Online Research @ Cardiff

This is an Open Access document downloaded from ORCA, Cardiff University's institutional repository: <https://orca.cardiff.ac.uk/id/eprint/106374/>

This is the author's version of a work that was submitted to / accepted for publication.

Citation for final published version:

Robertson, Simon ORCID: <https://orcid.org/0000-0001-8340-1317> 2019.
Reasons and practical possibility. Inquiry 62 (9-10) , pp. 1104-1135.
10.1080/0020174X.2018.1445026 file

Publishers page: <http://dx.doi.org/10.1080/0020174X.2018.1445026>
<<http://dx.doi.org/10.1080/0020174X.2018.1445026>>

Please note:

Changes made as a result of publishing processes such as copy-editing, formatting and page numbers may not be reflected in this version. For the definitive version of this publication, please refer to the published source. You are advised to consult the publisher's version if you wish to cite this paper.

This version is being made available in accordance with publisher policies.

See

<http://orca.cf.ac.uk/policies.html> for usage policies. Copyright and moral rights for publications made available in ORCA are retained by the copyright holders.

Reasons and Practical Possibility

Forthcoming in *Inquiry*

Simon Robertson

Cardiff University

robertsons3@cf.ac.uk

Abstract

According to a common thesis about normative reasons for action, you have a reason to perform a given action only if you can act for that reason. This thesis has long had broad appeal and is intended to capture the practical character of practical reasons. I'll call it the 'Practicality Thesis'. Recently, however, various writers have developed subtly different objections to it, each designed to show that there can be actions you have a reason to perform even though you could *not* act for that reason—because, were you aware of the reason-giving facts in the ways needed to act for the reason, it would no longer be a reason for you to so act. This article defends the Practicality Thesis against such objections. It considers some extant defenses but shows that these are inadequate. It then advances an alternative approach designed to counter any structurally similar objection.

Reasons and Practical Possibility

According to a common thesis about normative reasons for action, you have a reason to perform a given action only if you can act for that reason. This thesis has long had broad appeal. It figures as the first premise in what has been dubbed the ‘Classical Argument’ for internalist or neo-Humean theories of reasons. But it is also accepted by many externalists too—including many Kantians, Aristotelians, and others—and has significant ramifications not just for the substantive content of true reason claims, but also our underlying conceptions of autonomy, moral obligation, blame, and responsibility.¹ One rationale for it emerges from the thought that an account of practical reasons should indeed be *practical*, by being sensitive to the live practical possibilities of the agents it attributes reasons to: if a reason for an agent to act is to genuinely favour the performance of *that* action by *that* agent, and thus favour *that* agent’s *actually doing* what it recommends, it must be something which *that* agent *could* act in light of. The thesis is thereby intended to capture the practical character of practical reasons. I’ll call it the ‘Practicality Thesis’.

This thesis has come under increasing fire, however. For, according to a certain ‘finkish’ style of objection, there are some reasons to act you cannot act *for*—because, were you aware of the reason-giving facts in ways needed to act for those reasons, they would no longer be reasons. This article defends the Practicality Thesis against such objections. It focuses on two, subtly different, versions of the worry—one from Elijah Millgram, another from Mark Schroeder—though the response developed is designed to counter *any* structurally similar case. Officially, the paper defends the Practicality Thesis on behalf of an internalist model of reasons. Nonetheless, it should have broader appeal and import. For not only are the resources it draws on available to many externalists too, it addresses some fundamental yet widely neglected issues concerning how to individuate reasons and what is involved in acting for a reason. §1 outlines the basic internalist model and Classical Argument for it. §2 presents the focal objection, via Millgram and Schroeder. §3 considers, though rejects, three extant responses. §§4–6 advance my alternative approach.

¹Versions of the thesis are advanced by e.g. Williams 1995; Korsgaard 1986; McDowell 1995; Skorupski 2010 (esp. 252–9), the latter including illuminating explanation of its wider significance. The Classical Argument for internalism (so labeled by Schroeder [2007, 6–9]) is made by Williams (1995, 38–9).

1. An internalist model

There are many models of reasons we might aptly label ‘internalist’. Partly for sake of familiarity, I’ll start from Bernard Williams’ internalism. On his most basic formulation (1981, 101):

(RI) A has a reason to ϕ only if A has some motive which would be served by her ϕ ing.²

‘Motive’ is a blanket term Williams uses to denote any motivationally valent item within an agent’s existing psychological profile. These include (a) the agent’s *occurrent* desires, aims, ends, interests, evaluative commitments, and the like, which comprise or feature as elements in occurrently motivating states; (b) the agent’s *background* and *standing* desires, aims, ends, interests and evaluative commitments—items that need not currently, but could, motivate her; and (c) any motivationally priming *dispositions* of character, evaluation and emotion that shape (a) and (b).

RI presents a necessary condition for an agent’s having some reason or another. Often, though, we want to know whether an agent has some *specific* reason—whether some particular fact (or set of facts) gives an agent a reason. (The Practicality Thesis also concerns acting for a specific reason.) The concept of a reason is typically taken to be relational: reasons are relations holding between some fact (or set of facts), an agent, and an action.³ And we often individuate a specific reason via that fact (or set of facts) which gives the agent the reason.⁴ Precising RI accordingly:

(RI*) that p gives A a reason R to ϕ only if A has some motive which, given that p , would be served by her ϕ ing for R.

Two preliminary points about RI*. First, usually more than one fact contributes to a single reason. This will be important later; but I’ll meanwhile use ‘that p ’ to stand for some reason-

²Reasons externalism is here understood as the denial of RI.

³See e.g. Scanlon 2014; Schroeder 2007, 2011; Skorupski 2010.

⁴(a) More precisely, a specific reason is properly identified via the particular relation holding between its relata (this enables us to distinguish different reasons provided by the same facts). (b) The word ‘fact’ is not ideal here. Although some take a ‘reason-giving fact’ to be a true proposition, I here use this expression to denote *whatever is picked out by* a true proposition (be it an object, property, feature of a situation, a priori truth, etc.). There is no single term covering the full range of such items; but this approach allows us to say that one and the same item identified under different modes of presentation counts as a single ‘fact’—more on this later. (c) We can nonetheless treat *reasons* as either substantial or nominal properties denoted by true normative propositions. I’ll use the locutions ‘that p gives A a reason to ϕ ’ and ‘that p is a reason for A to ϕ ’ interchangeably, taking both to imply: that p has some property R of standing in a relation to A and ϕ , where R is the relational property of being a reason.

giving fact or set of facts. Second, if that p gives A a reason R to ϕ , and if A ϕ s for that reason, then A ϕ s under some description ‘implicating’ that p . (We return to some complications in §4.)

Williams also endorses an additional psychological thesis, which we may call the ‘Sentimentalist Thesis’ (attributable to Hume and other British sentimentalists, amongst others). According to it, what one is able to appreciate *as a reason* and be motivated by depends on—is ineradicably shaped, constrained and influenced by—one’s antecedent motives. A little more precisely:

(ST) A is able to appreciate that p is a reason R for her to ϕ , and to be motivated to ϕ for R , only if: A has some motive which, given that p , would be served by her ϕ ing for R .

The Classical Argument for internalism (henceforth CA) can now be stated as follows:⁵

- (1) that p gives A a reason R to ϕ only if it is possible that A ϕ s for R
- (2) it is possible that A ϕ s for R only if it is possible that A is motivated to ϕ for R
- (3) it is possible that A is motivated to ϕ for R only if A has some motive which, given that p , would be served by her ϕ ing for R

So,

(RI*) that p gives A a reason R to ϕ only if A has some motive which, given that p , would be served by her ϕ ing for R .

We’ve already encountered one rationale for the Practicality Thesis, represented here by (1): for a normative reason to favour the performance of the action by the particular person it is attributed to, it must be something which, given the way that person actually is, that person can act for. The Practicality Thesis will be our main focus. But it is worth explaining the other steps in CA, partly to ward off some potential misunderstandings about the Practicality Thesis and partly to show how externalists as well as internalists could accept it.

The following should be uncontroversial: acting *for a reason* involves being motivated to act (for that reason); and it is *possible* that a person acts for a given reason only if she *could* be

⁵This is a reconstruction of the argument at Williams 1995, 38–9. It is presented slightly differently by e.g. Millgram 1996, 198; Parfit 1997, 112; Schroeder 2007, 7; these differences won’t affect what follows. (For a very different interpretation of Williams’ argument, see Finlay 2009.)

motivated to act for that reason. This gets us (2).⁶ The rationale for (3) is twofold. If A were to ϕ for reason R, (3) is trivially true: the motivational state A would be in when A acts for R just *is* a motive. But (3) can also be understood as an application of ST: the considerations a person is able to appreciate as reason-giving and be motivated by depends on her antecedent motives. Many externalists who oppose CA object to (3) and ST, on grounds that we can appreciate reasons in ways utterly uninfluenced by our subjective motives (via pure practical reasoning or normative insight, say).⁷ Nonetheless, although CA, as an argument for internalism specifically, explicates (1) and then (2) in terms of (3), the defence of the Practicality Thesis developed here does not require accepting (3). This should give it broader appeal. For although I'll officially be defending the Practicality Thesis on behalf of an internalist model of reasons, the resources by which I do this are available to many externalists too.⁸

⁶Although (2) itself should be uncontroversial, Parfit (1997, 113–4) seems to suggest that moving from (1) to (2) in the context of an argument for internalism begs the question, by presupposing that normative reasons must be connected to our motivational possibilities. However, it is difficult to see what the objection really is. Part of his worry appears to be that one would accept (2) only if one already accepts the further view that normative facts *are* (or are *reducible* to) non-normative facts about motivation, where that further view is what needs to be established for internalism to be true (hence moving from (1) to (2) begs the question). But that is peculiar in various respects. For one thing, many internalists (including Parfit's express target: Williams) deny that normative reasons are (reducible to) psychological items and instead treat the consequent of RI merely as a *truth-condition* for reason-sentences (e.g. Williams 1981, 101). It is then hard to see how *moving* from (1) to (2) begs the question (Dancy 2000, 17–19 makes a similar point). Perhaps Parfit is really objecting to (1) itself. Alternatively, his thought might be that (2) overlooks a non-motivational reading of 'being able to act for a reason'; I return to this in n.8.

⁷E.g. Korsgaard 1986; McDowell 1995; Millgram 1996; Skorupski 2010. For defence of ST, see e.g. Prinz 2007, Ch.1; Katsafanas 2013, 115–32.

⁸Note, also, that although (2) and (3) explicate the 'possibility' of acting for a reason in motivational terms, this is entirely consistent with the *further* view that acting (and being able to act) for a normative reason involves a certain form of 'reason-responsiveness': an ability to respond to relevant reason-giving facts by grasping their normative significance or seeing them *as* reasons, where such reason-responsiveness is not itself a motive (see e.g. Millgram 1996; Skorupski 2010: 253). This claim about reason-responsiveness is often associated with externalist views of reasons. But is also something internalists can accept (and I think should). It just adds a *further* condition on being able to act for a reason. As noted, (2) should be uncontroversial and accepted by both internalists and externalists. Nonetheless, one can accept (2) and *also* hold that a person could be motivated to act for a given reason (and hence could act for that reason) only if she is able to appreciate the relevant reason-giving facts as reasons; thus, her being able to appreciate some fact as a reason is part of what enables her to be motivated to act for that reason (i.e. makes this possible). (3) and ST are likewise consistent with such claims about reason-responsiveness. Internalists can hold that it is our antecedent motives which enable us to appreciate things as reasons in the ways we do (and to then be motivated to act for those reasons)—our motives being a lens through which we can appreciate potential reasons in the ways needed to act for them. Disagreement over ST and (3), therefore, is not about reason-responsiveness as such, but how exactly reason-responsiveness works (concerning, e.g., the influences our antecedent motives do or do not play in it). The salient point about reason-responsiveness for present purposes, then, is that it makes a claim about the mechanisms involved in being able to act for a reason that is additional to but consistent with each premise in CA. This is significant for the later arguments: my defence of the Practicality Thesis does not depend on these further views about reason-responsiveness, though is compatible with them; it can therefore be accepted by both internalists *and* externalists, irrespective of whether one *also* thinks that being able to act for a reason requires these reason-responsive abilities. (I do not nonetheless adopt Millgram's own reason-responsive terminology of 'appreciating reasons' when explaining and addressing his objection. My response in §6, though, does not rely on a reading of this that his externalism is entitled to reject.) I'm grateful to an anonymous reviewer for encouraging me to clarify these matters.

But let's now focus on the Practicality Thesis itself. Many externalists accept it.⁹ But some writers, externalist and internalist, have doubts about it. Elijah Millgram (1996) and Mark Schroeder (2007) in particular have each raised serious objections to it. The next section introduces both objections in turn.¹⁰

2. Two objections

Millgram considers an insensitive chap Archie who, 'because he is insensitive, his life is worse than it might be' (1996, 203). Archie 'realizes that things are not going well... but his insensitivity prevents him from seeing why' (*ibid.*). 'Archie's insensitivity', Millgram suggests, 'is a reason to act: reason to avoid acquaintances in their hour of grief' (*ibid.*)—a reason to avoid the funeral, say. Millgram continues: 'Archie's insensitivity is a deliberative incapacity: it consists in being unable to appreciate certain reasons for action... Because he is insensitive, he cannot see that his own insensitivity gives him reason for action' (*ibid.*). Yet, if Archie could reason better, perhaps by saying to himself 'I had better stay away from the funeral; if I go, I'll only make things worse', then 'he would *ipso facto* be sensitive enough not to have these reasons' (*ibid.*). Therefore, although Archie's insensitivity gives him a reason to avoid the funeral, he cannot appreciate the reason in the ways needed to be motivated by it or hence act for it—contra (1) and the Practicality Thesis.

We need to clarify three points concerning the reason Archie has to avoid the funeral. First, Millgram suggests that the fact that *Archie is insensitive* is such a reason. This cannot be the full story, though. Usually more than one fact contributes to a single reason. In this example it seems correct to suppose that Archie's insensitivity contributes to a reason to avoid the funeral only given various other facts—that his insensitivity would manifest itself at the funeral in ways that upset people, for instance. Let's provisionally identify this reason by saying 'that *Archie will upset people* is a reason for him to avoid the funeral'. Second, though, Millgram *also* suggests that the fact that *by going to the funeral Archie will make things worse* is a reason for him to avoid it. However, that *Archie will make things worse* need not provide a different reason from the fact that

⁹E.g. Korsgaard 1986 (and other neo-Kantians); McDowell 1995; Skorupski 2010.

¹⁰*Both* because, although structurally similar, some crucial differences render existing responses to the one unable to deal with the other—see §3.

Archie will upset people. Indeed, the first of these specifications just seems an elliptical statement of the reason, to which we can add more fine-grained specifications—that Archie would make things worse *because* he would upset people, for instance. Third, Millgram suggests that by upsetting people Archie will make his own life go worse. Presumably, though, it will also make things worse for the people Archie upsets. Plausibly, that he will upset people here contributes to two different reasons: that he would upset other people, alongside the fact that upsetting those other people is bad *for them*, is one reason for him to avoid the funeral; and that he would upset others, alongside the fact that this will make *his own* life go worse, is another.¹¹ For simplicity, when identifying ‘the reason’ attributed to Archie via the fact that he will upset people, I’ll focus on that reason concerning the effects Archie’s behaviour has on others.

We can now re-present Millgram’s objection. That Archie will upset people is a reason for him to avoid the funeral. But, because he is insensitive, he is unable to appreciate that he has that reason in the ways needed to act for it. Hence, he could not avoid the funeral for that reason. Furthermore, were Archie not so insensitive, he would not upset people; but then the fact he will upset people would not give him a reason to avoid the funeral, since there is no such fact. Therefore, although the fact that insensitive Archie would upset people is a reason for him to avoid the funeral, he could not avoid the funeral for that reason.

Schroeder offers a second counterexample:

Nate loves successful surprise parties thrown in his honor, but can’t stand unsuccessful surprise parties. If there is an unsuspected surprise party waiting for Nate in the living room, then plausibly there is a reason for Nate to go into the living room. There is certainly something that God would put in the ‘pros’ column in listing pros and cons of Nate’s going into the living room. But it is simply impossible to motivate Nate to go into the living room for this reason—for as soon as you tell him about it, it will go away. (2007, 165)

Here, the fact there is a surprise party in the lounge is a reason for Nate to go to the lounge. The alleged impossibility of his acting for that reason emerges from two thoughts. One trades on the

¹¹If you think there is only one reason here, the later arguments can be made by assuming we have two ways to specify it. The inclusion of valoric concepts in these reason-specifications is merely shorthand; I leave open whether it is the badness, or the non-evaluative features this rests on, that contributes to the reason.

party being a *surprise*. Were Nate aware of the party, it wouldn't be a *surprise*; hence, he could not enter the lounge for the reason that there is a genuinely *surprise* party—a party which will surprise him. (As Schroeder puts it: 'as soon as you tell him about [the surprise party]', the reason 'will go away'.) This does not rule out the possibility that Nate enters the lounge for the reason that there is a 'surprise' party—a party in which the intended surprise element is unsuccessful. Nonetheless, secondly, Nate 'can't stand unsuccessful surprise parties': he hates them so much he couldn't be motivated to attend this one. So, contra (1): although the fact that there is a surprise party awaiting Nate in the lounge gives him a reason to go there, he could not go there for that reason (whether the surprise element is successful or not).

These counterexamples are not merely clever, in ways admitting a quick technical fix.¹² Dealing with them will take us into important, but neglected, territory concerning how to individuate reasons and the bearing this has on acting for a reason. But before turning to these matters, the next section considers three extant lines of response.

3. Three responses

Response I: A first response is to deny that these are counterexamples. One might suggest that, in Millgram's example, if Archie really *could not* avoid the funeral for the supposed reason that he will upset people, it is *not* actually a reason for him to do so—and to assume otherwise begs the question against (1). We can unpack this as follows.

Premise (1) supplies a constraint on correct reason-attributions, specifying a scope constraint that must be met if a given fact provides a particular agent with a reason. We can here distinguish *candidate* reasons and reason-giving facts from *actual* ones—a candidate reason-giving fact being a fact that is a candidate for providing a reason, though with no guarantee that it succeeds. We could thereby treat the fact that Archie will upset people as a *candidate* reason-giving fact signalling a candidate reason for him to avoid the funeral. Nonetheless, if he could not avoid

¹²It might be thought that the Archie and Nate scenarios are rare and/or trade on paradoxical issues posing problems for many areas of philosophy—whereby, even if they highlight exceptions to (1), they don't undermine its spirit. The deeper worry, though, is twofold. First, there may be *many* similarly structured cases—being unable to act for reasons expressing modesty, humility, automaticity, etc. Second, if (1) is not uniformly true, CA shows only that some reasons are tied to agential motives in ways that satisfy RI*; that threatens more than the spirit of (1) and CA.

the funeral for this candidate reason, (1) implies that it is not actually a reason. Before explaining what is *not* entirely adequate with this response, it is worth warding off two other doubts about it.

A first doubt about *Response I* emerges from the thought that the concept of *a reason* is that of a *consideration which counts in favour of*. For, it might be objected, it is true that the fact that Archie will upset people is a consideration which counts in favour of his avoiding the funeral—and that truth just is the truth that the fact that Archie will upset people is a reason for him to avoid the funeral. Granting this commonplace concept of a reason, there are two ways to reply. One denies that the fact that Archie will upset people favours *his* avoiding the funeral. A similar consideration might favour *my* avoiding it—if, say, I'm prone to some insensitive outbursts yet sensitive enough to appreciate reasons to avoid situations where my insensitivity will likely manifest itself; it's just not a consideration favouring *Archie's* staying away, given his insensitivity. This seems strained, though. For there *is*, in one obvious sense, a consideration favouring Archie's staying away: namely, *he will upset people*. A second and better reply grants this and also grants the following: if A has a reason to ϕ , there is some consideration favouring A's ϕ ing. But it *denies* the converse—that is, denies that: (necessarily) if there is a consideration counting in favour of A's ϕ ing, then A has a reason to ϕ . We can deny this on grounds that there must be some constraint on the scope of reasons (both reasons in general and specific reasons)—some constraint on who or what it is correct to attribute a reason to. Suppose that the fact we will go hungry is a consideration counting in favour of at least some x 's not eating our food. It is nonetheless overblown to insist for each and every x in a position to eat our food—including wild creatures like bears, rabbits, termites, say—that such a consideration entails a normative reason for them to refrain from doing so. There may be substantive disagreement over what exactly the correct scope constraint is. Nonetheless, that there must be *some* such constraint should be uncontroversial: inanimate objects, plants, at least some animals and humans don't have the full array of reasons the rest of us might have. Thus, even if the fact that Archie will upset people favours his avoiding the funeral, it at least remains an open question whether he falls within the scope of this candidate reason. And (1) presents a candidate scope constraint: if Archie cannot act for this candidate reason, it is not a reason for him.

A second doubt about *Response I* is this: surely it would be *good* if Archie avoids the funeral—good for those he would upset, say—whereby, the doubt runs, it simply follows that he has a reason to avoid it. Internalists can accept the claim about *value*. What they deny is that truths about the value of actions uniformly entail (atomic, non-counterfactual) truths about reasons to act for particular agents (Williams 2001, 96). In particular, they deny that true value claims of the form ‘A’s ϕ ing would be good’ or ‘that p would make A’s ϕ ing good’ entail true reason claims of the form ‘A has a reason to ϕ ’ or ‘that p is a reason for A to ϕ ’ (assuming that the consequents of RI and RI* are not also truth-conditions for value claims). One rationale for denying these entailments is that, whereas we can truly say of an action an agent cannot perform that it may nonetheless be good were she to do it, it is often (perhaps always) false to say that she has a reason to do something she cannot do. The truth of a value claim need not depend on the agent’s abilities—whereas the truth of a normative reason claim typically does, because it is sensitive to what one *can* do and to the considerations one *can* act in light of. Correct reason attributions, in short, are subject to certain scope constraints that true value claims are not. Even if it would be good were Archie to avoid the funeral, it doesn’t follow that he has a reason to.

All this is perfectly coherent—indeed, sensible internalist fare—as far as it goes. It supplies a way, consonant with (1), to handle the Archie-example. However, there remain two worries with *Response I*. First, even those who accept (1) might query whether we should deny *so quickly* that Archie has this reason. Whether he really does have the reason, I believe, will turn on further details about the nature of his insensitivity (details not filled in by Millgram). However, before considering those details we also need some additional apparatus; so I’ll postpone further discussion of this issue for later (§6). Second, though, even *if* it is correct to deny that Archie has this reason, there are other cases where *Response I* is less credible. Nate gives a case in point. For, as Schroeder writes (2007, 165–6):

If what Nate enjoyed immensely was playing poker, then the fact there is poker being played in the living room would be a reason for Nate to go in. If what Nate really enjoyed was watching TV, then the fact that there is a TV in the living room would be a reason for him to go in. So I see no obvious reason to insist that Nate’s enjoyment of

successful surprise parties must be different. [...] It is not as if such reasons don't matter, after all—they still play a role in determining what Nate ought to do—they still show up on God's list of pros and cons.

In short, given that Nate loves surprise parties, it looks ad hoc and otherwise implausible to deny that the fact there is a surprise party in the lounge is a reason for him to go there. We therefore require an alternative response. If we grant that the agents in question have the reasons attributed to them, we'll need to show that they *could* act for those reasons.

Response II: So let's suppose that the fact that Archie would upset people is a reason for him to avoid the funeral. Millgram, recall, presents Archie's insensitivity as a *deliberative* incapacity, rather than a *motivational impossibility*. Indeed, he grants that Archie has some motive which would be served by avoiding the funeral—a desire to do things that make his life go better, say—whereby, one might urge, it remains a motivational possibility that Archie avoids the funeral for the reason that he would upset people. Sure, given Archie's *deliberative* incapacity he *will not* act for this reason. But this does not show he *cannot* do so, so long as the modality in (1) is understood motivationally.¹³

Two difficulties attend this response, however. First, consider Nate again. He has a motive which would be served by going to the lounge. Unlike Archie, he suffers no deliberative incapacity. Yet he is still unable to go to the lounge for the reason there is a surprise party there. So, restricting possibility in (1) to motivational possibility is not a generalising solution. Second, such a restriction looks both question-begging and ad hoc. For given that Archie's insensitivity is a deliberative incapacity—a general incapacitating condition, not a merely localized or temporary ailment—this *also* represents an inability that makes it *impossible* for Archie to act for the reason. It is therefore ad hoc to restrict 'possibility' in (1) to include motivational possibility but exclude deliberative possibility. Further independent support for this modal restriction is therefore needed, even to deal with Archie.

Response III: Millgram anticipates a different response, which, if satisfactory, could also obviate the second difficulty with *Response II* (1997, 218, fn. 14). It trades on our *idealizing* Archie's

¹³This, in effect, is part of the response given by Joyce 2001, 114–5.

deliberative abilities: given that Archie has some motive which would be served by avoiding the funeral for the reason that he would upset people, then *were Archie a better deliberator* he would appreciate that he has this reason in the ways needed to be able act for it. To elucidate, consider a deliberatively ideal counterpart to Archie: Archie⁺. Archie⁺'s motives are identical to Archie's. But, because Archie⁺ does not suffer the deliberative incapacities Archie suffers, Archie⁺ would appreciate the effects of his insensitivity in the ways needed to not act insensitively. Archie⁺ could therefore avoid the funeral for the reason that his insensitivity will upset people. Hence, it is also possible that *Archie* avoids the funeral for that reason, in the sense that if he were a better deliberator he could appreciate the reason in the ways needed to act for it.

Millgram offers the following riposte. If Archie and Archie⁺ have identical motives, and if Archie's motives are 'not particularly focused on improving the predicament of others, then [Archie⁺] will be unlikely to notice the relevant features of [Archie's] circumstances, or be able to think them through helpfully' (*ibid.*). Therefore, since Archie⁺'s motives are just like Archie's, Archie⁺ would also be insufficiently receptive to the reason to avoid the funeral. This is inconclusive, though. Archie's deficiency, Millgram originally stipulates, lies in his deliberative incapacity, not his motives. But Millgram's rejoinder now explains Archie's inability (to avoid the funeral for the relevant reason) in terms of either the absence of a relevant motive or an inadequate focus of the motives he does have. This is problematic in two respects. First, if Archie lacks the relevant motive, Millgram's internalist opponents may find the example less intuitively compelling and hence less incumbent to address. Second, if instead we assume that Archie does have the relevant motive and we then treat his inability as deliberative (such that, e.g., better deliberation would improve the focus of his motives), then Archie⁺ will get right what Archie gets wrong, precisely because Archie⁺ does not suffer the deliberative incapacities that prevent Archie from appreciating the reason. Thus, the claim central to *Response III* remains in place: since Archie would appreciate the reason were he a better deliberator, he could act for that reason (even if he, as he actually is, doesn't so act).

Nonetheless, there remain two serious problems with *Response III*. First, insensitive Archie—that person, as *he* is, with the deliberative incapacity he *actually* has—*could not* act for the

relevant reason. It is only some non-actual, hypothetically improved, version of Archie who could act for this reason. Absent a justification for restricting possibility to motivational but not also deliberative possibility, the second worry with *Response II* resurfaces.¹⁴ Second, Schroeder's example is immune to *Response III*: Nate has a motive that would be served by going to the lounge; he suffers no deliberative incapacity to idealize away; yet he could not go to the lounge for the reason that there is a surprise party there. *Response III* does not generalize.

So the three responses canvassed are inadequate: there are doubts about each as applied to Archie; none look plausible with respect to Nate. We therefore need a different approach. It would, moreover, be pleasing if we could provide a single unifying response dealing with Archie and Nate—and any similarly structured case. The rest of the paper seeks just that.

4. Identifying reasons

My starting hypothesis: whether it is correct to say that a person *could (or could not) act for a specific reason* depends in part on how we *specify that reason*. This section outlines a framework for understanding what is involved in acting for a reason. §§5–6 then apply it to Nate and Archie. I'll begin, though, by motivating the general approach, first drawing attention to two basic 'adequacy requirements' on an account of acting for a reason and then raising some complications that a satisfactory model should also be able to handle.

First, we would expect any account of what is involved in *acting for a reason* to explain when a person counts as acting for a particular reason and what that reason is. This is crucial for adjudicating the Practicality Thesis. It is also important if we want to evaluate not just what people do but the reasons for which they do those things. More generally, we would simply expect a philosophical model of action and motivation, insofar as it talks about acting for reasons, to be able to distinguish the different reasons for which people do or might act. Given our interest in the Practicality Thesis, let's concentrate on cases of acting for a *normative* reason.

The following comprises a very basic adequacy requirement on any account of this:

¹⁴Joyce (2001, 115) supplements a version of *Response III* by restricting 'possibility' in (1) so that it 'extends to granting the agent full information and no deliberative flaws, and not beyond', concluding that Millgram's objection 'has not succeeded in refuting [premise (1)]'. Granted, *Response III* shows that Millgram's objection does not decisively *refute* (1). Nevertheless, this is at best a partial victory: the deeper issue remains why this is an appropriate (and non-question-begging) modal restriction.

(AR₁) An account of acting for a normative reason must be able to explain when and why a person counts as acting for a particular normative reason (rather than some other, or no, such reason).

When someone acts ‘for a reason’, she is in some psychological (motivating) state that helps explain why she so acts. So claims about acting for a reason involve claims about our psychological (motivating) states. People can perform a given action type for different reasons. To tell which reason a person acts for, and hence to satisfy AR₁, we need a way to identify and distinguish the reasons for which she does or might act. The most obvious, perhaps only, informative way to do that is in terms of the *content(s)* of the psychological state(s) the person is in when she acts for a reason. Furthermore, though, to count as acting for *a particular normative reason* there must also be some relevant connection between the psychological state the person is in when she acts and *that* normative reason. We also therefore need a way to identify the relevant normative reason(s) for which she is or might be acting. This gives us a second (subsidiary) adequacy requirement:

(AR₂) To satisfy AR₁, we need to show how the psychological state a person is in when she acts for a reason is connected to a particular normative reason in a way that explains why that is the normative reason she is acting for.

As noted earlier (§1), we identify and individuate a specific normative reason primarily via the reason-giving facts providing that reason. The only way to satisfy AR₂, as far as I see, is to show how the content of the psychological state the person is in when she acts for a reason refers in some appropriate way (in a sense to be explained in due course) to the normative reason and/or the reason-giving facts providing that normative reason. The framework developed in the rest of this section provides a neat way to discharge our two adequacy requirements. That counts in its favour and provides some independent justification for it.¹⁵

The basic idea, that we can and should identify the reasons people act for via their contents, rests on a very minimal claim: acting for a reason involves being in some psychological

¹⁵ One way to oppose my later defence of the Practicality Thesis, then, would be to contest the framework it rests on and supplant it with an alternative which not only precludes analogous defences but also discharges our basic adequacy requirements and accommodates the complications to be noted next (or else explains why these need not be discharged or accommodated). I cannot in advance rule out all possibility of such alternatives; but opponents would need to give the relevant arguments.

(motivating) state with an expressible content. It should likewise be uncontroversial that, when a person acts for a reason, the content of the motivating state she is thereby in helps explain why she is so acting. As we could put this: that content helps explain why she so acts by comprising the description under which she so acts.¹⁶ This allows us to say that, to count as acting for a particular normative reason, the description under which one acts must make appropriate reference to that normative reason and/or the reason-giving facts providing it. It is a commonplace that a single fact (thing, feature, property, state of affairs, etc.) can be specified or described in different ways. This is true also of the reason-giving facts by which we individuate reasons. Some descriptions of a particular reason-giving fact make it more obvious or more immediately intelligible why the obtaining of that fact provides the reason it does. Nonetheless, if the fact that p as described one way gives you a specific reason R to ϕ , then (holding all else constant) that p described a different way gives you that same reason. That should be uncontroversial. It raises several complications, however, when it comes to acting for a reason. First, if a reason or reason-giving fact can be correctly described in different ways, there is the possibility that someone who acts for a particular reason could act for that reason under different descriptions. Second, as noted in §§1–2, typically more than one fact contributes to the obtaining of a single normative reason. That opens up the possibility of a far wider range of descriptions of any such reason and, in turn, a wider range of descriptions under which one might be aptly said to act for that reason. Third, there may then be some descriptions of a reason under which one *can* act, but others under which one *cannot*. These are rarely acknowledged complications. The account to follow accommodates them; and that should again count in its favour.

Doing so also opens up a line of response to the Archie and Nate cases. The crucial idea is this: it is indeed possible, given their motives and/or reason-responsive capacities, that they act under some description, such that, by so acting under that description, they would count as

¹⁶The ‘minimal claim’ should be acceptable to views as otherwise diverse and opposed as, e.g., Davidson 1963, Alvarez 2010. Although widely agreed that *acting for a reason* marks a species of *intentional action*, to avoid proliferating terminology (as well as controversies about intentions not crucial here) I’ll avoid such talk. Note, also, that this (Anscombian-Davidsonian) talk of acting ‘under a description’ is merely shorthand for the idea that the motivating state a person is in when she acts for a reason (typically) has expressible content. I think this is the best way to explicate what is involved in acting for a reason, though the later arguments could be reformulated via alternative idioms. At any rate, this way of putting things does not beg the question against the Millgram/Schroeder cases: these require that there is some thought which Archie and Nate would have to act under in order to count as acting for the relevant reason, the objection being that they cannot have that thought or therefore act for the reason.

acting for the normative reasons attributed to them. Thus they can act for these reasons, as the Practicality Thesis requires. To show how this works, we need to present the basic framework in a little more detail. I'll develop it in four stages.

Stage I. More than one fact typically contributes to there being a single normative reason. Suppose that Ann has a reason *R* to pick up her umbrella. We might identify *R* via the fact that *it is raining*. Nonetheless, that *it is raining* is not by itself sufficient to give Ann this specific reason. Various other facts must obtain too: that *Ann is going outside—to a party*, say, that *being soaked to the skin at the party will be uncomfortable*, that *Ann's umbrella will keep her dry*, and so forth. Such facts could be conceptually independent of one another; but they each contribute to this reason *R*. A comprehensive characterization of *R* would cite all facts relevantly necessary¹⁷ and jointly sufficient for the obtaining of *R*. It is therefore more accurate to say 'the set of facts that $p_1...p_n$ contribute to there being a reason *R* for *A* to ϕ '.¹⁸ In practice, we often cite a single fact when picking out a reason. Which fact we cite can depend on various pragmatic and contextual factors: the pragmatics of communicating a reason claim, the context of knowledge we assume amongst those we are communicating with, and so on. Thus, we might explain why Ann has a reason to grab her umbrella succinctly by saying 'it is raining'; while to someone who knows *that*, but who does not see why Ann has a reason to get the umbrella, we might cite the fact that she is going outside. Nonetheless, although we often cite just one fact as 'the reason', usually more than one fact contributes to a single reason.

Stage II. Next, we need to elaborate a little on what is involved in acting *for a reason*. The following is far from exhaustive but will suffice for present purposes. When *A* acts for a reason, *A* is in some psychological state that explains (or, more accurately, contributes to the explanation of) *A*'s so acting. However, *A* could be in some psychological state that explains her acting, even

¹⁷I here bracket complications surrounding what counts as *relevantly* necessary—including how to circumscribe the set of reason-contributing facts in a suitably narrow, localized way. This remains a significant (and widely unresolved) issue.

¹⁸This is compatible with the following theses: (a) A single fact could contribute to more than one reason: that *it's raining* may contribute to *A*'s having a reason to take an umbrella and *A*'s having a reason to take a taxi. (b) Some facts that are not necessary for the *obtaining* of *R* may affect *R*'s *weight*: that taxis are expensive is not a necessary condition for Ann's having a reason to take an umbrella, but it might make it stronger. (c) Different facts can play different roles in how they contribute to a reason—as direct favorers, prerequisites and enablers for favorers, background conditions, weight intensifiers, and more (cp. Dancy 2004, Ch.3; Schroeder 2007, Ch.2).

though she is not acting for a reason.¹⁹ Acting for a reason is a narrower category: the psychological state A is in when A acts for a reason involves some attitude or disposition (consciously acknowledged or not) that casts the action in a favourable light. (This is sometimes marked as a distinction between A's being in a motivating *state* and A's having a motivating *reason* proper: when A acts for a reason the motivating state A is in is a motivating reason.) And, when A acts for a reason, the thought under which A acts has a content, such that A acts under a description. The description under which A acts need not involve explicit reference to 'reasons', as in 'that *p* is a reason for me to ϕ '. The description could just take the form '*p*' or 'I will ϕ because *p*'. The content of this description, as given by *p*, must nonetheless help explain or make intelligible why A ϕ s, by casting the action in a favourable light (at least given other facts about A's desires, dispositions, evaluative commitments, etc.). And when A acts for a reason R, an adequate specification of the description under which A acts must render sufficiently intelligible why A's acting counts as acting for R. One obvious way in which A's ϕ ing under description '*p*' renders sufficiently intelligible why A ϕ s for reason R would be that the content of that description denotes the same item in whose terms we individuate R: namely, the R-giving fact that *p*. Putting these points together:

If that *p* is a normative reason R for A to ϕ , and if A ϕ s for reason R, then: (i) A ϕ s under some description (such as '*p*')²⁰; (ii) A's ϕ ing under that description explains why A ϕ s by casting her ϕ ing in a favourable light; and (iii) A's ϕ ing under that description renders intelligible why this counts as ϕ ing for R.

Suppose, for illustration, that the fact that *it is raining* were the only fact contributing to a reason R for Ann to take an umbrella. Thus, if Ann takes an umbrella for R: (i) Ann takes an umbrella under some description, such as 'it is raining', where (ii) the content of that description explains why she is taking an umbrella, by casting that act in a favourable light (given her desire to stay

¹⁹Such cases are commonplace and familiar: you might be driven by a compulsion or addiction over which you have no control; you might act out of habit in ways you would not act were you paying greater attention; etc. To anticipate a different case I'll return to: your false belief that you are acting for a normative reason may explain your action; but if there is no such reason, you are not acting for what actually is a normative reason.

²⁰A more precise rendering is given in *Stages III* and *IV*, and §5.

dry, say), and (iii) her taking an umbrella under that description explains why this counts as acting for R , since ‘it is raining’ denotes the same fact by which we identify R .

Stage III. When more than one fact—some set of facts that $p_1...p_n$ —contributes to A ’s having reason R to ϕ , then if A ϕ s for R the content of the description under which A ϕ s must include appropriate reference to (e.g., by picking out) *at least one of* the salient R -contributing facts that $p_1...p_n$.²¹ Suppose that the facts that *it is raining* and that *being soaked will be uncomfortable for Ann* were the only two facts contributing to Ann’s reason R to take an umbrella. Then if she takes an umbrella for R , the content of the description under which she acts will make appropriate reference to at least one of those R -contributing facts: that *it is raining* or that *being soaked will be uncomfortable*.²² Ann might act under just one such description. She may think to herself ‘I’ll take an umbrella because it’s raining’, without entertaining all R -contributing facts (just as we, for pragmatic purposes, might cite only one R -contributing fact when specifying R). Nonetheless, whichever of these descriptions she acts under, in typical cases she counts as acting for R if, but only if, the content of the description under which she acts makes appropriate reference to at least one salient R -contributing fact.²³

Stage IV. Given this, to show that A *can* act for normative reason R , we need to show there is at least one R -contributing fact the description of which could figure in or as the content of a description under which A acts. More precisely: assuming that $p_1...p_n$ contribute to A ’s having a reason R to ϕ and that A ϕ s for that reason, it must be possible that the description under which A ϕ s makes appropriate reference to at least one of the R -contributing facts that $p_1...p_n$. It must be *possible* in that, even if it is not possible that A ϕ s under description ‘ p_1 ’, there is some other description (e.g. ‘ p_2 ’) that it is possible A ϕ s under. Thus:

²¹I say more about ‘appropriate reference’ in §5.

²²The description under which Ann acts might refer to herself: ‘I want to avoid being uncomfortable’. On some views, ‘I want to avoid being uncomfortable’ and ‘Ann wants to avoid being uncomfortable’ express different propositions. So the point is not that the content of the description *under which A acts* is identical to the content by which *we pick out* the reason-giving fact. Rather, the content of the description under which A acts picks out the same *item* as that denoted by a correct *specification* of the reason-giving fact.

²³The content of the description under which A acts could denote a fact contributing to more than one reason. Ann might act under the description ‘it is raining’, where (a) the fact that *it is raining*, alongside the fact that *being soaked is uncomfortable*, together contribute to reason R_1 to take an umbrella, and (b) the fact that *it is raining*, alongside the fact that *using an umbrella in rainy conditions is fashionable*, together contribute to reason R_2 to take an umbrella. In that case, to assess whether A acts for R_1 or R_2 (or both), we may need a fuller explanation of A ’s thoughts (though sometimes it might be indeterminate; see also §5).

If that $p_1...p_n$ contribute to A's having a reason R to ϕ , and if it is possible that A ϕ s for reason R , then (i) it must be possible that the description under which A ϕ s makes appropriate reference to at least one of the R -contributing facts that $p_1...p_n$, where (ii) A's ϕ ing under that description explains why A ϕ s (by casting ϕ ing in a favourable light) and renders intelligible why that counts as ϕ ing for reason R .

Accepting this opens up a line of response to the Archie and Nate examples.

5. Nate's reason

This section focuses on Nate. §5.1 provides two sample descriptions under which Nate could enter the lounge and then a 'Basic Argument' to show that, were he to do so, he thereby counts as acting for the reason Schroeder attributes to him. The remainder of the section wards off some possible objections, using these to further elucidate and consolidate the Basic Argument.

5.1 *The Basic Argument*. So the situation is this: that there is a surprise party awaiting Nate in the lounge contributes to a reason R for Nate to enter the lounge; but, Schroeder's objection runs, it is not possible that Nate enters the lounge for that reason R . A key assumption underlying the objection is that it is not possible that Nate enters the lounge for R because he could not enter the lounge under the description 'there is a surprise party' (or any description mentioning the surprise party). My strategy, then, will be twofold: to show, firstly, that there are other descriptions of R the content of which could figure in or as the content of a description under which Nate could enter the lounge, and, secondly, to explain why, were Nate to enter the lounge under such a description, he counts as acting for R .

Here are two descriptions under which it is possible that Nate enters the lounge:

D1: I have a reason to go to the lounge.

D2: I will enjoy myself.

It is *possible* that Nate enters the lounge under either description in the sense that, were he aware either that he has a reason to go there or that he would enjoy himself were he to go there, D1 and D2 are descriptions under which he could enter the lounge—such that, by so acting, the

reason identified by Schroeder in terms of there being a surprise party does not disappear.²⁴ There are various ways Nate could become aware of one of these things; but, for simplicity, let's suppose we tell him. The question now is whether, by acting under one of these descriptions, he could count as acting for *R*. Consider D1 and D2 in turn.

Firstly, suppose we tell Nate there is a reason for him to go to the lounge. We don't reveal any more details about what that reason is—we don't tell him there is a surprise party, for instance; and we convey things in such a way that prevents him from suspecting there's a surprise party. Nate trusts what we say and now trundles off to the lounge under D1. So: there is a normative reason *R* for Nate to go into the lounge; he goes into the lounge for a reason; the reason he goes into the lounge for is the reason we have told him (truly) there is (he is not acting for any other reason); the reason we have told him there is, *is* the normative reason *R* to do with the surprise party. Hence, Nate could act for that reason *R* which Schroeder identifies via the fact that *there is a surprise party*, even if he does not act under a description the content of which includes explicit²⁵ reference to (e.g. mention of) the surprise party.²⁶

Secondly, we might instead tell Nate that he has this reason *R* by citing some *R*-contributing fact (besides there being a surprise party), where the description we cite of that fact could figure within the content of the description under which Nate enters the lounge. The most obvious facts are to do with Nate's enjoyment. One such fact is that Nate would enjoy the surprise party. Nate could not enter the lounge under the description 'I will enjoy the surprise party', of course. Nevertheless, there are other facts concerning what Nate would enjoy that contribute to this reason *R* to enter the lounge, given that there is a surprise party there: for

²⁴What if Nate remains *unaware* that he has a reason to go to the lounge or would enjoy himself—would that undermine this line of argument? No, for two reasons. First, Schroeder and I are talking about 'objective reasons': reasons the obtaining of which does not depend on the agent's beliefs about or epistemic access to the reason-giving facts. We are therefore both allowing that a further necessary condition for being able to act for a reason *R* is that the agent is aware of the reason (under some description) or aware of at least some *R*-giving facts. (This condition can be incorporated into CA and assumed throughout.) Second, although there is a sense in which Nate *could not* enter the lounge for *R* were he unaware of *R*, it is a different sense (registering an *epistemic* modality) from that in virtue of which Schroeder's example gets going (which trades on an alleged *conceptual* or *motivational* impossibility). To see this, we need only note that, were Nate to become aware of *R* under description D1 or D2, *R* does not 'go away'—whereas, were Nate aware of the surprise party, *R* would go away.

²⁵We return to the significance of this in §5.3.

²⁶Schroeder's other views might commit him to accepting this. For he holds that (a) the fact that there is a reason for *A* to ϕ *is itself* a reason for *A* to ϕ (these are not independent reasons; 2007, 32–33)—thus allowing that (b) when Nate acts under D1, the reason Nate identifies *is* the surprise party reason. (This is merely *ad hominem*; I'm not committing to (a).)

example, that *he would enjoy himself*. And ‘I will enjoy myself’, i.e. D2, *is* a description under which Nate could enter the lounge. In which case, since the fact that Nate would enjoy himself contributes to R, and since he could enter the lounge under description D2, where the content of that description denotes one of the R-giving facts, Nate could go into the lounge for reason R. He is, after all, acting for a reason; but he is not acting for any reason other than R. Indeed, he is acting for the reason we have told him there is—namely, that reason R to which the fact that there is a surprise party contributes and that Schroeder identifies in terms of there being a surprise party. Hence, Nate can go to the lounge for the reason that there is a surprise party.

That is the basic argument. We can summarize it as follows. Granting that it is possible that Nate acts under D1 or D2:

Basic Argument: (i) When Nate acts under D1 or D2 in light of what we’re told him, he acts for a reason. (ii) The reason he acts for is the reason we’ve told him there is. (iii) The reason we’ve told him there is, is the reason Schroeder identifies via the fact that *there is a surprise party*. So, (iv) when Nate acts under D1 or D2 in light of what we’re told him, the reason he acts for is the reason Schroeder identifies via the fact that *there is a surprise party*.

Simple as it is, this Basic Argument seems to me plausible. Nonetheless, I foresee two main types of objection to it. One of these queries whether, when Nate acts under D1 or D2 in light of what we’ve told him, he is acting for a *normative* reason; the other queries whether, even if he is acting for a normative reason, that reason is the reason Schroeder identifies via the surprise party. Let’s consider these in turn.

5.2 Merely explanatory reasons? A first objection to the Basic Argument concerns the phrase ‘acts for a reason’ in (i). This, one might allege, is ambiguous between *explanatory* and *normative* reasons. For the Basic Argument to go through, it needs to be that Nate acts for a normative reason. But, the objection goes, ‘acts for a reason’ in (i) signals only an explanation of Nate’s action (an explanatory or motivating reason), not a normative reason. Thus, when he acts under D1 or D2, although this explains his acting it does not show that he is acting for the normative reason that

there is a surprise party—because he is not acting for a normative reason at all. There are two ways to develop this objection.

Version (a): Suppose that, having told Nate he has a reason to go to the lounge (and/or that he has this reason because he'll enjoy himself), he goes to the lounge because he *believes that* there is a reason to go there (or *believes that* he'll enjoy himself). In standard cases, though, a person's beliefs are *not* themselves normative reasons. So, when Nate enters the lounge because he *believes that* D1 or D2, he is not entering the lounge for a normative reason.

Version (b): Alternatively, one might suppose that what explains why Nate goes to the lounge is the fact that we've told him there is a normative reason to go there (and/or told him that he'll enjoy himself). But the fact that *we've told him* there is a normative reason (etc.) is not itself a normative reason for him to act. So, he does not go to the lounge for a normative reason.

The response to both (a) and (b) is similar. Regarding (a), one can grant both that Nate's belief that he has a reason to go to the lounge contributes to the explanation of his going there *and* that this belief is not a normative reason. Yet granting that does nothing to undermine the claim that, when he enters the lounge in light of his belief the *content* of which is D1 or D2, he is acting for the normative reason identified by that content. And since the normative reason he believes he has is the normative reason we have told him (truly) there is, he is acting for the reason we have told him there is. Similarly, regarding (b), one can grant that the normative reason for Nate to go into the lounge is not *that we tell him* there is a reason or he'll enjoy himself. Rather, the normative reason is the reason we tell him there is, i.e. the reason we pick out when we tell him that he has a reason or he'll enjoy himself. Hence, this leaves open that he can act for the normative reason we've told him (truly) there is, where that reason is denoted by the content of the description under which he acts.

To respond to the more general concern, the phrase 'acts for a reason' in (i) denotes *both* an explanatory and a normative reason. When Nate enters the lounge under D1 or D2, this explains his so acting. And the content of that description denotes a normative reason. Hence, the content of the description that explains his so acting expresses a truth (which Nate knows, or believes truly, with justification, etc.) about the normative reason there actually is.

5.3 *Too tenuous?* Suppose, then, that by entering the lounge under D1 or D2, Nate could be acting for a normative reason. Nonetheless, one might now object, it's not obvious that the normative reason he acts for is the reason *to do with the surprise party*. I can think of three ways to develop the objection. Each suggests that acting under D1 or D2 is too tenuously related to the surprise party for Nate to count as acting for the reason that there is a surprise party. I'll call this the Too Tenuous Objection.

A first version claims that, because the surprise party is an essential part of R, for Nate to count as acting for R the description under which he enters the lounge must refer to the surprise party by *explicitly mentioning* it. Hence, because neither D1 nor D2 mention the surprise party, when Nate acts under those descriptions he is not acting for R. Note that this objection rests on an unstated assumption:

(*) When there is more than one description of a reason R, there is a particular description under which one must act if one is to count as acting for R.

However, there are at least three reasons to doubt that (*) is true in the way needed for the Nate example to work against (1). First, it would commit us to denying more generally that a person could act for a normative reason on the basis of testimony which, though known to be true and reliable, tells one *that* but not *why* there is a reason. This does not seem correct in other cases. Suppose that I know that the fact that an old friend is in the lounge is a reason for you to go there; I now tell you there is a reason for you to go to the lounge but I don't say why. Insofar as you know (or believe with justification, in a non-deviant way, etc.) that I am reliable or telling the truth, you could come to know (etc.) there is a good reason to go to the lounge—even though you don't know why. You are then in a position to act for the normative reason I've told you there is. Denying this—and denying more generally that one can act for normative reasons on the basis of (true, reliable, etc.) testimony telling us that, but not why, there is a reason—is a big commitment that at least needs considerable support. Second, the earlier discussion of Ann's reason to take an umbrella was also intended to illustrate that (*) is false: *Stage III* suggested that Ann might take her umbrella under the description 'it is raining' or under the description 'being

soaked will be uncomfortable’, where by acting under *either* description she could be acting for one and the same reason. Such cases are commonplace. Third, it is a familiar idea that the kinds of thick features by which we identify actions as, e.g., virtuous need not be (and often are not) the descriptions under which people exemplifying those virtues perform those actions.²⁷ It can be apt to describe actions as *courageous*, *generous* or *modest*, say, where from a third-person perspective we view such thick features as reason-giving. But people acting courageously, generously or modestly need not act under descriptions the content of which refer explicitly to ‘courage’, ‘generosity’ or ‘modesty’—indeed, in many cases, acting under those descriptions indicates a distinctive lack of the relevant quality. So (*) looks false, not just in a few isolated cases, but quite widely. If it *has to* apply in Nate’s case specifically, we are owed an explanation why.

Now for a second version of the objection: it might be suggested that Nate isn’t really acting for R because, even if the description under which he acts need not explicitly mention the surprise party, he would at least have to be *aware* that there is a surprise party. This relies on a different assumption:

(**) For A to count as ϕ ing for reason R, there must be some specific R-contributing fact of which A is aware.

Again, though, (**) is false, and for similar reasons to (*). First, the case of acting for reasons from testimony suggests that one can indeed act for a reason R without being aware of *any* of the R-contributing facts. Second, even if one did need to be aware of some R-contributing facts to count as acting for R, it is not obvious that there must be one of these in particular which one must be aware of. Suppose that the following facts are each necessary for Ann’s having a reason to take her umbrella: that it is raining, that she doesn’t want to be uncomfortable at the party, and that her umbrella (which, unbeknown to Ann, had lost its waterproof qualities) has recently been waxed by a friend to make it waterproof again. Ann could be acting for R when she grabs her umbrella under the description ‘it is raining’, even if she is unaware that the umbrella has now been re-waterproofed. She might even be unaware that it is raining (or believe falsely that it is not raining); yet if we tell her that she has reason R to take an umbrella because doing so will make

²⁷See e.g. Williams 1985, 10.

her party experience more enjoyable (but we don't mention the rain), she could again count as acting for *R*. Third, even though from a third-person perspective it may be apt for us to identify a person's reasons via certain virtue-descriptions—that's the thing to do because it is courageous (etc.)—the person acting virtuously need not conceptualize her situation in a way displaying awareness of those same reason-indicating descriptions. These sorts of considerations suggest that (**) is generally false. If (**) *has to* apply in Nate's case, it again remains to be seen why.

There is, however, a third and perhaps more forceful variant of the Too Tenuous Objection. Nate could act under D1 or D2 *whether or not* there is a surprise party and hence *whether or not* there is this normative reason. Thus, the objection runs, the connection between being able to act under D1/D2 and being able to act for reason *R* is too tenuous, because the fact that there is a surprise party fails to explain why, or place any constraints on whether, by acting under D1/D2 Nate acts for the normative reason to which the surprise party contributes.

Note in response, however, the fact that there is a surprise party *does* explain, and place an important constraint on, whether by acting under D1/D2 Nate acts for *R*. For if there were no surprise party, then if Nate enters the lounge under D1/D2 he is not acting for normative reason *R*—since there is no such reason. There will be things that explain his action: his beliefs, his trusting our testimony, or whatever. But if there is no such normative reason, it is not possible (in the relevant sense) that he acts *for that normative reason*.²⁸ We can draw two important points from this.

First, even though when Nate acts under D1/D2 and there is *no* surprise party he does not act for reason *R*, this does not show that when he acts under D1/D2 and there *is* a surprise party he is not acting for *R*. That would be a straightforward fallacy. Second, more significantly, we can now see why, when Nate acts under D1/D2 and there is a surprise party, far from being tenuous the connection between his acting under those descriptions and reason *R* is actually rather strong. For when he acts under D1/D2, he acts for *R* *only if* and (in part) *because there is a*

²⁸In that situation, his belief that there is such a normative reason would therefore be false, as would his belief that he is acting for it. This should be uncontroversial. However, it is easy to be misled here. When we say 'A acts for a reason' this can conceal an ambiguity (see also Alvarez 2010, Ch.5). It could be true if 'a reason' denotes only an explanatory or motivating reason, whereby one means that A acts under some description that explains what A is doing. Yet the same statement could be false if 'a reason' implicates a normative reason and there is no such reason. Denying that implies infallibilism regarding reflexive beliefs about acting for normative reasons; but such an infallibilism looks implausible.

surprise party. That there is a surprise party thereby plays an explanatory role with respect to *whether* by acting under D1/D2 Nate counts as acting for the reason R Schroeder identifies via the surprise party. For if there were no surprise party, by acting under D1/D2 he would not count as acting for R, since there is no such reason; and when we tell Nate he has a reason to go to the lounge or that he has the reason because he'll enjoy himself, what we tell him is true *because there is a surprise party*. So, the fact there is a surprise party contributes to the explanation of whether his acting under one of these descriptions counts as acting for R. This implies that acting for a normative reason depends *not just* on the narrow content of the description under which one acts, but *also* on the content of that description being suitably related to how things actually are (even if one is not fully aware of how things are).²⁹

Hence the third version of the Too Tenuous objection fails. These points also generate a further positive rationale for the conclusion that by acting under D1/D2 Nate acts for the reason Schroeder identifies via the surprise party: Nate acts for that reason because the proposition picked out within the content of the description under which he acts is true (since it denotes either a reason or reason-contributing fact)—where its truth is *explained by the obtaining of the reason-contributing fact that there is a surprise party*. Hence, Nate counts as acting for reason R because: firstly, he is acting for the reason we have told him there is and that reason is R (that was the Basic Argument); and, secondly, the obtaining of R is explained by the fact that there is a surprise party. I'll later use these ideas to defend the Practicality Thesis against Millgram's objection. First, though, it will be useful to ward off some further possible qualms.

5.4 *Other objections*. First, an 'Indeterminacy Objection'. Suppose there are many things Nate would enjoy if he went into the lounge—watching TV perhaps—and hence many reasons for him to do so. Then, one might object, when he enters the lounge under a description like D2,

²⁹This might be thought contentious because it implicates a certain (not uncontroversial, albeit widespread) view about mental content. I suspect that, if you find it contentious in this context, that's due to an ambiguity in the phrase 'acting for a normative reason'. On a *factive* reading: A acts for normative reason R under a description mentioning an R-contributing fact that *p*, only if *p*. On a *non-factive* reading: it is possible that A acts for normative reason R under some description mentioning that *p*, even if not-*p*. All I need here is that the factive reading is legitimate. (One way to then reparse the non-factive reading is to say that, if the supposed R-contributing fact that *p* as mentioned in the description under which A acts fails to obtain, A is not acting for R even though he may believe he is.) Opening the door to a wide-content view in this context does not commit one to a more global externalism about mental content, of course.

this description is too course-grained to determinately conclude that he acts for *R*: it leaves it indeterminate whether he is acting for *R*, by failing to rule out the possibility that he enters the lounge for some other enjoyment-directed reason, *R**, such as the reason that he would enjoy watching TV. There are two things to say in response.

First, we can grant that it may indeed be indeterminate which reason Nate is acting for. However, that does not show that he *could not* act for *R*. To defend (1), all we need is that it is *possible* that he acts for reason *R*—not that he definitively *does* act for this reason, nor that we could tell he does. Second, even if D2 is too course-grained to definitively vindicate the conclusion that Nate acts for *R*, we can fine-grain our explanation of *R* when telling Nate that he has this reason in ways that make it increasingly determinate that this is the reason he is acting for. For instance, we could cite certain facts in virtue of which Nate would enjoy himself were he go to the lounge: that his friend Ann is there, that we've got some nice IPA, or so on. Such facts contribute to the obtaining of this specific surprise party—thus making it more determinate that when he acts under a description the narrow content of which mentions one of these facts, the reason he is acting for is the reason we have told him there is (namely, that there is a surprise party). Hence, so long as there is *some* such description suitably related to this surprise party under which Nate could go to the lounge, he can act for the reason that there is the surprise party, even though he need not thereby act under a description the (narrow) content of which mentions the party.

A further objection (implicitly anticipated by Schroeder, 2007: 165), which we could call the Too Easy Objection, is that by explaining how it is possible that agents (Nate included) act for a reason in terms of there being *some suitable description of that reason under which they could so act*, this makes it *too easy* to satisfy the modal requirement in (1). For, the worry goes, it is possible that an agent acts, or could be brought to act, under just about *any* description. (This might be thought problematic if the descriptions under which I'm suggesting Nate counts as acting for *R* are somewhat thin, as in D1 and D2.) In which case, if (1) could be satisfied by just about any agent for just about any reason, it threatens to be trivial.

One response to the Too Easy Objection is to say that, *if* (1) is trivial, then so be it. (That might be all to the good of the internalist's argument CA, at least insofar as the other premises by which CA arrives at the non-trivial conclusion RI* are not also trivial—if they were, something would have gone awry.) However, (1) is not trivial, and for three reasons.

Firstly, it requires not just that there is some description under which one could act, but that one can in fact perform the action specified. Secondly, being able to act for a reason (under some description) requires that one *could* be motivated to so act (as (2) has it); insofar as one's account of motivation allows that there are some descriptions under which a person *could not* be motivated to act, it allows that there are indeed some reasons a person could not act for. In the case of the internalist argument CA, it is premise (3) as an application of the Sentimentalist Thesis (ST) that generates a non-trivial restriction on what one can be motivated to do. For A to be motivated to ϕ (under a given description), there must be some route from the things A already cares about, as embodied in her existing motives, to her ϕ ing (under that description). There may be many ways A could be brought, given the actual motives she has, to be motivated to act under some given description—via deliberation, reflection, advice, conversion, coercion, manipulation, brainwashing, etc. Nonetheless, according to those internalists who endorse both RI* and ST, if A has no motive whatsoever that would be served by ϕ ing under a specified description, then A could not be (brought to be) motivated to ϕ under that description. And there will be many cases in which people could not be brought to be motivated ϕ under any description (not even a description like 'there is a reason for me to ϕ ')—cases where, for instance, ϕ ing is itself too far removed from (or incongruent with) one's existing motives and commitments. ST is controversial, of course; and a defence of it goes beyond the scope of this paper.³⁰ The immediate point, though, is this: so long as one's favoured model of motivation (be it ST or something else) generates a non-trivial restriction on what a person could be motivated to do and hence what reasons they can act for, (1) is not trivial and so does not fall prey to the Too Easy objection. Thirdly, to be able to act for a given normative reason (under some

³⁰ Schroeder's discussion of 'desire' (2007, Ch.8) suggests some sympathy for it. As an ad hominem point, Schroeder sometimes agrees that bringing a person to be motivated to do some things, even things she would enjoy, can require a radical transformation in her psychology more generally, including her motives (2007, 165)—whereby effecting such changes is not actually *too* easy.

description) there must actually be that normative reason. This might initially sound dialectically awkward (e.g. question-begging), at least when (1) also functions as a premise in an argument for an account specifying a condition under which a person has a normative reason. However, even when (1) functions as such a premise, as it does in the argument for RI*, it is a premise in an argument for *just one* such condition (since RI* specifies only a necessary condition). Other conditions might also need to obtain. Hence, even if one could be brought easily to act under some description, this does not count as acting for a given normative reason unless other relevant conditions are met and the reason itself thereby obtains. These three considerations should each block the triviality worry.

5.5 *An alternative strategy.* I therefore think premise (1) can be defended. In case you have remaining doubts, however, there is an alternative approach: to explicate (1) more precisely with reference to the full set of reason-contributing facts.

Rather than saying ‘that p gives A a reason R to ϕ only if it is possible that A ϕ s for R ’, we could say ‘that $p_1...p_n$ give A a reason R to ϕ only if it is possible that A ϕ s for R ’—where the consequent can be unpacked more fully by saying ‘it is possible that A ϕ s under some description the content of which denotes at least one of the R -contributing facts, that $p_1...p_n$ ’. For example, if the fact that there is a surprise party in the lounge and the fact that Nate would enjoy himself were he to go into the lounge were the only two facts contributing to his having reason R to go to the lounge, it must be possible that Nate enters the lounge under a description the content of which denotes at least one of those R -contributing facts. In the context of the Classical Argument for internalism, we could then amend the rest of CA accordingly:

(1*) that $p_1...p_n$ contribute to there being a reason R for A to ϕ only if: it is possible that A ϕ s under some description the content of which denotes at least one of the R -contributing facts that $p_1...p_n$

(2*) it is possible that A ϕ s under some description the content of which denotes (at least one of the) R -contributing facts that $p_1...p_n$ only if: it is possible that A is motivated to ϕ

under some description the content of which denotes at least one of the R-contributing facts that $p_1...p_n$

(3*) it is possible that A is motivated to ϕ under some description the content of which denotes at least one of the R-contributing facts that $p_1...p_n$ only if: A has some motive which, given that $p_1...p_n$, would be served by his ϕ ing under some description the content of which denotes at least one of the R-contributing facts that $p_1...p_n$

Therefore,

(RI**) that $p_1...p_n$ contribute to there being a reason R for A to ϕ only if: A has some motive which, given that $p_1...p_n$, would be served by his ϕ ing under some description the content of which denotes at least one of the R-contributing facts that $p_1...p_n$.

This is more complex, certainly. But the complications concern only how we formulate the argument in a suitably precise way. It is just a precisified argument for the simple idea that an agent's reasons are relativised to her motives. So, even if the resources appealed to in the earlier defence of (1) are too contentious to rely on, internalists could still propose a version of (1) like (1*) and get to internalism via a suitably amended version of CA. (For the same kinds of reasons explained in §1, one could accept (1*) but still reject internalism of course.)

6. Archie's reason

Can a similar response be developed to counteract the Archie-example? If so, it will need to show that there is some description under which Archie could act, where acting under that description counts as avoiding the funeral for the normative reason to which the fact that he will upset people contributes. Now Archie could avoid the funeral under the description 'there is a reason for me to stay away', akin to how D1 functions in Nate's case. If so, the Practicality Thesis is intact. However, I'll here consider some more interesting ways he might be able to act for the reason in question.

We've so far picked out 'the' reason R Archie has to avoid the funeral via the fact that he will upset people. Again, though, there may be a wider range of facts contributing to R. Millgram identifies two: that Archie is insensitive; and that his attending the funeral will be unpleasant (or

otherwise bad) for the people he will upset. (There may be many others: that there is a funeral about to take place; that people attending funerals are emotionally fragile; etc.) The question is whether any of these *R*-contributing facts could figure as the content of a description under which insensitive Archie might avoid the funeral. To answer this we need to know more about the nature and extent of Archie's insensitivity (something under-described by Millgram). I'll show there are at least three different versions of Archie who, although insensitive to some degree, could act under a description where (by so acting) they could count as acting for this reason.

First, one way to show that Archie—call him 'Archie₁'—has reason *R* is by appealing *directly* to one of the *R*-contributing facts he might not have noticed: that it's nasty or inconsiderate to upset people, say. Hence we might tell Archie₁ that he has reason *R* because he would upset people and that would be nasty. *If* Archie₁ grasps this—if, that is, despite his insensitivity he is able to appreciate that this fact gives him a reason to avoid the funeral, in whatever ways are needed for him to be motivated by that reason—then it is possible that Archie₁ acts under a description like 'I will avoid the funeral because, if I go, I'll upset people and that would be nasty'. If so, or so long as Archie₁ can appreciate the reason-giving force of *some* such candidate *R*-contributing fact, it is possible that he acts for *R*. Hence (1) would be safe. Of course, this relies on a big 'if': there may be other versions of Archie incapable of appreciating the reason-giving force of *any* of these candidate *R*-contributing facts. So this direct response may have limited purchase.

There are, nonetheless, several more *indirect* approaches to take. Suppose that Archie₂ finds it unpleasant when others upset him; and suppose he thinks this gives them reasons not to do things which upset him. Then, absent relevant differences between his upsetting others and their upsetting him, he is at least *committed* to accepting that he has a reason to not upset others—and hence committed to accepting that he has a reason to avoid the funeral, given that he would upset people there. Now suppose, further, that Archie₂ *is* able to appreciate this. Then, it seems, there is nothing to preclude his acting under the thought 'I will avoid the funeral because I'll upset people if I go and I know how unpleasant it is to be upset'. In which case, it is possible that Archie₂ acts for reason *R*.

It is worth noting an objection Millgram makes in anticipation to such a response: were Archie₂ able to appreciate this reason, he would *ipso facto* be sensitive enough not to have this reason, since he would then be sensitive enough that he would not upset people. This is a little quick, however. Insensitivity comes in *types* and *degrees*. Archie₂ could be sufficiently sensitive to appreciate that he has a reason not to act insensitively—and thereby sensitive enough to appreciate that he has a reason to avoid the funeral—yet insufficiently sensitive to ensure that he always acts sensitively. So, Archie₂ could remain prone to insensitive outbursts even though he is able to appreciate that he has reason not to act in the sorts of ways he knows he is prone to act (indeed, appreciating that we have reasons to not do things we are prone to do is a commonplace predicament). To serve Millgram's conclusion, then, Archie's insensitivity must be quite extreme.

So consider a third version of Archie: Archie₃. Even if he is unable to appreciate that the fact that upsetting others is unpleasant for them gives him a reason to avoid the funeral, there are other facts that give Archie₃ a reason to avoid it—notably, as Millgram suggests, that upsetting people at the funeral is *bad for Archie*, because it will be one more thing making his life worse. Suppose we explain to Archie₃ that things are going badly for him because everyone dislikes him and that people dislike him because, by acting in the insensitive ways he does, he upsets (riles, annoys, offends) them. Or we might say to him that his behaviour riles others because, in acting so insensitively, he fails to treat them with even a basic modicum of respect of the kind they believe is due. Perhaps this gives a route into Archie₃'s being able to appreciate why his life is going badly—and, moreover, why he has a reason to avoid situations like funerals where his insensitivity manifests itself.

It may be suggested here that, even if Archie₃ is able to appreciate this reason to avoid the funeral, it is the wrong kind of reason. It is a prudential reason; and this does not show that Archie₃ is able to act for the reason *R* we are after—namely, that his upsetting people is bad for them. Fair enough. But now we say to Archie₃ that he will only begin to make his own life go better if he treats people better (less insensitively, with more respect) *for their sake*. For, we explain, it is constitutive of showing people respect that one does this for their sake, not (just) one's own; and while he could just mimic such respect, treating them *as if* he respects them as a

means to making his own life better, this will likely prove an ineffective long-term policy. If Archie₃ can appreciate that he has *this* reason not to upset people, he could avoid the funeral for the reason R that he will upset people in ways bad for them.³¹ If so, premise (1) is saved.

But what if Archie₄ is unable to appreciate any such reason? At this point—if he is *so insensitive* that he is completely unable to appreciate any of the candidate R-giving facts we present to him, and if he is utterly incapable of acting under any such description—I think we *should* conclude that he does *not* have that reason. This is simply what (1) implies. And *this* is where the points raised under *Response I* in §3 have their place: these facts, which do not give Archie₄ a reason, are nonetheless *candidate* reason-giving facts that might provide reasons for *others* less insensitive than Archie₄; and it would be *good* for those he will upset, as well as himself perhaps, were he to avoid the funeral. Nonetheless, Archie₄'s insensitivity runs so deep that it is plausible to say that *he* really does *fall outside the scope* of *this* reason. Note, though, that if Archie₄ really is incapable of appreciating the reason-giving force of any of the considerations we present to him (pertaining to upsetting others, being nasty, showing people respect, etc.) he will be unable to appreciate a far wider range of basic ethical reasons—including, for instance, reasons to treat people well for their sake. Such a person would be extremely different from—somewhat alien to and estranged from—the rest of us, whereby it is no longer obvious that he does fall within the scope of the same sorts of reasons as us (any more than a wild creature, say). Furthermore, attributing the normative reason to Archie₄ will have no practical effect on him, since it is not a reason he could act in light of. And, internalists will add, if it is not a consideration he could act in light of, this putative reason is not action-guiding when it comes to Archie₄ in the ways we expect an account of reasons for action to be; hence, as the Practicality Thesis predicts, it is unclear why we should think of it as a reason applying to Archie₄. Here Williams (2001, 95) might be right: insisting that certain considerations, which we regard as reasons, really are reasons for those incapable of appreciating or acting in light of them, amounts to brow-beating bluff (see also Manne 2014). Nor, moreover, would attributing the reason to Archie₄ have any correlative

³¹It may in practice take some time for Archie₃ to appreciate that he has reasons to treat people well for *their* sake, much as it takes time to cultivate such ethical dispositions in children. The internalist stance implies that Archie₃ (like young children) might thereby acquire reasons he previously lacked through the process of becoming sensitive to the sorts of considerations (candidate reasons) others appreciate more readily as reasons. This seems to me both psychologically and normatively realistic, though a full defense is too wide-ranging to begin here.

benefit for us. Archie₄ may just be someone aptly labelled *sociopathic*—someone we have good reason to avoid or protect ourselves against. Even so, given the extent of his insensitivity and how different this makes him from the rest of us—someone *incapable of appreciating the most basic ethical reasons*—there is no more reason to think he falls within the scope of these normative reasons than there is to think other reason-insensitive creatures do.³² Thus, when Archie is construed in such extreme terms as Archie₄, it becomes increasingly unobvious that this represents a clear-cut objection to (1).

In short, then: either Archie is sufficiently sensitive to at least some reasons that he will be able to appreciate that the fact that he will upset people is a reason for him to avoid the funeral (as with Archie₁₋₃); or his insensitivity is sufficiently great that it looks more plausible to say he falls outside the scope of this reason (as with Archie₄). If he falls into the first category, he could act for the relevant reason. If he falls into the second category, that he could not for the reason does not undermine (1). Either way, (1) remains intact.

7. Concluding remarks

We started with two counterexamples to the Practicality Thesis. Three extant responses were inadequate. The rest of the paper therefore sought a unified approach capable of dealing with both scenarios—and any structurally similar case. I've developed this defence mainly on behalf of an internalist model of reasons. If cogent, and if other elements in the Classical Argument can likewise be defended, we have a persuasive case for internalism. And that has significant implications for substantive issues concerning who has what reasons, and for foundational issues about the categoricity and authority of moral obligation. Nevertheless, the resources deployed in defence of the Practicality Thesis are available to many externalists too and should therefore have wider appeal. Several outstanding issues remain, including what positive arguments there are for

³²Attributing such reasons to him may facilitate certain practices of *blame* (see Williams 1995). However, if Archie could not act for reason *R*, then on most views he is not *blameworthy* for failing to do so. Skorupski also argues that we need a constraint on the scope of reasons. According to his 'cognitive internalism': that *p* is a reason for *A* to ϕ only if *A* has the ability to recognize, *de se*, that were it to be the case that *p* that would be a reason for *A* to ϕ (2010, 73–6, 253–6). Skorupski denies both RI and ST, though note that cognitive internalism plus ST entails RI.

the Practicality Thesis. The aim here, though, has been to defend it against a particular style of objection.³³

References

- Alvarez, Maria. 2010. *Kinds of Reasons*. Oxford: Oxford University Press.
- Dancy, Jonathan. 2000. *Practical Reality*. Oxford: Oxford University Press.
- Dancy, Jonathan. 2004. *Ethics Without Principles*. Oxford: Oxford University Press.
- Davidson, Donald. 1980. "Actions, Reasons, and Causes". Chap. 1 in *Essays on Actions and Events*. Oxford: Oxford University Press.
- Finlay, Stephen. 2009. "The Obscurity of Internal Reasons". *Philosophers' Imprint* 9 (7).
- Joyce, Richard. 2001. *The Myth of Morality*. Cambridge: Cambridge University Press.
- Katsafanas, Paul. 2013. *Agency and the Foundations of Ethics: Nietzschean Constitutivism*. Oxford: Oxford University Press.
- Korsgaard, Christine. 1986. "Skepticism about Practical Reason". *Journal of Philosophy* 83 (1): 5–25.
- Manne, Kate. 2014. "Internalism about Reasons: Sad but True?". *Philosophical Studies* 167 (1): 89–117.
- McDowell, John. 1995. "Might there be External Reasons?", in *World, Mind, and Ethics*, edited by J. Altham & R. Harrison, 68–85. Cambridge: Cambridge University Press.
- Millgram, Elijah. 1996. "Williams' Argument Against External Reasons". *Nous* 30 (2): 197–220.
- Parfit, Derek. 1997. "Reasons and motivation". *Proceedings of the Aristotelian Society* LXXI: 99–130.
- Scanlon, T. M. 2014. *Being Realistic About Reasons*. Oxford: Oxford University Press.
- Schroeder, Mark. 2007. *Slaves of the Passions*. Oxford: Oxford University Press.
- Schroeder, Mark. 2011. "Ought, Agents, and Actions". *Philosophical Review* 120 (1): 1–41.
- Skorupski, John. 2010. *The Domain of Reasons*. Oxford: Oxford University Press.
- Williams, Bernard. 1981. "Internal and External Reasons". Chap. 8 in *Moral Luck*. Cambridge: Cambridge University Press.

³³ Many thanks to the various audiences I've presented parts of this work to, to everyone who has read it, and to the many people I've discussed these issues with—especially Alexander Bird, Rory Drennen, Philip Ebert, Charis Marwick, Christian Pillar, Nick Shackel, and John Skorupski. I'm also grateful to an anonymous referee for very useful comments and suggestions on the penultimate version, which have helped me clarify and otherwise improve some key aspects.

Williams, Bernard. 1985. *Ethics and the Limits of Philosophy*. London: Fontana Press.

Williams, Bernard. 1995. "Internal Reasons and the Obscurity of Blame". Chap. 3 in *Making Sense of Humanity*. Cambridge: Cambridge University Press.

Williams, Bernard. 2001. "Postscript. Some Further Notes on Internal and External Reasons", in *Varieties of Practical Reasoning*, edited by E. Millgram, 91–7. Cambridge: MIT Press.