

REASONS AND ABILITY

Bart Streumer

b.streumer@rug.nl

17 December 2014

For Daniel Star (ed.), *The Oxford Handbook of Reasons and Normativity*

Abstract:

I argue that there can only be a reason for a person to perform an action if this person can perform this action. I begin by giving three arguments for this claim. I then reject several alternatives to it, argue that four counterexamples to it fail, and argue that a similar claim is true of reasons for belief. I end by showing that these claims can help us to distinguish deontic from evaluative judgements.

REASONS AND ABILITY

Most philosophers think that ‘ought’ implies ‘can’: they think that

- (O) It can only be the case that a person ought to perform an action if this person can perform this action.

Many philosophers also think that there is a close connection between what we ought to do and what there is reason for us to do. This suggests that ‘reason’ may also imply ‘can’: it suggests that

- (R) There can only be a reason for a person to perform an action if this person can perform this action.

In this chapter, I will defend (R). I will begin by giving three arguments for (R). I will then reject several alternatives to (R), argue that four counterexamples to (R) fail, and argue that a similar claim is true of reasons for belief. I will end by showing that these claims can help us to distinguish deontic from evaluative judgements.

1. The interpretation of (R)

There are different uses of the term ‘reason’. I will use this term to mean normative reason. I will take normative reasons to be considerations that count in favour of, or count against, having an attitude or performing an action. We should take these considerations into account in rational deliberation.

There are also different uses of the term ‘can’. When we say that there can only be a reason for a person to perform an action if this person can perform this action, this could mean that

- (R1) There can only be a reason for a person to perform an action if this person

would perform this action if he or she tried to perform it.

But it could also mean that

- (R2) There can only be a reason for a person to perform an action if there is a historically and nomologically accessible possible world in which this person performs this action,

where a possible world is historically and nomologically accessible if and only if it has the same past and laws of nature as the actual world. Or our claim could be weaker, and mean merely that

- (R3) There can only be a reason for a person to perform an action if there is a historically and nomologically accessible *or close* possible world in which this person performs this action,

where a possible world is historically and nomologically close if and only if its past and laws of nature are similar to those of the actual world.¹

There is an important difference between (R2) on the one hand and (R1) and (R3) on the other. If determinism is true, the only historically and nomologically accessible world is the actual world. This means that if (R2) is true, determinism implies that there can only be reasons for us to perform actions that we are in fact going to perform. That is hard to believe. By contrast, (R1) and (R3) do not have this implication. Determinists will therefore find these claims easier to accept.

Since (R3) is a rather weak claim, I will not interpret (R) as (R3). But I will not choose between (R1) and (R2). My arguments for (R) will therefore be arguments for the disjunction of (R1) and (R2). Since this disjunction entails (R3), these arguments will also support (R3).

¹ In (R2) and (R3), I use 'this person' to mean 'this person or his or her counterpart'.

2. The argument from crazy reasons

I will now give three arguments for (R). The first is what we can call *the argument from crazy reasons*.² Suppose that Fred is a normal human being, and consider the following claims:

- (1) There is a reason for Fred to travel back in time to prevent the crusades, slavery and the two world wars.
- (2) There is a reason for Fred to travel to the other side of the world within thirty seconds to prevent an immanent plane crash.
- (3) There is a reason for Fred to develop medicines against all known diseases by the end of the week.

There clearly are no such crazy reasons. (R) explains why not: because Fred cannot perform the actions (1), (2) and (3) say there are reasons for him to perform. You may nevertheless reject (R). But if so, you will need to propose other restrictions on the existence of reasons in order to make it the case that there are no such crazy reasons. You could, for example, propose the following restrictions:

There can only be a reason for a person to perform an action if this person does not have to travel back in time in order to perform this action.

There can only be a reason for a person to perform an action if this person is physically capable of performing this action.

There can only be a reason for a person to perform an action if this person is intellectually capable of performing this action.

But you will then face two problems. The first is that I could keep making further claims about crazy reasons. You will then have to keep proposing further restrictions on the

² I gave earlier versions of this argument and my next two arguments in Streumer 2007.

existence of reasons, without there being an obvious stopping point. And (R) is clearly a simpler and less *ad hoc* explanation of the non-existence of crazy reasons than your ever-expanding list of restrictions. The second problem is that for each restriction you propose, I can ask:

“*Why* can there only be a reason for a person to perform an action if this person does not have to travel back in time in order to perform this action?”

“*Why* can there only be a reason for a person to perform an action if this person is physically capable of performing this action?”

“*Why* can there only be a reason for a person to perform an action if this person is intellectually capable of performing this action?”

To which the obvious answers seem to be:

“Because a person *cannot* travel back in time.”

“Because a person *cannot* perform an action that he or she is physically incapable of performing.”

“Because a person *cannot* perform an action that he or she is intellectually incapable of performing.”

What explains and unifies your ever-expanding list of restrictions therefore seems to be (R). This suggests that instead of endorsing this ever-expanding list, you should simply endorse (R).

Some opponents of (R) deny that claims like (1), (2) and (3) are false. For example, Roger Crisp thinks that there can be reasons to perform “actions which cannot in an ordinary sense be performed, such as my flying several metres through the air to avoid a charging tiger”.³ Crisp defends this view in two ways. First, he suggests that we think that ‘ought’ implies ‘can’ because of “the close relation between ‘ought’ and emotional responses such as

³ Crisp 2006, p. 43.

guilt or blame”, and he claims that “the notion of a reason is quite comprehensible independently of the emotions”.⁴ But the argument from crazy reasons does not appeal to guilt or blame, and neither will my next two arguments for (R). Second, Crisp admits that his reason to fly several metres through the air “is of course largely irrelevant to deliberation”, but he claims that this does not matter, since there can be reasons of which we are unaware which are similarly irrelevant to deliberation.⁵ I will come back to this when I give my third argument for (R).

Other opponents of (R) agree that claims like (1), (2) and (3) *seem* false, but claim that this appearance is misleading. For example, Mark Schroeder defends a Humean theory of reasons, according to which

For R to be a reason for X to do A is for there to be some *p* such that X has a desire whose object is *p*, and the truth of R is part of what explains why X’s doing A promotes *p*.⁶

This theory entails that there can be some very crazy reasons. For example, it entails that it can be the case that

(4) There is a reason for you to eat your car,

as long as there is some *p* such that you have a desire whose object is *p*, and the fact that is this reason is part of what explains why your eating your car promotes *p*.⁷ I am told that some people can eat cars.⁸ But Schroeder’s theory implies that (4) can be true of you even if you are not one of these people.

⁴ Crisp 2006, pp. 42-3.

⁵ Crisp 2006, pp. 43, 39.

⁶ Schroeder 2007, p. 59.

⁷ Schroeder 2007, pp. 92-7.

⁸ The French entertainer Michel Lotito apparently ate a Cessna 150 airplane by cutting it up into very small pieces. Presumably he could have eaten a car as well.

Schroeder admits that (4) seems false, but he suggests that this claim seems false merely because it conversationally implicates that your reason to eat your car is weighty. He thinks we can cancel this implicature by saying what your reason for you to eat your car is, and by adding that this reason has very little weight. For example, we can say that

(4*) The fact that your car contains your recommended daily intake of iron is a reason for you to eat your car, but this reason has very little weight.

If so, we could make similar claims about (1), (2) and (3). For example, we could cancel (1)'s implicature that Fred's reason to travel back in time is weighty by saying that

(1*) The fact that the crusades, slavery and the two world wars caused an enormous amount of suffering is a reason for Fred to travel back in time to prevent these events, but this reason has very little weight.

I agree that (1*) sounds less indefensible than (1). But the difference between (1) and (1*) seems to me too small to make (1*) defensible. Schroeder may think that he needs to accept the existence of certain crazy reasons because their existence is implied by his theory. But he could easily prevent his theory from having this implication by taking A to range only over the actions that X can perform.⁹ I therefore continue to think that (1), (2) and (3) are false.

3. The argument from tables and chairs

A second argument for (R) is what we can call *the argument from tables and chairs*. There are no reasons for tables or chairs to perform any action whatsoever, since inanimate objects cannot perform actions. When a person cannot perform an action, this person is in the same

⁹ If Schroeder took A to range only over the actions that X can perform, however, this would not fully solve what he calls the 'too many reasons problem' (2007, pp. 84-7). He would still need to solve this problem by saying that the claim that there is a reason for X to do A conversationally implicates that this reason is weighty.

position with regard to *this* action that a table or a chair is in with regard to *all* actions. This suggests that just as there are no reasons for a table or a chair to perform any action whatsoever, there is no reason for this person to perform this action. In other words, it suggests that (R) is true.

If you reject (R), you could offer a different explanation of why there are no reasons for tables or chairs to perform any action whatsoever. What explains this, you could say, is not that

- (1) There can only be a reason for an entity to perform an action if *this entity* can perform *this action*,

but is instead that

- (2) There can only be a reason for an entity to perform an action if *entities of this kind* can perform *actions*.¹⁰

And you could then say that the argument does not show that (R) is true, since (2) does not entail (R).

But suppose that Kate cannot perform any action whatsoever. Since persons are entities of a kind that can perform actions, if (2) is true but (1) is false, there can then be reasons for Kate to perform very many actions. But Kate herself cannot perform any of these actions. What makes it true that persons are entities of a kind that can perform actions is only that other people can perform these actions. And why would whether there are reasons *for Kate* to perform these actions depend on whether *other people* can perform them? It is clearly more defensible to say that it depends on whether Kate *herself* can perform these actions. It is therefore clearly more defensible to endorse (1).

If you reject (R), you could also say that what explains why there are no reasons for tables or chairs to perform any action whatsoever is that

¹⁰ I use the term 'entity' to cover both inanimate objects and persons.

- (3) There can only be a reason for an entity to perform an action if this entity can perform *actions*, regardless of whether this entity can perform *this* action.

As before, you could then say that the argument does not show that (R) is true, since (3) does not entail (R).

Suppose again that Kate cannot perform any action whatsoever. In that case, (1) and (3) both entail that there is no reason for her to perform any action whatsoever. But suppose next that there is a very small change in Kate's situation, such that there are now exactly two actions that she can perform. If (3) is true but (1) is false, there can then all of a sudden be reasons for Kate to perform very many actions, all but two of which Kate cannot perform. Why would this very small change in Kate's situation lead to this enormous change in the number of reasons there can be for her? It is clearly more defensible to say that whether there is a reason for Kate to perform an action depends on whether she can perform *this* action. As before, it is therefore clearly more defensible to endorse (1).

Finally, if you reject (R), you could also say that what explains why there are no reasons for tables or chairs to perform any action whatsoever is that

- (4) There can only be a reason for an entity to perform an action if this entity can respond to reasons.

Once again, you could then say that the argument does not show that (R) is true, since (4) does not entail (R).

We respond to a reason when we perform an action that there is a reason for us to perform because we believe that there is this reason for us to perform it. What (4) adds to (3) is therefore that

- (5) There can only be a reason for an entity to perform an action if this entity can have beliefs about reasons.

But suppose that Kate is irreversibly paralyzed but can still have beliefs about reasons. If (5) is true but (1) is false, there can then be reasons for her to perform very many actions. If Bob asks Kate to pick him up from the train station, there can be a reason for her to pick him up,

even though she cannot do this. If Susan has had an accident outside Kate's house, there can be a reason for Kate to go outside and help Susan, even though she cannot do this. Why would whether there can be reasons for Kate to perform these actions depend merely on whether she can have beliefs about reasons? It is clearly more defensible to say that it also depends on whether Kate can actually perform these actions. Once again, it is therefore clearly more defensible to endorse (1).

I therefore conclude that (1) is the best explanation of why there are no reasons for tables or chairs to perform any action whatsoever. Since (1) entails (R), this explanation supports (R).¹¹

4. The argument from deliberation

A third argument for (R) is what we can call *the argument from deliberation*. Suppose that (R) is false. In that case, when you engage in rational deliberation about what to do, you need to take into account not only reasons to perform actions that you *can* perform, but also reasons to perform actions that you *cannot* perform. Since you know that the crusades, slavery and the two world wars caused an enormous amount of suffering, you would then almost always have to conclude that there is most reason for you to travel back in time and prevent the crusades, slavery and the two world wars. And you would then have to try to travel back in time to prevent the crusades, slavery and the two world wars. This shows that if (R) were false, rational deliberation would almost always result in your trying to perform actions that you cannot perform. But rational deliberation should not have such pointless results. This suggests that (R) is true.

You may think that this argument assumes the truth of consequentialism. But it does

¹¹ According to Ulrike Heuer, the argument from tables and chairs is ultimately “just the argument [from crazy reasons] under a different guise” (2010, p. 240). But there is a clear difference between these arguments: whereas the argument from crazy reasons assumes that there are no reasons of a certain kind and shows that (R) is the best explanation of the non-existence of these reasons, the argument from tables and chairs assumes that there is a difference between agents and non-agents and shows that (R) is the best explanation of this difference.

not. The argument merely assumes that rational deliberation should take into account the amount of suffering that an action would prevent, and that the larger this amount of suffering is, the more likely it should be that rational deliberation results in your trying to perform this action. That is true according to any defensible moral view.

If you reject (R), you could propose a different restriction on rational deliberation that prevents it from having such pointless results. You could, for example, propose one of the following restrictions:

- (1) When judging which reasons there are for or against performing an action, a person should only take into account reasons to perform actions that he or she can perform.
- (2) When judging which action there is most reason for him or her to perform, a person should only take into account reasons to perform actions that he or she can perform.
- (3) A person should only try to perform actions that he or she can perform.

But if (1), (2), or (3) are true, the only reasons that can make a difference to the result of rational deliberation are reasons to perform actions that a person can perform.¹² In that case, why should we think that there are also reasons to perform actions that a person *cannot* perform? Part of what makes a consideration a reason for a person is that this person should take this consideration into account in rational deliberation. If a consideration cannot make a difference to the result of rational deliberation, therefore, this consideration does not seem to be a reason at all.

As we have seen, Roger Crisp denies this. He admits that reasons to perform actions that we cannot perform are “largely irrelevant to deliberation”, but he claims that this does

¹² Alternatively, you may make the more general claim that the less likely it is that a person will successfully perform an action, the less likely it should be that this person’s rational deliberation results in his or her trying to perform this action. Like (1), (2) and (3), this claim also entails that the only reasons that can make a difference to the result of rational deliberation are reasons to perform actions that a person can perform.

not matter, since there can be reasons of which we are unaware which are similarly irrelevant to deliberation.¹³ But reasons of which we are unaware merely *do not* make a difference to the result of rational deliberation. These reasons would make a difference to this result if we were aware of them. We therefore should not deny that they exist. By contrast, if (1), (2) or (3) is true, reasons to perform actions that we cannot perform *cannot* make a difference to the result of rational deliberation. That suggests that these reasons do not exist.

John Gardner also denies this. According to Gardner, if you see a man drowning in a stormy sea, there is a reason for you to rescue this man even if you cannot rescue him. But Gardner endorses a variant of (3) by saying that there is then *no* reason for you to *try* to rescue this man.¹⁴ Why should we think that there is a reason for you to actually rescue this man, even though there is no reason for you to try to rescue him? According to Gardner, we should think this because this reason explains why you should feel “horrified at the realisation that it would be so utterly futile” to try to rescue this man, and explains why you should not “walk past without compunction”.¹⁵ But if that is what this reason explains, it seems to be a reason to feel horrified and to refrain from walking past without compunction rather than a reason to actually rescue this man.

Of course, rational deliberation should sometimes result in a person’s trying to perform an action that this person cannot perform: namely, if this person does not know that he or she cannot perform this action and has the justified belief that there is most reason for him or her to perform this action.¹⁶ For example, if you do not know that you cannot rescue this drowning man and you have the justified belief that there is most reason for you to rescue him, then rational deliberation should result in your trying to rescue him. But this is

¹³ Crisp 2006, pp. 43, 39.

¹⁴ Gardner 2004, p. 55.

¹⁵ Gardner 2004, pp. 55-6.

¹⁶ Kearns and Star 2009, p. 236, give an example of this kind. They suggest that in such cases there is a reason for you to perform the action that you cannot perform, since when you realise afterwards that you could not perform this action it still seems natural to say that there was a reason for you to perform it. I doubt this. It seems more natural to say that you were right to *think* that there was a reason for you to perform this action, but that in fact there was no such reason because you could not perform this action.

compatible with (R): (R) allows that if you do not know that you cannot rescue this man, you can be justified in believing that there is most reason for you to rescue him. By contrast, if (R) were false, rational deliberation would not merely sometimes but *almost always* result in your trying to perform actions that you cannot perform.

We can next note that, in addition to (R), it also seems true that

(RA) There can only be a reason *against* a person's performing an action if this person can perform this action.

For example, consider the action of producing a round square. Since inconsistent propositions entail everything, the production of a round square entails the destruction of the world. If (RA) was false, there would therefore be an extremely weighty reason against your producing a round square.¹⁷ But there clearly is no such reason. For just as reasons to perform actions that you cannot perform cannot make a difference to the result of rational deliberation, neither can reasons against performing actions that you cannot perform. The relation between reasons and rational deliberation therefore not only supports (R), but also supports (RA).

5. Alternatives to (R)?

Some philosophers reject (R) because they endorse an alternative to (R) that they take to be more defensible. For example, you may reject (R) because you think that

(1) There can only be a reason for a person to perform an action if this person *believes* that he or she can perform this action.

But suppose that Bob has the deluded belief that he is Napoleon. If he thinks he can win the Battle of Waterloo tomorrow, can there be a reason for him to win this battle? Of course not. Bob may believe that there is a reason for him to win this battle, but that does not mean there

¹⁷ I owe this point to Roy Sorensen.

actually is a reason for him to do this. This shows that (1) is not a defensible alternative to (R).

You may also reject (R) because you think that

- (2) There can only be a reason for a person to perform an action if this person has the *justified* belief that he or she can perform this action.

Suppose again that Bob believes that he is Napoleon. But now suppose that his belief is justified, since he has unwittingly taken a powerful drug that makes his surroundings look like a muddy field in Belgium in 1815 and that makes the people around him look like members of the Imperial Guard. If Bob now thinks he can win the Battle of Waterloo tomorrow, can there be a reason for him to win this battle?¹⁸ You may think there can be. But suppose that Bob asks you what there is reason for him to do tomorrow. You clearly should not answer: “There is a reason for you to win the Battle of Waterloo.” Instead, you should tell him that there is a reason for him to seek urgent medical help. This shows that (3) is not a defensible alternative to (R) either.¹⁹

Finally, you may reject (R) because you think that

- (3) There can only be a reason for a person to perform an action if this person has *sufficient evidence* that he or she can perform this action.

¹⁸ You may object that since Bob is hallucinating, when he uses the phrase ‘the Battle of Waterloo’ he does not refer to the actual Battle of Waterloo. But this does not matter. For the purpose of the example, we can stipulate that we are using the phrase ‘the Battle of Waterloo’ to refer to whatever it refers to when Bob uses it.

¹⁹ You may think that you should not answer “There is a reason for you to win the Battle of Waterloo” because this conversationally implicates that this is Bob’s strongest reason, whereas in fact there is a much stronger reason for him to seek urgent medical help. But it would clearly also be strange to answer: “There is some reason for you to win the Battle of Waterloo, but there is more reason for you to seek urgent medical help.”

But (3) clearly faces the same problem as (2).²⁰

More generally, we should ask the following question: why is whether a person believes, has the justified belief or has sufficient evidence that he or she can perform an action relevant to whether there is a reason for this person to perform this action? The answer seems to be: because whether this person can *actually* perform this action is relevant to whether there is a reason for him or her to perform this action. That suggests that instead of endorsing (1), (2) or (3), we should endorse (R).

6. Counterexamples to (R)?

Other philosophers reject (R) because they think there are counterexamples to (R).²¹ Two of these examples have been given by Ulrike Heuer. In Heuer's first example, someone cannot perform an action at time t_1 , but can perform this action at a later time t_2 after gaining a new ability. Heuer writes:

I cannot play the piano, say. Is there, therefore, no reason for me to do it? If there was no reason to play the piano for someone who can't play it already, there would presumably be no reason to learn to play it either. Reasons for learning something require that there is a reason for doing what (as yet) one cannot do.²²

²⁰ If you follow Kearns and Star 2009 in thinking that reasons for a person to perform an action are evidence that this person ought to perform this action, you may say that there can be a reason for Bob to win the Battle of Waterloo, since if he is sufficiently deluded there can be (misleading) evidence that he ought to win this battle even though he cannot actually win this battle. But this would commit you to telling Bob that there is a reason for him to win the Battle of Waterloo, which strikes me as implausible.

²¹ If these examples were counterexamples to (R), they would also be counterexamples to the claim that 'ought' implies 'can'. For exhaustive and compelling answers to purported counterexamples and other objections to the claim that 'ought' implies 'can', see Vranas 2007.

²² Heuer 2010, p. 237. Heuer draws a general distinction between what she calls 'derivative reasons', which exist only because of the existence of other reasons, and what she calls 'non-derivative reasons' (p. 236). She claims that her reason to learn to play the piano is a derivative reason. I discuss Heuer's arguments in Streumer 2010.

But this example merely shows that (R) should be indexed to time.²³ If we take 't₁' and 't₂' to refer to any two subsequent points in time, we should take (R) to say that

(R*) There can only be a reason at t₁ for a person to perform an action at t₂ if this person can at t₁ perform this action at t₂.

What does this mean? If we interpret (R) as (R1), it means that

(R1*) There can only be a reason at t₁ for a person to perform an action at t₂ if this person would perform this action at t₂ if he or she tried to perform a certain series of actions between t₁ and t₂.

And if we interpret (R) as (R2), it means that

(R2*) There can only be a reason at t₁ for a person to perform an action at t₂ if there is at t₁ a historically and nomologically accessible possible world in which this person performs this action at t₂.

Suppose that it would take me two years to learn to play the piano. In that case, there can be a reason for me today to play the piano in two years' time, since I can start taking lessons today that will enable me to play the piano in two years' time. And this reason can give rise to a reason for me to start taking lessons today. If we take (R) to be indexed to time, this example therefore does not show that (R) is false.

In Heuer's second example, there is a reason at t₁ for a person to perform an action at t₂, but this person then makes it impossible for him- or herself to perform this action at t₂. She writes:

²³ It has often been pointed out that the claim that 'ought' implies 'can' should similarly be indexed to time. See Zimmerman 1996 and Vranas 2007, p. 171.

A person, call her Lilly, has a reason to attend a meeting which she dreads, but she can make it impossible that she will attend by, say, locking herself into a room and throwing the key away. Is it now true that she doesn't have a reason to go to the meeting? After all, she can't. If so, it would be unclear why she has a reason not to disable herself, or to overcome the self-inflicted obstacle once it exists.²⁴

Suppose that the meeting will start at 10.00 and finish at 10.30, and that Lilly locks herself in and throws away the key at 9.00. In that case, (R) says that after 9.00 it is no longer true that there is a reason for Lilly to attend the meeting. But if you reject (R), you cannot defensibly say that it *never* ceases to be true that there is a reason for Lilly to attend the meeting.²⁵ Instead, you will presumably say that this ceases to be true at 10.30, since after 10.30 the meeting has finished, which means that after 10.30 Lilly can only attend the meeting by changing the past. But since Lilly locked herself in and threw away the key at 9.00, it is *already true after 9.00* that she can only attend the meeting by changing the past. This suggests that the time at which it ceases to be true that there is a reason for Lilly to attend the meeting is 9.00 rather than 10.30, exactly as (R) says.

According to Heuer, if there is no longer a reason for Lilly to attend the meeting once she has made it impossible for herself to attend it, it is “unclear why she has a reason not to disable herself”.²⁶ But I think this is clear: if a consideration is a reason for a person to perform an action, this very same consideration is also a reason for this person not to make it impossible for him- or herself to perform this action. For example, suppose that the reason for Lilly to attend the meeting is the fact that an important decision will be taken at the meeting.

²⁴ Heuer 2010, p. 237. Such examples have often been put forward as counterexamples to the claim that ‘ought’ implies ‘can’. See, for example, White 1975, p. 149, and Sinnott-Armstrong 1984. For responses, see Zimmerman 1996, pp. 97-100, Haji 2002, pp. 47-9, Howard-Snyder 2006, pp. 235-6, and Vranas 2007, pp. 175-82.

²⁵ See Vranas 2007, pp. 176-7, and p. 201 n. 10. Of course, if you want to reject (R), you could dig in your heels and say that this never ceases to be true. But that strikes me as very implausible. Does it remain true hundreds of years from now, after Lilly has died, that there is a reason for her to attend the meeting?

²⁶ Heuer 2010, p. 237.

If so, this very same fact is also a reason for her not to make it impossible for herself to attend the meeting.²⁷

Moreover, (R) allows that after 9.00 the following claims are still true:

There *was* a reason for Lilly to attend the meeting.

Because there was a reason for Lilly to attend the meeting, there was a reason for Lilly not to make it impossible for herself to attend the meeting.

Because there was a reason for Lilly to attend the meeting, she can be blamed for not having attended the meeting.

Because Lilly did not attend the meeting, there is a reason for her to apologise and to make sure that she attends the next meeting.

Since (R) allows that after 9:00 these claims are still true, and since any defensible view has to say that there is *some* time at which it ceases to be true that there is a reason for Lilly to attend the meeting, we should conclude that after 9.00 there is no longer a reason for her to attend the meeting. I therefore think that this example does not show that (R) is false either.

So-called Frankfurt cases can also be regarded as counterexamples to (R).²⁸ Suppose that Black is a powerful neurosurgeon who has implanted a device in Susan's brain that enables him to manipulate her decisions, without Susan being aware of this. Black wants Susan to kill Fred. But Black can accurately predict what Susan will decide to do, and he predicts that Susan will kill Fred all by herself, without any intervention from Black being necessary. This is in fact what happens. Since Black did not intervene, Susan seems blameworthy for having killed Fred. But she could not have avoided killing Fred, since if Black had predicted that Susan was not going to kill Fred all by herself, he would have used

²⁷ Heuer also writes that if Lilly no longer has a reason to attend the meeting once she has made it impossible for herself to attend it, it is "unclear why she has a reason . . . to overcome the self-inflicted obstacle once it exists." But this example can only be a counterexample to (R) if Lilly cannot overcome this self-inflicted obstacle. If (R) is true, there is therefore no reason for Lilly to overcome this obstacle.

²⁸ Examples of this kind were first given by Frankfurt 1969. For a brief overview of the extensive literature about them, see Fischer 2011.

the device in her brain to make her decide to kill Fred. This seems to show that

- (A) A person can be blameworthy for having performed an action even if this person could not have avoided performing this action.²⁹

And it is often thought that

- (B) A person can be blameworthy for having performed an action only if there was a reason for this person not to perform this action.

If (B) is true and if Susan is blameworthy for having killed Fred, this implies that there was a reason for Susan not to kill Fred. But if (R) is true, there was no reason for Susan not to kill Fred, since she could not have avoided killing him. This suggests that Frankfurt cases are counterexamples to (R).³⁰

But instead of rejecting (R), I think we should reject (A) or (B). We can reject (B) by endorsing a different claim about blame instead, such as the claim that

- (B*) A person can only be blameworthy for having performed an action if this person *should have believed* that there was a reason for him or her not to perform this action.³¹

Since Susan was unaware of Black's ability to manipulate her decisions, she should have believed that there was a reason for her not to kill Fred. (B*) therefore allows that Susan is blameworthy for having killed Fred.

²⁹ This is the negation of what Copp 1997 calls the 'blameworthiness reading' of Frankfurt's principle of alternate possibilities.

³⁰ Widerker 1991 and Copp 1997, p. 445, show that (A), the claim that 'ought' implies 'can', and a claim that is analogous to (B) are an inconsistent triad. The same is true of (A), (B) and (R). For discussion, see Haji 2002, pp. 36-58.

³¹ Haji 1998 proposes a view along these lines. Copp 1997, p. 448, criticises it.

Alternatively, we can also reject (A). Since Susan would not have been blameworthy for having killed Fred if Black had made her decide to kill Fred, we can say that what Susan is really blameworthy for is not simply *having killed Fred*, but is instead *having killed Fred all by herself*.³² If that is so, Frankfurt cases do not show that (A) is false. For Susan *could* have avoided killing Fred all by herself, by not deciding all by herself to kill Fred and thereby forcing Black to use the device in her brain to make her decide to kill Fred. You may object that it makes no sense to say that Susan is responsible for having killed Fred all by herself without thereby being responsible for having killed Fred.³³ If ‘responsible’ means causally responsible, that is true, since we cannot cause X in a particular way without thereby causing X. But (A) is a claim about *blameworthiness*, and we *can* be blameworthy for having done X in a particular way without thereby being blameworthy for having done X. For example, suppose that I promise to read your book closely, but that I then break my promise and only read your book superficially. In that case, I am not blameworthy for having read your book, but I am blameworthy for having read your book superficially. Similarly, Susan may not be blameworthy for having killed Fred, but she may be blameworthy for having killed Fred all by herself. I therefore think that Frankfurt examples do not show that (R) is false either.

Finally, R. M. Hare gives an apparent counterexample to the claim that ‘ought’ implies ‘can’ that may also seem to be a counterexample to (R). Suppose that my friend Bob has fallen ill and is in hospital. I could then say:

I ought to visit Bob in hospital, but unfortunately I cannot.

Similarly, I could say:

³² This point is made by Naylor 1984, though she puts it in terms of responsibility rather than in terms of blameworthiness. Copp 1997, p. 444, similarly denies that the person in a Frankfurt example is blameworthy for the action he or she performs, and adds that this person “may deserve blame for something else”, such as “permitting herself to be the kind of person” who would perform an action of this kind. For critical discussion of such views, see Fischer 2011.

³³ See Kane 2005, p. 85, and Kane 1998, pp. 41-2.

The fact that Bob is my friend is a reason for me to visit him in hospital, but unfortunately I cannot.

Does this show that (R) is false? Hare argues that it does not show that 'ought' does not imply 'can'. He writes:

When I say "I ought but I can't", I am prescribing in general for cases like mine; I certainly think that a man in my situation ought, *if he can*, to do the act in question; but the prescription fails to apply in my case because of the impossibility of acting on it. It is as if I said "If I were able, it would be the case that I ought (full force); but since I am not able, that lets me out".³⁴

Similarly, when I say that the fact that Bob is my friend is a reason for me to visit him in hospital, I am saying that this consideration is a reason for Bob's friends to visit him in hospital *if they can*. Since I cannot visit Bob in hospital, that "lets me out", as Hare puts it. You may object that this leaves it unclear why I should apologise to Bob for not visiting him in hospital. But if Bob and I both know that I cannot visit him in hospital, such an apology is mere politeness. A real apology is only required if I failed to do something I could have done, such as phoning Bob in hospital or sending him a get well card. I therefore think that this example also does not show that (R) is false.

6. Extending (R) to reasons for belief

If the arguments I have given show that (R) is true, similar arguments may show that

(RB) There can only be a reason for a person to have a belief if this person can have this belief.

³⁴ Hare 1963, p. 53.

For if (R) is the simplest and least *ad hoc* explanation of the non-existence of crazy reasons for action, (RB) is likewise the simplest and least *ad hoc* explanation of the non-existence of crazy reasons for belief, such as reasons for people with limited mathematical abilities to believe complex mathematical theorems. If a generalised version of (R) is the best explanation of the fact that there are no reasons for inanimate objects like tables and chairs to perform actions, a generalised version of (RB) is likewise the best explanation of the fact that there are no reasons for inanimate objects to have beliefs. And if (R) is the best explanation of the fact that rational practical deliberation should not result in a person pointlessly trying to do what he or she cannot do, (RB) is likewise the best explanation of the fact that rational epistemic deliberation should not result in a person pointlessly trying to believe what he or she cannot believe.

If you think that (RB) is false, this may be because you think that a reason for a belief is a consideration that is evidence for this belief, and that a consideration can be evidence for a belief regardless of whether anyone can have this belief. For example, you may think that a reason to believe a complex mathematical theorem is a consideration that is evidence for this theorem, and that this consideration can be evidence for this theorem regardless of whether anyone can believe this theorem. But I think that reasons for belief are not the same thing as evidence. Evidence for a belief often counts in favour of having this belief, I think, but not if we are not currently attending to this evidence and forming this belief would be a waste of our cognitive resources.³⁵ Similarly, if (RB) is true, evidence for a belief often counts in favour of this belief, but not if we cannot have this belief.

You may also think that (RB) is false because our beliefs are not under our voluntary control. But ‘can’ in (RB) need not mean ‘can merely by trying to do so’. In other words, we need not interpret (RB) as saying that

(RB1) There can only be a reason for a person to have a belief if this person would have this belief if he or she tried to have it.

³⁵ See Harman 1986, p. 12. Kelly 2003, p. 625, claims that when I “stumble upon strong evidence” for a trivial or useless belief, there is a reason for me to have this belief. I think this is plausible only because I am then attending to this evidence.

That is clearly false. Instead, we should interpret (RB) as saying that

(RB2) There can only be a reason for a person to have a belief if there is a historically and nomologically accessible possible world in which this person has this belief,

or perhaps as saying that

(RB3) There can only be a reason for a person to have a belief if this person would have this belief if he or she was aware of sufficient evidence for this belief.

You may object to (RB3) that someone could be so irrational that this person would not believe that *p* if he or she was aware of sufficient evidence for the belief that *p*, even though it may nevertheless be true that this person *can* believe that *p*. This objection may show that we should not interpret (RB) as (RB3). But in a similar way, someone could be so weak-willed that this person would not perform an action if he or she tried to perform this action, even though it may nevertheless be true that this person *can* perform this action. This may similarly show that we should not interpret (R) as saying that

(R1) There can only be a reason for a person to perform an action if this person would perform this action if he or she tried to perform it.

If this objection is sound, therefore, it does not drive a wedge between reasons for action and reasons for belief, but instead identifies a problem for conditional analyses of ‘can’ of the kind that (R1) and (RB3) incorporate.

7. Deontic and evaluative judgements

Many philosophers make a distinction between deontic and evaluative normative judgements. Deontic judgements include judgements about an action or attitude’s being obligatory, prohibited, right, wrong, someone’s duty, or such that someone ought or ought not to perform or have it. Evaluative judgements include judgements about something’s being good, bad, or

such that it ought or ought not to obtain.³⁶

Philosophers who think that ‘ought’ implies ‘can’ usually make analogous claims about other deontic judgements. For example, they also think that

A person can only have an obligation to perform an action if this person can perform this action.

A person can only be prohibited from performing an action if this person can refrain from performing this action.

A person can only have a duty to have an attitude if this person can have this attitude.

By contrast, most philosophers would not make such claims about evaluative judgements. For example, they do not think that

A state of affairs can only be good if someone can make this state of affairs obtain.

It can only be true that a state of affairs ought to obtain if someone can make this state of affairs obtain.

A character trait is bad only if its possessor can avoid having this character trait.

Why is this? If (R) and (RB) are true, it may more generally be true that

(RG) There can only be a reason for a person to perform an action or to have an attitude if this person can perform this action or have this attitude.

Now suppose that

(D) A judgement is deontic if and only if this judgement entails the claim that there is a reason for the person it is about to perform the action or have the

³⁶ See Tappolet 2013.

attitude it is about.³⁷

For example, the judgement that Bob has an obligation to keep his promise is deontic, since it entails the claim that there is a reason for Bob to keep his promise. But the judgement that it would be good if Bob kept his promise is evaluative, since it does not entail the claim that there is a reason for Bob to keep his promise.³⁸

If (D) is true, deontic judgements entail the claim that there is a reason for the person they are about to perform the action or have the attitude they are about. And if (RG) is true, there can only be such a reason if this person can perform this action or have this attitude. By contrast, evaluative judgements do not entail this. This may explain why claims that are analogous to the claim that ‘ought’ implies ‘can’ seem true of deontic judgements but not of evaluative judgements.

This explanation faces two problem cases. The first is that many deontic judgements, such as the judgement that keeping one’s promises is obligatory, are not about a specific person. But we can perhaps take such judgements to be about everyone *who can keep his or her promises*. These judgements could then be true even if no one could keep his or her promises, but they would then be vacuously true, in the sense that the set of people they are about is empty.

The second problem case are judgements about permissibility. Philosophers often take such judgements to be deontic, but (D) is not true of these judgements: the judgement that it is permissible for a person to perform an action or to have an attitude does not entail that there is a reason for this person to perform this action or to have this attitude. But perhaps we should say that judgements about permissibility are evaluative rather than deontic. Philosophers may take these judgements to be deontic because they take them to be

³⁷ This is compatible with the existence of false deontic judgements: deontic judgements entail the claim that there is a reason for the person it is about to perform the action or have the attitude it is about, but if the judgement is false this claim may be false as well.

³⁸ Of course, the judgement that it would be good if Bob kept his promise may conversationally implicate that there is a reason for him to keep his promise, especially if we express it in Bob’s presence. But unlike a relation of entailment, this conversational implicature can be cancelled.

equivalent to judgements about obligations: they take an action to be permissible if and only if not performing this action is not obligatory. But these judgements may not in fact be equivalent: whereas the first judgement ascribes a normative property to an object, the second judgement may merely say that the *absence* of an object *lacks* a normative property.³⁹ If so, judgements about permissibility are not interdefinable with judgements about obligations.

8. Conclusion

I conclude that there can only be a reason for a person to perform an action if this person can perform this action. I also conclude, more tentatively, that there can only be a reason for a person to have a belief if this person can have this belief. And I also conclude, even more tentatively, that the truth of these claims can help us to distinguish deontic from evaluative normative judgements.⁴⁰

References

- Copp, David. 1997. "Defending the Principle of Alternate Possibilities: Blameworthiness and Moral Responsibility". *Noûs* 31: 441-456.
- Crisp, Roger. 2006. *Reasons and the Good*. Oxford: Oxford University Press.
- Fischer, John Martin. 2011. "Frankfurt-Type Examples and Semicompatibilism: New Work". In Robert Kane, ed., *The Oxford Handbook of Free Will, Second Edition*. Oxford: Oxford University Press.
- Frankfurt, Harry. 1969. "Alternate Possibilities and Moral Responsibility". *Journal of Philosophy* 66: 829-839.

³⁹ In that case, the judgement that not performing an action is not obligatory conversationally implicates that this action is permissible. I say more about this in Streumer forthcoming.

⁴⁰ For helpful comments on earlier versions of this chapter, I am grateful to Samuel Asarnow, Alex Gregory, the participants in my Spring 2013 graduate class at the University of Reading, and audiences at Eindhoven University of Technology, the University of Groningen, the Fourth Annual Dutch Conference on Practical Philosophy, and the 2013 Saint Louis Annual Conference on Reasons and Rationality. I am especially grateful to Daniel Star for discussion and comments on this chapter.

- Gardner, John. 2004. "The Wrongdoing that Gets Results". *Philosophical Perspectives* 18: 53-88.
- Haji, Ishtiyaque. 1998. *Moral Appraisability*. Oxford: Oxford University Press.
- Haji, Ishtiyaque. 2002. *Deontic Morality and Control*. Cambridge: Cambridge University Press.
- Hare, R. M. 1963. *Freedom and Reason*. Oxford: Oxford University Press.
- Harman, Gilbert. 1986. *Change in View*. Cambridge, Mass.: MIT Press.
- Heuer, Ulrike. 2010. "Reasons and Impossibility". *Philosophical Studies* 147: 235-246.
- Howard-Snyder, Frances. 2006. "'Cannot' Implies 'Not Ought'". *Philosophical Studies* 130: 233-246.
- Kane, Robert. 1998. *The Significance of Free Will*. Oxford: Oxford University Press.
- Kane, Robert. 2005. *A Contemporary Introduction to Free Will*. Oxford: Oxford University Press.
- Kearns, Stephen, and Daniel Star. 2009. "Reasons as Evidence". *Oxford Studies in Metaethics* 4: 215-42.
- Kelly, Thomas. 2003. "Epistemic Rationality as Instrumental Rationality: A Critique". *Philosophy and Phenomenological Research* 66: 612-40.
- Kelly, Thomas. 2003. "Epistemic Rationality as Instrumental Rationality: A Critique". *Philosophy and Phenomenological Research* 66: 612-40.
- Naylor, Margery Bedford. 1984. "Frankfurt on the Principle of Alternate Possibilities". *Philosophical Studies* 46: 249-258.
- Schroeder, Mark. 2007. *Slaves of the Passions*. Oxford: Oxford University Press.
- Shah, Nishi. 2002. "Clearing Space for Doxastic Voluntarism". *The Monist* 85: 436-45.
- Sinnott-Armstrong, Walter. 1984. "'Ought' Conversationally Implies 'Can'". *Philosophical Review* 93: 249-61.
- Streumer, Bart. 2007. "Reasons and Impossibility". *Philosophical Studies* 136: 351-384.
- Streumer, Bart. 2010. "Reasons, Impossibility and Efficient Steps: Reply to Heuer". *Philosophical Studies* 151: 79-86.
- Streumer, Bart. Forthcoming. *Unbelievable Errors: An Error Theory about All Normative Judgements*.
- Tappolet, Christine. 2013. "Evaluative vs. Deontic Concepts". In Hugh LaFollete, ed., *The International Encyclopedia of Ethics*. Oxford: Wiley-Blackwell.

- Vranas, Peter. 2007. "I Ought, Therefore I Can". *Philosophical Studies* 136: 167-216.
- White, Alan R. 1975. *Modal Thinking*. Oxford: Blackwell.
- Widerker, David. 1991. "Frankfurt on 'Ought Implies Can' and Alternative Possibilities".
Analysis 51: 222-4.
- Williams, Bernard. 1970. "Deciding to Believe". Reprinted in his *Problems of the Self*.
Cambridge: Cambridge University Press, 1973.
- Zimmerman, Michael. 1996. *The Concept of Moral Obligation*. Cambridge: Cambridge
University Press.