You ought to ϕ only if you may believe that you ought to ϕ

Benjamin Kiesewetter

Published in: The Philosophical Quarterly 66 (265): 760-82, (2016).

https://doi.org/10.1093/pq/pqw012

This is a pre-print. Please cite the published version.

Abstract: In this paper I present an argument for the claim that you ought to do something only if you may believe that you ought to do it. More exactly, I defend the following principle about normative reasons: An agent A has decisive reason to ϕ only if she also has sufficient reason to believe that she has decisive reason to ϕ . I argue that this principle follows from the plausible assumption that it must be possible for an agent to respond correctly to her reasons. In conclusion, I discuss some implications of this argument (given that some other standard assumptions about reasons hold). One such implication is that we are always in a position to be justified in believing all truths about what we have decisive reason (or ought) to do.

In this paper I present an argument for an interesting claim about normative reasons:

The Principle of Decisive Reasons (PDR): Necessarily, if A has decisive reason to ϕ , then A has sufficient reason to believe that she herself has decisive reason to ϕ .

This claim is interesting because it has substantial consequences for some theories about reasons and for some issues that have recently been the focus of the philosophical debate on normativity. To anticipate one of these consequences, consider the widely shared view

that a person's reasons for belief are given by this person's evidence. On this assumption, PDR entails some kind of evidence requirement on the existence of decisive reasons:

Evidence Requirement: A has decisive reason to ϕ only if A has sufficient evidence that she herself has decisive reason to ϕ .

This is a significant constraint on reasons. An example put forward by John Broome (2007a, 352) serves to illustrate this point. Suppose that the fish on the plate in front of you contains salmonella, but you have no evidence that this is the case. According to Broome, the fact that the fish contains salmonella could still provide you with a decisive reason against eating it. It is on the basis of this assumption that Broome rejects the view that rationality consists in responding correctly to reasons: since you cannot rationally be expected to avoid eating the fish, but you might have decisive reason to avoid eating the fish, requirements of rationality and requirements of reasons must come apart. This argument is blocked by the evidence requirement. If you have no evidence that the fish contains salmonella, then (as far as the example tells us) you lack sufficient evidence that you have decisive reason to avoid eating the fish. Hence the evidence requirement entails that you do not have decisive reason to avoid eating the fish in Broome's example. This is a controversial implication, and (as I try to demonstrate at the end of this paper) it might have important impacts on the debate on reasons and rationality, as well as on other philosophical disputes.

I shall argue that PDR follows from the conceptual possibility of responding correctly to reasons. I will start with some clarificatory remarks about the content of PDR in §I. In §§II-III, I explain and defend the crucial assumption of the argument, namely that it must be possible for an agent to respond correctly to her decisive reasons. Subsequently, I present an argument that this assumption, together with a relatively uncontroversial claim about reasons, entails PDR in §IV. The final section §V concludes by highlighting some of the consequences of PDR. I provide a formal version of the argument in the appendix.

I. Preliminaries

Let me start with some clarificatory remarks. First, the reasons that PDR refers to are normative reasons. Second, I shall assume a conception of normative reasons (nowadays standard), according to which such reasons are facts (or true propositions) that stand in a relation to an agent A such that they count in favour of some kind of response ϕ of the agent.¹ By "response", I mean something that can be done or held for a reason, such as an action or a belief. I focus on reasons for action and belief, but plausibly other attitudes can be favoured by reasons, and reasons for such attitudes will be subject to PDR as well. Thus, PDR connects truths about normative reasons of all kinds with truths about normative reasons for beliefs.

Third, I mean A's having *sufficient* reason to ϕ to imply that A's reasons – all things considered – normatively *permit* A to ϕ , while I mean A's having *decisive* reason to ϕ to imply that A's reasons – all things considered – normatively *require* A to ϕ .² In the latter case, one might just as well say that A *ought* to ϕ : I take it that "A has decisive reason to ϕ " is the thought that A expresses when she concludes after deliberation "I should ϕ " or "I ought to ϕ ".³ PDR can thus be rephrased as the claim that you ought to ϕ only if you may believe that you ought to ϕ . It may be understood as a version of Francisco Suárez' principle that a doubtful law is not binding (*lex dubia non obligat*), which was accepted by many early modern scholastics and arguably also by Immanuel Kant.⁴

¹ See e.g. Dancy (2000), Parfit (2011), Raz (2011), Scanlon (1998), Skorupski (2010). I here understand the standard conception broadly as to include views according to which the notion of a reason is definable in other normative terms, such as Broome's (2013, Ch. 4), or reducible to natural facts, such as Schroeder's (2007) and Smith's (1994).

² By a "normative requirement", I mean a requirement that is normative; I am not referring to those specific wide-scope principles that Broome used to call "normative requirements" (Broome 1999), but no longer does (Broome 2013, 27).

 $^{^3}$ I also think that it is the thought that others express when they say, "You ought to ϕ " in giving A advice. Advice, however, raises specific questions about how an agent's evidence is related to this agent's reasons that I have to bracket here. I have discussed some of these questions elsewhere (in Kiesewetter 2011 and forthcoming) and must leave the discussion of the particular implications that PDR has for the truth conditions of statements of advice for another occasion.

⁴ Suárez, "De bonitate", disp. 12, sec. 6, n. 8. For a discussion of the *lex dubia* principle in early modern scholasticism and Kant, see Schuessler (2006, Ch. 2 and Ch. 5). PDR is also closely related to some claims that have been defended in the recent literature on normativity. On the natural assumption that what we may believe corresponds to what we can justifiably believe, PDR is the reverse of the principle that if one

Finally, I have to address a complication. Some authors use the expression "A has reason to ϕ " to mean something like "There is a reason for A to ϕ and A is aware of, or has epistemic access to, this reason", and others think that the term 'reason' is systematically ambiguous between different normative senses, e.g. one sense that carries the implication that A has epistemic access to the reason and another sense that does not carry this implication.⁵ So let me clarify. Like most other authors, I take the expressions "A has a reason to ϕ " and "There is a reason for A to ϕ " to be synonymous.⁶ When using the first of these expressions, I do not thereby wish to suggest that the relevant reason is epistemically available to A. That is, my *use* of this expression leaves it open whether reasons, in the sense I am concerned with, have to satisfy some kind of access condition or not. As I have already indicated, and will further elaborate later on, PDR can be used as premise in an argument to the effect that we must have epistemic access to our decisive reasons (the *evidence requirement*). But this is a (possible) conclusion of this paper, not something that I presuppose.

If you think that the expressions 'decisive reason' and 'sufficient reason' are ambiguous between different readings, then I hope you will agree that one of these readings is pertinent to deliberation about what one ought to, or may, do or believe. The notions of 'decisive reason' and 'sufficient reason' that I use are the ones that settle the

can justifiably believe that one ought to ϕ , one ought to ϕ - a thesis defended by Gibbons (2009, 172–3) and Way and Whiting (forthcoming), although in terms of doxastic (or *ex post*) justification rather than propositional (or *ex ante*) justification. I believe that this thesis is falsified by the principle that 'ought' implies 'can', but I am sympathetic to a qualified version that accommodates this concern (cf. Kiesewetter 2013, Ch. 7.8). Similarly, PDR may be seen as the reverse of a version of Titelbaum's "fixed point thesis" (2015, 261), although Titelbaum restricts his claim to rational requirements and permissions, which he sees as at least "potentially" distinct from the notions of sufficient and decisive reasons (cf. pp. 265–7).

⁵ See Lord (2010) and Skorupski (2010, 113) for the former of these views, and Schroeder (2008) for the latter

⁶ See e.g. Williams (1979, 101). Similarly, I am not differentiating between a person's evidence, the evidence that *there is* for this person, and the evidence that this person *has*. However, while I do not think that relating a person to reasons in the above-mentioned ways presupposes that the reason is available to the person (as I go on to explain), I do indeed think that relating a person to evidence in the above-mentioned ways presupposes that the evidence is available to the person. This is because in the case of evidence, there does not seem to be any relation that these expressions could refer to other than the availability relation itself. What this means is that the step from "A has reason to believe p" to "A has evidence that p" involves not only the assumption that reasons for belief are provided by evidence, but, in addition, the assumption that the relevant evidence must be available to the agent. I address these assumptions in §5.

conclusions of such deliberation.⁷ The common view, I take it, is that PDR is false for these notions. For the standard account of the relevant notion of 'sufficient reason to believe' is that one has sufficient reason to believe p only if one has sufficient evidence that p. Yet many philosophers working on these issues would concur in Broome's opinion that just as with any other beliefs, the correct conclusion of deliberation may be that one ought to ϕ , even if one has insufficient evidence for this claim.⁸ If you already believe that this common view is false, then the conclusions of this paper may not surprise you that much, but the paper may still provide you with a new argument for what you already believe.

II. The Possibility of Responding Correctly to Reasons

The argument that I shall put forward in favour of PDR rests on the assumption that it is always possible for an agent to respond correctly to her decisive reasons. In this section, I will discuss the notion of 'responding correctly to reasons', and argue that the assumption is correct. The thesis that I seek to defend is this:

The Possibility of Responding Correctly to Reasons (PRR): Necessarily, if A has decisive reason to ϕ_1 , ..., and A has decisive reason to ϕ_n , then it is possible that A has decisive reason to ϕ_1 , ..., A has decisive reason to ϕ_n , and A responds correctly to her decisive reasons.

Let us call a claim to the effect that A has decisive reason to ϕ a *decisive reason claim*. According to PRR, it is a conceptual truth that if a decisive reason claim – or a conjunction of such claims – is true, then it is possible that this claim or conjunction of claims is true while A responds correctly to her decisive reasons.

-

⁷ PDR might thus be expressed as the following claim: the correct conclusion of deliberation about whether to ϕ is that one ought (or has decisive reason) to ϕ , only if the correct conclusion of deliberation about whether to believe that one ought to ϕ is that one may (or has sufficient reason to) believe this.

⁸ See e.g. Williams (1979, 102–3).

As I will explain presently, I understand the relevant notion of 'responding to reasons' such that it involves conforming to one's decisive reasons while being aware of the fact that one has these reasons. PRR thus requires that it is possible to conform to all one's decisive reasons together while believing that they exist. But note that by 'possible', I mean no more than 'conceptually possible'. It is compatible with PRR that it might not be psychologically possible for you to respond correctly to your reasons given the evidence available to you. Thus, PRR does not beg the question against those who want to deny the evidence requirement. It is also compatible with PRR that you might be unable to respond correctly to your reasons because you lack the relevant capacities. PRR is neutral on the question of whether beings that lack the capacity to have beliefs about reasons could be subject to decisive reason claims. What motivates PRR is that responding correctly to reasons should at least in principle be open to those who have the relevant reason-responding capacities. If it were conceptually impossible for an agent to respond correctly to her decisive reasons, then this agent could not respond to them correctly, no matter what capacities or psychology she has. Denying PRR commits one to this possibility, which strikes me as quite implausible.9

In what follows, I shall elucidate the notion of 'responding correctly to reasons' that I adopt, and say why I think that PRR is plausible for that notion. To begin with, I will assume that responding correctly to reasons entails *conformity* with decisive reasons:

The conformity condition: Necessarily, if A has decisive reason to ϕ , and A responds correctly to her decisive reasons, then A ϕ -s.

This might be questioned in the following way. Suppose you have decisive reason to ϕ , and you try your best to ϕ , but due to some coincidence in the world, you end up not ϕ -ing. In such a case, you might want to say that you responded correctly to your reasons

⁹ It is plausible that a stronger version of PRR, which requires e.g. nomological and not merely conceptual possibility, can be justified on the same ground. Since the argument for PDR requires only the weaker claim about conceptual possibility, I restrict myself to it here.

even though you did not conform to them. In my view, this case is better described as one in which the claim that you had decisive reason to ϕ is undermined by the fact that even trying your best to ϕ did not lead to your ϕ -ing – at least this seems to be implied by a plausible interpretation of the principle that 'ought' implies 'can'.

However, let us grant for the sake of the argument that you fail to conform to your decisive reasons in the case under consideration and that there is a legitimate use of the phrase 'responding correctly to reasons' that applies nevertheless and thus does not comprise the conformity condition. This is compatible with there also being a legitimate use of that phrase that does comprise the conformity condition. Up to a certain extent, we are free to define the meaning of a technical phrase like 'responding correctly to reasons' in ways that fit our purposes. No matter whether one can fail to conform to a decisive reason despite trying one's best, there is a natural and useful notion of 'responding correctly to reasons' that entails conformity, and this is the notion that I adopt.¹⁰

Accordingly, the interesting question is not whether responding correctly to reasons entails conformity, but whether PRR is true if 'responding correctly to reasons' is understood as entailing conformity. Certainly, everyone who accepts *some* version of 'ought' implies 'can' thereby also accepts that it must be conceptually possible for us to conform to each decisive reason claim. But should we likewise accept that it must be possible for us to conform to all decisive reason claims together? This conclusion immediately follows if we accept that decisive reason claims agglomerate:

Agglomeration: Necessarily, if A has decisive reason to ϕ , and A has decisive reason to ψ , then A has decisive reason to $[\phi \text{ and } \psi]$.

¹⁰ In addition, I would like to note that the cited objection is ineffective against a version of the conformity condition that is restricted to decisive reasons for beliefs, and that the argument for PDR in §4 requires only the truth of this restricted condition (i.e. the second premise of the argument could be replaced by the restricted condition).

¹¹ Note that 'A has decisive reason to ϕ ' does not entail that A has one particular reason to ϕ that is decisive, but only that A's reasons taken together decisively favour ϕ -ing. For this reason, the consequent of the agglomeration principle – 'A has decisive reason to $[\phi$ and ψ]' – should not be taken to imply that there is any particular reason to $[\phi$ and ψ]; it merely means that A's reasons taken together decisively favour $[\phi$ -ing and ψ -ing].

For if it must be possible for A to conform to each decisive reason claim, and decisive reason claims agglomerate, then it must also be possible for us to conform to all of our decisive reason claims together. Moreover, not only does the agglomeration principle seem intuitively plausible, the principle that decisive reason claims should be jointly satisfiable is also independently quite compelling. If reasons conflict because the responses they favour are incompatible, then either one of the reasons outweighs the other, or the reasons are equally strong, or neither is the case and the reasons are incommensurable. Outweighed reasons are not decisive reasons. If we have (otherwise undefeated) equally strong or incommensurable reasons for incompatible responses, then we have decisive reason to give one of these responses, but we do not have decisive reason to give each of them. So in no case of conflicting reasons do we have decisive reasons for incompatible responses.¹²

Famously, Bernard Williams rejected both the agglomeration principle and the requirement of joint satisfiability with respect to *moral obligations* in order to make room for the possibility of tragic moral dilemmas.¹³ Note, however, that I am concerned not with the moral 'ought', but with the all-things-considered 'ought' of deliberation. Even if normative conflicts are granted on the level of qualified 'oughts' such as the moral 'ought' (and I wish to express no opinion on this matter here), there is no reason to suppose that such conflicts could occur on the level of decisive reasons. Proponents of tragic moral dilemmas typically agree with this; they deny that moral obligations necessarily provide decisive reasons rather than maintaining that there can be decisive reasons for incompatible responses. Williams, for example, explicitly embraces the view that there

¹² I give an independent argument to the effect that the deliberative 'ought' obeys a principle of joint satisfiability in Kiesewetter (2015, 930–4).

¹³ See Williams (1965, 117–23). By a *tragic* moral dilemma I mean a dilemma that involves moral obligations to perform incompatible actions. Such dilemmas are sometimes called "genuine" moral dilemmas. But I do not think that genuine dilemmas need to be tragic. Conflicts of equally strong or incommensurable moral reasons can provide genuine moral dilemmas that are not tragic. So in my view, one need not allow for tragic moral dilemmas (thus denying the joint satisfiability of moral obligations) in order to allow for genuine moral dilemmas.

cannot be conflicting 'oughts' on the all-things-considered level of deliberation. ¹⁴ I shall here assume that this is indeed the case.

I have been assuming that responding correctly to decisive reasons entails conforming to them, but clearly responding to reasons goes beyond conformity: you can conform to a reason by mere accident, while responding is non-accidental. What else does it involve? A natural idea is that one responds correctly to one's decisive reasons if one's conformity is guided by an awareness of the fact that one has these reasons. This entails:

The normative belief condition: Necessarily, if A has decisive reason to ϕ , and A responds correctly to her decisive reasons, then A believes that she herself has decisive reason to ϕ .

Two points should be noted about this claim. First, the standard for what counts as having the relevant belief should not be set too high. That an agent believes that she has decisive reason to \$\phi\$ does not imply that she would use the term 'decisive reason' to express her belief; it primarily implies that her having thoughts that can correctly be described by using that terminology.\(^{15}\) Moreover, the normative belief condition should not be taken to imply that whenever one responds correctly to one's decisive reasons, one has to make an explicit and conscious judgment before giving the response one's reasons call for. Such a requirement could indeed give rise to a "one thought too many" objection.\(^{16}\) A response has to be carried out against the background of an understanding of one's reasons as reasons, indeed as decisive reasons, but that does not necessarily mean that one has the thought presently in mind that one has decisive reason for the response. The relevant belief may be implicit rather than explicit.

Second, and quite importantly, the normative belief condition does not imply that we are always required to conform to our reasons in the reflective manner that is involved in a full-blown response. What I say is compatible with the plausible claim that

¹⁴ Cf. Williams (1965, 123–4; 1980, 119), see also Horty (2003, 588–9).

¹⁵ It also should not be taken to imply her having knowledge of the principles that are constitutive of the concept of a (decisive) reason (see Raz 2011, 32–3).

¹⁶ Compare Williams (1976, 18).

we are often permitted to conform to our decisive reasons in more automatic ways, without engaging in the activity of assessing reasons at all.¹⁷ Recall that PRR merely requires the *conceptual possibility* of responding correctly to reasons.

Yet the normative belief condition might be questioned. Consider the case of Huckleberry Finn. Huck, famously, had a guilty conscience because he did not turn in his friend Jim, who escaped from slavery. We may suppose that even though Huck had decisive reason not to turn in Jim, he did not believe this to be the case. Since Huck did not turn in Jim, he conformed to this reason. But since he did not believe himself to have this reason, the normative belief condition entails that he did not respond correctly to it. On the one hand, this appears to be the correct verdict, because Huck's decision was not guided by recognition of his reason, at least not as a reason. On the other hand, the verdict may also be questioned, because Huck's decision was not merely accidental but involved a sensitive reaction to Jim's humanity and the value of their friendship. Thus, even though Huck did not assess his reasons correctly during reflection, he was in some sense sensitive to the reasons to which he conformed, and this might be taken to be sufficient for it to be the case that he has responded correctly to his reasons.

Again, I shall grant that there is a weak notion of 'responding correctly to reasons', according to which Huck responded correctly to his reasons. The notion that I adopt, however, is a more restrictive one, according to which Huck did not respond correctly to his reasons. Karen Jones (2003, esp. 189–90) makes a useful distinction between "tracking" reasons via a reliable non-reflective mechanism, and "responding" to reasons via reflective assessment of reasons as reasons. Drawing on this distinction, I shall say that Huck *tracks* his reasons, but does not respond to them correctly.

This is, of course, a merely terminological issue. The substantial question is whether PRR should be accepted when it is interpreted along the lines of the normative belief condition. And it seems to me that it should. Consider first the claim that it must

¹⁷ This point is familiar from the work of Joseph Raz (esp. Raz 1990, 178–82). What I call "responding correctly to reasons" seems to me equivalent to what Raz calls "compliance with reasons". Cf. Raz (1990, 178): "If the need to give Jane moral support […] is a reason for Derek to stay at home, then he conforms with that reason if he does stay at home. […] If Derek not only stays at home but does so because he realises Jane's need and that it is a reason for him to so act, then we would say that he complies with the reason."

be possible for us to respond, in the sense specified by the normative belief condition, to each particular decisive reason claim. It is part of the idea of a reason that it is a consideration that can be taken to count in favour of a certain kind of response, a response that can be carried out from an understanding of that consideration as counting in favour of that very response. That is why there are reasons to have beliefs, but no reasons, for example, to have hiccups. The point is not that having hiccups is not "up to us" as it is sometimes put, for beliefs are not up to us either. The difference is that hiccups cannot possibly be sensible reactions to the recognition of reasons for them. And if reasons are, as T.M. Scanlon (1998, 20), among others, points out, essentially related to responses that are sensitive to judgements about reasons, then it would be very astonishing if we could have decisive reasons that we could not possibly respond to in the way that the normative belief condition requires. So for example, it would be very astonishing if Huck had decisive reason against turning in Jim, even though it were in principle impossible for him (or any other agent) not to turn Jim in out of a recognition of the fact that there are decisive reasons against doing this.

Granted that it must be possible to respond to each decisive reason claim in particular, it might still be asked why it should be possible to respond to all decisive reason claims together. Why is it not enough to say that it must be possible to respond to each particular decisive reason claim, and possible to *track* or *conform to* all decisive reason claims together? One objection to this contender for PRR is that it appears to conflict with the agglomeration principle.²⁰ But once again, the idea that it should be possible to respond not only to each individual decisive reason claim, but to all of them together, is also independently plausible. Reasons should not put us under demands that we cannot

¹⁸ As Williams (1970), among others, has emphasized.

¹⁹ Note that if there were reasons for and against having hiccups, it would, at least in principle, seem to be possible for us to *track* such reasons. Actors report, for example, that they stop having hiccups once they enter the stage. So if the possibility of *being tracked* via reliable unreflective mechanisms were enough for considerations to provide reasons, then we could well have reasons for or against having hiccups.

²⁰ Strictly speaking, the agglomeration principle and the possibility of responding to each particular decisive reason claim only entail the possibility of conforming to all decisive reason claims while believing that one has decisive reason for the conjunction of the relevant responses, and not, as PRR requires, the possibility of conforming to all decisive reason claims while believing each of these claims. It is difficult to see, however, how one could motivate rejecting the latter possibility while accepting the former.

at the same time satisfy and understand. If it were impossible for us to meet *both* the conformity condition *and* the normative belief condition for all of our decisive reasons, then we could conform to or track the decisive reasons that we have at a particular time *only if* we did not believe that they existed, and we could believe that they existed *only if* we did not conform to or track them. That is, we could conform to some of our decisive reasons only accidentally or via an unreflective mechanism – it would be impossible for us to be aware of our overall conformity with reasons, let alone for our overall conformity with reasons to be guided by reflection. This strikes me as implausible. A notion of 'decisive reason' or 'ought' that has its natural place in deliberative conclusions should be taken to imply the possibility that one consciously conforms to all of one's decisive reasons.

III. An alleged counterexample

Mark Schroeder has recently put forward a case that might be regarded as a counterexample to PRR.²¹ Suppose that Nate really enjoys surprise parties thrown in his honour, and that there is a surprise party waiting for him at home. Schroeder maintains that under such circumstances the fact that there is a surprise party waiting for him is a reason for Nate to go home. Similarly, he might claim that Nate has reason to go to the surprise party. Assuming that there are no defeaters present, Nate's reason might be decisive. But then it seems that it is indeed conceptually impossible for Nate to respond correctly to his reasons. For as soon as he responded to his reasons by forming the belief that he has decisive reason to go to the surprise party, he could no longer conform to this reason, since there would no longer be a *surprise* party he could go to.

Similar examples have been discussed in the context of accounts of reasons that build on the role that reasons play in good reasoning, such as Bernard Williams' (1979) and Kieran Setiya's (2007b).²² I agree with Setiya (2009, 538) that they are inconclusive. Nate might have decisive reason to go home, the reason being (for example) that there is

²¹ Schroeder (2007, 33). Schroeder (2004, 363) credits the example to Nate Williams.

²² See e.g. Millgram (1996, esp. §3), Smith (2009, 523–4).

13

a surprise waiting for him. This is a reason that he could respond to correctly. We need not assume that Nate literally has reason to go to the surprise party (or that the fact that there is a surprise party is a reason for him), in order to make sense of the situation.

While Schroeder would conclude that some reasons cannot be responded to from the possibility of surprise party reasons, I would conclude that there cannot be surprise party reasons since such reasons could not be responded to. How can we decide the issue? On the one hand, we should all agree that it would be good for Nate to go to the surprise party, and that other things equal persons have reason to \$\phi\$ if \$\phi\$-ing would be good for them.²³ On the other hand, we should also agree that other things aren't always equal. It would often be good for us to travel back in time to correct our past mistakes, but presumably we have no reason to do this. It can even be good for us to have hiccups, but that provides no reason for having them. Neither do we have reasons against having headaches or nightmares even though it would be good for us not to have them. When does the fact that ϕ -ing would be good for us provide a reason for us to ϕ and when not? A very natural answer is that it provides such a reason just when awareness of this fact can guide our ϕ -ing, when it is possible for us to ϕ in light of that fact. As Joseph Raz (2011, 27) puts it: "to say of a fact that it is a reason for action is not merely to say that it shows the action to have some good, some point to it. It is to say something like that it can rationally guide an agent towards that action." Or, in the words of John Searle (2001, 104): "you have to be able to reason with reasons."

In an important recent article, Schroeder himself acknowledges this distinction between the evaluative and the deliberative. He maintains that the deliberative sense of 'ought', i.e. the 'ought' of decisive reasons, must be distinguished from an "evaluative sense, on which it means, roughly, that were things ideal, some proposition would be the case" (2011, 1). Sceptics about surprise party reasons agree with Schroeder that Nate ought to go to the surprise party in the evaluative sense; that is, they agree that it would be good or ideal if he went. What is at issue is whether Nate ought also in the deliberative

²³ Similarly, we can agree that the fact that there is a surprise party waiting for him is a good-making feature (going home is good for Nate *because* of that fact), and that *ceteris paribus* good-making features are reasons.

sense to go to the surprise party. Schroeder (2011, 8–11) himself suggests a number of criteria for distinguishing the deliberative from the evaluative 'ought'. I shall argue that accepting Schroeder's own hallmarks of the deliberative 'ought' should lead us to reject the claim that Nate ought, in the deliberative sense, to go to the surprise party, and adopt something like the alternative description that I suggested, according to which Nate ought to go home because there is a surprise waiting for him.

Schroeder's first criterion is that the deliberative 'ought' is "the right kind of thing to close deliberation" (2011, 9). Note that while Nate can correctly conclude in deliberation that he ought to go home, he cannot correctly conclude in deliberation that he ought to go to his own surprise party: whenever he draws this conclusion, it is actually false. Schroeder's second criterion is that "knowing what someone ought to do, in the deliberative sense, settles the question of what is advisable for them to do" (ibid.). Note that while it seems good advice to tell Nate to go home, it does not seem to be good advice at all to tell him to go to the surprise party. Schroeder's third criterion is that "it is legitimate criticism of someone that he or she does not do what he or she ought to have done" in the deliberative sense (ibid.). Note that while Nate might legitimately be criticized for not going home (if people told him that there was a surprise waiting for him), it would seem inappropriate to criticise him for not performing an action described as 'going to the surprise party'. Finally, the fourth criterion is that claims about the deliberative 'ought' presuppose claims about what is in the agent's "power to do" (2011, 10). And note that while it is certainly in Nate's power to go home, it is at least questionable that it is in his power to go to his own surprise party: it is nothing that he can intentionally do under this description.²⁴

I conclude that according to Schroeder's own criteria, we should reject the assumption that Nate ought, in the deliberative sense, to go to the surprise party. Consequently, we should reject the assumption that he has decisive reason to do so. The

²⁴ Schroeder mentions a fifth hallmark, according to which "the deliberative 'ought' is more closely connected to the notion of obligation than the evaluative 'ought'" (2011, 10). This hallmark also does not support the assumption that Nate ought, in the deliberative sense, to go to the surprise party.

correct analysis of Schroeder's example is compatible with the claim that it must be possible for us to respond correctly to our decisive reasons.

IV. The Argument

So far I have been arguing for the Possibility of Responding Correctly to Reasons (PRR), and defended it against an alleged counterexample. PRR entails the first premise of the argument for PDR:²⁵

(1) If it is possible that A has decisive reason to ϕ_1 , ..., and A has decisive reason to ϕ_n , then it is possible that A has decisive reason to ϕ_1 , ..., A has decisive reason to ϕ_n , and A responds correctly to her decisive reasons.

In the discussion of PRR, I have stipulated a notion of 'responding to reasons' that entails the conformity condition and the normative belief condition. This makes premises (2) and (3) true by definition:

- (2) Necessarily, if A has decisive reason to ϕ , and A responds correctly to her decisive reasons, then A ϕ -s.
- (3) Necessarily, if A has decisive reason to ϕ , and A responds correctly to her decisive reasons, then A believes that she herself has decisive reason to ϕ .

The fourth and final premise of the argument is as follows:

(4) Necessarily, if A has no sufficient reason to believe that p, then A has decisive reason not to believe that p.

²⁵ The logical form of PRR is $\Box(p\rightarrow\Diamond q)$, which contraposes to $\Box(\neg\Diamond q\rightarrow\neg p)$. By the K axiom, this entails $(\Box\neg\Diamond q)\rightarrow(\Box\neg p)$, which contraposes to $(\neg\Box\neg p)\rightarrow(\neg\Box\neg\Diamond q)$ and is thus equivalent to $\Diamond p\rightarrow\Diamond \Diamond q$. By the S4 axiom, which seems uncontroversial for conceptual necessity, we get $\Diamond p\rightarrow\Diamond q$, which is the logical form of premise (1).

This claim seems to me fairly uncontroversial. We are not normatively permitted to have beliefs for which we lack sufficient reason. Instead, if we have insufficient reason to believe p, and insufficient reason to believe not-p, then we are normatively required to *withhold* belief with respect to p.²⁶

It might be objected that this assumption presupposes that the old controversy between W.K. Clifford (1877) and Williams James (1896) on whether "it is wrong always, everywhere, and for anyone, to believe anything upon insufficient evidence" (Clifford 1877, 77) be decided in favour of Clifford. James, however, does not claim that we are permitted to have beliefs for which we have insufficient *reasons*, but instead that we sometimes have sufficient reasons to believe something for which we do not have sufficient evidence. In other words, James does not object to (4) but to evidentialism about reasons for belief. While some of the consequences that can be drawn from PDR do indeed depend on whether evidentialism is true, premise (4), and all other premises of the argument for PDR, can be accepted by evidentialists and non-evidentialists alike.

With premises (1)-(4) at hand, we can now run the following argument for PDR:

- (5) Assume for *reductio*: It is possible that A has decisive reason to ϕ , but no sufficient reason to believe that she herself has decisive reason to ϕ .
- (6) Hence, it is possible that A has decisive reason to φ, and A has decisive reason not to believe that she herself has decisive reason to φ (from 4 and 5).
- (7) Hence, it is possible that A has decisive reason to φ, A has decisive reason not to believe that she herself has decisive reason to φ, and A responds correctly to her decisive reasons (from 6 and 1).

²⁶ Some authors would go further and claim that we are never permitted to ϕ if we have insufficient reason to ϕ . Note that I am not making this stronger claim; it seems to me debatable in the case of action. If it is conceivable that an agent has no reason to ϕ , but also no reason not to ϕ , we might plausibly hold that this agent is normatively permitted to ϕ , even if she does not have sufficient reason to ϕ . But while it is debatable whether *actions* always have to be licensed by a sufficient reason, it seems to me uncontroversial that *beliefs* indeed carry with them a commitment to sufficient reasons. If there is such a disanalogy, it might be because actions are not (as it is sometimes claimed) necessarily performed 'under the guise of the good', while there really is no question that we believe things 'under the guise of the true', since beliefs simply *are* attitudes of taking something to be true.

- (8) Hence, it is possible that A believes that she herself has decisive reasons to φ, A has decisive reason not to believe that she herself has decisive reason to φ, and A responds correctly to her decisive reasons (from 3 and 7).
- (9) Hence, it is possible that A believes that she herself has decisive reasons to φ and A does not believe that she herself has decisive reason to φ, and A responds correctly to her decisive reasons (from 2 and 8).
- (10) It is impossible that A believes that she herself has decisive reasons to ϕ and A does not believe that she herself has decisive reason to ϕ (*logical truth*).
- (11) It is impossible that A has decisive reason to ϕ , but no sufficient reason to believe that she herself has decisive reason to ϕ (reductio of 5, from 9 and 10).
- (12) Necessarily, if A has decisive reason to ϕ , then A has sufficient reason to believe that she herself has decisive reason to ϕ (from 11).

In (5), I assume for *reductio* what PDR denies: that it is possible to have decisive reason to ϕ , while lacking sufficient reason to believe that one has this reason. Since (4) tells us that insufficient reasons to believe entail decisive reasons not to believe, I conclude in (6) that it is possible to have a decisive reason to ϕ and a decisive reason not to believe in this reason. From premise (1), we can conclude that it is possible to have both of these decisive reasons *and* respond correctly to one's decisive reasons, as I record in (7). But responding correctly to the first reason entails (by the normative belief condition) believing that one has decisive reason to ϕ , while responding correctly to the second reason entails (by the conformity condition) *not* believing that one has decisive reason to ϕ . So we are forced to conclude that it is possible for us to believe *and* not to believe that we have decisive reason to ϕ (9), which runs counter to the logical truth that contradictions are impossible (10). This constitutes a *reductio* of the assumption made in (5), which allows us to confirm its negation (11) that in turn is equivalent to PDR (12).

V. Implications

If it is always conceptually possible for an agent A to respond correctly to her decisive reasons, as I have argued is plausible, then there is always sufficient reason for A to believe what there is decisive reason for A to ϕ . We should therefore accept what I have called the "Principle of Decisive Reasons" (PDR). The aim of this final section is to consider some possible implications of this result.

While the argument for PDR does not presuppose evidentialism or any other substantial conception about reasons for belief, what conclusions to draw from PDR essentially depends on what having sufficient reason to believe p amounts to. This question, however, is itself subject to considerable controversy and cannot be satisfactorily discussed here. I will restrict myself to providing a sketch of what seems to me the most plausible, and also most widely accepted, account of sufficient reasons for belief, say something about why it is plausible, and then consider the implications of PDR given that this account holds true.

The account that I have in mind holds that A has sufficient reason to believe p if, and only if, the evidence available to A sufficiently supports p, where I take it that a body of evidence sufficiently supports p just when believing p on the basis of this evidence is epistemically permissible. Instead of saying that the evidence available to A sufficiently supports p, I will use the simpler formulation that A has sufficient evidence that p. This gives us:

The standard account: A has sufficient reason to believe that p iff A has sufficient evidence that p.

The standard account, as I understand it here, leaves open a number of controversial questions. It leaves open, for example, how strong sufficient evidence has to be. It also leaves open how exactly we should conceive of the availability relation between A and the relevant body of evidence. Following Timothy Williamson (2000, ch. 9), we might hold that the relevant evidence has to be in A's actual possession, for example in virtue of being

known by A. On a more liberal conception, the relevant evidence only has to be in A's potential possession; according to John Gibbons (2013, ch. 7), for example, it consists in what A is in a position to know. These matters are for discussion on another occasion. What is important for present purposes is merely that the relevant evidence bears some kind of availability relation to A, such that A is in a position to form the relevant beliefs on the basis of that evidence by using her relevant capacities.

What are the alternatives to the standard account? First, one might deny the evidentialist assumption (implicit in the standard account) that reasons for belief are only provided by evidence for the belief's content, and maintain some kind of pragmatism, according to which at least some reasons for belief are provided by the benefits of having the belief. My main reason for rejecting this view is that it does not fit with the phenomenology of deliberation. As Richard Moran (1988, 143), among others, has pointed out, deliberation about whether to believe p is "transparent" on deliberation about whether p is the case. From a deliberative standpoint, then, we do not treat (or are not even able to treat) reasons for believing that p as different from evidence for p. Allegedly non-evidential reasons for belief are better understood as *practical* reasons for bringing some belief about (if one can) or reasons to wish that one had those beliefs, rather than as reasons for those beliefs themselves.²⁷ In any case, evidentialism is clearly the prevailing view about reasons for belief, and I shall assume in this section that it is true.

Second, one might accept some form of evidentialism and yet deny that the relevant body of evidence must be available to A. According to one alternative to the standard account, the relevant body of evidence that determines what an agent has sufficient reason to believe is just the set of all facts or truths, whether available to this agent or not. This view threatens to trivialize PDR. Since every truth is entailed by all truths (and thus certain, conditional on the set of all truths), it would follow that we always have sufficient reason to believe every truth. On this conception of sufficient

²⁷ See Kolodny (2005, 547–51) and Shah (2006) for arguments supporting Moran's inference from transparency to evidentialism, and Parfit (2011, App. A) for an argument in favour of understanding allegedly non-evidential reasons for beliefs as reasons for action.

reasons for belief, PDR would amount to no more than a trivial consequence of the more general principle 'If p, then A has sufficient reason to believe p'.

However, the view that we always have sufficient reason to believe each truth also strikes me as very implausible. Recall that the notions of 'decisive reason' and 'sufficient reason' that I am concerned with are the notions that figure in deliberation about what to do or to believe. I take it that at least sometimes, the correct conclusion of deliberation about whether to believe p is that one ought to withhold belief with respect to p, because one neither has sufficient reason to believe p, nor sufficient reason to believe not-p. The notion of 'sufficient reason' relevant in deliberation about what to believe thus cannot be determined by the body of evidence that consists of all facts, since in this case one could never lack sufficient reason to believe p and lack sufficient reason to believe not-p. Moreover, it also seems clear that it cannot be determined by any other body of evidence that includes information unavailable to the agent. For it seems clear that if the available evidence is insufficient for believing p or for believing not-p, then it is correct for an agent to conclude that she ought to withhold belief, no matter whether some unavailable body of evidence would be sufficient.

The standard account thus seems to capture the notion of 'sufficient reason for belief' that is relevant for deliberation quite well. In what follows, I will consider some possible consequences of PDR, assuming that the standard account is correct. I will do so in a rather explorative manner, in some cases drawing on further common assumptions without offering much argument, despite being aware that some of them are subject to controversy. The point here is simply to show how PDR can be engaged in different discursive contexts, provided that some common assumptions can be maintained.

Evidence and justification requirements

According to the standard account of sufficient reasons for belief, A has sufficient reason to believe p iff A has sufficient evidence that p. Thus, on the standard account PDR entails the following evidence requirement:

Evidence Requirement: A has decisive reason to ϕ only if A has sufficient evidence that she herself has decisive reason to ϕ .

As I pointed out in the beginning already, this is a significant restriction on the truth of decisive reason claims, and it is in tension with much of the literature on this topic.

According to what might be called the standard account of *justification*, A is in a position to be justified in believing p iff A has sufficient evidence that p.²⁸ Together with the standard account of reasons for belief, this entails that A has sufficient reason to believe that p iff A is in a position to be justified in believing p - a claim that can also be regarded as independently plausible. Together with this claim, PDR entails:

Justification requirement: A has decisive reason to ϕ only if A is in a position to be justified in believing that she herself has decisive reason to ϕ .

This means, in effect, that we are always in a position to be justified in believing all truths about what we ought to do. Again, this is an interesting and potentially far-reaching result.

Objectivism vs. perspectivism

Recent philosophical discussion shows a renewed interest in the question of whether what an agent ought to do depends on this agent's epistemic perspective or circumstances (as perspectivists claim), or on the totality of objective fact (as objectivists claim).²⁹ Most authors take reasons for belief to depend on the limited body of evidence that is available to the agent. Moreover, as I have tried to make plausible in the beginning of this section, they are right about that. But many still resist the view that practical reasons and 'oughts' depend on epistemic circumstances in the very same way.³⁰ A closer look at PDR,

²⁸ See, e.g., Conee and Feldman (2004, 1) for the claim that "facts about whether or not a person is justified in believing a proposition supervene on facts describing the evidence that the person has".

²⁹ For objectivism, see e.g. Thomson (1990, 229–234) and Graham (2010); for perspectivism, see e.g. Zimmerman (2008) and Kiesewetter (2011).

³⁰ Feldman (1988) and Skorupski (2010, 41f.) explicitly affirm a disanalogy between theoretical and practical reasons, maintaining that the former are "perspectival" or "relative to an epistemic field" in a way

however, reveals that such a disanalogy between the theoretical and the practical does not stand scrutiny, at least on the level of decisive reasons or 'oughts'. For according to PDR, decisive reasons of all kinds depend on sufficient reasons to believe. So if reasons for belief depend on epistemic circumstances, all (decisive) reasons do.

The relevance of normative uncertainty

Suppose that you ought to ϕ , but you do not know it. This possibility gives rise to a number of disputes among philosophers. I shall briefly indicate the impact that PDR has on two of these disputes. The first concerns the question of whether normative ignorance exculpates under certain circumstances. In a much-discussed article, Gideon Rosen (2004) has argued for a form of scepticism about moral responsibility on the basis of the assumption that normative ignorance can be exculpatory. Others have tried to resist this conclusion by denying that assumption.³¹ Although I cannot do justice to the complexity of the debate here, it seems to me worth pointing out that PDR may be taken to support the view that normative ignorance does not exculpate. This is because according to the evidence requirement (which, as I have argued, is implied by PDR), we always have sufficient evidence for truths about what we ought to do. It follows that ignorance about what one ought to do never results from a lack of sufficient evidence, but is always due to the fact that the agent does not form a belief despite the fact that the available evidence supports it. Thus, PDR entails that one can always avoid normative ignorance by forming beliefs in accordance with one's evidence. It seems clear that this makes it much harder to argue that normative ignorance exculpates.

A further dispute is concerned with the question of whether normative ignorance or uncertainty poses further practical questions about how to deal with the risk of wrongdoing. Suppose, for example, that you are inclined to believe that you ought to perform an act of blasphemy, but since you are not fully certain that God does not exist you also give some credence to the claim that you actually ought not to perform it. You

the latter are not. Gibbons (2010, 335), in contrast, questions whether there could be any good explanation for such a disanalogy. PDR provides an explanation for why there can be no such explanation.

³¹ See e.g. Harman (2011).

also believe that exercising blasphemy when you ought not to do so is considerably worse than refraining from exercising blasphemy when you actually ought to exercise blasphemy. According to some authors, such circumstances pose a further, substantial deliberative question about what one ought to do when one does not know what one ought to do.³² This assumption, however, creates a lot of problems. The relevant question can be given a non-trivial answer only if the two 'oughts' occurring in it have different meanings. Yet once we accept that first-order normative uncertainty poses the need for a second-order sense of 'ought', we should likewise accept that second-order normative uncertainty poses the need for a third-order sense of 'ought', and so on.³³ We seem faced with a potentially infinite regress of questions of the form "What ought_{n+1} I to do when I don't know what I ought, to do?" and a potentially infinite number of senses of 'ought' issuing conflicting demands on our behaviour. In the light of these problems, Michael Zimmerman has drawn the sceptical conclusion that the question "What ought I to do when I don't know what I ought to do?" is either incoherent or can be given only the trivial answer that one ought to do whatever it is that one ought to do, whether one knows what that is or not.34

While the arguments supporting this sceptical view strike me as convincing, it is also natural to complain that it provides a disappointing and unhelpful response to the apparently legitimate concern of an agent who is looking for guidance in a case of normative uncertainty. At this point, however, PDR may be brought into play as a response to this concern. For as I have argued, a plausible implication of PDR is that we are always in a position to justifiably believe all truths about what we ought to do. Thus, if you don't know what you ought to do, PDR suggests that you are in a position to find it out by considering your evidence – you need not seek guidance in considering a further practical question about what to do when you don't know what you ought to do.

_

³² See e.g. Sepielli (2009).

³³ Compare Sepielli (2014).

³⁴ Zimmerman (2008, 70). Harman (2015) also seems to embrace this sceptical view, although for different reasons.

Internal vs. external reasons

At least since Bernard Williams' seminal paper "Internal and external reasons", there has been an ongoing dispute among philosophers working on practical reasons as to whether statements to the effect that an agent has a reason for an action presuppose the existence of what Williams calls an element of this agent's "subjective motivational set". Williams' own argument for this kind of internalism has been contested by many. PDR may shed a new light on the issue. Recall that Williams includes dispositions of evaluation among the elements of the motivational set (cf. 1979, 105). And recall that on the standard account, PDR entails that one has decisive reason for some response only if one has sufficient evidence that one has this reason. Williams' internalist may claim that having sufficient evidence for a decisive reason claim involves having some disposition of evaluation. I am not endorsing this claim here, I am only employing it as an interesting hypothesis that the internalist might endorse and the externalist has to deny. For given that PDR and the standard account of reasons for belief are true, this hypothesis will indeed entail that truths to the effect that A has decisive reason to φ presuppose the existence of some element in A's subjective motivational set.

The normativity of rationality

Many philosophers suggest that rationality requires us to believe in accordance with our evidence.³⁶ Assuming the standard account once again, this entails that rationality requires us to believe in accordance with our reasons for belief. According to PDR, we have sufficient reasons to believe all truths about our decisive reasons. It follows that rationality at least requires us not to have false beliefs about our decisive reasons.

³⁵ See e.g. Hooker (1987), Korsgaard (1986), McDowell (1995), and Millgram (1996). For a good overview of the debate, see Finlay and Schroeder (2008).

³⁶ See e.g. the entry on "evidence" in the *Stanford Enyclopedia of Philosophy*: "insofar as one is rational, one is disposed to respond appropriately to one's evidence" (Kelly 2006, §2). Other authors make similar claims: "If you have no clear evidence that *p* is true […], you are rationally required not to believe *p*" (Cullity 2008, 79). "It is irrational for one to believe that p just in case one lacks sufficient evidence for p" (Ross 2012, 173). "Rational subjects have the capacity to grasp the evidence available to them and to form beliefs that accord with that evidence, the justification of subjects' beliefs turns on whether or not they have exercised such evidence-grasping and belief-forming capacities as they have" (Smith 2004, 167). "Rationality requires one to conform one's beliefs to one's evidence" (Williamson 2000, 12).

This is an interesting result in itself and yet also provides us with a new perspective on the current debate over the normative status of rational requirements.³⁷ Since irrationality is generally seen as worthy of criticism, it seems that we must have reasons, plausibly decisive reasons, to satisfy the requirements of rationality. This assumption, however, leads to serious difficulties. A consideration of the rational connection between intentions and beliefs about decisive reasons brings such difficulties to the surface. As is well known from the literature, a 'narrow scope' account of this connection allows for unacceptable forms of bootstrapping – it entails that whenever we believe we have decisive reasons to ϕ , we have decisive reasons to intend to ϕ . According to a 'wide scope' account, in contrast, we have decisive reasons either to intend to ϕ or not to believe we have decisive reasons to φ. But such symmetric requirements seem incapable of providing us with rational guidance for our decisions, and in special cases they also seem to license unacceptable forms of bootstrapping.³⁹ Note, however, that if rationality requires us not to have false beliefs about our decisive reasons, as suggested above, then there would be no objectionable bootstrapping involved in an asymmetric normative requirement to intend what we rationally believe about our decisive reasons. So PDR provides us with a promising starting point to account for the normativity of rationality. 40

The ideas mentioned in this section obviously need much more elaboration, and do not amount to anything more than a sketch. What I hope to have here shown is only the potential impact that the Principle of Decisive Reasons could have on a number of controversial issues that are discussed in the current literature.⁴¹

³⁷ See e.g. Broome (2007b), Kolodny (2005). For a recent survey of the debate, see Way (2010).

³⁸ See Broome (1999), Kolodny (2005).

³⁹ See Kolodny (2005) and Schroeder (2004) for the former point, and Raz (2005) and Setiya (2007a) for the latter.

 $^{^{40}}$ For an account of the normativity of rationality that makes use of the idea that rationality consists in responding correctly to evidence, see Kiesewetter (2013).

⁴¹ Earlier versions of this essay have been presented at the Australian National University, Canberra; Humboldt University of Berlin; Torcuato Di Tella University, Buenos Aires; the University of Sydney; the conference "Reasons: Action, Belief, Perception" at Saarland University, Saarbrücken, in October 2013; the 8th international conference of the Society for Analytical Philosophy (GAP) at the University of Konstanz in September 2012; and the earliest version was presented at the Humboldt-Princeton Graduate

Appendix

In this appendix, I formalize the argument for PDR, and show that it is valid. Let \square represent conceptual necessity, and let the letters on the left represent the following propositions:

C: A responds correctly to her decisive reasons.

 $R\phi$: A has decisive reason to ϕ .

r ϕ : A has sufficient reason to ϕ .

Bp: A believes that p.

φ: A φ-s.

PDR, the demonstrandum, thus reads: $\Box(R\phi \rightarrow rBR\phi)$.

The first premise of the argument states: "If it is possible that A has decisive reason to ϕ_1 , ..., and A has decisive reason to ϕ_n , then it is possible that A has decisive reason to ϕ_1 , ..., A has decisive reason to ϕ_n , and A responds correctly to her decisive reasons." To represent decisive reason claims, I use " Ψ_1 " as a variable for propositions of the form "A has decisive reason to ϕ_1 ". Premise (1) thus reads:

(1)
$$\Diamond(\Psi_1 \wedge ... \wedge \Psi_n) \rightarrow \Diamond(\Psi_1 \wedge ... \wedge \Psi_n \wedge C)$$
, premise

Premise (2) states: "Necessarily, if A has decisive reason to ϕ , and A responds correctly to her decisive reasons, then A ϕ -s." And Premise (3) states: "Necessarily, if A has decisive

Workshop in Philosophy, in July 2011, where I was so fortunate as to have Jack Woods as a commentator, who also helped me in subsequent correspondence. I would like to thank the participants for their valuable feedback. I am especially grateful to Anne Burkard, Andreas Carlsson, Ryan Cox, Daniel Friedrich, Jan Gertken, Felix Koch, Conor McHugh, Thomas Schmidt, Jonathan Way, and two anonymous referees for written comments, and Christian Barry, Eleonora Cresto, Christoph Fehige, Seth Lazar, Susanne Mantel, Elijah Millgram, Anders Nes, Andreas Pittrich, Stephan Schmid, Rudolf Schuessler, Nicholas Southwood, and Horacio Spector for discussing this essay with me or helping me with it in other ways. Work on this essay has been supported by the ARC Discovery Grant DP120101507 and the Deutsche Forschungsgemeinschaft (DFG project "Principles of the Deliberative Ought").

reason to ϕ , and A responds correctly to her decisive reasons, then A believes that she herself has decisive reason to ϕ ." This gives us:

- (2) $\Box((R\phi \land C) \rightarrow \phi)$, true by definition
- (3) $\Box((R\phi \land C) \rightarrow BR\phi)$, true by definition

Premise (4) states: "Necessarily, if A has no sufficient reason to believe that p, then A has decisive reason not to believe that p."

(4)
$$\Box(\neg rBp \rightarrow R \neg Bp)$$
, premise

From (1)-(4), which are the premises of the argument, PDR can be derived as follows:

- (5) $\Diamond(R\phi \land \neg rBR\phi)$, assumption for reductio
- (6) $\Diamond (R \phi \land R \neg BR \phi)$, from (4) and (5)
- (7) $\Diamond (R \phi \land R \neg BR \phi \land C)$, from (1) and (6)
- (8) $\Diamond (BR \phi \land R \neg BR \phi \land C)$, from (3) and (7)
- (9) $\Diamond (BR\phi \land \neg BR\phi \land C)$, from (2) and (8)
- (10) $\neg \Diamond (BR \phi \land \neg BR \phi)$, logical truth
- (11) $\neg \Diamond (R \phi \land \neg rBR \phi)$, reductio of (5), from (9) and (10)
- (12) $\square(R\phi \rightarrow rBR\phi)$, from (11)

Quod erat demonstrandum.

References

- Broome, John. 1999. "Normative Requirements." Ratio 12 (4): 398-419.
- ———. 2007a. "Does Rationality Consist in Responding Correctly to Reasons?" *Journal of Moral Philosophy* 4 (3): 349–74.
- ——. 2007b. "Is Rationality Normative?" *Disputatio* 2 (23): 161–78.
- Clifford, W.K. 1877. "The Ethics of Belief." Reprinted in: *The Ethics of Belief and Other Essays*, 70–96. Amherst, MA: Prometheus (1999).
- Conee, Earl, and Richard Feldman. 2004. *Evidentialism. Essays in Epistemology*. Oxford: Oxford University Press.
- Cullity, Garrett. 2008. "Decisions, Reasons, and Rationality." Ethics 119 (1): 57–95.
- Dancy, Jonathan. 2000. Practical Reality. Oxford: Oxford University Press.
- Feldman, Richard. 1988. "Subjective and Objective Justification in Ethics and Epistemology." *The Monist* 71 (3): 405–19.
- Finlay, Stephen, and Mark Schroeder. 2008. "Reasons for Action: Internal vs. External." In *Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta. URL = http://plato.stanford.edu/archives/spr2009/entries/reasons-internal-external/.
- Gibbons, John. 2009. "You Gotta Do What You Gotta Do." Noûs 43 (1): 157-77.
- ——. 2010. "Things That Make Things Reasonable." *Philosophy and Phenomenological Research* 81 (2): 335–61.
- ——. 2013. *The Norm of Belief*. Oxford: Oxford University Press.
- Graham, Peter A. 2010. "In Defense of Objectivism about Moral Obligation." *Ethics* 121 (1): 88–115.
- Harman, Elizabeth. 2011. "Does Moral Ignorance Exculpate?" Ratio 24 (4): 443-68.
- ——. 2015. "The Irrelevance of Moral Uncertainty." Oxford Studies in Metaethics 10: 53–79.
- Hooker, Brad. 1987. "Williams' Argument Against External Reasons." *Analysis* 47 (1): 42–44.
- Horty, John F. 2003. "Reasoning with Moral Conflicts." Noûs 37 (4): 557–605.
- James, William. 1896. "The Will to Believe." Reprinted in: *The Will to Believe and Other Essays in Popular Philosophy*, 1–31. New York: Dover (1956).
- Jones, Karen. 2003. "Emotion, Weakness of Will, and the Normative Conception of Agency." *Royal Institute of Philosophy Supplement* 52: 181–200.
- Kelly, Thomas. 2006. "Evidence." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta.
 - URL = http://plato.stanford.edu/archives/fall2008/entries/evidence/.
- Kiesewetter, Benjamin. 2011. "'Ought' and the Perspective of the Agent." *Journal of Ethics and Social Philosophy* 5 (3): 1–24.
- ———. 2013. *The Normativity of Rationality*. Ph.D. thesis, Humboldt-Universität zu Berlin.
- ———. 2015. "Instrumental Normativity: In Defense of the Transmission Principle." *Ethics* 125 (4): 921–46.

- ——. "How Reasons Are Sensitive to Available Evidence." In *Normativity: Epistemic and Practical*, edited by Conor McHugh, Jonathan Way, and Daniel Whiting. Oxford: Oxford University Press (forthcoming).
- Kolodny, Niko. 2005. "Why Be Rational?" Mind 114 (455): 509–63.
- Korsgaard, Christine M. 1986. "Skepticism about Practical Reason." *The Journal of Philosophy* 83 (1): 5–25.
- Lord, Errol. 2010. "Having Reasons and the Factoring Account." *Philosophical Studies* 149 (3): 283–96.
- McDowell, John. 1995. "Might There Be External Reasons?" Reprinted in: *Mind, Value, and Reality*, 95–111. Cambridge, MA: Harvard University Press (1998).
- Millgram, Elijah. 1996. "Williams' Argument Against External Reasons." *Noûs* 30 (2): 197–220.
- Moran, Richard. 1988. "Making Up Your Mind." Ratio 1 (2): 135–51.
- Parfit, Derek. 2011. On What Matters. Vol. 1. Oxford: Oxford University Press.
- Raz, Joseph. 1990. "Rethinking Exclusionary Reasons." Postscript to the Second Edition of: *Practical Reason and Norms*, 178–99. Oxford: Oxford University Press (repr. 1999).
- ——. 2005. "The Myth of Instrumental Rationality." *Journal of Ethics and Social Philosophy* 1 (1): 2–28.
- ——. 2011. From Normativity to Responsibility. Oxford: Oxford University Press.
- Rosen, G. 2004. "Skepticism about Moral Responsibility." *Philosophical Perspectives* 18 (1): 295–313.
- Ross, Jacob. 2012. "Rationality, Normativity, and Commitment." Oxford Studies in Metaethics 7: 138–81.
- Scanlon, T.M. 1998. What We Owe to Each Other. Cambridge, MA: Belknap Press.
- Schroeder, Mark. 2004. "The Scope of Instrumental Reason." *Philosophical Perspectives* 18 (1): 337–64.
- ——. 2007. *Slaves of the Passions*. Oxford: Oxford University Press.
- ——. 2008. "Having Reasons." *Philosophical Studies* 139 (1): 57–71.
- ——. 2011. "Ought, Agents, and Actions." The Philosophical Review 120 (1): 1–41.
- Schuessler, Rudolf. 2006. Moral im Zweifel. Vol. 2. Paderborn: Mentis.
- Searle, John R. 2001. Rationality in Action. Cambridge, MA: MIT Press.
- Sepielli, Andrew. 2009. "What to Do When You Don't Know What to Do." Oxford Studies in Metaethics 4: 5–28.
- ——. 2014. "What to Do When You Don't Know What to Do When You Don't Know What to Do…." *Noûs* 48 (3): 521–44.
- Setiya, Kieran. 2007a. "Cognitivism about Instrumental Reason." *Ethics* 117 (4): 649–73.
- ———. 2007b. *Reasons without Rationalism*. Princeton: Princeton University Press.
- ——. 2009. "Reply to Bratman and Smith." *Analysis* 69 (3): 531–40.
- Shah, Nishi. 2006. "A New Argument for Evidentialism." *The Philosophical Quarterly* 56 (225): 481–98.
- Skorupski, John. 2010. The Domain of Reasons. Oxford: Oxford University Press.
- Smith, Michael. 1994. The Moral Problem. Oxford: Blackwell.

- ——. 2004. "The Structure of Orthonomy." *Royal Institute of Philosophy Supplement* 55: 165–93.
- ——. 2009. "Reasons With Rationalism After All." Analysis 69 (3): 521–30.
- Suárez, Francisco. "De Bonitate et Malitia Humanorum Actuum." In: *Opera Omnia*, Vol. 4. Paris 1856.
- Thomson, Judith Jarvis. 1990. *The Realm of Rights*. Cambridge, MA: Harvard University Press.
- Titelbaum, Michael G. 2015. "Rationality's Fixed Point (or: In Defense of Right Reason)." Oxford Studies in Epistemology 5: 253–94.
- Way, Jonathan. 2010. "The Normativity of Rationality." *Philosophy Compass* 5 (12): 1057–68.
- Way, Jonathan, and Daniel Whiting. "If You Justifiably Believe That You Ought to Φ, You Ought to Φ." Forthcoming in: *Philosophical Studies*. doi:10.1007/s11098-015-0582-2.
- Williams, Bernard. 1965. "Ethical Consistency." *Proceedings of the Aristotelian Society, Supplementary Volumes* 39: 103–38.
- ——. 1970. "Deciding to Believe." Reprinted in: *Problems of the Self*, 136–51. Cambridge: Cambridge University Press (1973).
- ——. 1976. "Persons, Character and Morality." Reprinted in: *Moral Luck*, 1–19. Cambridge: Cambridge University Press (1981).
- ——. 1979. "Internal and External Reasons." Reprinted in: *Moral Luck*, 101–13. Cambridge: Cambridge University Press (1981).
- ——. 1980. "Ought and Moral Obligation." Reprinted in: Moral Luck, 114–23. Cambridge: Cambridge University Press (1981).
- Williamson, Timothy. 2000. Knowledge and Its Limits. Oxford: Oxford University Press.
- Zimmerman, Michael J. 2008. Living with Uncertainty. The Moral Significance of Ignorance. Cambridge: Cambridge University Press.