

12

Reasoning with Reasons

Daniel Star

When considering the nature of normative reasons, it is natural to start by attempting to identify truisms or platitudes that may inform our efforts to develop a theory of reasons. One such truism is that reasons are at least typically *pro tanto* considerations (this is how they are to be distinguished from facts about what we ought to do). Another is that they have the potential to make right or wrong. Yet another is that they play a role in deliberation or reasoning (broadly construed), in guiding agents to various conclusions or endpoints (beliefs, intentions, actions, etc.). Some accounts of reasons, although they aim to respect such truisms—at least in general terms (it is not *prima facie* unreasonable to assert of any or all of these claims that they are generalizations that admit of exceptions)—also aim to go beyond them. Other accounts take one of these claims to be fundamental, and attempt to bring the truism in question sharply into focus, in order to make it the basis of a more precise analysis. Here I will compare one of each kind of account. The focus throughout will be on the third truism mentioned above: whatever else they are, reasons are things we reason with. An important feature that both of the accounts discussed here share is that the authors who defend them are interested, from the get go, in attempting to provide a unified account of normative reasons that applies to both practical and epistemic reasons, and not just an account of reasons for action. This distinguishes these theorists from a number of earlier theorists.

Conor McHugh and Jonathan Way (2016, Forthcoming), Kieran Setiya (2014), and Jonathan Way (2017) all contend that normative reasons, whether reasons for belief or reasons for action, are premises of good reasoning. I agree that one important *role* that reasons play is that they are premises of good reasoning, but I do not think that in saying this we get at the essence of normative reasons. It makes good sense to ask whether there is some more fundamental claim we might make about the nature of reasons. I will argue that when we raise this

question in a certain way we are led to accept the view that reasons are facts that are evidence concerning what one ought to do (in the case of reasons for action), or ought to believe (in the case of reasons for belief). I focus, in particular, on the precise account of reasons provided by Way (2017). I do this both in order to simplify the discussion—it is difficult to deal with several versions of this approach to reasons at the same time—and because Way raises a problem for his account of reasons that I take to be illuminating more generally. He proposes an interesting solution to this problem, but I argue that his solution is not satisfactory. It would appear that this problem cannot be satisfactorily solved by a *reasons as premises of good reasoning* account of reasons. However, it can be solved by supplementing Way’s account with my own preferred account of reasons, *reasons as evidence*.

1 Reasons and Reasoning

I have defended the following analysis of what it is for a fact to be a reason, often with a coauthor (Kearns and Star 2008, 2009, 2013, 2015; Star 2010, 2011, 2015).¹

R: Necessarily, a fact F is a reason for an agent A to ϕ *if and only if* F is evidence that A ought to ϕ .

This analysis, which is at the heart of *reasons as evidence*, explains what it is for a fact (or true proposition) to be a reason in terms of that fact being evidence for the truth of an ought proposition. In the case of reasons for action, “ ϕ ” is to be replaced by a particular type of act (e.g. “keep a promise”). In the case of reasons for belief, “ ϕ ” is to be replaced by “believe p ” (or “suspend belief in proposition p ”), for some p . One might also say that any fact that is evidence that I ought to believe some proposition p is evidence that I ought to believe p *because* it is evidence that p ; if there are no pragmatic reasons for belief, an explanation of this kind will always be appropriate.

A simple example will help in understanding R, as it applies to reasons for action. It will be useful to return to this example later.

You are hurrying to meet a friend, Jed, whom you promised to meet in a few minutes for coffee, and you come across a very sick stranger who needs your assistance to get to a hospital. It is apparent to you that providing such assistance will take quite a bit of time, so it is not possible to both help the

¹ In the first half of this section, some material from section 3 of chapter 1 of Star 2015 is repeated.

stranger and keep your promise to Jed. On the basis of considering the reasons that apply to you, you decide to help the stranger.²

A natural way of describing your decision to help the stranger in this situation is to say that you recognize that there is a reason to meet your friend, as well as a reason to help the stranger, and you judge that the reason to help the stranger is a stronger or weightier reason than the reason to meet your friend. Another way of describing your situation is to say that you recognize that there is evidence that you ought to meet your friend and evidence that you ought to help the stranger, and you judge that the evidence that you ought to help the stranger is stronger or weightier than the evidence that you ought to meet your friend. *Reasons as evidence* has it that the relevant claims about evidence and reasons are extensionally equivalent. Furthermore, it is an advantage of *reasons as evidence* that it can explain something that is mysterious on many other accounts of reasons—the *strength* or *weight* of reasons—in terms of something significantly less mysterious, viz. the strength of evidence, which is itself perhaps best understood in terms of epistemic probability (Kearns and Star 2008, 44–5; 2009, 231–2).³

We can speak of it being the case that there “is” a reason, just as we can speak of it being the case that there “is” evidence. Similarly, we can speak of it being the case that a person “has” a reason, just as we can speak of it being the case that a person “has” evidence. If a stranger outside my front door urgently needs my help, but I am inside and unaware of that fact, there is still a reason for me to help the stranger (there is evidence that I ought to help the stranger), but this is not a reason (evidence) that I possess. There may well be general constraints on which facts can count as unpossessed evidence, and such constraints as there are will also be constraints on unpossessed reasons. Plausibly, a fact *F* can only be evidence for me that *p* if an *idealized* counterpart of myself would possess this evidence (that is, be aware of *F* and recognize that it is evidence that *p*).

Although for most purposes it is sufficient to focus on **R** as a claim about a necessarily true *biconditional*, I think that what, in fact, explains the truth of **R** is

² This example is very similar to one provided by Ross in his classic discussion of *prima facie* duties (2002, 18).

³ Our precise claims were “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***” (Kearns and Star 2009, 232).

not some further truth; rather, I take it that *reasons as evidence* is ultimately best understood as an account of the *real definition* of a normative reason.⁴ According to *reasons as evidence*, thus understood, all reasons for action *are* evidence that I ought to do something in particular.⁵ Going beyond this core analysis, we might additionally suppose that some of the facts that are reasons *also* play a fundamental right-making role (or wrong-making role); they *make it the case that* particular acts are right (or wrong). These *fundamental* reasons would be the reasons that we aim to specify in normative ethical theories, along with universal principles concerning what we *ought* to do.

Way defends the following “constitutive account” of what it is for a fact to be a reason.

R1: For the fact that p to be a reason for S to ϕ is for there to be a good pattern of reasoning from the [correct] belief that p , perhaps together with other correct attitudes which S has, to ϕ -ing. (Way 2017, 254)

In order to understand **R1**, as Way intends it to be interpreted, a few key things need to be said. Regarding *reasoning*, it is important that we not restrict ourselves to a narrow conception of reasoning, such that only conscious, calculative thinking processes count as instances of reasoning; rather, reasoning is to be understood as “any psychological process by which we come to form, revise, or sustain an attitude . . . in light of some consideration” (Way 2017, 252).⁶ **R1** does

⁴ As we said in Kearns and Star (2009, 219), “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.”

⁵ Gert (2016) suggests an addendum to *reason as evidence*. He thinks that it is important to distinguish between *justifying* reasons and *requiring* reasons, since, on some occasions, reasons may justify my doing any of a range of permissible acts, none of which I am required to do. He suggests that the defender of *reasons as evidence* can very naturally capture this distinction by claiming that facts that are *evidence that it is not the case that one ought not ϕ* are also reasons, of the justifying type. I think it might well be fine for defenders of *reasons as evidence* to accept this friendly suggestion. However, I am not sure that we need to. So far as supererogatory acts are concerned, it is possible to distinguish between what is morally *required* and what one ought to do, and specify that an act is supererogatory if the total evidence that it is morally required is weaker than the total evidence (with respect to morally relevant considerations alone) that one ought to do it (all evidence that an act is morally required is also evidence that one ought to do it, but the reverse does not hold). And, so far as cases where it might seem like it is not true to say that one ought to do any particular act are concerned, we might hope to capture the relevant intuitions by focusing on rational appropriateness or correctness, rather than reasons; often one does not have sufficient evidence that one ought to do anything, but it is still rationally appropriate (or correct) to do the act that the strongest evidence that one ought favors (of course, this act may be an act of gathering more evidence).

⁶ Way includes, in this definition, “for a reason” (he also adds “in light of some consideration”). Given that he wishes to understand reasons in terms of reasoning, this might suggest a vicious circularity. However, he is careful to avoid circularity by indicating he means anything one takes, perhaps wrongly, to be a reason (Way 2017, 252).

not say that reasons are premises in just any process that might count as reasoning (we are only too familiar with fallacious forms of reasoning, after all, and the objectivity of reasons cannot be secured without excluding them); rather, the focus is on *good* patterns of reasoning. Patterns of good reasoning are abstract structures, not instantiated by agents, so **R1** alone will not tell us when a reason is a reason *for an agent*; in order to do this, it will be necessary to relate patterns of reasoning to psychological features of particular agents (253). The way to do this is to restrict the account to reasoning to and from *correct* attitudes. If a glass has petrol in it, then my false belief that it contains gin does not correspond to a reason (fact); and, if I ought not murder, then my intention to do so neither provides nor corresponds to a reason to buy a weapon, as a means to the end of murdering someone (253–4).

As these examples bring out, we are talking about an account of *objective* reasons, and not subjective or motivating reasons. It is worth underscoring that *goodness* and *correctness* are both concepts essential to this analysis (much as the concepts of *evidence* and *ought* are essential to the analysis provided by *reasons as evidence*). Finally, Way asserts that “**R1** . . . takes the correctness of attitudes and the goodness of reasoning to be prior to reasons” (Way 2017, 254). This is, of course, essential if we are to take **R1** to be providing us with an analysis of what reasons are.

This *reasons as premises of good reasoning* account of reasons is very general in scope—it is not meant to be restricted merely to reasons for action. If you correctly believe that *p*, and you could use modus ponens (a good pattern of reasoning, if ever there was one) to conclude *q* on the basis of *p* and *if p then q*, then *p* is a reason to believe *q*. If you correctly believe that a person in front of you is in urgent need of medical assistance, and there is a good pattern of reasoning that will lead from this belief to an intention to do what you can to save their life, then the fact that the person in front of you is in urgent need of medical assistance is a reason to form the said intention. And if you correctly intend to do what you can to save someone’s life, and you correctly believe that phoning for an ambulance is a necessary means to saving their life, then you have a reason to phone for an ambulance. These examples are provided merely for the sake of illustration. The defender of *reasons as premises of good reasoning* will ultimately need to say more about what correctness and good patterns of reasoning amount to, with respect to different attitudes (as Way 2017 and McHugh and Way Forthcoming endeavor to do).

Notice that both *reasons as evidence* and Way’s account of reasons share a very interesting property, a property that we should all take to be a virtue in an account of reasons. They are both highly general accounts of reasons that purport

to provide a unified *and* informative analysis of what normative reasons are. Some accounts of normative reasons are unified and some are informative, but few are both. Reason primitivists, like Scanlon and Parfit, do *not* think it is possible to provide a true, informative analysis of normative reasons. Since they think of the concept of a reason as basic, they are in a fine position to think of reasons for belief and reasons for action as all normative reasons, in the exact same sense. They need not posit any disunity to reasons in general. Humean accounts of *practical* reasons are, on the other hand, informative. The Humean may explain practical reasons in terms of the desires of agents. However, it is relatively implausible to claim that reasons for belief are to be explained in terms of desires, so the Humean is likely to try to explain reasons for belief in a quite different way, e.g. in terms of evidence. This suggests there is an ambiguity in the term “reason” as used by Humeans, since it is now unclear why reasons for action and reasons for belief are all, well, reasons. The Humean, so construed, gives up on unity in return for informativeness. Way and I, in our different ways, need sacrifice neither unity or informativeness, since our accounts of reasons clearly explain what reasons are in other terms, yet they are also designed to apply to reasons for belief and reasons for action (as well as other types of normative reasons, such as reasons to care, reasons to intend, etc.).

I will focus on a problem with the *reasons as premises of good reasoning* view below. Before I do that, let us consider the possibility that **R**, before it is interpreted as a statement providing a real definition of normative reasons, and **R1** (or something close to it) might be compatible. If they are compatible, we can ask which of them provides the more fundamental analysis of reasons. I suggest that if **R** and **R1** are both true, then it is **R** that provides the more fundamental analysis of reasons.

The following seems like a perfectly intelligible and sensible question to ask: what *type* of fact or proposition is best suited to be a premise of good reasoning, quite generally? In pursuing an answer to this question, we might perhaps find that no one type of fact plays this role, but we should not begin by assuming that this is the case.⁷ I contend that there is one thing that plays the role: facts (true propositions) that are *evidence that one ought*.

⁷ Way (2017, 254) makes use of a relevant distinction here: a distinction between a *constitutive* account of the reason-relation (which he takes himself to be providing), and *substantive* accounts of the different types of considerations that might stand in this relation. He suggests that a substantive account of reasons for action might be very different than a substantive account of reasons for belief (e.g. it might, in general, be good reasoning to form conclusions about what one ought to do on the basis of considerations of utility, but not good reasoning to form conclusions about what one ought to believe on the basis of considerations of utility). So while there may be unity on one level, there

2 The Options

Suppose that we ask: what sort of thing must reasons be for them to be well suited to feature as premises in practical reasoning (or, in other words, for it to be appropriate to rely on them when engaging in practical reasoning)? I think they must be facts that are evidence concerning what one ought to do. My strategy here will be to consider a number of alternatives to *evidence that one ought*,⁸ and argue, in relation to each of them, that *evidence that one ought* is a better candidate for realizing the role in question.⁹ I will make use of a substantial, but, I think, thoroughly realistic assumption: many ordinary, conscientious agents are ignorant of the fundamental right-makers and wrong-makers (*qua* reasons or right-makers) that it is the job of normative ethics to specify (see Star 2015).

It might be said that it is nonetheless rationally appropriate for all agents to respond to (1) whatever the fundamental right-makers and wrong-makers (and permissibility-makers) are that are specified in the correct normative ethical theory.¹⁰ If one were in a position to know or rely on such a theory, one might use it to work out what one should do, then do it (or work out what it is permissible to do, then pick one of the things it is permissible to do). But if one is ignorant of this theory—a position we seem to all be in—then one doesn't know what the fundamental right-makers and wrong-makers are, and one doesn't know the principles by which one might otherwise work out what one

will be disunity on the other. I think there is a *deeper unity* to reasons than this, since both practical and epistemic reasons are evidence that one ought to ϕ .

⁸ In all cases where I use the phrase “evidence that one ought” in this way, I intend to be talking about a class of *facts*, understood as true propositions.

⁹ In this section, I repeat some material from section 6 of chapter 2 of my monograph (Star 2015); this material is put to a different use here.

¹⁰ The reader might wonder why I don't consider the option that ordinary agents are able to respond to *non-fundamental* right-makers. This option is not obviously false, but it is wholly unhelpful in the present context, as one might harmlessly redescribe the aim of this section as being one of deciding what facts get to count as non-fundamental right-makers, in the absence of knowledge of fundamental right-makers (agents who lack such knowledge are not in a position to rationally determine what the non-fundamental right-makers are by way of an inference that begins with the fundamental right-makers). I sometimes discuss fundamental right-makers (and wrong-makers and permissibility-makers) without referring to them as reasons. I actually take them to also be reasons (in Star 2015, I call them *fundamental reasons*), in virtue of the fact that they too are *evidence that one ought*. What I wish to emphasize here is that, assuming normative ethics is a viable research program and the fundamental right-makers are likely to be quite surprising to ordinary reasoners, this evidence—the evidence provided by fundamental right-makers—is not normally possessed by any agents (or, at least, much of it isn't).

ought to do. It seems very odd to suppose that when, due to ignorance, one is not able to respond to fundamental right-makers, it is rationally appropriate, in any case, to respond to them, but I suppose one might be misled into thinking that this is true *if* one were to only pay attention to reasons for action, and ignore reasons for belief. It might sometimes very luckily be true that it is possible for agents to act in conformity with fundamental right-makers in a condition of ignorance (and without taking up the alternative of responding to evidence concerning what they ought to do) *if* they were to constantly act in ways that, by their own lights, they don't have good reason to. However, this is clearly to act in ways that are epistemically irrational.

Generally speaking, intentionally acting in ways that one does not take oneself to have reason to act will violate requirements of rationality with respect to the beliefs and intentions that precede one's acts—roughly speaking, our intentional acts are explained by our intentions, and deliberately forming intentions that are contrary to what one believes one ought to do (or believes it is permissible to do) will violate a requirement of rationality. This is enough to take this alternative off the table. The discussion of this option has, however, brought out one important *desideratum* for the rest of our survey of the alternatives to *evidence that one ought*, and that is that we must keep in mind that, very generally, agents will be in a condition of ignorance of the correct ethical theory and the fundamental right-makers and wrong-makers specified therein.

What might some other alternatives be to responding to evidence concerning what one ought to do as providing reasons, in a situation where one is not able to respond directly to fundamental right-makers (due to ignorance of them *qua* reasons)? Here are some: (2) an agent could respond to non-natural non-analyzable “reasons”; (3) an agent could respond to her desires; (4) an agent could respond to what is good; (5) an agent could do nothing; and (6) an agent could form, and act on, whatever beliefs about reasons will effectively lead her to respond to the fundamental right-makers (independently of whether such beliefs are true or well justified). No doubt this list is not exhaustive, but these seem to be the clearest alternatives. I will now argue that none of them are good alternatives. Of course, this strategy leaves me vulnerable to critics who might argue that some further alternative is preferable to any of these options, but that does not mean the strategy is a bad one to employ here.

Imagine an ethically conscientious agent attempting to determine what to do, on some particular occasion. I suggest we do this not because we are trying to provide an account of reasons for conscientious agents in particular, but just because if the options we consider are not good alternatives for conscientious agents to evidence concerning what one ought to do, then that will be enough to rule them out as good

alternatives in general. We should imagine that this agent is also either very uncertain that she has cognitive access to any relevant fundamental right-makers, appreciated as such (during the window of time that is open for deliberation), or even justifiably believes that she does not have such access. This is necessary to avoid the agent confusing derivative and fundamental right-makers (of course, I don't think that ordinary agents usually classify reasons into these categories, but that doesn't matter for our purposes). And we should imagine that it is in fact true that this agent doesn't have any straightforward cognitive access to any relevant fundamental right-makers, understood as such (or understood to be fundamental reasons). The focus here is not simply on ignorance of ethically relevant non-ethical facts, but on ignorance of *which* facts are fundamental right-makers (reasons). Let me further specify that the agent in question is aware that one option on the table is for her to respond to evidence concerning what she ought to do—she is to compare the alternatives to this option.

So, now our agent considers (2) reasons taken to be non-natural, unanalyzable entities (reason relations). These are the reasons of the reason primitivists (especially Scanlon and Parfit). I will call them reasons* (doing this does not beg the question, since, merely referring to such reasons, if they exist, as reasons* does nothing to rule out the possibility that reasons* just are reasons). Suppose that when our agent is deliberating she can either pay attention to some reasons* *or* to evidence concerning what she ought to do (the reasons* being non-identical to the evidence concerning what she ought to do, we are assuming for the moment). Would she not reason as follows? It seems clear that if the evidence that what I ought to do comes apart from some purported reasons (reasons*), I should attend to the first before I attend to the latter (and if I can only attend to one, I should attend to the first). What could be a more reliable guide to acting as I ought than evidence concerning what I ought to do? Why instead trust a primitive metaphysical entity or relation that one supposedly is in a position to respond to, if it is *not* the best candidate for a guide that will lead one to act as one should? Now, one might have direct cognitive access to these primitive metaphysical entities (relations), but in this case one would have evidence concerning what one ought to do; reasons* would not then appear to be a practical alternative to *reasons as evidence*. In any case, reason primitivists should avoid claiming (and usually do avoid claiming) that *evidence one ought* always coincides with reasons*, since this will undermine the claim, essential to this view, that it is not possible to provide an informative analysis of reasons.¹¹

¹¹ Admittedly, there is room in logical space for the view that reasons* and *evidence one ought* are necessarily coextensive, but what would be the motivation at this point for claiming that no

Let us next consider option (3), the possibility of an agent responding to her desires in a way where she can draw the inference that she ought to F (or, at least, that by F-ing she will be getting as close as she can to acting as she ought) on the basis of her desires to achieve the end that F will lead to her achieving (which, we may suppose, are stronger, as a set, than her desires for any other end). The problem is that, as a conscientious agent, she will be aware that desires are often irrational or biased states, and certainly not always as they ought to be—it would be one thing if she were to respond to desires that she ought to have, but she is aware that, since she doesn't have access to any fundamental right-makers (i.e. fundamental reasons), she doesn't have access to any fundamental reasons to desire.

Two possible thoughts might next occur to this agent: (i) she might think she is a virtuous or fully rational person, so her desires are reliable (having been shaped to avoid bad biases, etc.); or (ii) she might consider what she would desire *if* she were virtuous or fully rational. The first of these options is certainly not going to provide a general alternative to *evidence that one ought*, since it will only apply to the virtuous or fully rational agent (assuming that those who know they are virtuous are allowed to infer that they ought to F on the basis of desiring to F). In any case, it collapses into the view that she should respond to evidence concerning what she ought to do. If she asks herself why she should follow her desires, her thought that she should do so because she is virtuous or fully rational can be understood as resting on the thought that facts about the desires of the virtuous or the fully rational are good *reliable indicators* (i.e. evidence) concerning ought-facts—if they were not reliable in this way, why ought they be followed?

Alternatively, if the agent we are imagining is not virtuous, but is considering trying to aim for what she would desire if she *were* virtuous, or fully rational, she will need to examine *evidence* for thinking that her idealized self would desire the end that requires her to F. Now she might think of such evidence in either of two ways: (a) she might think her idealized self would be in a position to do as she ought by responding appropriately to fundamental right-makers, perhaps partly in virtue of having consistent and coherently unified desires (Smith 1994), or (b) she might think her idealized self would not be in such a position, but she would at least be fully rational (her non-ethical beliefs would be impeccable and her desires would be consistent, coherently unified, etc.). If she accepts option (a),

satisfactory, informative analysis of what it is for something to be a reason is available (once an obvious candidate is thought to be always present)? One motivation might be to hold on to the thought that there is something distinctively non-natural about the normative realm, but *reasons as evidence* is consistent with the claim that *ought* facts are non-natural normative facts, and consistent with the claim that fundamental right-makers are non-natural normative relations, so the non-naturalist desideratum could still be satisfied in other ways.

it seems that, in her present position of ignorance of the fundamental right-makers, she should think that evidence concerning what her idealized self would want her now to do *just is* evidence concerning what she now ought to do (so this alternative option collapses into *evidence that one ought*). We have a ready explanation available here: the reason why one might think one should defer to one's ideal self is most plausibly thought of as being that one thinks one's ideal self is in a better position to *recognize* what it is important to do; that is, that she is better placed to respond well to the relevant evidence.

If, on the other hand, she accepts option (b), she is now left comparing evidence concerning what she ought to do (the option that, for the sake of comparison, we put on the table at the beginning) and evidence concerning what she would desire if virtuous or fully rational, where this, *ex hypothesi*, can come apart from being in a position to recognize fundamental right-makers. In this case, the option of acting on merely more consistent and coherent desires (or more virtuously shaped desires) will not seem like a good alternative to acting on evidence concerning what she ought to do, since such desires (or virtuous dispositions) could not be *more* reliable (in relation to the relevant ought-facts) than all her relevant evidence concerning what she ought to do.

The next option is to suppose that the conscientious agent who is ignorant of fundamental right-makers might reason about and act on the basis of (4) facts about what is good or best. The impetus for saying that agents are generally ignorant of fundamental right-makers came from the idea that ordinary agents are ignorant of normative ethical theory. This means we shouldn't take them to be committed to a view that all reasons are evaluative at base. They will, if appropriately epistemically cautious, be open-minded about issues that divide consequentialists and deontologists. Hence, if they have a choice between acting on evidence concerning what they ought to do and evidence concerning what is good (or best) in some context, if this evidence points in different directions, it seems appropriate for them to act on the evidence concerning what they ought to do. Having said this, I don't need to deny that evidence that an option is good (or will realize some good) is also evidence that one ought to select this option.¹² My point is that in the cases where these come apart it appears appropriate to act on *evidence that one ought*, and that this suffices to establish that *evidence that one ought* plays a more primary role so far as practical reasoning is concerned.

¹² I also do not need to deny that there are more general senses of "good," "best," and "better than" that we use when talking about deontic matters, e.g. when we say that an act that conforms with duty is better than an act that does not conform with duty.

Alternative (5) is to do nothing. It seems clear that, in general, doing something that is based on evidence concerning what one ought to do is rationally preferable to doing nothing. It *must* be preferable because it is not as though doing nothing is itself something that the evidence concerning ought needs to rule out—such evidence may, in fact, direct one to do nothing (or, more precisely, to wait until one has gathered more evidence, or until a certain event occurs, or to simply sit tight in order to play it safe). However, if a rational assessment of such evidence does *not* direct one to do nothing, surely it is preferable to do as the evidence suggests one ought.

There is now just alternative (6) to consider. This might appear to be an attractive alternative to some consequentialists, in particular. Some consequentialists (some traditional utilitarians, and some so-called “global” consequentialists) might think that agents should form, and act on, whatever beliefs about reasons will effectively lead them to respond to fundamental right-makers, independently of whether such beliefs are true or well justified.¹³ For our purposes here, we need not deny that this is sometimes true (it is not difficult to imagine cases of great moral import, for instance, where some degree of rational belief formation seems worth sacrificing, from the point of view of an agent considering making such a sacrifice). It is enough to point out that doing one’s best when it comes to being guided by reasons is not the same thing as doing what is actually best: conforming with reason is fundamentally about us doing our rational best to follow the reasons, both epistemic and practical, that apply to us, in order to get as close as we can to always doing and believing what we should do and believe, but there is no guarantee that such guides to right action and belief—despite being the most reliable guides that we have—will always lead us to the promised land.¹⁴

3 Weighing Reasons

Let me now return to Way’s role-based account of reasons. Recall that the main statement of this view is the following.

R1: For the fact that p to be a reason for S to ϕ is for there to be a good pattern of reasoning from the [correct] belief that p , perhaps together with other correct attitudes which S has, to ϕ -ing.

¹³ Plausibly, *most* of the time consequentialism would direct people to form beliefs in a standard rational fashion, since it is extremely plausible that our rational belief-forming capacities produce more good in the long term than attempts to circumvent or undermine them would (especially since too much circumvention will lead to some long-term undermining of these capacities).

¹⁴ And if there are situations where we can and should decide to be irrational for a good end, the leap into irrationality is a leap where we at least risk leaving these guides (reasons) behind.

As Way (2017, 257–61) recognizes, the *reasons as premises of good reasoning* account faces a problem: many reasons are “outweighed” (it is relevant that he uses this term himself), and when they are outweighed it will *not* be good reasoning to reason from the relevant premise to the action that this premise (purportedly a reason) favors. Recall the earlier example. On the assumption that when faced with a conflict between a promise to meet a friend for a casual conversation over coffee and a duty to assist someone in urgent need of help, one ought to do the latter, a person in this situation who reasons from the fact that he promised to meet his friend to the action of meeting his friend is intuitively engaging in *bad* reasoning. So, according to R1, there is no reason at all in this situation to keep the promise in question. But this is clearly false. All cases of this general kind appear to be counterexamples to R1.

As a solution to this problem, Way (262–3) proposes that good patterns of reasoning be understood to be *defeasible*. It is good reasoning to reason from your belief that you promised to meet your friend to doing so when there is no countervailing consideration, but it is not good reasoning to reason from your belief that you promised to meet your friend to doing so when you also believe that you will thereby leave someone to die; in the second case, it is instead good reasoning to reason from your belief that someone nearby is going to die without your help to a conclusion that is an act of saving them. This is an interesting solution, but I’m afraid it is not satisfactory. Let me discuss four problems with the solution.¹⁵

The first problem might be termed the problem of *weak reasons*. There may be some weak reasons that, even when they are not outweighed, it is not appropriate to act on. Way might deny that there are any such reasons for action. I grant that it is not clear that such reasons for action exist. Never mind, since it seems clear that there are weak reasons for belief that, even when they are not outweighed, it is not appropriate to form beliefs on the basis of. I can have very weak evidence that *p* and no or little evidence that not *p*. A friend with a poor, but still not completely unreliable memory says she seems to remember a stranger named Fred was at a particular party that I was unable to attend last year. This is clearly a weak reason to believe a stranger named Fred was at the party. Suppose I have not heard anything else about Fred and the party; in this case, I have no reason to believe a stranger named Fred was *not* at the party. Still, it is inappropriate for me

¹⁵ One relatively superficial problem I will not discuss is that not all good theoretical reasoning is *defeasible* reasoning; clearly some theoretical reasoning—deductive reasoning—is not. I take it Way’s response here would be to simply agree that not all good patterns of reasoning are defeasible, so that a complete account of good reasoning would include both defeasible and indefeasible patterns of reasoning.

to believe that this Fred was at the party merely on the basis of the weak reason I have to believe he was at the party. Way might reply that all cases where one might form beliefs on the basis of very weak evidence would be cases where bad forms of reasoning would be involved, but one would like to hear more about *why* we should think that this is true. It is natural to think that what explains *why* it would be bad reasoning to form a belief that Fred was at the party on the basis of my friend's testimony is that the evidence is too weak for it to be appropriate to reason in this way. Way cannot avail himself of this explanation.¹⁶

The second problem might be termed the problem of *bruteness*. In the light of the problem Way raises and his purported solution, one might ask: is there just a large number of *brute* facts about which particular forms of reasoning are good and which are bad? It seems natural to want more of an explanation of what all the good forms of reasoning have in common. It could be said that this is not so much a problem, as a request for more information (and one that McHugh and Way Forthcoming aim to address). In the case of practical reasoning, which we might take to involve the weighing of practical reasons, one natural way of providing more information here would be to say that good reasoning is simply reasoning that conforms with the correct substantive ethical theory. Conformity with the correct ethical theory would here involve good patterns of practical reasoning being understood to be patterns that lead from facts that the correct ethical theory would have us think of as basic reasons (or duties, or goods) to conclusions about what we ought to do that this theory would endorse. This response seems fairly natural, but it leads straight to a third problem.

The third problem is a problem of *needing normative ethical theory to understand good reasoning*. As I just suggested, one might think that the question of what good practical reasoning amounts to is an ethically *substantive* matter. This seems to suggest that it is the job of normative ethical theory to give us an account of good forms of practical reasoning. However, it seems preferable to many of us to have an account of good reasoning that is more general and does not require particular substantive ethical commitments—or far fewer of them, anyway.¹⁷

¹⁶ As a reader pointed out, Way could claim that there are, in fact, no weak reasons for outright belief, but that, in cases where it is tempting to say that there are, we should say instead that there are really just reasons for weaker credence states. I'm not sure why this view would need to talk of *reasons* for belief at all, rather than simply fittingness conditions for particular credence states, as the notion of reasons having weights would seem unnecessary at this point. In any case, this view would be revisionary, since we ordinarily take ourselves to have weak and strong reasons for outright belief: when the evidence is weak we have weak reasons, and when the evidence is strong, we have strong reasons.

¹⁷ McHugh and Way (Forthcoming) do provide an account of good reasoning: roughly speaking, it is reasoning that takes us from fitting attitudes to fitting attitudes. I do not mean to overlook that

Consider a conscientious person, ignorant of the correct ethical theory, trying to work out what she ought to do in a morally difficult situation (not the example we have been discussing, where it seems clear that one reason is stronger than the other, but a situation where it is difficult to judge which of two reasons is stronger). Suppose she reasons about a number of relevant facts. That is, she thinks very carefully about what she ought to do. As it happens she does the objectively wrong thing, failing, in fact, to act on the strongest set of reasons (not because she lacks some relevant non-ethical information, but because she doesn't know which facts provide her with fundamental normative reasons). It seems appealing to say that such an agent has engaged in good ethical reasoning, yet Way's solution to the problem he brings up seems to rule this out. We need an account of good ethical reasoning that is somewhat independent of ethical theory, and whatever the objective reasons turn out to be. One might also think that something like this is true of non-deductive theoretical reasoning concerning evidence: good reasoning of this kind can't be conceived of as simply taking us from correct inputs to correct outputs, where correctness is understood in terms of truth (since it is not strictly truth preserving).

Finally, there is a problem of *capturing ranges of weights*. We generally take reasons to have weights (or strengths), and it does not appear that this talk—or, indeed, phenomenology—of weights is *exhausted*, or fully explained by, discrete patterns of good reasoning.¹⁸ Consider a series of examples where one can either help A or help B relieve pain (and not help both A and B), and we are able to assign numbers to the amount of pain (perhaps directly measuring C-fibers firing) each person experiences. The weight of the reason to help A and the reason to help B seems to vary with, and be, at least in part, a function of, the degree of pain each person is experiencing. And this seems true even when we restrict our attention to a region where one clearly ought to help A or clearly

we are provided with this very general, informative claim about good reasoning (and I fully admit I have not explored their account in detail here): my point is that in order to go beyond this characterization of good reasoning and identify which exact forms of ethical reasoning count as good (i.e. actually preserve fittingness, on their approach), it will be necessary to do substantive normative ethics.

¹⁸ McHugh and Way (2016, 593, fn. 39) write, “we understand outweighing in terms of patterns of reasoning: roughly, where p is a reason to A and q is a reason not to A, p outweighs q if it is good reasoning to move from the beliefs that p and q to A-ing and not good reasoning to move from these beliefs to not-A-ing.” This approach to understanding outweighing is susceptible to the present line of criticism, precisely because it does not provide any story about degrees of closeness to the point where a reason of one kind will outweigh a reason of another kind (for an intuitive sense of what is meant by degrees of closeness to outweighing here, think about gradually increasing the physical weight of one item on a scale, while leaving the weight of the initially heavier object on the other side of the scale constant).

ought to help B. Reasons can increase in strength without this leading to a difference in what one ought to do, and Way's purported solution does nothing to explain this. *Reasons as evidence*, on the other hand, has an account of the weight of reasons that will apply here—as A's pain increases, the evidence that one ought to help him or her (rather than B) increases.

4 Conclusion: Knowledge as a Standard of Correctness

Let me conclude by responding to a concern the reader might have regarding my view that both Way's claim about the relation of reasons and reasoning and my own favored account of reasons are compatible, in such a way that it is really *reasons as evidence* that accounts for what reasons are. **R1** refers to reasoning both from and to *correct* attitudes (the "other correct attitudes" in **R1** intentionally implies that the belief that *p* is a correct attitude; Way takes correctness for beliefs to be a matter of truth). **R** does not rely on a concept of *correctness*, and I did not say anything about correct attitudes when sketching the *reasons as evidence* account of reasons, so one might worry that it is just not clear, one way or the other, whether **R** and **R1** are compatible. I will not attempt to provide a general account of correct attitudes here, for all attitudes that are used when reasoning,¹⁹ but I want to end by saying something about correctness for the true belief that *p* in **R1**. First, I need to introduce **R1***.

R1*: For the fact that *p* to be a reason for S to ϕ is for there to be a good pattern of reasoning from *knowledge* that *p*, perhaps together with other correct attitudes which S has, to ϕ -ing.

The only change here from **R1** is that the standard of correctness for beliefs is no longer truth, but (that the beliefs be) knowledge. I think it is actually **R1*** that is best combined with **R**, but **R1*** is very close in content to **R1**. It is certainly close enough to **R1** for us to assume that were Way to object to it, this would be because of a substantive disagreement about whether it is better to think of truth or knowledge as providing the right criterion for the correctness of beliefs, and not for some reason to do with the project of analyzing reasons in terms of premises of good reasoning *per se*.

¹⁹ However, it is predictable that I will use **R** as widely as possible. For example, a reason to form a particular intention is, on my view, evidence that one ought to form the intention (or, perhaps, evidence that it is permissible to form the intention, if one thinks Way is right to focus on permissibility when thinking about the correctness of intentions). An account of *correctness* for intentions will need to respect this thought.

Williamson (2000, 2005) and others (especially Hawthorne and Stanley 2008) have argued that knowledge is the norm, or standard of correctness for reasoning and action. In a chapter of my book, I argue that important benefits come from combining *reasons as evidence* with the view that knowledge is the standard of correctness for reasoning and action, and I argue that these two theses are natural partners (Star 2015). The standard or norm in question can be expressed as follows.

The Knowledge Norm: It is correct for *A* to treat *p* as a reason for action just in case *A* knows *p*.

This statement of the knowledge norm is close to identical to a number of formulations of the norm that one finds in the literature.²⁰ Notably, these formulations fail to distinguish between reasons that it is correct or appropriate to act on (or form intentions on the basis of), and facts that it is correct or appropriate to take to be reasons but also correct or appropriate to take to be outweighed by other reasons.²¹ It is notable that this weakness in the knowledge norm, as it is normally stated, is the exact same weakness that Way locates in his account of reasons. I think *reasons as evidence* provides the missing ingredient here.²²

Simply in order to make it clear how *reasons as evidence* can avoid the problem Way raises for his account of reasons, I will end with the contentions concerning correctness (or rational appropriateness) and practical reasons that I provide in my book (Star 2015, chapter 4). First, I take it to be correct to treat a fact merely as a reason (outweighed or not) just in case it is known *and* it is evidence that one possesses that one ought to do a particular act. Second, I take it to be correct to act on a reason to Φ just in case (a) one knows one ought to Φ , *or*, when one does not know one ought to Φ , (b) the evidence that one possesses concerning what one ought to do makes it most probable that one ought to Φ , relative to one's other options (or no less probable that one ought to Φ than do any other act).²³ Cases

²⁰ Williamson, for example, states the norm as "One knows *q* iff *q* is an appropriate premise for one's practical reasoning" (2005, 231).

²¹ Ichikawa (2012) rightly emphasizes that the knowledge norm does not tell us *which* bits of knowledge it is appropriate to act on and that this is why it seems vulnerable to certain putative counterexamples discussed by Brown (2008).

²² More precisely, I take it that I can provide *two* missing ingredients: I suggested in Section 1 that *reasons as evidence* has its own account of the weight of reasons; here I am focusing on the issue of simply distinguishing between when it is correct or rationally appropriate to take something to be a reason and when it is correct or rationally appropriate to act on a reason.

²³ This is a claim about when it is appropriate to act, and not about when it is appropriate to form a belief, since the option of suspending belief must be taken into account when considering the appropriateness, or otherwise, of belief formation, but the option of suspending belief has no analogue in the case of action.

of type (a) are those where the evidence that one possesses concerning what one ought to do will warrant drawing the conclusion that one ought to Φ , while cases of type (b) will be those where one is not warranted in concluding that one (objectively) ought to Φ , but where action is nonetheless required and one must thus do the best one can to follow the goal of acting as one ought to.

In all likelihood, this account of rational correctness will need to be improved upon or supplemented—this being what one should generally expect in philosophy—but it at least suggests that it may be possible to analyze *correctness* (using knowledge, evidence, and ought propositions, and with respect to rationality), rather than take it to be a basic concept. Like Way, I take notions of correctness and fittingness to be of genuine importance to ethics and epistemology. Unlike Way, I doubt any such notion corresponds to a fundamental normative property (cf. McHugh and Way 2016). One reason to doubt this is that fittingness relations can be truthfully referred to by someone who completely repudiates a practice or character state (e.g. certain sexist acts are fitting with respect to certain sexist practices or rules that we should in no way endorse; or, overeating is fitting with respect to the vice of greed). But this is a topic for another day.^{24, 25}

References

- Brown, Jessica (2008). “Knowledge and Practical Reason.” *Philosophy Compass* 3: 1135–52.
- Gert, Joshua (2016). “The Distinction between Justifying and Requiring: Nothing to Fear.” In Barry Maguire and Errol Lord (eds), *Weighing Reasons*. Oxford: Oxford University Press, 157–72.

²⁴ After recently considering Clayton Littlejohn’s critical response to my book, I have become convinced that the defender of *reasons as evidence* who is also a friend of the knowledge norm cannot actually accept R1* in the form stated above (see Littlejohn 2016 and Star 2016, for clarification). The closest claim of this kind that I believe he or she can accept is actually R1**: if there is a good pattern of reasoning from *knowledge* that p , then p is a reason for S to ϕ . Correctness with respect to the norms that govern beliefs (perhaps because they are constitutive of belief) might provide reasons, but, on this view, it isn’t the case that all reasons for belief align with such correctness conditions. An alternative strategy for the defender of *reasons as evidence* would be to give up on the knowledge norm and claim instead that correctness conditions specify that we should follow true justified beliefs. One might then accept R2*: for the fact that p to be a reason for S to ϕ is for there to be a good pattern of reasoning from a true justified belief that p , perhaps together with other correct attitudes which S has, to ϕ -ing.

²⁵ I would like to thank the editors for their helpful comments, as well as the philosophers who pressed me on various issues when I presented earlier versions of this chapter at the University of Southampton, New York University Abu Dhabi, Dartmouth College, Southern Methodist University, and the Central European University. Since I focus herein on a paper by one of the editors, let me note that it was never suggested to me that I do anything of this kind, and that the editors were not aware that I was interested in doing so when they invited me to contribute to this volume.

- Hawthorne, John and Stanley, John (2008). "Knowledge and Action." *Journal of Philosophy* 105: 571–90.
- Ichikawa, Jonathan Jenkins (2012). "Knowledge Norms and Acting Well." *Thought* 1: 49–55.
- Kearns, Stephen and Star, Daniel (2008). "Reasons: Explanations or Evidence?" *Ethics* 119 (1): 31–56.
- Kearns, Stephen and Star, Daniel (2009). "Reason as Evidence." *Oxford Studies in Metaethics* 4: 215–42.
- Kearns, Stephen and Star, Daniel (2013). "Weighing Reasons." *Journal of Moral Philosophy* 10 (1): 70–86.
- Kearns, Stephen and Star, Daniel (2015). "Weighing Explanations." In Iwao Hirose and Andrew Reisner (eds), *Weighing and Reasoning: A Festschrift for John Broome*. Oxford: Oxford University Press.
- Littlejohn, Clayton (2016). "Do Reasons and Evidence Share the Same Residence?" *Philosophy and Phenomenological Research* 93: 720–7.
- McHugh, Conor and Way, Jonathan (2016). "Fittingness First." *Ethics* 126: 575–606.
- McHugh, Conor and Way, Jonathan (Forthcoming). "What Is Good Reasoning?" *Philosophy and Phenomenological Research*.
- Ross, W. D. (2002). *The Right and the Good*. Oxford: Oxford University Press.
- Setiya, Kieran (2014). "What Is a Reason to Act?" *Philosophical Studies* 167: 221–35.
- Smith, Michael (1994). *The Moral Problem*. Malden, MA: Blackwell.
- Star, Daniel (2010). "Moral Skepticism for Foxes." *Boston University Law Review* 90: 497–508.
- Star, Daniel (2011). "Two Levels of Moral Thinking." *Oxford Studies in Normative Ethics* 1: 75–96.
- Star, Daniel (2015). *Knowing Better*. Oxford: Oxford University Press.
- Star, Daniel (2016). "Replies to Cuneo, Driver, and Littlejohn." *Philosophy and Phenomenological Research* 93: 728–44.
- Way, Jonathan (2017). "Reasons as Premises of Good Reasoning." *Pacific Philosophical Quarterly* 98: 251–70.
- Williamson, Timothy (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.
- Williamson, Timothy (2005). "Contextualism, Subject-Sensitive Invariantism and Knowledge of Knowledge." *Philosophical Quarterly* 55: 213–35.