

## What is a reason to act?

Kieran Setiya

© Springer Science+Business Media Dordrecht 2012

**Abstract** Argues for a conception of reasons as premises of practical reasoning. This conception is applied to questions about ignorance, advice, enabling conditions, “ought,” and evidence.

**Keywords** Action · Advice · Ought · Reasons · Reasoning · Evidence

A reason for action is a premise of practical reasoning. When someone acts on the ground that  $p$ , she reasons to action or intention from the proposition that  $p$ . Some authors reserve the use of “reasoning” for calculative activity, for thought that invokes normative concepts, or for the kind of deliberation that is itself intentional. None of these restrictions will operate here. In our artificially inclusive sense, any instance of doing something for a reason counts as reasoning. More broadly, practical reasoning incorporates any form of thought to which assessments of practical rationality apply.

These observations speak to what some call “explanatory” or “motivating” reasons.<sup>1</sup> This essay is not primarily concerned with them, but with normative or justifying reasons, considerations that count in favour of action. Normative reasons bear a different relation to practical reasoning. When a fact is a reason for  $A$  to  $\phi$  in the normative or justifying sense, it need not be a reason for which she acts; she may not even be aware of it. But the fact is a premise for sound reasoning to a desire or motivation to  $\phi$  whose further premises are available to  $A$ .<sup>2</sup>

---

<sup>1</sup> For conflicting theories, see Davidson 1963 on “primary reasons”; Smith 1987, Dancy 2000, and Alvarez 2010.

<sup>2</sup> I assume that normative reasons are facts in that they are true propositions, not worldly states that correspond to them: normative reasons must be as finely individuated as the premises of practical reasoning.

---

K. Setiya (✉)  
Department of Philosophy, University of Pittsburgh, Pittsburgh, PA 15260, USA  
e-mail: kis23@pitt.edu

This picture has been developed in different ways and in different idioms. Introducing a collection of essays on practical reason published in 1978, Raz appealed to “practical inference”: “The premises of a valid practical inference if they are all true (or justified) state a reason” (p. 5).<sup>3</sup> For Williams (1979), the pivot was “deliberation,” which he tied to the agent’s “subjective motivational set.” We can filter off the “internalist” and instrumentalist themes in Williams’ essay to state an abstract principle anyone can accept:

The fact that  $p$  is a reason for A to  $\phi$  if and only if there is a sound deliberative route from A’s psychological states, together with the belief that  $p$ , to the desire to  $\phi$ .<sup>4</sup>

In *Reasons without Rationalism* (Setiya 2007), the framing concept is that of practical thought and the dispositions that govern it. Assuming that a good disposition of practical thought is one that is good, as such—a disposition to engage in good practical thought—the idea of a reason to act as a premise of sound reasoning comes to this:

*Reasons*: The fact that  $p$  is a reason for A to  $\phi$  just in case A has a collection of psychological states, C, such that the disposition to be moved to  $\phi$  by C-and-the-belief-that- $p$  is a good disposition of practical thought, and C contains no false beliefs. (Setiya 2007, p. 12)

Some notes of clarification. First, this principle is concerned with *pro tanto* reasons, reasons that can be outweighed; accordingly, it speaks of being moved, not of acting or intending.<sup>5</sup> Second, the exclusion of false beliefs echoes Williams on the soundness of deliberative routes and Raz on the need for true premises in practical inference that corresponds to reasons. A related point: the relevant collection of psychological states typically will not include the whole array of A’s beliefs. If one has false beliefs, one can still have reasons, but the reasoning to which they correspond will be free from error. Third, good practical thought does not include redundant premises, ones that make no difference to the case for doing  $\phi$ . If it did, irrelevant facts would count as reasons. Fourth, apart from non-redundancy, *Reasons* does not support any particular view of what is good in practical thought. It is not implicitly instrumentalist. Nor does it claim that when the fact that  $p$  is a reason for A to  $\phi$ , A is capable of being moved to  $\phi$  by the belief that  $p$ . It is thus consistent with, though it does not entail, “externalism”

<sup>3</sup> Raz treats the conclusion of practical inference as a deontic proposition—that one ought to  $\phi$  relative to these considerations, or that there is reason to  $\phi$ —not as action, intention, or desire (Raz 1978, pp. 5ff). The formulation in the text abstracts from this dispute.

<sup>4</sup> Like Raz, Williams takes practical inference to involve the belief that one has reason to  $\phi$  (Williams 1979, pp. 104, 107–108), though his principal concern is motivation.

<sup>5</sup> It is thus immune to the objections of Broome (2007, pp. 368–369, 371–373), who notes that practical rationality is consistent with failing to act or intend to act on the ground that  $p$ , when the fact that  $p$  is a reason for me to  $\phi$ . This reason may be outweighed, or I may have misleading evidence to that effect. Despite this, rationality does require a correct response to reasons, in the form of motivation or desire.

about reasons to act.<sup>6</sup> Finally, *Reasons* claims no priority for either side of its equation.

The truth, I think, is that *Reasons* says very little. It is a harmlessly illuminating principle that connects two things which surely must be connected: facts' being reasons for action, on one side, and the process of practical thinking, inference, deliberation, on the other. Whether a given fact is a practical reason has to do with the soundness of practical reasoning from that fact. It has to do both with the state of the world and with the standards of practical rationality. *Reasons* tells us how.<sup>7</sup>

Innocent though it is, *Reasons* is denied by some, ignored by others. In what follows, I defend this principle, and in doing so resolve disputes about the nature of practical reason: about full information, advice, enabling conditions, holism, the weight of reasons, "ought," and evidence.

## 1 Examples and advice

The plausibility of *Reasons* comes out most effectively by comparison. Suppose we want to relate facts about reasons to facts about practical rationality, knowing that there must be some connection here. Our first thought might be that reasons for action are considerations by which an agent would be moved if she were practically rational, conforming to the ideal standards of practical thought. We thus arrive at what Michael Smith (1995) has called the "example model" of reasons, which treats an idealized version of the agent as an example for her to follow. This view is decisively flawed. As Smith points out, it goes wrong when my reasons would be different if I were more fully rational than in fact I am. He cites an example by Watson (1975, p. 210), in which I am angry with my opponent after a humiliating game of squash. If I were practically rational, and so not gripped by irrational anger, I would want to shake hands with him, but in my fury I would probably lose my cool; there is no reason for me to take that risk. The same example shows that there are reasons for me to act in ways I would not care to act if I were fully rational. For instance, there is reason for me to hit the showers at once, which I would not want to do if I were sufficiently rational to ignore or not to feel such anger. The example model gets this wrong.

A second objection is more mundane. When I am mistaken about my circumstance in ways that practical rationality would not correct, I may be moved by considerations that are not good reasons. In Williams' well-known example, I falsely believe that the glass in front of me contains not petrol but gin (Williams 1979, p. 102). The fact that I am thirsty is not a reason to mix it with tonic and drink it, but I would be moved to do so if I were fully rational, so long as my false belief persists. The example model could be revised so as to identify reasons with the facts by which I would be moved if I were fully rational and had no false beliefs. But this

---

<sup>6</sup> On internalism about reasons, see the essays collected in Setiya and Paakkunainen (2012).

<sup>7</sup> Throughout this essay, "practical rationality" means the excellence of practical reason, not just freedom from "practical irrationality" in a narrow sense. On the restricted use of "irrationality" as a constraint on reasons, see "Against Internalism" (Setiya 2004).

would only intensify the problem above: if I had no false beliefs, let alone all relevant true beliefs, my reasons would be different yet again.

These difficulties do not arise for the principle of *Reasons*. Because it looks to patterns of good reasoning, or good dispositions of practical thought, not to ideal rationality, it can allow for reasons that depend on rational imperfection and for reasons we would have only if we were more rational than in fact we are. Since it is a formal principle, it does not specify the content of these reasons, but there is nothing to prevent its being a good disposition of practical thought to be moved to hit the showers, not to shake my opponents hand, when I am irrationally enraged. *Reasons* can also explain why the fact that I am thirsty is not a reason to drink the petrol I believe to be gin: good reasoning corresponds to reasons only when it does not rest on false beliefs.

This explanation prompts a natural question: why not correct for ignorance as well as error? But that “correction” would be a mistake. Reasons for action may correspond to practical thought that depends on ignorance of fact. So, for instance, if I do not know that my team is going to lose, the fact that the odds are good may be a reason to place a bet.<sup>8</sup> The reasoning by which I am moved to do so turns on lack of knowledge, but it corresponds to reasons nonetheless. This point is not specific to ignorance of the future. The odds may give me reason to buy a ticket in a lottery that has already taken place, even if, as a matter of unknown fact, the winning ticket has been drawn by someone else. In general, we need to allow for reasons by which we are moved in conditions of uncertainty. We could make this explicit by adapting *Reasons* to partial belief:

The fact that  $p$  is a reason for A to  $\phi$  just in case A has a collection of psychological states, C, such that the disposition to be moved to  $\phi$  by C-and-the-belief-that- $p$  is a good disposition of practical thought, and C contains no false beliefs or mistaken degrees of belief.

This formulation may seem problematic. How often are we in a position to reason in ways that are not contaminated by errors of confidence? Won't this principle give us fewer reasons than we actually have? But the problem disappears when we acknowledge the fine texture of practical thought. Even if am wrong to be certain that  $p$ , it may be a good disposition of practical thought to be moved to  $\phi$  by confidence above some threshold,  $n$ , along with the belief that  $q$ , where I am right to be confident at least to that degree. The fact that  $q$  is thereby vindicated as a reason for me to  $\phi$  despite my over-confidence. If we equate degrees of belief with beliefs about epistemic probability—“On my evidence, it is more or less likely that  $p$ ”—these complexities fall under *Reasons* itself. At any rate, they will be left implicit from now on.<sup>9</sup>

The asymmetrical treatment of ignorance and false belief provides a point of contrast with Smith's own response to the problems posed above. He treats an idealized version of the agent—one who deliberates correctly and has all relevant

<sup>8</sup> For an example of this kind, attributed to Frank Jackson, see Dancy (2000, pp. 65–66).

<sup>9</sup> If degrees of belief are correct when they match epistemic probabilities, beliefs may be correct only when they are both true and justified, or epistemically sound. For instance, we might restrict the premises of sound reasoning to what the agent knows. For a view in this spirit, see Hawthorne and Stanley (2008).

true beliefs—as an advisor, not an example. More carefully, he takes reasons to match the desires one would have for one’s actual circumstance if one were fully rational and fully informed.<sup>10</sup> What the idealized agent wants herself to do is, in effect, the advice she would give herself to follow. If I were fully rational and fully informed, I would want myself to hit the showers, not shake my opponent’s hand, in a fit of irrational anger. Nor would I want myself to mix petrol with tonic and take a drink. In each case, the advice model works fine. It does less well with reasons that depend on partial belief. Because the ideal in question is one of full information, I would not want myself to bet on the losing team, no matter what the odds. The model wrongly predicts that there is no reason to do so.

This problem might be solved.<sup>11</sup> More worrying is that the “relative” version of the advice model—“there is reason for A to  $\phi$  in C just in case A would want himself to  $\phi$  in C if he were fully rational and fully informed”—conflicts with the universality of reasons—“if there is reason for A to  $\phi$  in C, there is reason for anyone to  $\phi$  in C”—unless there is convergence in the relevant desires of rational agents. Hence the “non-relative” version, proposed by Smith, on which there is reason for A to  $\phi$  in C just in case *everyone* would want themselves to  $\phi$  in C if they were fully rational and fully informed.<sup>12</sup> Agents who meet these conditions must want themselves to act in the same way, in a given circumstance, if there is reason for anyone to act in that way. The requirements of rationality must fix what we want for ourselves, with full information, regardless of our initial desires. As others have complained, this demand is difficult to meet: in its non-relative version, the advice model threatens to eliminate reasons.<sup>13</sup> What has not been sufficiently emphasized is that it does so even if it is possible to reason well. If convergence fails for the desire to  $\phi$  in C, the non-relative advice model implies that an agent who finds herself in C and is moved to  $\phi$  on the basis of fully rational deliberation from a correct conception of her circumstance has no reason whatever to  $\phi$ . A fact by which she is rationally moved, without false belief, is not a reason! This consequence is incredible. And nothing compels it. We can meet the pressures that

<sup>10</sup> Smith counts full information as a condition of full rationality, along with correct deliberation (Smith 1994, p. 156). My terminology separates the two.

<sup>11</sup> Compare Smith 2006, in which “should” is indexed to contextually relevant expectations and values. This move could be adapted to reasons. In each case, the effect is to build disputed claims about the rational treatment of risk and uncertainty into the nature of normative properties. *Reasons* is less tendentious.

<sup>12</sup> Smith (1994, pp. 151–152, 164–174).

<sup>13</sup> The strength of Smith’s condition can be hard to make out. After all, if the circumstance includes psychological facts about the agent, we are bound to want the same things when our circumstances are the same. This is, however, irrelevant to the advice model, which asks for the desires we would have *about* our behaviour in C, not what desires we would have in C itself. Suppose, for instance, that practical rationality is purely instrumental, a matter of putting means to ends; and suppose that A is altruistic, desiring happiness for all, while B is utterly selfish. If A were fully rational and fully informed, what desires would he have about his behaviour in the unfortunate circumstance in which he becomes like B? Being altruistic, instrumentally rational, and well-informed, he wants even those without altruistic desires to act in ways that benefit others: that is what he wants himself to do in the circumstance described. In contrast, B would want himself to act only in ways that promote his own interests. Even in conditions of full rationality and full information, A and B would differ in the relevant desires. These issues are further discussed in Sobel (1999).

motivate Smith by accepting *Reasons*.<sup>14</sup> It is free from the defects of the example model, and so long as we include in an agent's circumstance the facts of his psychology, it is not only consistent with but entails the universality of reasons: if the fact that  $p$  is a reason for A to  $\phi$ , it is a reason to  $\phi$  for anyone who shares the relevant collection of psychological states, and for whom it contains no false beliefs.

## 2 Enabling conditions

In *Ethics without Principles*, Dancy (2004, Part I) draws a contrast between reasons for doing  $\phi$  and mere conditions for the existence of a reason. This distinction is no doubt real. Among the necessary conditions of something's being a reason for A to  $\phi$  is the fact that A exists. This need not be a premise of the practical reasoning by which A should be moved to  $\phi$ . It is a condition, not a reason. On some accounts, desire plays a "background" role in practical thought.<sup>15</sup> It is rational to be moved to  $\phi$  by the desire that  $p$  and the belief that doing  $\phi$  is an effective means to that end. That I have this desire is relevant to what there is reason for me to do without being a reason to  $\phi$ : I should be moved by the desire itself, not by a belief about that desire. *Reasons* allows for and explains these possibilities: among the facts in virtue of which there is reason for me to  $\phi$ , only the contents of beliefs that figure in C will count as reasons in their own right.

Dancy's more contentious claim is that even some of these considerations, which are premises of sound reasoning, are mere enabling conditions, not reasons to act. In his example, the fact that I have promised to do something is a reason to do it, while the absence of duress, possession of ability, and lack of competing reasons, though relevant to practical reasoning, are said to be mere conditions (Dancy 2004, pp. 38–41). Here *Reasons* disagrees. It counts every premise of sound reasoning as a reason to act. There is a pragmatic explanation of our tendency to focus on the promise, not on the facts of duress, ability, and so on, when giving a reason. Ability is arguably a condition of every practical reason. Duress is uncommon and we assume its absence. Citing the promise is thus a more informative and more natural way to bring out the pattern of practical reasoning under which the relevant motivation falls. In the right context, however, any fact that is a premise of sound reasoning can be given as a reason. If I have been pressured into making various promises, so that duress is salient, the fact that a particular promise was made

<sup>14</sup> In addition to theoretical arguments for the advice model, there are cases in which it seems to give a better verdict than *Reasons*, as when it is impossible to act on a putative reason without defeating its purpose. Is the fact that I have forgotten a meeting at noon a reason for me to check my calendar? It is not a reason by which it would be rational to be moved. As I have argued elsewhere (in Setiya 2009, p. 538), facts like this are not reasons to act but reasons to respond to action in related ways: the fact that I have forgotten my appointment is a reason to be glad if I check my calendar and dismayed if I do not.

<sup>15</sup> See Pettit and Smith (1990).

without duress will be a sensible thing to cite as a reason for keeping it. The same goes for ability. Imagine the following conversation:

“A, B, and C are all in need.”

“But I can’t help A or B.”

“Then the fact that you *can* help C is a reason to do so.”

Exchanges of this kind seem awkward, perhaps, but not incorrect. The situation is similar to that evinced by causal talk, where “the cause” is picked out from background conditions by explanatory salience—the striking of the match, not the presence of oxygen—though in a suitable context, any condition could be “the cause.” If oxygen is scarce, it makes sense to say “The match lit because oxygen was present.”

Appeal to what we ordinarily call “the reason” is therefore inconclusive. And Dancy’s theoretical arguments are weak. According to Dancy (2004, p. 39), “the fact that my promise was not given under duress is not a second reason for [keeping] it, to be set alongside the first one.” About the fact that there is no greater reason to do anything else: “If [this] was itself a further reason over and above those on which it passes judgement, we would be forced to reconsider the balance of reasons once we had asserted [it], in a way that would continue ad infinitum.” (Dancy 2004, p. 40) What these objections assume is that, if the fact that  $p$  is a reason to  $\phi$  and the fact that  $q$  is a reason to  $\phi$ , the conjunctive fact that  $p \& q$  must be a stronger reason to  $\phi$  than either  $p$  or  $q$  alone. We thus generate the absurdity that, by citing further premises of the practical argument for keeping a promise, we build an ever-stronger case for doing it, on the assumption that each premise is a reason. The proper response, however, is not to deny that some premises are reasons, but to reject the assumption that, by adding reasons to  $\phi$ , one inevitably finds a weightier reason. When the relevant facts are premises of the same sound reasoning, as in the present case, this will not be so.<sup>16</sup>

A final argument turns on a subjunctive test for reasons.

That she is in trouble and needs help is a consideration that favours my helping her. That I am the only person around does not seem to be another reason, on top of the first one [even though it is a premise of sound practical reasoning]. It is not as if, even if she were not in trouble, that I am the only person around would still favour my helping her. (Dancy 2004, pp. 41–42)

But as Dancy elsewhere notes, that something would not be a reason in a different circumstance is consistent with its being a reason here and now (Dancy 2004, pp. 18–21, 73–74). So, again, there is no objection to the verdict of *Reasons*.

As this fact reveals, one form of “holism” is readily explained by the picture of reasons as premises of sound reasoning. Except in the limiting case where C is empty, *Reasons* predicts that a fact that provides a reason in one circumstance can be disabled in others: where other conditions of C do not obtain, or where a belief

---

<sup>16</sup> Ironically, Dancy himself rejects the problematic assumption (Dancy 2004, p. 15). I come back to it in Sect. 3.

involved is false.<sup>17</sup> What is a reason in a given circumstance may not be a reason elsewhere. Since practical reasoning typically rests on more than a single belief, the contrasting “atomism” is implausible. Nor can it be saved by conjoining the premises of practical thought to make a “complete” or “inclusive” reason. If sound reasoning from this conjunctive premise turns on anything but belief—doubt, ignorance, non-cognitive psychology—the reason in question will not apply when these conditions fail. Understood in this way, the doctrine of holism is boring but true.<sup>18</sup>

### 3 Weighing reasons

In refusing to distinguish reasons from enabling conditions among the premises of sound reasoning, I was led to deny the following claim:

If the fact that  $p$  is a reason for  $A$  to  $\phi$  and the fact that  $q$  is a reason for  $A$  to  $\phi$ , the conjunctive fact that  $p \ \& \ q$  is a stronger reason for  $A$  to  $\phi$  than  $p$  or  $q$  alone.

When  $p$  and  $q$  are premises of the same practical reasoning, this principle fails. It follows that reasons for action are not “quasi-additive” in the sense defined by Selim Berker:

(i) individual reasons always make discernible individual contributions to the overall rightness or wrongness of a given action and (ii) the individual contribution made by a reason of positive valence always positively affects the total reason in favor of the action in question, and the individual contribution made by a reason of negative valence always negatively affects the total reason in favor of the action. (Berker 2007, p. 130)

Berker doubts that any coherent conception of reasons could deny that they are quasi-additive. If reasons count in favour of action, how could adding reasons fail to count for more? (Berker 2007, §V) In my view, *Reasons* refutes this argument. A reason to  $\phi$  counts in favour of doing so by contributing a premise for sound reasoning to the desire or motivation to  $\phi$ . Conjoining this reason with another premise of the same reasoning adds nothing to its strength.

<sup>17</sup> In this way, *Reasons* restricts the phenomenon of “silencing.” If the fact that doing something will cause pleasure is a reason for me to do it, but not when the pleasure is guilty, the possibility of guilt must be excluded by the starting points of sound practical thought. My reasoning must include beliefs whose truth implies that the pleasure is innocent. In other words: reasons can be silenced by a change in circumstance only if sound reasoning anticipates the relevance of that change. This may seem to count against *Reasons*, but it does not. We can see this by reflecting on practical thought in conditions of ignorance. If it is good practical thought to be moved to  $\phi$  by C-and-the-belief-that-doing- $\phi$ -will-cause-pleasure, even though C does not contain beliefs whose truth is incompatible with guilt, it is, in effect, practically rational to bet on the pleasure’s being innocent. That reason survives even if the gamble fails, like my reason to bet on the losing horse. Note that, even in this case, C had better include the absence of beliefs which imply that the pleasure is guilty, or there would be reason to pursue it even when I believe it is.

<sup>18</sup> Note that holism about reasons does not entail, or even support, the more radical “particularism” of Dancy 2004. As Dancy (2004, pp. 81–82) concedes, holism is consistent with the codifiability of practical reason; and Dancy’s challenge to codifiability—how does it follow from the nature or “logic” of reasons to act?—does not depend on it.



In order to confirm this diagnosis, we need to say more about the relative weight of reasons to act. In doing so, we provide a theory of “ought” or “should, all things considered.” The basic thought is this: if reasons correspond to sound reasoning that issues in motivation or desire, we can measure the relative weight of reasons by the relative strength of motivation. In *Reasons without Rationalism*, I proposed the following:

Reasons correspond to collections of psychological states that fuel good practical thought. One reason is *stronger* than another just in case it is a good disposition of practical thought to be *more strongly moved* by the collection of states that corresponds to it, than by the collection that corresponds to the other. (Setiya 2007, p. 13)

When the fact that  $p$  is a reason for A to  $\phi$ , A has a collection of psychological states, C, such that the disposition to be moved to  $\phi$  by C-and-the-belief-that- $p$  is a good disposition of practical thought. When we compare the weight of reasons, we focus on the most limited collection of states for which this condition holds, the minimal pattern of reasoning that takes us from the premise to the relevant motivation or desire. If this collection contains beliefs besides the belief that  $p$ , their contents will also be reasons to  $\phi$ . Since they contribute to the same pattern of reasoning, however, and since this pattern is minimal, the conjunction of these reasons is no stronger than any of them alone: they are not quasi-additive. The collection of states that corresponds to each reason is the same as that which corresponds to their conjunction, so one cannot be more strongly moved by one collection than the other.

With this account of the relative weight of reasons, “ought” or “should” can be reductively explained. When the fact that  $p$  is a reason for A to  $\phi$  and the fact that  $q$  is a reason for A to  $\phi$  then the fact that  $p \ \& \ q$  is a reason for A to  $\phi$ . Let the conjunction of all reasons for A to  $\phi$  be the *total reason* for A to  $\phi$ . A ought to  $\phi$ , all things considered, just in case the total reason for him to  $\phi$  is stronger than the total reason to do anything else.<sup>19</sup>

Although it is true as far as it goes, this claim hides serious complications. To begin with, reasons may be incomparable. In the formula for the relative strength of reasons, “good disposition” must be read exclusively: of the conflicting dispositions we might have, only one can be good in the relevant sense. It is, in effect, the right disposition, while alternatives are wrong. We otherwise get the absurd result that two reasons can be stronger than each other, since the corresponding dispositions are both equally and adequately good. The proper verdict here is that neither reason has more weight. At the same time, these reasons may not be equally strong, since it may be not be good, as practical reasoning, to treat them equally. (To say that reasons are of equal weight is to say that the right disposition is the disposition to be moved by each of them to the same degree.) In that case, the reasons will be incomparable: not equally strong, but also not such that one is stronger than the

<sup>19</sup> This formula ignores sound reasoning that does not involve belief, since it does not correspond to reasons. In my view, this is harmless, since there is no such thing. When it is valid, practical reasoning always involves beliefs about the circumstance of action, even if it also involves non-cognitive states.

other.<sup>20</sup> In the face of normative diversity, when different patterns of practical reasoning make sense (or are themselves incomparable), the reasons to which they correspond cannot be determinately compared.<sup>21</sup>

A more surprising source of incomparability is epistemic. Think back to the possibility of reasons that depend on ignorance or doubt, as when a fact about the odds is a reason for me to place a bet on a particular team. At the same time, the fact that the team is going to lose, despite the odds, is a reason not to do so, a reason of which I am sadly unaware. Which reason is stronger? The question is perplexing. It has no comfortable answer, and the formula above explains why. The practical reasoning by which I would be moved to place the bet rests on a belief about the odds and lack of knowledge about the winner. The practical reasoning by which I would be moved to refrain turns on knowing that my team will lose. Since one cannot have both sets of psychological states, it cannot be practically rational to be more strongly moved by one set than the other, or to be moved by them to just the same degree. Reasons that depend on ignorance of fact cannot be compared with reasons that the fact itself provides. More generally, reasons can be compared only when the cognitive states involved in the reasoning that corresponds to them are compatible with one another.

What this predicts is that “ought” and “should, all things considered” can be evaluated only relative to a set of comparable reasons, and so only relative to a cognitive or informational state. And this is exactly what we find. In a now-familiar scenario, ten miners are trapped in shaft A or shaft B, though we do not know which one.<sup>22</sup> Water threatens to flood both shafts. If we block one shaft, the miners in the other shaft will die. If we block neither, both shafts will fill part way with water and just one miner, the furthest down, will die. Faced with this predicament, we announce our conclusion:

We ought to block neither shaft.

The problem is that the miners must be in one shaft or the other. Either way, there will be a fact—that they are in shaft A, or that they are in shaft B—that is a compelling reason to do otherwise. Although we know this in advance, our apparently conflicting assertion seems apt. What is going on? The answer, in part, is that the reason to block neither shaft corresponds to sound reasoning from a condition of ignorance about the location of the miners, while the reason to block shaft A or B corresponds to reasoning from knowledge of where they are. These reasons are incomparable: no-one can instantiate both conditions at once, so no-one can be more strongly moved by one pattern of reasoning than the other. In asking what we ought to do, we have to focus on a set of comparable reasons, and since we do not know where the miners are, and do not expect to find out, we focus on practical reasoning that is compatible with ignorance. When we gather and weigh

<sup>20</sup> This possibility is elaborated, in connection with a “virtue theory” of practical reason, in Setiya (2007, pp. 77–79). On the general question of evaluative incomparability, see the essays collected in Chang (1997).

<sup>21</sup> It follows that even complete or inclusive reasons may be incomparable and so not quasi-additive.

<sup>22</sup> For this version of the example, see Kolodny and MacFarlane (2010); they take the case from Derek Parfit, who credits it in turn to Donald Regan.

the corresponding reasons, we conclude that the total reason to block neither shaft is stronger than the total reason to block shaft A or the total reason to block shaft B. Relative to these considerations, that is what we ought to do.

In writing this description, I have been careful to skirt a contested issue about the information-relativity of “ought” and “should”: whether it is best conceived as subject-, context-, or assessment-sensitivity. According to the first approach, the truth of an “ought”-sentence is fixed by the set of reasons relevant to the subject of that sentence, which is fixed in turn by the body of information accessible to her. The problem is that an observer who knows the location of the miners can aptly assert, “They should block shaft B,” even if the crucial information is completely inaccessible to us. According to the second approach, the truth of an “ought”-sentence at a context of utterance is fixed by the set of reasons relevant to that context, which is fixed in turn by a contextually specified body of information. When I utter the sentence above, what I say is true, since it fits the balance of contextually salient reasons. Meanwhile, what the observer says is also true, since the context is different, and the salient reasons include the fact that the miners are in shaft B. Despite appearances, we do not contradict each other: we can both be right. (Note that the context that determines the relevant information need not be the informational state of the speaker: it may turn in part on the speaker’s intentions or on what is conversationally presupposed.) According to the third and final approach, the truth of an “ought”-sentence at a context of utterance is further relative to a context of assessment, by which the relevant information and reasons are fixed. An apparent advantage of such “relativism” is that it can explain why the observer seems to contradict my view.<sup>23</sup> For the relativist, the very proposition I assert is one

<sup>23</sup> Kolodny and MacFarlane (2010, §2.2). For what it is worth, I am not convinced by their arguments. If we think of “ought” as being evaluated with respect to a set of reasons fixed by a body of information, which is fixed in turn by the discursive context, and we assume that the deliberator intends to focus on the information that will be available before he acts, we can explain how advisors and deliberators disagree. The advisor’s information, however unexpected, is relevant to the deliberator’s claim. Nor is that claim excessively precarious or speculative, since the deliberator may have good reason to doubt that new information will come to light. What about distant observers? If they mean to rely on inaccessible information, their claims do not conflict with those we make in the context of deliberation. On the other hand, an observer may want to evaluate an action from our point of view. It will then be natural to say, of the rescuers in the mineshaft case, “They did what they should in blocking neither shaft,” and in doing so, one speaks the truth. The position of the deliberator is equally flexible, as we can see by looking at conditionals like these:

- If the miners are in shaft A, we ought to block shaft A.
- If the miners are in shaft B, we ought to block shaft B.

Such conditionals may seem true. But ask yourself: suppose the miners are in shaft A, but you don’t know it; should you block shaft A? Surely not! Even if the miners are in shaft A, we ought to block neither shaft. What is happening here is a shift in the contextually salient body of information. On one reading, we should block neither shaft, not knowing where the miners are. The mere fact that the miners are in shaft A, or B, does nothing to change this, and the conditionals come out false. (This point is more vivid the more obscure the relevant facts.) On another reading, we accommodate the assertion of the conditionals by evaluating “ought”-claims with respect to a different body of information, including facts about the location of the miners. Then the conditionals come out true; and it is a mistake to say, as we originally did, “We ought to block neither shaft.”

that he denies—correctly, since it is false at the context from which it is being assessed.

The dispute between these theories turns on subtle questions about the nature and scope of normative disagreement. Is the conflict between my judgement and that of the observer merely apparent? Does this appearance fade when our intentions coincide? And can it really be salvaged by assessment-sensitivity? Although it is not my purpose to resolve these questions here, there is no reason to doubt that their resolution is consistent with the present theory of reasons, reasoning, “ought” and “should.”

#### 4 Evidence

It may, however, conflict with the theory of reasons proposed by Kearns and Star (2008, 2009), on which reasons are evidence of what one ought to do. More carefully:

R: Necessarily, a fact  $X$  is a reason for an agent  $N$  to  $\Phi$  if and only if [...]  $X$  is evidence that  $N$  ought to  $\Phi$ . (Kearns and Star 2008, p. 37; see also Kearns and Star 2009, pp. 216–217)

In taking “ought” as primitive, Kearns and Star do not explain its information-relativity. Nor can they do so through the information-relativity of reasons as evidence,<sup>24</sup> since on their account ought-facts are not constituted, but made probable, by reasons.

This points to a more basic difficulty, which Kearns and Star acknowledge, that it “seems impossible for it to be the case that one ought to perform an action if there is no reason at all to do so.” (Kearns and Star 2008, pp. 51–52) This follows directly from the account of “ought” as fixed by the balance of reasons, which tracks the intuitive idea that what one ought to do is what there is most reason to do. How can it be explained by the conception of reasons as evidence? Kearns and Star reply that it is “unattractive” to suppose that one ought to  $\phi$  when one cannot know that one ought to  $\phi$ . Since “very plausibly, one can know a fact only if one is able to have evidence of this fact,” it follows that one ought to  $\phi$  only if there is evidence that one ought to  $\phi$ , and thus, by  $R$ , only if there is reason to  $\phi$  (Kearns and Star 2008, p. 53). Neither principle employed in this explanation is secure. The possibility of knowing without evidence is not incoherent: it figures in discussions of self-knowledge and, more relevantly, in some versions of moral epistemology.<sup>25</sup> Setting this point aside, the impossibility of unknowable *oughts*, however plausible, must also be explained. How does it follow from the metaphysics of *ought* that it cannot be epistemically inaccessible? The problem for reasons as evidence is not that the principle, “if one ought to  $\Phi$ , there is evidence that one ought to  $\Phi$  is [...] implausible,” but that it is left as a brute necessity, not derived from the nature of

<sup>24</sup> As on the theory of evidence as raising epistemic probability; Kearns and Star (2009, pp. 231–232).

<sup>25</sup> For forms of intuitionism that allow for non-evidential knowledge, see Audi (2004) and Shafer-Landau (2003).

reasons, *ought*, and evidence (Kearns and Star 2008, p. 53). Appealing to other brute necessities does not help.

There are, in any case, clear exceptions to *R*. While reasons for *A* to  $\phi$  may always count as evidence that he should, evidence that *A* should  $\phi$  need not amount to reasons. Three illustrations. First, when it rests on false belief, the premises of good reasoning are not reasons to act. They may nonetheless be evidence of what one ought to do. That I am thirsty is evidence that I ought to drink what is in the glass in front of me; though if I am acting on the false belief that the glass contains not petrol but gin, it is not a reason to do so. Second, the distinction between reasons and enabling conditions is ignored by *R*. If desires play a background role in practical reasoning, the fact that I desire that *p* may be evidence that I ought to  $\phi$ , where doing  $\phi$  is an effective means to that end, even though it is not a premise of practical thought, and therefore not a reason to act. I should be moved by the desire itself, not by a belief about that desire. Finally, the fact that *p* may be evidence that I ought to  $\phi$  of an abstruse or esoteric kind. An epistemic genius might see the connection by prodigious theoretical inference. But if it shows no defect of *practical* thought to miss the point, the fact that *p* is not a practical reason.

In each case, *Reasons* gives the proper verdict where *R* does not. The comparison is instructive. Principle *R* is what we get if we begin with the idea of reasons as premises of practical reasoning, an idea that is codified by *Reasons*, and assume that reasoning that concludes “I ought to  $\phi$ ” is necessarily practical. This, I conjecture, is where Kearns and Star go wrong. One of their principal arguments for the doctrine of reasons as evidence draws on the link between reasons and reasoning (Kearns and Star 2009, § 2.3). But they equate “reliable practical reasoning” with “reasoning that is generally successful in terms of issuing in correct judgements concerning what it is one ought to do” (Kearns and Star 2009, p. 224). If one thinks of practical reasoning as theoretical reasoning about a normative, practical subject-matter, one will expect the premises of sound reasoning about this topic—which are evidence that one ought to  $\phi$ —to constitute practical reasons. Hence the theory of reasons as evidence.

The problem with this argument is that, while it makes sense to call such reasoning “practical” in virtue of its object (what one ought to do), this is not the sense of “practical reasoning” that figures in our guiding thought. As I said at the very beginning of this essay, in interpreting the idea of reasons for action as premises of practical reasoning, we should count as practical reasoning any form of thought to which assessments of practical rationality apply. This includes being moved by a reason directly, without the mediation of normative beliefs—being moved by the belief that *p* without believing that the fact that *p* is a reason to  $\phi$ —and acting for a reason without concluding that one should perform the relevant act.<sup>26</sup> There is more to practical reasoning than forming and responding to such beliefs. More significantly, there is also less. It follows from the failure of principle *R*, on grounds described above, that theoretical reasoning about whether one ought to  $\phi$  may not amount to practical reasoning, in this sense: it cannot always be assessed for practical rationality. It may be part of being practically rational that one

---

<sup>26</sup> I defend these possibilities in Setiya (2010).

recognize reasons as such and be able to determine what one ought to do when one knows the reasons on each side. But as we have seen, a fact can be evidence that one ought to  $\phi$  without being a reason to  $\phi$ —as in the presence of false beliefs, when the evidence that one ought to  $\phi$  is a mere enabling condition, or when its epistemic relevance is obscure. Being indifferent to such facts is not, in itself, a defect of practical reason.<sup>27</sup>

The contrast between practical and theoretical rationality points towards a final virtue claimed by Kearns and Star: that the conception of reasons as evidence alone permits a unified treatment of reasons for action and belief (Kearns and Star 2009, § 2.1).<sup>28</sup> Reasons for action are evidence that one ought to  $\phi$ ; reasons for belief are evidence that one ought to believe that  $p$ . In my view, such unity is not enough to overcome the flaws considered so far. Nor does the unity run as deep as Kearns and Star suppose. They contend that “we can weigh reasons to act and epistemic reasons to believe against each other.” (Kearns and Star 2009, p. 220) Sometimes, the harm of believing  $p$  outweighs the otherwise decisive evidence of its truth: one should not believe that  $p$ , all things considered. I am sceptical of this description. The situation is rather one in which I ought, epistemically, to believe that  $p$ , even though there is reason to wish for the absence of that belief.<sup>29</sup> In any case, the principal question is comparative: does *R* provide the only plausible account of what is common to practical and epistemic reasons? Not if *Reasons* can be adapted to belief, as in fact it can.

*Epistemic Reasons:* The fact that  $p$  is a reason for A to believe that  $q$  just in case A has a collection of psychological states, C, such that the disposition to be more confident that  $q$  in light of C-and-the-belief-that- $p$  than in light of C alone is epistemically rational, and C contains no false beliefs.

Reasons for belief stand to degrees of belief as reasons for action stand to degrees of motivation: the kind of motivation which, when decisive, both occasions and informs intentional action.

An adequate defence of this principle would be another project. But so long as something in the vicinity holds, the conception of reasons as premises of sound reasoning finds sufficient unity in the application of reasons to action and belief. At the same time, it accommodates reasons that depend on ignorance, though not on error; it explains the distinction between reasons and enabling conditions; it makes sense of the relative weight or strength of reasons; and it predicts the information-relativity of “ought” and “should.” Though it does not answer the most important questions of ethics, *Reasons* provides a framework in which these questions can be most plainly and profitably asked.

<sup>27</sup> Nothing I have said here provides a test for when evidence that one ought to  $\phi$  is or is not a reason to  $\phi$ . That is a question about the content of practical rationality. I am trying to state the relationship between its standards, whatever they are, and facts about what there is reason to do.

<sup>28</sup> They also claim that *R* best explains how reasons can be weighed (Kearns and Star 2009, §2.6). But they do not consider *Reasons* and the corresponding theory of relative strength, developed in the text above. Their complaint is more persuasive against the theory of reasons as explaining *oughts* in Broome 2004; see Kearns and Star (2008, pp. 42–5).

<sup>29</sup> Here I agree with Parfit (2001), among others.

**Acknowledgments** For discussion of these topics, I am grateful to Samuel Asarnow, Jonathan Dancy, Cian Dorr, Alida Liberman, Richard Moran, Joseph Raz, Michael Smith, Jonathan Way, and Ralph Wedgwood; a reader for the journal gave very helpful comments on earlier drafts.

## References

- Alvarez, M. (2010). *Kinds of reasons*. Oxford: Oxford University Press.
- Audi, R. (2004). *The right in the good*. Princeton, NJ: Princeton University Press.
- Berker, S. (2007). Particular reasons. *Ethics*, 108(1), 109–139.
- Broome, J. (2004). Reasons. In R. J. Wallace, M. Smith, S. Scheffler, & P. Pettit (Eds.), *Reason and value: Themes from the moral philosophy of Joseph Raz* (pp. 28–55). Oxford: Oxford University Press.
- Broome, J. (2007). Does rationality consist in responding correctly to reasons? *Journal of Moral Philosophy*, 4(3), 349–374.
- Chang, R. (1997). *Incommensurability, incomparability, and practical reason*. Cambridge, MA: Harvard University Press.
- Dancy, J. (2000). *Practical reality*. Oxford: Oxford University Press.
- Dancy, J. (2004). *Ethics without principles*. Oxford: Oxford University Press.
- Davidson, D. (1963). Actions, reasons, and causes. Reprinted in *Essays on actions and events* (pp. 3–19). Oxford: Oxford University Press, 1980.
- Hawthorne, J., & Stanley, J. (2008). Knowledge and action. *Journal of Philosophy*, 105(10), 571–590.
- Kearns, S., & Star, D. (2008). Reasons: Explanations or evidence? *Ethics*, 119(1), 31–56.
- Kearns, S., & Star, D. (2009). Reasons as evidence. *Oxford Studies in Metaethics*, 4, 215–242.
- Kolodny, N., & MacFarlane, J. (2010). Ifs and oughts. *Journal of Philosophy*, 107(3), 115–143.
- Parfit, D. (2001). Reasons and rationality. In D. Egonsson, J. Josefsson, B. Petersson, & T. Rønnow-Rasmussen (Eds.), *Exploring practical philosophy* (pp. 17–39). Aldershot: Ashgate.
- Pettit, P., & Smith, M. (1990). Backgrounding desire. *Philosophical Review*, 99(4), 565–592.
- Raz, J. (1978). Introduction. In J. Raz (Ed.), *Practical reasoning* (pp. 1–17). Oxford: Oxford University Press.
- Setiya, K. (2004). Against internalism. *Nouûs*, 38(2), 266–298.
- Setiya, K. (2007). *Reasons without rationalism*. Princeton, NJ: Princeton University Press.
- Setiya, K. (2009). Reply to Bratman and Smith. *Analysis*, 69(3), 531–540.
- Setiya, K. (2010). Sympathy for the devil. In S. Tenenbaum (Ed.), *Desire, practical reason, and the good* (pp. 82–110). Oxford: Oxford University Press.
- Setiya, K., & Paakkunainen, H. (2012). *Internal reasons: Contemporary readings*. Cambridge, MA: MIT Press.
- Shafer-Landau, R. (2003). *Moral realism*. Oxford: Oxford University Press.
- Smith, M. (1987). The Humean theory of motivation. *Mind*, 96(381), 36–61.
- Smith, M. (1994). *The moral problem*. Oxford: Blackwell.
- Smith, M. (1995). Internal reasons. *Philosophy and Phenomenological Research*, 55(1), 109–131.
- Smith, M. (2006). Moore on the right, the good, and uncertainty. In T. Horgan & M. Timmons (Eds.), *Metaethics after Moore* (pp. 133–148). Oxford: Oxford University Press.
- Sobel, D. (1999). Do the desires of rational agents converge? *Analysis*, 59(3), 137–147.
- Watson, G. (1975). Free agency. *Journal of Philosophy*, 72(8), 205–220.
- Williams, B. (1979). Internal and external reasons. Reprinted in *Moral luck* (pp. 101–113). Cambridge, England: Cambridge University Press, 1981.