objection, which is examined in the section 4.) In these circumstances, how can your belief normatively require you to intend to sleep?

The same objection may be raised against the instrumental reasoning I described in section 1. I mention this because that type of instrumental reasoning seems to me incontrovertibly correct. So, if the misdirection objection applies to it, this raises a doubt about the misdirection objection. Suppose you intend an end, and believe a particular means is necessary to this end. I claimed in section 1 that correct reasoning will lead you from this intention and belief to intending the means. Consequently, your intention and belief normatively require you to intend the means. But you believe the way to achieve your end is to take the means, not to intend the means, so why should you be normatively required to intend the means? Once again, this seems to be a misdirection of the normative requirement. You may not believe that intending the means will help you achieve the end. Indeed, in some circumstances, intending the means may actually prevent you from achieving the end. How, then, can you be normatively required to have this intention?

Suppose you intend to perform at your best in the exam tomorrow, and believe you will do so only if you sleep now. I claim that this intention and belief normatively require you to intend to sleep now. But suppose that intending to sleep now will actually prevent you from doing so. (Assume once again that you do not know this.) Can you really be normatively required to have this intention?

In each sleeping example, I claim that correct reasoning sets out from one or more premise-states and arrives at the conclusion-state of intending to sleep. The premise-states normatively require the conclusion-state. Yet in both examples, this conclusion-state defeats the premise-states in a particular way – a different way in each case. In the normative example, the premise-state is your belief that you ought to go to sleep, and the conclusion-state prevents you from doing this thing that you believe you ought to do. That is one sort of defeat. In the instrumental example, one of your premise-states is your intention to perform at your best, and the other is your belief that sleeping is a necessary means of doing so. The conclusion-state prevents you from taking what you believe is a necessary means to achieve your intended end. That is another sort of defeat. We would not expect the conclusion-state of a reasoning process to defeat the premise-states in this way. This is an aspect of the misdirection objection.

Response.

First, remember my claim is that the premise-states normatively require the conclusion-state. It is not that you ought to enter the conclusion-state. Plausibly, in each example, you ought not to enter the conclusion state. That is consistent with my claim.

Whether a piece of reasoning is correct is not determined by whether it puts you into a state you ought to be in, or even by whether it serves its premise-states in some sense. It is determined by the formal relations that hold amongst the propositions that constitute the contents of the premise-states and conclusion-state. For instance, in the example of section 1, these contents form a valid syllogism, and that is enough to make the reasoning correct. Reasoning is a process that takes place amongst mental states, but it is validated by the formal relations that hold amongst the contents of those states. Normative requirements amongst the states are determined by the demands of correct reasoning. If the contents stand in the relations that make the reasoning correct, the mental states are normatively required to stand in the corresponding relations.

Reasoning is validated by its contents, and the mental states have to go along willy-nilly. This explains why, in the reasoning set out in (ϵ), you end up with an intention to ϕ, and why this intention is normatively required by your belief that you ought to ϕ. The content of your intention is that you ϕ. That is precisely what you believe you ought to do, so the content of your intention stands in the appropriate relation to the content of your normative belief. At the level of content, there is no hint of misdirection. The mental state that has this content is an intention; this intention is therefore normatively required.

Separately, the mental states that constitute a piece of reasoning may have various other properties, and stand in various other relations. These properties and relations may cause trouble. For example, it may sometimes happen that the conclusion-state of a correct reasoning process is a state you ought not to be in, perhaps because it prevents you from winning a prize. It may also happen that the conclusion-state defeats the premise-states in some sense, as the sleeping examples illustrate. This is unfortunate when it happens, but it cannot affect the correctness of the reasoning process, because that is determined by the contents. Consequently, it cannot affect the normative requirements that hold among the mental states.

Reasoning holds a potential curse. Reasoning correctly, you say to yourself 'I ought to sleep, so I'll sleep'. Then you find yourself stuck with the intention of going to sleep. It may not be an intention you ought to have, and it may keep you awake. The content of your intention is the proposition 'I'll go to sleep', and going to sleep is exactly what you believe you ought to do. So there is nothing wrong at the level of the content. But tragically, you find yourself with the intention of going to sleep, which you ought not to have. Correct reasoning can be cursed, but it remains correct reasoning. It still determines normative requirements for your mental states.

That is my response to the misdirection objection. The objection is to the requirement claim, but my response comes from the direction of the reasoning claim as a whole. Only if (δ) is reasoning can I insist it is validated by its contents. The requirement claim is logically weaker than the reasoning claim, but it cannot be defended independently. You can only be normatively required to move from a normative belief to an intention if this movement is actually a process of reasoning. Otherwise, we could not explain why you are normatively required to end up with an intention, when you do not believe you ought to have this intention.

4. The inability objection.

It may happen that you cannot help believing you ought to ϕ, and at the same time cannot intend to ϕ. You cannot either intend to ϕ or not believe you ought to ϕ. But the requirement claim says you ought, if you believe you ought to ϕ, to intend to ϕ. Equivalently: you ought

either to intend to ϕ or not believe you ought to ϕ. So if ought implies can, the requirement claim must be false.

How might you be unable to intend to ϕ? When you intend to ϕ, a causal process is in train that you believe will result in your ϕing. (I do not mean you believe the process is in train. I mean you believe the conditional: If that process is in train, it will result in my ϕing.) Normally, much of this process takes place in your own mind, and it often depends on your making plans and decisions along the way. To form the intention to ϕ, you must set such a process in train. If there is no process you can set in train that you believe will result in your ϕing, you cannot form the intention. In that case, it is not within your ability to intend to ϕ.

Here are some examples. First, you may not be able to set in train a process that you believe will make you more self-confident, so you cannot intend to be more self-confident. However, you may also be unable to help believing that you ought to be more self-confident, perhaps because you have convincing evidence for this fact.

Second, on the basis of long experience you may believe you will not get up early tomorrow. You cannot intend to get up early, because you do not know how to set in train a process that you believe well get you up early. However, you may also be unable to help believing that you ought to get up early.

A third example.[9] In your pill-box are two pills. One cures headaches; the other tummy-aches. You have a headache, and you also know it is risky to take both pills at once. Consequently, you cannot help believing you ought to take the headache pill only, and not the other one. However, it is dark, and you cannot tell the pills apart. Consequently, you cannot form the intention of taking the headache pill only.

The inability objection is ostensibly directed at the requirement claim, but it is actually more effective against the reasoning claim. If we were troubled by the objection, we could easily adjust the requirement claim to accommodate it, by adding an extra condition. For example, we could say that believing you ought to ϕ requires you to intend to ϕ, provided you can intend to ϕ. But the reasoning claim cannot accommodate the objection so easily. Correct reasoning has to be validated by the formal relations among its contents. So any further condition would have to be included within the content of the reasoning. To accommodate the objection, we would have to take your reasoning to be something like:

I ought to tack
and I can intend to tack,
so I shall tack,

for example. But this cannot be correct reasoning. Correct reasoning must be validated by appropriate relations amongst its contents. The content of the second premise-state is ‘I can intend to tack’. This is a proposition about the conclusion-state itself; it is not appropriately related to the *content* of the conclusion-state. So the reasoning claim cannot accommodate the inability objection.

Response

The examples seem possible to me, so I can respond to this objection only by questioning the principle that ought implies can, on which it rests.

This principle has exceptions. For one thing, a psychological failing is often not enough to excuse you from an ought. Even if you cannot bring yourself to give up cream cakes, still it may be the case that you ought to give them up. Even if you cannot bring yourself to believe you are mortal, still you ought to believe it.

Take a case of theoretical reasoning. Suppose you believe a proposition *P* (for instance,

that all men are mortal) from which another proposition *Q* (for instance, that you are mortal) follows by an immediate and obvious inference. I claim that believing *P* normatively requires you to believe *Q*; you ought either to believe *Q* or not believe *P*. Now suppose you cannot help but believe *P*, because you have convincing evidence for it. But suppose that for some reason or other you cannot bring yourself to believe *Q*. So you cannot either believe *Q* or not believe *P*. Does this excuse you from the normative requirement? I would say not. Rationality demands that you believe the immediate and obvious consequences of your beliefs. If you are prevented from satisfying this demand by some failing in your psychology, that does not excuse you.

Can we take a similar line with normative practical reasoning? For some cases at least, I think so. In the getting-up example, I claim you ought either to intend to get up early or not believe you ought to get up early. Your inability to conform to this requirement stems from a psychological failing on your part, namely your inability to intend to get up early. Plausibly, this failing is not enough to excuse you from the requirement. So, plausibly, the requirement still stands.

In other cases your inability seems a better excuse, because it is not so easy to blame it on a psychological failing. In the pill example, it is not your fault that darkness prevents you from intending to take the headache pill only. Because of that, you seem to have a good excuse for not satisfying the normative requirement I claim you are under. It might therefore not really be a normative requirement after all.

Nevertheless, the pill example may succumb to a different strategy. The inability objection relies on the principle that ought implies can. The idea that underlies this principle is that 'you ought' can apply only to things you have control over. But this idea presumably licenses another principle too: that ought implies can intend. If you cannot intend to do something, you do not truly have control over your doing it, and so it cannot be the case that you ought to do it. That seems plausible; at any rate it should seem plausible to anyone who believes ought implies can. The strategy is to use an instance of 'ought implies can intend' against an instance of 'ought implies can'.

Suppose you find you cannot intend to take the headache pill only. According to the principle that ought implies can intend, it follows that it is not the case that you ought to take the headache pill only. Moreover, when you find you cannot intend to take the headache pill only, you acquire good evidence that it is not the case that you ought to take the headache pill only. Therefore, you ought not to believe you ought to take the headache pill only. If you nevertheless cannot help but believe this, that is a psychological failing on your part.

The requirement claim is that you ought either to intend to take the headache pill only, or else not believe you ought to take the headache pill only. But you are unable to do this. Your inability implies the claim is incorrect, provided ought implies can. However, we now have an argument that says your inability results from a psychological failing on your part. A psychological failing may not be enough to excuse you from an ought; ought may not imply can in this case. So the requirement claim may still be correct.

In so far as it works, this strategy for handling the pill example can be generalized to handle any example where you believe you ought to ϕ yet cannot intend to ϕ. So it could supply a general response to the inability objection.

However, I am not convinced this strategy does work, because I am not convinced that ought implies can intend. In the pill example, I am not convinced it is not the case that you ought to take the headache pill only. You have a headache, and it is risky to take both pills at once. You cannot distinguish the pills because it is dark. But until you have convinced

yourself that you cannot distinguish the pills, I think you ought at least to try and think of a way to distinguish them. Why ought you to do this? So far as I can see, only because it is a possible means towards something you ought to do. That something must be taking the headache pill only. So it must be the case that you ought to take the headache pill only. This argument casts doubt on my strategy for handling this example.

But then, a doubt about 'ought implies can intend' should also be a doubt about 'ought implies can', since the two principles reflect one general idea. And 'ought implies can' is the basis of the inability objection. So my discussion of this objection has reached an impasse. I have no conclusive argument either for it or against it. I think this objection is the most serious threat to the reasoning claim.

5. The Humean objection

David Hume said that 'reason alone can never be a motive to any action of the will.'[10] I assume that forming an intention is an action of the will. Given that, Hume's remark can be interpreted in various ways. On the strongest interpretation, it says that reasoning cannot conclude in an intention. This appears to have been Hume's own view. It means that practical reasoning as I defined it does not exist, not even instrumental reasoning. My example of instrumental reasoning contained in section 1 is a counterexample to this strong view, and in section 1 I identified a flaw in Hume's defence of it.

According to a weaker interpretation of Hume's remark, reasoning cannot generate an intention out of anything apart from a prior intention, so if the conclusion-state of a piece of reasoning is an intention, the premise-states must include an intention. This makes room for instrumental reasoning. According to a still weaker interpretation, if the conclusion-state is an intention, the premise-states must include a 'motivated state', which may be a desire or an intention or perhaps something else, but not a belief. My putative pattern of reasoning (δ) violates even this condition, so even the weakest Humean view is an objection to my reasoning claim. I shall therefore concentrate on this one.

A friendly Humean need not flatly reject (δ) as reasoning. Instead she might say that a normative belief is covertly a motivated state. If so, (δ) would conform to the weakest Humean condition, because its premise-state would be a motivated state. If I felt the need, I should not much mind accepting this friendly offer. (It would have to be explained how normative beliefs, understood this way, participate in theoretical reasoning just like any other belief.[11]) But as yet I do not feel the need, since I know no good argument for even the weakest Humean view.

I cannot review every Humean argument, and in any case I have little to add to the many existing critiques of them.[12] I criticized an argument of Hume's in section 1, and here I shall criticize a prominent modern one.[13] This one concentrates on the explanation of intentional action. Suppose you tack. If the reasoning claim is correct, the explanation of your tacking might be: you believed you ought to tack, through correct reasoning this belief caused you to intend to tack, and this intention in turn caused you to tack. But Humeans think your tacking could not be explained this way (unless your belief is a covertly motivated state), so the reasoning claim cannot be correct.

They do not doubt that your intention to tack could cause you to tack; their objection is focussed on the first part of the explanation. They think your belief could not cause your intention through correct reasoning. In general they think that no genuine belief can cause an intention through correct reasoning; only a motivated state can do that.

They should not doubt that a genuine belief can cause an intention. Intentions can be

caused in all sorts of ways. You can wake up with a new intention, and you can get one by hypnosis or a knock on the head. It is easy to construct a story in which a genuine belief causes an intention. Moreover, it is easy to construct a story in which this happens through a mental process. Here is one. In your long-forgotten childhood, you had a trying encounter with a dog named 'Wilhelm'. You now acquire the belief that the dog behind you is named 'Wilhelm'. This belief causes you to intend to leave its presence.

So Humeans should recognize that a genuine belief can cause an intention through a mental process. Their argument can only be that any mental process through which this happens cannot be correct reasoning. In particular, (δ) cannot be correct reasoning. But they still need an argument to show why not, and this brings them back to where they started. Their worry about explaining action is a diversion; it leads to no new argument that (δ) is not correct reasoning.

True, I have suppressed a part of the explanation in the dog example. Your belief causes your intention, but only because you have a particular disposition: you are disposed to intend to leave the presence of any dog you believe is named 'Wilhelm'. This disposition might be called a motivated state. Similarly, your normative belief that you ought to tack causes you to intend to tack only because you have a particular disposition: you are disposed to intend to ϕ whenever you believe you ought to ϕ. This disposition might be called a motivated state. But it is obviously not a premise-state in your reasoning. According to the reasoning claim, it is simply a disposition to reason correctly, so it does not count against the reasoning claim at all.

6. Conclusion

The reasoning claim is neither proved nor disproved. It is more threatened by the inability objection than by the other objections.

Notes

I am grateful to Jonathan Dancy, Brad Hooker and Derek Parfit, Mozaffar Qizilbash, Joseph Raz and Thomas Scanlon for very useful comments. Also to audiences at Rutgers, MIT, Berkeley and the Joint Session at York. This is a slightly improved version of the paper that appears in the *Proceedings of the Aristotelian Society*, Supplementary Volume 75 (2001), pp. 175–93.

1. John Broome, 'Practical reasoning', in *Reason in Nature: New Essays in the Theory of Rationality*, edited by José Bermùdez and Alan Millar, forthcoming.
2. David Hume, *A Treatise of Human Nature*, bk 3, pt 1, sect 1.
3. There is a fuller account in the paper mentioned in note 1.
4. This point was made to me by Christian Piller, commenting on this paper in his 'Normative practical reasoning', *Proceedings of the Aristotelian Society*, Supplementary Volume 75, pp. 195–216.
5. The content of your belief is *de se*, to use David Lewis's term in his 'Attitudes *de dicto* and *de se*', *Philosophical Review*, 88 (1979), pp. 513–43, reprinted in his *Philosophical Papers, Volume 1*, Oxford University Press, 1983 pp. 133–60. This means it cannot be a proposition, if a proposition is a set of possible worlds. In this paper I use 'proposition' more generally.
6. In *Mind and World* (Harvard University Press, 1994), John McDowell argues that perception has a conceptual content, which he takes to be propositional. If he is right, mine is a bad example.
7. There is an example in my paper 'Reason and motivation', *Proceedings of the Aristotelian Society*, Supplementary Volume 71 (1997), pp. 131–46.

8. This example comes from Gregory Kavka's 'The toxin puzzle', *Analysis*, 43 (1983), pp. 33–6.
9. This is a version of a problem presented to me by Gijs van Donselaar.
10. *Treatise*, bk 2, pt 3, sect 3.
11. This task is undertaken by (amongst many others) Allan Gibbard in *Wise Choices, Apt Feelings: A Theory of Normative Judgment*, Oxford University Press, 1990.
12. It seems invidious to mention just one of the many excellent critiques, but I shall mention the one that prompted my interest in the subject: Derek Parfit's 'Reasons and motivation', *Proceedings of the Aristotelian Society*, Supplementary Volume 71 (1997).
13. The leading source is Bernard Williams's 'Internal and external reasons', in his *Moral Luck*, Cambridge University Press, 1981, pp. 101–13.