



8

Reasons as Evidence

Stephen Kearns and Daniel Star

Normative reasons are strange beasts. On the one hand, we are all intimately familiar with them. We couldn't live for long without the guidance they continually offer us when we are trying to work out what to believe and what to do. At the same time, they seem to resist being analyzed in other terms. We can say that they "count in favor of..." (acts, beliefs, intentions, desires, emotions, etc.), and some have thought that this rather uninformative characterization is pretty much all we will ever be able to come up with when attempting to answer the question, what are reasons? (Scanlon 1998: 17; Parfit 2007).

Philosophers have distinguished between species of reasons in a number of ways (moral/prudential/aesthetic, practical/theoretical etc.), but it is commonly thought that no unified and informative analysis of the genus is possible.¹ Some think *reasons for action* can be analyzed in terms of (ideal) desires, but most of the very same philosophers would be unhappy with the idea that *reasons for belief* could also be analyzed in terms of desires (ideal or otherwise). No unified analysis of reasons seems possible.

We would like to thank the many participants at the *Fourth Annual Metaethics Workshop* who helped us with their comments, as well as a number of people who offered us similarly excellent comments elsewhere. In particular, we appreciate the feedback we received from Wylie Breckenridge, John Broome, Krister Bykvist, Fabrizio Cariani, David Chalmers, Roger Crisp, Terence Cuneo, James Dreier, Andy Egan, Geoffrey Ferrari, Peter Graham, Gilbert Harman, Karen Jones, Clayton Littlejohn, Errol Lord, Ofra Magidor, Julia Markovits, Elinor Mason, Sean McKeever, James Morauta, Adam Pautz, Wlodek Rabinowicz, Andrew Reisner, Jacob Ross, Jonathan Schaffer, Mark Schroeder, Wolfgang Schwarz, Russ Shafer-Landau, Michael Smith, Nicholas Southwood, Ralph Wedgwood, and two anonymous referees for *Oxford Studies in Metaethics*.

¹ The 'counting in favor of' account of reasons subscribed to by Scanlon and Parfit might be thought to be unified, but it is not informative, since these authors deny that it is possible to provide an analysis in other terms of what it is to be a reason. Other accounts of reasons are informative, but not unified.

The purpose of the present chapter is to suggest that despair on this front is premature. We believe it is possible to give an informative and unified analysis of reasons. A reason to ϕ is simply evidence that one ought to ϕ , where ϕ is either a belief or an action.² Of course, this will seem far from obvious to many readers, especially in the case of reasons for action. We are fortunate that this is the case, because it provides us with a good rationale for trying to convince readers that our analysis is correct. In the last part of the chapter, we consider the views of opponents (real and imagined). Our main claim may seem obviously true to some other readers. If so, we are also fortunate that this is the case, because we think our view is intuitively appealing, and because we need all the allies we can get.

The chapter has three main parts. To begin with, we quickly run through our new analysis of what it is to be a reason, without providing any substantive arguments for this analysis. We discuss both reasons for belief and reasons for action, but the most controversial of the principles we provide in the first part of the chapter is a principle connecting reasons for action and evidence. In the second part of the chapter, we provide arguments for this principle, and in the last part of the chapter we examine some important objections to it.

1. REASONS

Here are several principles, each of which specifies necessary and sufficient conditions for being a reason of a particular kind (or for being a reason of the most general kind, in the first case):

Reasons

R Necessarily, a fact F is a reason for an agent A to ϕ **iff** F is evidence that A ought to ϕ (where ϕ is either a belief or an action).

Reasons for Belief

RB Necessarily, a fact F is a reason for an agent A to believe that P **iff** F is evidence that A ought to believe that P .

Reasons for Action

RA Necessarily, a fact F is a reason for an agent A to ϕ **iff** F is evidence that A ought to ϕ (where ϕ is an action).

² We believe this analysis can be extended to reasons to intend, to desire, etc., but we will not discuss such reasons.

Epistemic Reasons for Belief

ER Necessarily, a fact F is an *epistemic* reason for an agent A to believe that P **iff** F is evidence that A ought to believe that P, and F is evidence that A ought to believe that P because F is evidence that P.

Pragmatic Reasons for Belief

PR Necessarily, a fact F is a *pragmatic* reason for an agent A to believe that P **iff** F is evidence that A ought to believe that P, and this is not just because F is evidence that P (and it may be the case that F is not evidence that P).

R is our main claim. It says that a fact F is a reason just in case it is evidence that one ought to ϕ , for some particular ϕ , where this ϕ is either a belief or an action. RB and RA are belief-specific and action-specific versions of R. The reader should note that it is *not* a feature of our account of reasons that we collapse the distinction between reasons for belief and reasons for action, or attempt to reduce one kind of reason to the other kind of reason. We have found this to be a common misunderstanding of our view. There is nothing in our general view that rules out the possibility that the oughts that govern belief are very different than the oughts that govern action.

In the case of reasons for belief, we claim that a fact F is a reason to believe a proposition P just in case F is evidence that one ought to believe P. ER and PR are both versions of RB, which we have provided simply in order to demonstrate that our account of reasons leaves room for there to be both epistemic reasons for belief and pragmatic reasons for belief.³ We claim that in the case of epistemic reasons for belief, the very same fact F that is evidence that one ought to believe a particular proposition P is itself evidence for the truth of this proposition, and it is *because* it is evidence that P that the fact F is evidence that one ought to believe P. In the case of pragmatic reasons for belief, on the other hand, the fact F that is evidence that one ought to believe a proposition P is evidence that one ought to believe P in a way that is not merely explained by F being evidence for the truth of P, and F may well completely fail to be evidence that P.⁴

³ It is an interesting feature of our account of reasons that it seems evidentialist in one sense, but non-evidentialist in another sense. Our account clearly has an evidentialist flavor about it because we say that all reasons to believe rest on evidence, but our account also leaves open a possibility that is typically taken to be incompatible with evidentialism, i.e. the possibility that some reasons to believe are pragmatic, rather than epistemic.

⁴ Why do we say that “F *may* fail to be evidence that P,” rather than the simpler “F fails to be evidence that P”? In most cases of pragmatic reasons to believe that one might ordinarily consider, the pragmatic reasons to believe are not also epistemic reasons to

It might seem to be a substantial objection to RB that there is a simpler account of what it is to be a reason for belief readily available to us. Perhaps a reason to believe a proposition just *is* evidence for the truth of that proposition (in all cases). This would either rule out pragmatic reasons to believe altogether, or make all reasons to believe pragmatic in nature (assuming one could provide a plausible pragmatist account of evidence). Nonetheless, it might be thought that this simpler claim still has an advantage over our claim that a reason to believe a proposition is evidence that one ought to believe that proposition, just because it is more simple, or because it does not contain any mention of oughts. In any case, we believe this objection to RB can be largely defused merely by pointing out that it can still be true on our account that:

EE Necessarily, for all agents A, F is an epistemic reason for A to believe P iff F is evidence that P.

EE is not inconsistent with any of the principles listed above. In fact, EE follows from ER in conjunction with two very reasonable principles:

EO Necessarily, for all agents A, if F is evidence that P then F is evidence that A ought to believe P.

EC Necessarily, for all agents A, F is evidence that A ought to believe P and (F is evidence that A ought to believe P because F is evidence that P) iff F is evidence that A ought to believe P and F is evidence that P.

If the reader happens to feel compelled to deny that there are any pragmatic reasons to believe, he or she should still be happy with ER and EE and may feel free to scratch out the word “epistemic” (not because its inclusion makes either of these principles false, but rather because it will then play no important role).

The attentive reader who compares RB and EE may well wonder which statement is more fundamental. This raises an important issue that we do not intend to resolve here. All of the claims we have listed above are *if and*

believe, so the relevant F does fail to be evidence for the truth of the relevant P. However, it seems that one and the same fact can be both a pragmatic reason for belief and an epistemic reason for belief, assuming one accepts that there are any pragmatic reasons for belief at all (and it doesn't really matter to us, so far as the main arguments in this paper are concerned, if the reader doesn't accept that there are such reasons). Consider the complex fact that *water is thirst-quenching and believing that water is thirst-quenching will make my life go much better than if I don't believe that water is thirst-quenching*. This complex fact seems to be both a pragmatic reason for belief and an epistemic reason for belief. It is evidence that I ought to believe water is thirst-quenching, but this is not only because it is evidence that water is thirst-quenching, for the same fact is also evidence that I ought to believe water is thirst-quenching in virtue of the good practical consequences that will follow from doing so.

only if claims. As such, they are weaker than identity claims, and are only analyses in a *broad* sense of the term. We have found it easier to argue for the truth of these *if and only if* claims than to argue for corresponding claims that state that the properties of being a reason and being evidence of an ought are identical. We will be more than satisfied if the reader accepts our arguments for the *if and only if* principles alone. Nonetheless, we also believe that the *best* explanation of the truth of all these principles is that the property of being a reason and the property of being evidence of an ought are identical.

Although we have described R as our main claim, we will actually spend most of this chapter discussing RA. This is because we take RA to be the most controversial of the claims listed above, RA and RB together entail R, R is not true if RA is not true, and perhaps the best thing that can be said for R (apart from what can be said for RA) is that it provides the basis for a *unified* account of reasons for belief and reasons for action.

2. ARGUMENTS FOR RA

We will present six arguments for RA. Two of the arguments take the form of an inference to the best explanation, one argument is based on induction, and three of the arguments are deductively valid. Each argument on its own *might* be considered only weakly persuasive; in combination, the six arguments provide us with a strong case for the new analysis of reasons.

2.1. The Simplicity Argument

The first argument runs as follows:

- (1) Epistemic and practical reasons are of a kind.
- (2) RA provides the only plausible account of reasons according to which (1) is so.
- (3) Therefore, RA is true (*inference to the best explanation*).

Let us consider premises (1) and (2) in turn. The first premise simply states the position that we consider to be the natural *default* position concerning the relationship between epistemic reasons (i.e. epistemic reasons for belief) and practical reasons (i.e. reasons for action and pragmatic reasons for belief). Epistemic and practical reasons should be thought of as being of the same basic kind prior to the presentation of good arguments to the contrary.

What entitles us to think of (1) as stating a default position? We believe we are entitled to do so on the basis of linguistic evidence. Here are three relevant linguistic observations. First, and most obviously, the same word, “reason,” is used when people talk about epistemic reasons for belief as when people talk about reasons for action. If epistemic reasons for belief alone concern evidence and reasons for action alone concern right-makers and wrong-makers (for instance), then it is surprising that we use the word “reason” in cases involving evidence as well as cases involving right-makers or wrong-makers. One might have thought we would have different words for these different kinds of entities. It is certainly difficult to detect any ambiguity in our ordinary normative reasons talk.

Secondly, the word “reason” behaves the same way, grammatically speaking, in both reasons for action and reasons for belief talk, so long as we restrict our attention to normative reasons. We talk about there being a “reason *to act*” when concerned with a normative reason for action, and we similarly talk about there being a “reason *to believe*” when concerned with a normative epistemic reason to believe. Of course, it is true that there are also non-normative uses of “reason” in English, but it is interesting to note that “reason” does not normally take the infinitive when used in a non-normative fashion. We say “the reason *that* it is raining is because the clouds are heavy,” or “the reason *why* it is raining is that the clouds are heavy,” but we do not say “the reason *to* it is raining is because the clouds are heavy.”

Thirdly, it is possible to construct grammatically correct sentences of the following form: F is a reason to believe P and to ϕ (where ϕ is an action). An example of a sentence that has this form is: “That the ground is wet is a reason to believe it is raining and to take an umbrella.” Note that the word “reason” only appears once in this well-formed sentence. One would not expect this sentence to be grammatically correct if epistemic reasons for belief and reasons for action were reasons of two very different kinds. Of course, one *can* say “That the ground is wet is a reason to believe it is raining and a reason to take an umbrella” but the second occurrence of “a reason” seems unnecessary.

Apart from these reasons to accept that (1) states a position that should be considered to be the default position, and to accept that (1) is true, there are additional reasons to accept (1). Most notably, it seems that we can weigh up reasons to act and epistemic reasons to believe against each other. This would not be possible if they were not of a kind. Here is an example of what we have in mind: imagine a professional high-jumper who must jump higher and higher every time he jumps in a particular prestigious competition if he is to succeed in winning a medal. The high-jumper is aware that he could think carefully about the height of the jump on each

occasion he jumps and thus come to form a well-justified belief about whether or not he is going to succeed on each particular occasion. However, he is also aware that paying attention to this kind of reason for belief may be something that will guarantee that he fails to successfully make certain jumps. He has a reason to ignore the evidence concerning whether or not the high-jump is too high for him, given his track-record, and this is the reason he has to win the competition. A reason for action trumps a reason for belief, over and over again.

Similarly, a distraught mother might ignore evidence that her kidnapped child has been murdered, and continue searching for him, and this extra effort might actually be required if she is to find her child. Her reason to believe that continuing to search for her child is futile would be outweighed by the very weighty reason that speaks in favor of continuing to search for her child.

The second premise says that RA provides the only plausible account of reasons according to which epistemic and practical reasons are of a kind.⁵ Being an epistemic reason for belief is simply being evidence for something. (On the simplest view, an epistemic reason for P is evidence that P. On our view, an epistemic reason for P is both evidence that P and that one ought to believe that P.) Epistemic reasons for belief are the same kind of thing as reasons for action (and pragmatic reasons for belief). Therefore, one very natural view is that reasons for action are also evidence for something. Evidence for what, however? Given that practical reasons ultimately concern what one ought to do (or believe), then we should infer that practical reasons are evidence for propositions concerning what one ought to do (or believe). In short, we extend a plausible analysis of epistemic reasons to all normative reasons.

In order to capture the idea that epistemic reasons and practical reasons are of a kind, other analyses of reasons must find something else that they have in common. If we should not analyze practical reasons in terms of evidence, then how should we analyze them (keeping in mind that epistemic reasons should be so analyzed)? One view is that reasons of both kinds 'count in favor' of beliefs or actions. This is no *analysis*, since the notion of counting in favor is one that cannot be understood separately from the notion of being a reason (as we have already noted).

One might think that reasons for action and reasons for belief are both types of explanations. John Broome argues that normative reasons for action are (parts of) explanations of why one ought (or ought not) to perform actions (see Broome 2004 and Kearns and Star 2008). However, epistemic

⁵ Thank you to Sean McKeever for the written comments that helped us improve this section of the chapter.

reasons are not explanations, but evidence. (This can be brought out with the following case. A piece of undiscovered evidence, such as the fact that Bob's fingerprints are on the murder weapon, is a reason to believe that Bob is the murderer. It is not any part of an explanation concerning what we ought to believe, however, because we do not *have* this evidence/reason.⁶) If epistemic reasons are to be analyzed in terms of evidence and nothing more, and if epistemic and practical reasons are of a kind, then practical reasons must also be analyzed in the same way.

2.2. The Standard Cases Argument

Our second argument is fairly straightforward. It runs as follows:

- (1) Standard cases of practical reasons to ϕ are cases of evidence that one ought to ϕ , and vice versa.
- (2) Therefore, RA is true (*argument from induction*).

In order to provide support for (1) we could catalogue a very large number of examples, but we hope the reader will in fact be content to be provided with just a couple of carefully chosen examples. It should be apparent from the general features of these examples that not much work would be required to generate many more suitable cases.

Here is a fairly simple case to start with. Imagine that I (Stephen or Daniel) am passing by my friend John's house and I find him standing in the street next to his car. He is wincing and crying out for help. I notice his foot is stuck under his car wheel. I see, or at least infer, that he is in pain. Clearly, the fact that John is in pain is a reason to help him and to believe that I ought to help him. It is also evidence that I ought to help him. Perhaps I had promised to meet another friend in two minutes' time for a conversation about his new metaethics paper. I remember that fact as I rush over to the car to find a way to help John, and I feel a touch of guilt that I may not get to fulfill my promise. I ask myself (even as I try lifting the car), ought I to be helping John? After all, the fact that I promised to meet my other friend very soon is a reason to rush to meet the other friend. It is also evidence that I should now be rushing to meet the other friend. Clearly John's pain provides much stronger evidence that I ought

⁶ It might be objected that there are cases where one ought to believe a proposition even though one doesn't *have* evidence for that proposition. Such cases always seem to involve some form or other of epistemic irresponsibility on the part of the relevant believer which explains why it is that they ought to believe a particular proposition (since they ought to *have* the evidence that is readily available to them). Thus we merely need to stipulate that the above example involves no epistemic irresponsibility.

to be doing exactly what I am now doing (by now, I am attempting to lift the car off his foot using a crowbar). The pain also provides a reason to help John that is much stronger than the reason I have to meet the other friend (i.e. the reason that springs from the fact that I have promised to meet him).

This first case involved moral reasons. Let us now consider a case that concerns prudential reasons for action. Jack works hard for a good charity relief organization during the day. He lives alone. In the evenings he likes to either carefully read excellent science books, in order to come to a better understanding of the world, or simply relax in front of his television. He generally finds that he is unable to enjoy the relatively petty shows on television any evening when he has been reading one of his books, and he is unable to concentrate on reading one of his books any evening when he has already watched some television, due to its mind-numbing properties. One night after arriving home he is wondering what to do (sometimes he doesn't bother wondering what to do, but sometimes he does), having already determined that he wants to do one of these two things. The fact that reading a particular book would help him better understand the world is a reason for him to read that book. The same fact is evidence that he ought to read the book. The fact that watching television would give him a lot of pleasure (certainly much more than he would get from reading the book) is a reason for him to watch television. The same fact is evidence that he ought to watch television. He weighs these reasons against each other, and judges that he ought to read a book on the basis of these considerations. This is, on our view, just another way of saying that he weighs the evidence that he ought to read a book against the evidence that he ought to watch television, and judges that he ought to read a book on this basis. However, even if you do not accept the last claim, it seems that you should accept that the normative reasons that figure in this case correspond, on a one-to-one basis, to examples of evidence that he ought to act in the ways he is considering acting.

Cases such as these seem to be examples where it is not controversial to claim that the facts that provide practical reasons to ϕ also provide evidence that one ought to ϕ . They are examples where the relevant facts are transparent to the agent, that is, where there are no false beliefs playing any role in deliberation and there is no misleading evidence around clouding the water. As such, they are *standard* cases, for we all encounter such cases over and over again during the course of a day.

However, it must be admitted that examples that lack the transparency or simplicity of these cases may well be thought to tell against RA. Consider a case of someone who is thirsty and who, in a state of ignorance (which he is in through no fault of his own), is about to drink from a glass that

contains deadly poison. It is far from uncontroversial to say that this person has a normative reason to drink from the glass, even though he clearly has evidence that he ought to drink from the glass (e.g. it looks like it contains an ordinary thirst-quenching liquid). Niko Kolodny (2005) has argued that such a person has no real reason to drink the liquid he has the option of drinking, although it may well be rational for him to form an intention to do so. We will not be assessing Kolodny's interesting arguments here, but we wish to note that basic intuitions concerning the matter of whether or not such an unlucky person has a reason to drink from such a glass seem to differ from philosopher to philosopher (Kolodny himself does not claim otherwise). We return to supposed counterexamples to RA below, in the third part of the chapter.

2.3. The Deliberation Argument

The third of our arguments is deductively valid. We have split it up into an argument for the left to right of RA (i.e. the claim that reasons to ϕ are evidence that one ought to ϕ) and the right to left of RA (i.e. the claim that evidence that one ought to ϕ is a reason to ϕ). We will consider the former first. The argument runs as follows:

- (1) Practical reasons to ϕ can play an important role in reliable practical reasoning concerning whether or not one ought to ϕ .
- (2) If practical reasons to ϕ are not evidence that one ought to ϕ then they cannot play this important role in reliable practical reasoning.
- (3) Therefore, reasons to ϕ are evidence that one ought to ϕ .⁷

By “reliable practical reasoning” we mean practical reasoning that is generally successful in terms of issuing in correct judgments concerning what it is one ought to do (or, at least, judgments that get as close to being correct as is normally possible). Bearing this in mind, the first premise seems true. It may even be true by definition, because there is clearly a very tight connection between reasons and reasoning. It is plausible to suppose that, whatever else they are, reasons are facts that can be correctly used in reasoning about what to do or believe. This is suggested by the fact that the words “reason” and “reasoning” have a similar etymology. In any case, (1) seems obvious as a matter of empirical fact. Over and over

⁷ This conclusion may seem to be stating the strong identity thesis that we said we would not directly argue for in this chapter. However, this conclusion does not state that the *property* of being a reason is identical to the property of being evidence of an ought. All it says is that facts that are reasons are evidence of oughts. Similar considerations apply to similarly worded conclusions below.

again, we each consider practical reasons when we are engaged in practical reasoning.

Premise (2) also seems true, although it is by no means as obvious as (1). Why do we think (2) is true? Reasons are often successfully used by rational people to determine what they ought to do, but we think this would be *miraculous* if such reasons are not evidence. Consider a person who is trying to work out what it is she ought to do. It seems natural to describe her as aiming to hit a target. Being a sensible person, she realizes that she may fail to hit the target, that is, despite her best efforts, she might fail to do what it is she really ought to do. Being a reasonable person, she will try her best to hit the target, given relevant time constraints (and, for all we say on the topic, it may be reasonable for her to be content with getting very close to hitting the target). Now, ask yourself: how could the person we are considering ever properly aim to hit the target, or, indeed, ever succeed in hitting the target (other than through sheer luck), if reasons to ϕ are *not* evidence that one ought to ϕ ? On the other hand, if reasons to ϕ are evidence that one ought to ϕ , then it is possible to see how agents are able to reliably engage in practical reasoning, in order to work out what they ought to do. We can conclude that (3) reasons to ϕ are evidence that one ought to ϕ .

Let us now turn to the right to left direction of RA, which the following argument is designed to defend:

- (4) If a fact F is evidence that one ought to ϕ , then F can play an appropriate role in one's reliably concluding that one ought to ϕ .
- (5) If a fact F can play this role, then F is a reason to ϕ .
- (6) Therefore, evidence that one ought to ϕ is a reason to ϕ .

Again, this argument is valid. Let us now assess premises (4) and (5).

Premise (4) says that evidence that an agent ought to ϕ can help this agent conclude that she ought to ϕ . The plausibility of this idea stems from the very notion of what it is for a fact to be evidence for something. We use evidence precisely to work out which propositions are true. If a fact is evidence that one ought to ϕ , then such a fact is able to help an agent conclude that she ought to ϕ . This fact *raises the probability* of the proposition that she ought to ϕ (see Section 2.6 below for more details). If the agent is reasoning well, she can use this fact to conclude that she ought to ϕ on those occasions she ought to ϕ .

Premise (5) says that if a fact can play an appropriate role in an agent's reliably concluding that she ought to ϕ , then this fact is a reason to ϕ . This follows from the idea that it is sufficient for a fact to be a reason that it plays such a role in reasoning. Practical reasons simply *are* whatever facts are used

in reasoning in this way. Reasons are facts that can help us determine what we ought to do. Reasons to ϕ are facts that can help us determine that we ought to ϕ .

One way to show that premise (5) is true to assume it is false. On the assumption that (5) is false, it is possible that a person may come to know what she ought to do without considering or even knowing any of the reasons she has for acting. Thus, if there is a fact that is evidence that one ought to ϕ without being a reason to ϕ , then one may use this evidence to work out that one ought to ϕ without being aware of any reason to ϕ at all. This result is unattractive as it implies that reasons need not play any important role in our deliberations about what we ought to do. (We expand upon this point in our paper ‘Reasons: Explanations or Evidence?’ 2008.)

We conclude, then, that reasons to ϕ are evidence that one ought to ϕ and that evidence that one ought to ϕ is a reason to ϕ . This is equivalent to RA. Because a fact’s being able to play a certain important role in practical reasoning is both necessary and sufficient for this fact to be a reason and for it to be evidence, we conclude that reasons are evidence.

2.4. The Public Role Argument

Our fourth argument complements the third argument. Again, we have split it up into a defense of the left to right of RA and the right to left of RA. Here is the argument for the left to right of RA:

- (1) Practical reasons to ϕ can play an important public role in rationally convincing people (by providing justifications) that they ought to ϕ , and rationally convincing people that when one ϕ -ed one did what one ought to have done.
- (2) If practical reasons to ϕ are not evidence that one ought to ϕ then (1) is not true.
- (3) Therefore, reasons to ϕ are evidence that one ought to ϕ .

Consider premise (1) to begin with. When asked why one did something, there is a vast difference between a reply that offers up a causal reason (e.g. “I was pushed onto you”) and a reply that offers up a justifying reason (e.g. “I realized that pushing you out of the way of the bus would be the only way to save your life”). In cases where one is *only* causally explaining what happened one is typically implicitly admitting that one had insufficient justification for what one did, or that one was failing to respond appropriately to reasons at all (as might be the case if one had lost control of one’s body due to having

been pushed by someone else). In contrast, when one appeals to normative reasons, one can justify (or at least partly justify) one's actions. Similarly, one can appeal to normative reasons in order to rationally convince others about what they ought to do.

Now consider premise (2). When considering how we go about convincing others to act one way or the other, or convincing them that we were right to have acted one way or the other, people generally recognize that there is all the difference in the world between efforts to convince others that depend on violence or propaganda (subtle or otherwise), and efforts that consist in providing evidence. But if practical reasons to ϕ are not evidence that one ought to ϕ then this distinction effectively collapses and there is no genuine trading in rational justifications to be had, so (1) is not true. In this context, it is worth reminding ourselves of an important and common criticism of Ayer and Stevenson's simple versions of non-cognitivism. It is often said that Ayer and Stevenson effectively collapse the distinction between providing a normative reason and causing to believe, and our view both endorses this objection to simple non-cognitivism and further explains the basis of the objection. It is only because (3) reasons to ϕ are evidence that one ought to ϕ that we are able to rationally and reasonably persuade each other that there are various things that we ought to do.

The argument for the right-to-left reading of RA runs as follows.

- (4) If a fact F is evidence that someone ought to ϕ , then F can play an appropriate public role in rationally convincing that person that she ought to ϕ and in rationally convincing other people that she ought to ϕ .
- (5) If a fact F can play this role, then F is a reason to ϕ .
- (6) Therefore, evidence that one ought to ϕ is a reason to ϕ .

Again, it should be clear how this argument parallels that given in Section 2.3 above. Our defense of these premises also parallels our defense of the related premises in Section 2.3. If a fact is evidence that one ought to ϕ , then this fact can be used to justify one's own ϕ -ing. If a fact is evidence that someone ought to ϕ , informing them of this fact can rationally convince this person that she ought to ϕ . Thus premise (4) is true. Furthermore, this role that facts can play is sufficient to ensure that such facts are reasons. It is the *reasons* one has to ϕ which one appeals to in order to justify one's own ϕ -ing. It is the *reasons* to ϕ to which one appeals in order to convince others that they ought to ϕ . Thus premise (5) is true. We conclude, therefore, that RA is true.

2.5. The Normative Principles Argument

In this section, we present another argument for RA. The basic argument runs as follows:

- (1) A fact F is a reason to ϕ if and only if it is normally the case that if a fact relevantly similar to F obtains, then one ought to do something relevantly similar to ϕ -ing.
- (2) A fact F is evidence that one ought to ϕ if and only if it is normally the case that if a fact relevantly similar to F obtains, then one ought to do something relevantly similar to ϕ -ing.
- (3) Therefore, a fact F is a reason to ϕ if and only if F is evidence that one ought to ϕ (that is, RA is true).

It should be clear why the conclusion follows from (1) and (2).⁸ We shall now discuss each premise in turn.

The first premise says that a fact is a reason to perform a certain action if and only if it is normally the case that if something like this fact obtains, then one ought to perform something like this action. In order to see why we should accept this premise, consider the following typical fact that can act as a reason to perform an action: John is in pain. This fact is, like most typical reasons, a non-normative fact that simply describes how the world is. Such a fact can, however, act as a normative reason to try to help John.

Thus we have, on the one hand, the non-normative fact that John is in pain and, on the other, we have the normative fact that one should try to help John. How are the two linked? A natural answer to this is that there is a (hedged) normative principle that states that, if someone is in pain, then one ought to try to help that person. That is, it is generally the case that if a certain type of fact obtains (e.g. a fact of the form 'x is in pain'), then a certain type of normative fact obtains (e.g. a fact of the form 'one ought to

⁸ One might worry that what counts as *relevantly similar* to F and to ϕ -ing is something different in the first premise than what it is in the second (i.e. what counts as relevantly similar with regards to reasons is different from what counts as relevantly similar with regards to evidence), in which case the argument would not be valid. In fact, we intend the occurrences of the phrase 'relevantly similar' not to be read with reference to the notions of evidence or reason. We think it is straightforwardly true that, for instance, it is normally the case that if a fact relevantly similar to the proposition that John is in pain obtains, then one ought to do something relevantly similar to helping John. The phrase 'relevantly similar' is simply meant to ensure that, whatever similarity the pertinent fact has to F, the associated action has a relevant similarity to ϕ -ing, and vice versa. Thank you to Elinor Mason for prompting us to think carefully about this matter.

help x'). Thus the fact that John is in pain is a reason to help John because, normally, when relevantly similar facts obtain (e.g. 'Mary is in pain,' 'Henry is in pain'), then one ought to perform relevantly similar actions (e.g. help Mary, or help Henry). What bridges the gap between non-normative facts and unconditional normative facts are conditional normative principles that state, in effect, that if something like this non-normative fact obtains, then so does something like the unconditional normative fact.

A slightly different account of the link between the fact that John is in pain and the fact that one ought to try to help him is that the former *explains* the latter (see Broome 2004). That is, one ought to try to help John *because* he is pain. But what makes *this* true? Why does John's being in pain explain why one ought to help him? The most natural answer to this question is again that there is a (hedged) normative principle that says that if someone is in pain, then one ought to try to help that person. It is in virtue of this principle that we may truly say that one ought to help John because he is in pain.

A non-normative fact is a reason to perform a certain action, then, whenever there is a normative principle that says that if a relevantly similar fact obtains, then one ought to perform a relevantly similar action. Such a principle need not be strict or necessary. It can be hedged and contingent. However, such a principle needs to be more robust than a simple material conditional. It is not sufficient for F to be a reason to ϕ that the material conditional 'if a fact relevantly similar to F obtains, then one ought to do something relevantly similar to ϕ -ing' is true. Such a material conditional would be true if it simply happened that its consequent were true. The normative principle expressed in (1) needs, at the very least, to be a generic truth that links F-like facts to relevant normative facts. If the normative principle is true, it does not simply *happen* to be the case that if a fact relevantly similar to F obtains, then one ought to do something relevantly similar to ϕ -ing. Rather, this is *normally* the case. Such a principle links non-normative facts to unconditional normative facts.

The second premise is very similar to the first. It says that a fact is evidence that one ought to ϕ if and only if it is normally the case that if a relevantly similar fact obtains, then one ought to do something relevantly similar to ϕ -ing. To see why this is true, again consider the fact that John is in pain. When this fact is evidence that one ought to help John, as it surely often is, what makes it evidence?

An initially plausible answer to this question is that the fact that John is in pain is evidence that one ought to help John if and only if this fact *reliably indicates* that one ought to help John. Though we think something like this is correct, it is too simple as it stands. A generalization of this idea would produce the following principle. A fact is evidence for a proposition if and

only if this fact reliably indicates this proposition. If this were right, then there could be no misleading evidence for necessarily false propositions, as no fact could reliably indicate something necessarily false. Consider also the following case: a normally reliable telephone book says that John's number is 123456. As it happens, the book is wrong. The fact that the book says that John's number is 123456 does not reliably indicate that John's number is 123456. Still, the fact that the book says John's number is 123456 is still *evidence* that his number is 123456. Therefore, a fact can be evidence for a proposition without reliably indicating this proposition.

The obvious fix for this problem is to say that, though the fact that the book says John's number is 123456 does not reliably indicate that his number is 123456, there are many relevantly similar facts that do reliably indicate relevantly similar propositions (e.g. the book does reliably give the correct numbers for Mary and Henry and many others). We may therefore conclude that a fact is evidence for a proposition if and only if relevantly similar facts reliably indicate relevantly similar propositions. In the normative case, then, we can say that a fact F is evidence that one ought to ϕ if and only if facts relevantly similar to F reliably indicate propositions relevantly similar to the proposition that one ought to ϕ . Furthermore, if facts relevantly similar to F reliably indicate propositions relevantly similar to the proposition that one ought to ϕ , then it is normally the case that if a fact relevantly similar to F obtains, then one ought to perform an action relevantly similar to ϕ -ing. That is, a generic conditional is true whenever the antecedent of this conditional reliably indicates the consequent. From these ideas, premise (2) follows. From (1) and (2), RA follows.

In short, the argument presented in this section is the following. A fact, F , is a *reason* to perform a certain action if and only if a principle holds that links F to the proposition that one ought to perform this action. A fact, F , is *evidence* that one ought to ϕ if and only if this same principle holds. Therefore, F is a reason to ϕ if and only if it is evidence that one ought to ϕ .

2.6. The Strength of Reasons Argument

In this section we present a final argument for RA. This argument goes as follows:

- (1) Reasons can have different strengths.
- (2) RA is the best explanation of how this is possible.
- (3) Therefore, RA is true. (*inference to the best explanation*).

This argument takes the form of an inference to the best explanation. Such an inference may not be deductively valid, but it provides good reason to believe RA. Let us examine each of the premises of the argument.

Premise (1) is uncontroversial. Some reasons are stronger than others. If John is in excruciating pain, then this is a strong reason to help him. If he offers you five pounds to help him, this is a less strong reason to help him. Reasons can outweigh other reasons. If Mary offers you five pounds not to help John, this reason not to help John is outweighed by the fact that John is in excruciating pain. Reasons can combine their strengths. The fact that John is in excruciating pain and the fact that he offers you five pounds to help him combine to make an even stronger reason to help him. All these facts need explaining. What exactly does it mean to say that one reason is stronger than another? How do reasons outweigh other reasons? How do reasons combine their strengths?

Premise (2) says that RA suggests very attractive answers to these questions. Evidence also comes in different strengths. If John is in excruciating pain, then this is strong evidence that one ought to help him. If he offers you five pounds to help him, this is weaker evidence that one ought to help him. Evidence can outweigh other evidence. If Mary offers you five pounds not to help John, this evidence that one ought not to help him is outweighed by the fact that John is in excruciating pain. Pieces of evidence can combine their strengths. The fact that John is in excruciating pain and the fact that he offers you five pounds to help him combine to form stronger evidence that one ought to help him.

These parallels between reasons and evidence are suggestive enough in themselves to give us good reason to believe RA. Furthermore, these parallels provide us with an explanation of what the strength of a reason comes to. Without appealing to RA, the strength of reasons is somewhat mysterious. Some who appeal to a primitive 'counts in favor of' relation between a reason and an action may say that one reason is stronger than another when this reason counts more strongly in favor of this action, where such a relation is also primitive. This is not very satisfying.⁹

When we appeal to RA, however, accounting for the strength of reasons is unproblematic. We have a good grasp of the idea of the strength of evidence. The strength of a piece of evidence E for a proposition P depends on the

⁹ John Broome (2004) appeals to the notion of a weighted explanation to explain the way in which the strengths of reasons can be compared. We criticize this idea, as well as his general account of reasons, in a separate paper that compares his account of reasons with our own, 'Reasons: Explanations or Evidence?' (2008).

degree to which E increases the probability of P.¹⁰ The more probable P is given E, the stronger evidence E is that P is true. E is stronger evidence than another piece of evidence E* for P if and only if E makes P more probable than E* makes P. E outweighs a piece of evidence E* if and only if E is evidence for P, E* is evidence for \sim P and E makes P more probable than E* makes \sim P. Two pieces of evidence, E and E* can combine to form stronger evidence if the probability of P given the conjunction of E and E* is greater than both the probability of P given E and the probability of P given E*. The strength of evidence for the truth of a proposition, then, can be accounted for in terms of the probability of the proposition being true given this evidence.

This idea translates very nicely into an account of the strength of a reason. The strength of a reason to ϕ , R, depends on the degree to which R increases the probability that one ought to ϕ . The more probable it is that one ought to ϕ given R, the stronger reason to ϕ R is. R is a stronger reason to ϕ than another reason R* if and only if R makes the proposition that one ought to ϕ more probable than R* makes it. R outweighs R* if and only if R is a reason to ϕ , R* is a reason not to ϕ , and R makes the proposition that one ought to ϕ more probable than R* makes the proposition that one ought not to ϕ . Two reasons R and R* can combine to create a stronger reason to ϕ if the probability that one ought to ϕ given the conjunction of R and R* is greater than the probability that one ought to ϕ given R and the probability of the same proposition given R*. The strength of a reason to perform an action can be accounted for in exactly the same way as the strength of evidence can.

We claim, then, that if we understand normative reasons to be evidence for oughts, we are able to give a very attractive account of what it is for such reasons to have strengths. A reason to ϕ makes the proposition that one ought to ϕ more probable. The stronger the reason is, the more probable it is that one ought to ϕ . RA best explains how reasons can have strengths.

¹⁰ This raises two questions. First, what kind of probability is increased by evidence? It is evidential or epistemic probability. For an extended discussion of this see Timothy Williamson (2002: ch. 10). Second, what is the probability increased relative to? There are two natural answers to this that are not quite right. One is that evidence increases the probability of the truth of a proposition from a position of total ignorance. This is not correct because a fact's being evidence for a proposition can sometimes depend on the existence of other evidence. The other is that evidence increases the probability of the truth of a proposition given one's entire body of other evidence. This is not correct because it is possible for a fact to be evidence for a proposition even if one's entire body of other evidence already makes the probability of this proposition equal to 1 (one's evidence may be overdetermined). The correct answer is that evidence increases the probability of a proposition relative to some salient relevant subset of one's total body of evidence. Thank you to Jacob Ross for prompting us to address this point.

3. OBJECTIONS TO RA

3.1

In this part of the chapter, we consider various objections to RA and furnish replies. The first objection is that there are counterexamples to the thesis that all pieces of evidence that one ought to ϕ are reasons to ϕ . Consider the following example. A newspaper says that there are people starving in Africa.¹¹ This is evidence that one ought to give money to Oxfam. However, the fact that the newspaper says that people are starving in Africa is not a *reason* to send money to Oxfam. This may seem most obvious when the newspaper (which is, on the whole, reliable) actually incorrectly reports that there are people starving in Africa. Rather, it is the fact (if it is a fact) that there are people starving in Africa that is a reason to send money to Oxfam. Thus it seems possible for a fact to be evidence that one ought to ϕ without being a reason to ϕ .

Why exactly might one think that the fact that the newspaper says that there are people starving in Africa is not a reason to send money to Oxfam? A plausible diagnosis is that such a fact, though evidence that one ought to send money to Oxfam, does not *make it right* to send money to Oxfam. What makes it right to send the money is that there are people starving in Africa. Evidence that one ought to ϕ may not be a reason to ϕ if this evidence is not a right-maker for ϕ -ing. Given that there are cases in which a fact is evidence for a normative proposition but where this fact does not make this proposition true, not all evidence that one ought to do something is a reason to do it.

How should we reply? First, it is plausible that a fact can act as a right-maker for an action simply by being good evidence that one ought to perform this action. The newspaper case can illustrate this point. Let us say that, in fact, though the newspaper is extremely reliable, in this case they got it wrong. Still, someone who reads this newspaper may be acting immorally if they do not send money to Oxfam (or do something similar). The fact that the newspaper says that people are starving in Africa may in itself be enough to create an obligation to send money to Oxfam.

Secondly, those who reject the idea that the fact that the newspaper says there are people starving in Africa is a reason to send money to Oxfam may also need to reject the idea that the fact that there are people starving in Africa is a reason to send money to Oxfam. After all, the fact that people

¹¹ James Morauta may have been the first to suggest that cases like this might be construed as counterexamples.

are starving *need* not mean that they are badly off. In some distant possible worlds, starving might be extremely pleasant and not life threatening at all. Being starving certainly *indicates* that a person is badly off, but this is just to say that it is evidence that one ought to send money to Oxfam. Indeed, *most* of the facts we cite as reasons merely indicate what we ought to do. Therefore, unless one wishes to deny that most facts that we think of as reasons really are reasons, the fact that the newspaper says there are people starving in Africa is really a reason to send money to Oxfam.

3.2

A second objection attacks the claim that if F is a reason to ϕ , it is evidence that one ought to ϕ . Consider the following case. There is a time-bomb in a building in which John resides. This time-bomb is counting down and, when it explodes, it will destroy the building and kill everyone in it. It seems, then, that John has a reason (a very good reason) to flee the building. This reason is that there is a time-bomb in the building. However, John has no evidence whatsoever that there is a time-bomb in the building. It is extremely well hidden. Therefore, it seems that the fact that there is a time-bomb in the building is a reason for John to flee the building, but it is not *evidence* that John ought to flee the building, as he (and, we may assume, everyone else) is oblivious to this fact. Therefore, RA is false.

This objection relies on two claims. The fact that there is a time-bomb in the building is a reason for John to flee. The fact that there is a time-bomb in the building is not evidence that John should flee. Both of these claims are controversial. If John really has *no idea* that there is a time-bomb in the building, does he really have a reason to flee? And isn't it possible that a fact could be evidence even if no one is aware of this fact (how could we look for evidence if all the evidence we have is known to us)? It seems that in some moods we think that the fact that there is a time-bomb in the building is a reason to flee even when no one knows this fact, while in some moods we do not. Similarly, in some moods we think that the fact that there is a time-bomb in the building is evidence that one ought to flee even when no one knows this fact, while in some moods we do not.

We think that to account for this, it is necessary to distinguish there *being* a reason/evidence from someone's *having* this reason/evidence. The fact that there is a time-bomb in the building is indeed a reason for John to flee the building. However, John does not *have* this reason. It is in this sense that John has no reason to flee the building. Similarly, the fact that there is a time-bomb in the building is indeed evidence that John ought to

flee the building. However, John does not have this evidence. It is in this sense that John has no evidence that he ought to flee. Once this distinction is pointed out, the objection to RA loses its force. The above case is not one in which a fact is a reason to ϕ without being evidence that one ought to ϕ . It is simply a case in which a fact is a reason to ϕ without anyone *having* it as evidence that one ought to ϕ . This is perfectly consistent with RA.

Not only is the distinction between there *being* a reason/evidence and *having* a reason/evidence a distinction that provides us with a response to the above objection; it is also a distinction that provides us with additional evidence for our view. That is to say, we take it that one thing that speaks in favor of our view is precisely the fact that the being/having distinction in the case of reasons maps very nicely onto the being/having distinction in the case of evidence. To say that there *is* evidence that one ought to do something in particular is to say that there *is* a reason to do that thing, while to say that one *has* evidence that one ought to do something in particular is to say that one *has* a reason to do that thing.¹²

3.3

The third objection is simple. If F is a reason for an agent A to ϕ , then A *can* ϕ . However, it is not the case that if F is evidence that A ought to ϕ then A can ϕ . Therefore RA is false. In short, 'reason' implies 'can,' while 'evidence' does not.

We certainly admit that, even if a fact is evidence that one ought to ϕ , it need not be the case that one can ϕ . Therefore, we must reject the idea that 'reason' implies 'can.' This idea is obviously derived from the more familiar thesis that 'ought' implies 'can.' Some philosophers reject this thesis. Such philosophers will have little sympathy with 'reason' implies 'can.' What about those people who accept 'ought' implies 'can'?

RA does retain some connection between reasons to ϕ and the ability to ϕ . If it is true that a person ought to ϕ only if she can ϕ , then it is also true that she has *conclusive* reason to ϕ only if she can ϕ . This is because a person who has conclusive reason to ϕ has conclusive evidence that she ought to ϕ . Furthermore, if someone has conclusive evidence that she ought to ϕ ,

¹² We could thus list principles that concern *having* reasons, in order to complement our original list of principles. For instance, RA* would correspond to RA.

RA* Necessarily, F is a reason for an agent A to ϕ and A has this reason to ϕ iff F is evidence that A ought to ϕ and A has this evidence that he/she ought to ϕ .

then she ought to ϕ . RA, then, does not break the link between ‘reason’ and ‘can’ completely.

Still, in those cases in which a person cannot ϕ , but has (non-conclusive) evidence that she ought to ϕ , RA claims that she has a (non-conclusive) reason to ϕ . This is a rather weak principle. Consider a surgeon who, unbeknownst to her, is performing surgery on someone whom she cannot save. After her patient dies, she realizes this. Would it be right of this surgeon to conclude that she had no reason at all to save her patient? That absolutely nothing counted in favor of her saving the patient? It seems not. She had a reason to save the patient, but she was simply unable to do so.

3.4

Another objection runs as follows. One has reason to perform supererogatory actions. One does not have evidence that one ought to perform supererogatory actions, since supererogatory actions are essentially actions that are beyond the call of duty. Therefore not all reasons for action are evidence that one ought to perform these actions. Of course, this objection depends on there really being supererogatory actions. While the “beyond the call of duty” thought is present in moral common sense, it has proven notoriously difficult to provide an account of supererogatory actions that succeeds in making good sense of their existence. Either as a result of this, or for other reasons, many moral philosophers doubt that there are any supererogatory actions.

One way to understand supererogatory actions is to depend on a distinction that is sometimes made between obligations and oughts and claim that supererogatory actions ought to be done but are not obligatory (though this need not be thought of as a sufficient condition for an action to be supererogatory). This is suggested by the phrase “beyond the call of duty.” It should be obvious that this conception of the supererogatory is entirely compatible with our reasons as evidence thesis.

However, it might be claimed that this is an incorrect conception of the supererogatory (perhaps because there is no relevant distinction to be made between obligations and oughts). An alternative interpretation of “beyond the call of duty” is the view that it is not the case that one ought to perform supererogatory acts (this certainly isn’t a sufficient condition). This conception seems to conflict with RA. However, we have three responses available to us at this point. First, we could simply accept this view but deny that we have reasons to perform supererogatory actions. Though it may be good, kind, or nice to perform such actions, one really has no reason to perform them and there is thus no evidence that one ought to perform them. Secondly, we could instead claim that there are reasons to perform

supererogatory actions and also evidence that one ought to perform these actions, albeit evidence that is always *misleading*. Finally, one could modify RA in the following way:

Necessarily, F is a reason for an agent A to ϕ iff F is evidence that A ought to ϕ or F is evidence that is supererogatory for A to ϕ .

This is our least favorite solution to the problem, but is still one that is very much in the spirit of the account of reasons defended in this chapter.

3.5

We sometimes receive the following objection to RA. One fact F can be both a reason to ϕ and a reason not to ϕ , but no fact can be both evidence for P and evidence for not P. Therefore, not all reasons to ϕ are evidence that one ought to ϕ . Consider the following example, provided by Andy Egan when pressing this objection.¹³ A runner in a competition realizes that if he sticks his leg out at a particular time he will trip over his main competitor without others observing him do so. The fact that if he were to do this his main competitor would trip over is a reason to stick his leg out (because it will increase his chances of winning the race) *and* it is a reason not to stick his leg out (either because it would provide an unfair advantage, or because it would be likely to injure the other person). However, this fact is *not* both evidence that he ought to stick his leg out and evidence that it is not the case that he ought to stick his leg out. This is because it is impossible for a fact F to be both evidence that P and evidence that not P. After all, no fact can increase the probability that P and increase the probability that not P. Therefore, RA is false.

This objection relies on conflating two distinct notions—that of evidence that one ought not to ϕ and that of evidence that it is not the case that one ought to ϕ . Look carefully at RA. There is no negation before any ought. According to RA, a reason to ϕ is evidence that one ought to ϕ and a reason not to ϕ is evidence that one ought not to ϕ (rather than evidence that it is not the case that one ought to ϕ). It should be uncontroversial that if a fact is evidence that one ought not to ϕ this does not entail that this fact is evidence that it is not the case one ought to ϕ .¹⁴ To see that this is so, just consider any fact that is evidence that an action is not merely

¹³ The same objection was pressed on us by James Dreier, who provided his own colorful example.

¹⁴ To be fair to our critics, it should be said that their worries may have sprung from the highly plausible thought that any true proposition that one ought not to ϕ will entail a true proposition that it is not the case that one ought to ϕ (this seems correct, on the plausible assumption that there are no deontic conflicts). This is a different claim and we

permissible. Such a fact increases the probability both of the truth of the proposition that one ought to ϕ and of the truth of the proposition that one ought not to ϕ .¹⁵ It will not increase the probability of the truth of the proposition that it is not the case that one ought to ϕ . Thus, evidence that one ought to not ϕ is not always evidence that it is not the case that one ought to ϕ .

In the above example, then, the fact that if the runner were to stick his leg out, it would trip his main competitor is both evidence that he ought to do so, and evidence that he ought not to do so. It is not, however, evidence that it is not the case that he ought to do so. It is not the case, then, that on the view that reasons are evidence, we must accept the implausible idea that one fact can be both evidence for a proposition and evidence against the same proposition.

3.6

The following objection was suggested by Ralph Wedgwood. Consider Buridan's Ass. An ass stands between two equidistant identical bales of hay. There is reason for him to eat the left bale (e.g. it is nutritious) and there is reason for him to eat the right bale (it is just as nutritious). However, there is no evidence that the ass *ought* to eat the left bale or that he *ought* to eat the right bale. After all, there is absolutely nothing to choose between the two bales. It is thus simply obvious that it is not the case that he ought to eat the left bale and that it is not the case that he ought to eat the right bale. Rather, he ought to eat one or the other, but which one he eats does not matter. Therefore, there are reasons that are not evidence. RA is false.

We have two mutually incompatible replies to this objection. We wish to remain agnostic as to which reply is correct. First, it is not obvious to us that the ass lacks evidence that he ought to eat the left bale of hay. The left bale of hay is, after all, nutritious. Is this not then evidence that he ought to eat it? Of course, this evidence is not conclusive, but it is no part of our

are not opposed to it. The example that follows establishes that evidence is not closed under this entailment.

¹⁵ To see why this must be so, notice that the relevant probability space can be completely divided into three non-overlapping sectors: the probability of the truth of the proposition that the action is merely permissible, the probability of the truth of the proposition that one ought to do the act, and the probability of the truth of the proposition that one ought not to do the act. These three probabilities must add up to one, so when the first probability decreases, as it does when one receives new evidence that the act is not merely permissible, then the other two probabilities must both increase (assuming this new evidence is not also evidence that the act is obligatory or evidence that it is forbidden).

view that reasons must be conclusive evidence for oughts. Secondly, it is not obvious to us that the ass has a reason to eat the left bale of hay. The ass has a reason to eat either the left bale of hay or the right bale of hay, but does he have a reason to eat the left bale in particular? It is not clear.

Though we are undecided on these matters, we are committed to the idea that *if* the ass has a reason to eat the left bale then he has evidence that he ought to eat the left bale. In light of the fact that intuitions about this case are not firm, RA is not threatened.

3.7

Suppose you are a good person and you trust yourself to make good decisions.¹⁶ You know that at a future point in time you will do a particular act. The fact that you will do this act is evidence that you ought to do it (because you are a person who makes good decisions), but it doesn't seem like a reason to do it. How could the fact that you will do something be a reason to do it?

How should we respond to this objection? There are two kinds of case that would fit the preceding description: there are cases where one is aware one will do a particular act and one is aware of possessing sufficient reasons for performing the act that are separate from the fact that one will do the act; and there are cases where one is aware one will do a particular act but one is not aware of (sufficient) reasons for performing the act that are separate from the fact that one will do the act. About each kind of case we wish to say different things. If one is aware of the sufficient reasons of the standard kind to do an act, then the fact that one will do the act in the future is not evidence that one can *have*, even though it *is* evidence that one ought to perform the act, and (thus) the fact that one will do the act in the future is not a reason that one can *have*, even though it *is* a reason to perform the act. On the other hand, if one is not aware of having sufficient reasons of the standard kind that one will do an act, then the fact that one will do the act in the future is evidence that one can *have*, and (thus) the fact that one will do the act in the future is a reason that one can *have*.

A different example that will quickly illustrate why one should accept this view is a case where you provide your name (John, say) to another person then hear that person utter your name soon afterwards. Clearly the later uttering of the name is evidence that you are John and it is a reason for a bystander who has never met you to believe that you are John. Would you

¹⁶ We owe this example to Peter Graham. The objection that follows was also pressed upon us, in one form or another, by a number of other people with other examples. Adam Pautz was the first to raise an objection of this kind.

take it to be a reason to believe you are John? Not ordinarily. Ordinarily, your particular epistemic state of knowing to a high degree of certainty that your name is John would stand in the way of being able to think of the utterance of your name as evidence for yourself (i.e. evidence that you only now *have*) that your name is John, or as a new reason to believe that your name is John. A situation in which you could accept this utterance as new evidence and a new reason to believe is one where you are suffering from amnesia. What explains why it is not possible to have evidence/a reason in one case, but it is possible to have evidence/a reason in the other case? We believe it is the apparent *independence* of pieces of evidence in one case and not in the other. If a person immediately repeats your own name to you without anything strange having happened to your memory in the interim, the relevant utterance is blocked from being considered as a reason because it is not viewed as evidence independent of evidence that you already possess.

If we now return to the original example we can see that the very same story that we just told about reasons for belief can be told about reasons for action. If one is aware of the ordinary reasons on the basis of which one will do a particular act at some point in the future, then the mere fact that one will do this act is not in itself evidence concerning what one ought to do that one can *have*, and it is (therefore) not a reason that one can *have*. It is still a reason to act. Is this a bullet that we are biting? That this question should be answered in the negative becomes clear if one imagines a situation where one suddenly forgets all the ordinary reasons that speak in favor of doing a particular act but where one does not give up an intention to do that act and hence still believes that one will do it. Remember that we are assuming that one believes oneself to be trustworthy (if this is not granted, then there is no reason to accept that the fact that one will do a particular act is evidence that one ought to do it). It seems quite right that one would therefore take the fact that one is going to do a particular act as evidence that one ought to do the act, and as a reason to do the act, and we hope that the reader will share this judgment on reflection.¹⁷

3.8

Sean McKeever and Michael Ridge effectively deny RA. They argue that “For a fact to be a reason for action cannot plausibly be understood in

¹⁷ Consider the following advice: Do what the virtuous person would do in your situation. The fact that a virtuous person would do a certain act is in fact a reason to do this act. This is similar to what is going on in the above case. If one knows oneself to be virtuous, then the fact that one will do a certain act is a reason to do it.

terms of that fact's reliably *indicating* anything, not even reliably indicating that the action ought to be done" (McKeever and Ridge 2006: 127), and they take it that this implies that reasons for action are fundamentally different from reasons for belief, because they can not be understood in terms of evidence. They provide an example that is meant to establish this conclusion (also from McKeever and Ridge 2006: 127):

Let us suppose that the fact that an action would give someone pleasure is typically a reason in favour of the action. This in no way entails that the fact that an action would give someone pleasure reliably indicates the rightness of the action. To see that this entailment does not hold, simply imagine a world in which whenever one person gives another person some pleasure a perverse demon causes (and is known to cause) great pain to some innocent person. Here we have an illustration of how being a reason for action and being a reliable indication of the rightness of an action can dramatically come apart.

The example is supposed to work like this. In a certain possible world, an action, ϕ -ing, would give somebody pleasure. This is a reason to ϕ . However, it is not evidence that one ought to ϕ in this world because it does not reliably indicate that one ought to ϕ . This is because whenever someone gives another person pleasure, a demon causes someone else great pain. Thus it is never the case in this world that one ought to give someone pleasure. Therefore, the fact that ϕ -ing would give someone pleasure does not reliably indicate that one ought to ϕ .

We have two responses to this objection. The first point has been mentioned before. McKeever and Ridge's argument relies on the idea that a fact is evidence for the truth of a (contingent?) proposition only if it reliably indicates this proposition. In our discussion of the Normative Principles Argument, however, we mentioned that a fact can be evidence for the truth of a proposition even if this fact does not reliably indicate this proposition. Even McKeever and Ridge say that "the only obvious and serious difficulty here is with reasons for believing necessary truths, which arguably are equally well reliably indicated by everything" (2006: 127). As we have seen above, facts can be evidence even for contingent propositions without reliably indicating these propositions. The phonebook case is such an example. We thus reject this premise of their argument.

McKeever and Ridge's argument also relies on the idea that the fact that ϕ -ing would cause pleasure is a reason to ϕ . However, one may dispute this. As McKeever and Ridge say, this is merely 'typically' a reason to ϕ . This implies that it is merely contingently a reason to ϕ . But if it is merely contingently a reason to ϕ , then surely the demon world is one world at least in which it is no reason at all to ϕ (if not that world, which?). We conclude, therefore, that McKeever and Ridge's argument fails.

This brings our discussion of arguments against RA to an end. We hope to have persuaded you that RA and R are plausible statements of how reasons and evidence are related. Like many other philosophers, McKeever and Ridge believe that no unified account of reasons for action and reasons for belief can be provided. They claim that while it is fine to view the world as containing a very large number of reasons for belief, it would be very odd to think that we are surrounded by a similarly large number of reasons for action (2006: 124–37). We do not think this is odd at all.¹⁸ Rather, we think it is very odd that the idea that there may be an underlying unity to reasons for action and reasons for belief—a unity that is capable of being analyzed (and seems to us to be best explained by the thesis that reasons *are* evidence)—has not been given more of a run for its money.

REFERENCES

- Broome, John (2004) 'Reasons,' in J. Wallace, M. Smith, S. Scheffler, and P. Pettit (eds.), *Reason and Value: Themes from the Moral Philosophy of Joseph Raz* (Oxford: Oxford University Press), 28–55.
- Kearns, Stephen, and Star, Daniel (2008) 'Reasons: Explanations or Evidence?' *Ethics* 119(1).
- Kolodny, Niko (2005) 'Why be Rational?' *Mind* 114: 509–63.
- McKeever, Sean, and Ridge, Michael (2006) *Principled Ethics* (Oxford: Oxford University Press).
- Parfit, Derek (2007) *Climbing the Mountain*, unpublished MS.
- Scanlon, T. M. (1998) *What We Owe to Each Other* (Cambridge, MA: Harvard University Press).
- Star, Daniel (2007) 'Review of S. McKeever and M. Ridge, *Principled Ethics*' *Notre Dame Philosophical Reviews*, <http://ndpr.nd.edu/review.cfm?id=9203>.
- Williamson, Timothy (2002) *Knowledge and its Limits* (Oxford: Oxford University Press).

¹⁸ The view that this claim is not really that odd is discussed in Star (2007).