

# Reasons, Reason, and Context<sup>1</sup>

Daniel Fogal  
Uppsala University

Language [is] one of the principal instruments or helps of thought; and any imperfection in the instrument, or in the mode of employing it, is confessedly liable, still more than in almost any other art, to confuse and impede the process, and destroy all ground of confidence in the result. For a mind not previously versed in the meaning and right use of the various kinds of words, to attempt the study of methods of philosophizing, would be as if some one should attempt to become an astronomical observer, having never learned to adjust the focal distance of his optical instruments so as to see distinctly.

— J.S. Mill, *Logic*

## 1 Introduction

The notion of a normative reason has played an increasingly prominent role in recent theorizing. This focus on reasons has been especially pronounced in ethics and metaethics, but more recently has taken hold in epistemology as well.

Although I'm critical of this recent trend, I'm not going to be arguing directly against it in what follows. Instead, my goal will be to point out various subtleties in how we ordinarily think and talk about reasons—subtleties which, if taken seriously, have various upshots, both substantive and methodological. I'll focus on two subtleties in particular. The first concerns the use of *reason* (in its normative sense) as both a count noun and as a mass noun, and the second concerns the context-sensitivity of ordinary reasons-claims. The more carefully we look at the language of reasons, I'll argue, the clearer its limitations and liabilities become. The cumulative upshot is that although talk of reasons is intelligible and useful for the purposes of communication, we should be wary of placing much weight on it when engaging in substantive normative inquiry. By way of illustration, I'll consider some potential pitfalls of taking our talk of reasons too seriously, explaining how careful attention to the language of reasons undermines the main argument for moral particularism, Mark Schroeder's recent defense of Humeanism about practical reasons, and the "reasons-first" program in metanormativity.

The paper is structured as follows. In §2 I consider the various ways in which we standardly use the common noun *reason*, focusing on the relationship between the use of *reason* (in its normative sense) as a count noun and as a mass noun. I argue that there's a strong case to be made in favor of understanding reasons (count) in terms of reason (mass), rather than vice versa. In §3 I explore some notable consequences of

---

<sup>1</sup> A slightly shorter version of this paper is forthcoming in *Weighing Reasons* (Oxford University Press), eds. Errol Lord and Barry Maguire.

understanding reasons in terms of reason, and briefly assess its overall philosophical significance. In §4 I present new data concerning the contextual variability of reasons-claims that reveals a tension between the theoretical role that reasons are supposed to play and the more practical role that they—or at least our claims about them—in fact play. I conclude in §5 by exploring some consequences for particular issues in (meta)ethics. It turns out that claims about reasons are only ever the beginning, and never the end, of normative inquiry.

## 2 Reasons and Reason

### 2.1 The language of reasons

Talk of reasons can be confusing—at least to the theorist—despite being utterly commonplace. That’s because there are several distinct strains in such talk, and failing to be sensitive to their differences can easily lead one astray. To start with, like many others I think it’s worth distinguishing between the reasons *why* something is the case (these are commonly called “explanatory reasons”), the reasons *why*—or *for which*—someone does something (commonly called “motivating reasons”<sup>2</sup>), and the reasons for someone *to* do something (commonly called “normative reasons”).<sup>3</sup> Although this three-way distinction is widely recognized, there is considerable disagreement over the proper characterization of each class as well as the relationship between them. I’ll consider the prospects of providing a unified analysis below. For now, however, I’ll briefly mention two complications before focusing on a third.

The first complication concerns the class of motivating reasons—those which help rationalize or explain (at the so-called “personal” level) why some agent *S*  $\varphi$ -s or is disposed to  $\varphi$ , where  $\varphi$  is a verb phrase denoting some action or attitude. For within the class of motivating reasons there’s a further distinction to be drawn between (what I’ll call) “factual reasons” and “teleological reasons”, on the one hand, and what Davidson (1963) calls “primary reasons”, on the other. The basic form of factual reasons-claims is ‘*S*’s reason for  $\varphi$ -ing is (was) that *p*’ while the basic form of teleological reasons-claims is ‘*S*’s reason for  $\varphi$ -ing is (was) to  $\psi$ ’. In both cases, however, there are other ways of saying more or less the same thing. What’s crucial is that in both cases the reason specified by the relevant clause (‘that’-clause in the former and ‘to’-infinitival in the latter) is the content of one of the agent’s action-guiding attitudes (e.g. belief or knowledge in the former and intention or action-guiding preference in the latter). Roughly put, a factual reason is a fact the awareness of which helps explain why the agent  $\varphi$ -s, whereas a teleological reason is a goal or purpose that an agent aims to achieve or promote by  $\varphi$ -ing and which thereby helps explain her doing so. Hence we

---

<sup>2</sup> Despite being standard, this label is misleading. We can presumably have “motivating” reasons for belief, for instance, but it’s doubtful that we’re actually motivated to have beliefs—at least in general. The same goes for other “reasons-responsive” attitudes (intentions, etc.).

<sup>3</sup> I’m ignoring the use of *reason* to mean ‘faculty of reason’ and the use of *reason* as a verb.

say things like ‘Victor’s reason for leaving early was that he was tired’ or ‘Jen’s reason for going into law was that it pays well’ as well as ‘Sasha’s reason for becoming vegetarian was to reduce animal suffering’ or ‘The only reason (why) Karl exercised regularly was to lose weight’.

A primary reason, on the other hand, is a psychological state (e.g. belief, desire), or combination of such states, that helps explain—in the right sort of way—why an agent  $\varphi$ -s. Hence we say things like ‘Jorge realized he was tired, and that’s the reason (why) he left early’ as well as ‘Sally wanted to reduce animal suffering, and that’s the reason (why) she became vegetarian’ or ‘Bob thinks he’s Superman, and that’s the reason (why) he thinks he can fly’. As the Victor/Jorge and Sasha/Sally examples suggest, there’s a close connection between factual/teleological reasons and primary reasons, and they are best seen as providing complimentary explanations rather than competing ones. Primary reasons-claims highlight the agent’s contentful attitude(s)—those that play a certain explanatory role—whereas factual/teleological reasons-claims highlight the content of those attitudes.

The second complication concerns the class of normative reasons—the considerations which “count in favor of” performing certain actions and/or having certain attitudes. For many have insisted on the need to distinguish between the reasons *there are* for some agent to  $\varphi$  and the reasons that agent *has* to  $\varphi$ , where only the latter are of direct relevance to rationality and its ilk.<sup>4</sup> (A similar distinction arises with talk of evidence.) Suppose, for example, that you have skin cancer but it has yet to be detected. Although the fact that you have skin cancer is a reason for you to visit the doctor, it’s not a reason you intuitively *have*—it’s not something you could be expected to take into account in deliberation, nor anything else that plausibly makes a difference concerning what it’s reasonable for you to do. So although your failing to go to the doctor may be unfortunate, it wouldn’t be *irrational* if you remain ignorant of your condition.

I myself am happy granting something like the distinction between reasons there are and reasons had. But there are a lot of complications concerning how such a distinction is to be understood, and how exactly (if at all) it manifests itself in ordinary thought and talk. Since these additional complications won’t matter in what follows, I’ll simply help myself to the intuitive distinction without negotiating the details. Although for convenience I’ll mostly focus on claims concerning reasons there are rather than reasons had, both should be kept in mind.

Unlike the first two complications, the one I’m most concerned with has been largely (though not entirely) overlooked. It’s the distinction between the use of *reason*—in its normative, but not motivating or explanatory, sense—as a count noun (‘Julie has many reasons to lie’) and as a mass noun (‘Julie has lots of reason to lie’),

---

<sup>4</sup> The former are often called “objective reasons” and the latter “subjective reasons”—see, e.g. Parfit (2011) and Schroeder (2007)—though I think those labels do more harm than good. Broome’s (2013) distinction between “owned” and “unowned” reasons is slightly better.

as well as between the facts that such nouns are used to report.<sup>5</sup> Intuitively, count nouns denote (classes of) “things” that are countable, and hence can occur with cardinal numerals (*one, two, three...*) and take plural form (-s), while mass nouns denote “stuff” that’s not countable, and hence do not occur with cardinal numerals and are generally singular or unmarked. However, the term ‘stuff’ can be misleading, and is generally far less apt for so-called “abstract” mass nouns (‘information’, ‘freedom’, ‘reason’, ‘advice’, ...) than it is for more “concrete”, substance-denoting mass nouns (‘water’, ‘cheese’, ‘sand’, ‘beef’, ...). The same is true of ‘thing’, since count nouns vary considerably in the clarity and precision with which their referents are individuated or denotations demarcated, with abstract count nouns oftentimes being particularly poorly individuated—contrast ‘cat’, ‘chair’, and ‘microscope’ with ‘cloud’, ‘detail(s)’, and ‘explanation’. This reinforces the point that the mass/count noun distinction is a grammatical one, having to do with morphosyntactic and semantic properties, and contrary to what the intuitive gloss above might suggest it doesn’t (by itself) have metaphysical implications.<sup>6</sup> For instance, although ‘bean(s)’ is count and ‘rice’ is mass, they both denote granular substances, and although ‘knowledge’ is mass and ‘belief’ is count, they both denote mental states. In general, we shouldn’t draw metaphysical conclusions solely from linguistic data.

Note that some determiners combine with both count and mass nouns, including *any*:<sup>7</sup>

- |                                       |         |
|---------------------------------------|---------|
| (1a) I don’t have any water.          | (mass)  |
| (1b) I don’t have any reason to lie.  |         |
| (2a) I don’t have any dogs.           | (count) |
| (2b) I don’t have any reasons to lie. |         |

Similarly promiscuous determiners include *some, the, no, this, that, what, more*, and genitives such as *my*. Other determiners, however, are wholly or predominantly restricted to one class of nouns, and hence more discriminating. For example, *each, every, several*,

---

<sup>5</sup> As Payne and Huddleston (2002) note, the mass/count distinction best understood as an instance of polysemy (more than one semantically related sense for a single word), rather than homonymy (distinct lexical items that happen to be pronounced and spelled alike).

<sup>6</sup> This point is further reinforced by the fact that there are nouns that straddle the mass/count divide, including so-called “fake” or “collective” mass nouns (‘furniture’, ‘silverware’, ‘luggage’, ...) and “plural” mass nouns (‘earnings’, ‘belongings’, ‘dues’, ...). Morphosyntactically, fake mass nouns are mass nouns, but they are interpreted in much the same way as count nouns (see Barner and Snedeker (2005) for experimental evidence, and McCawley (1975), Rothstein (2010), and Pelletier (2012), among others, for more general discussion). Plural mass nouns take plural form (often obligatorily), but otherwise behave like mass nouns (for discussion, see Gillon (1992), Ojeda (2005), and Schwarzschild (2009), among others).

<sup>7</sup> I use the term *determiner* broadly to include quantifying expressions as well as definite and indefinite determiners.

*few, a/an, many, either, neither, and one* all select for count nouns, while *much, little, enough*, and the lack of a determiner typically select for mass nouns. Hence the difference between:

- (3a) There is water, but not much/#many. (mass)  
(3b) There is reason to lie, but not much/#many.  
  
(4a) There are dogs, but not many/#much. (count)  
(4b) There are reasons to lie, but not many/#much.<sup>8</sup>

Although a lot more could be said about the mass/count distinction, the important point is just that the common noun *reason* in its normative—but not explanatory or motivating—sense is standardly used both ways. The relationship between the two uses has nonetheless been neglected, with the vast majority of theorists focusing on the count noun. This is surprising because many common nouns in English pattern both ways, and when they do there is usually a story to tell about the relationship between the two uses—one which makes clear which is to be understood in terms of the other. And it's important because once we notice that a robust mass/count distinction arises with respect to *reason*, we're faced with the question of which, if either, use should be taken as basic, as well as which facts—count-y facts about reasons or mass-y facts about reason—are most normatively significant.

## 2.2 Count and mass

There are a variety of ways in which mass/count noun pairs can be related.<sup>9</sup> Consider *cheese*, for instance. Although it is standardly used as a mass noun to denote quantities of cheese ('How much cheese do you want?'), it is also used as a count noun to denote kinds of cheese ('What cheeses do you like?'). Same goes for *bread, coffee, virtue*, and many others. Call this class of mass/count pairs **K**, for *kind*.

Next consider *beer*. Like *cheese* it is standardly used as a mass noun to denote quantities ('How much beer do you want?') and as a count noun to denote kinds ('How many beers are on tap?'). So it belongs to **K**. But as a count noun it can also be used to denote conventional units or individual servings of beer ('How many beers do you want?'). So *beer* also belongs to what I'll call category **U**, for *units*.<sup>10</sup> And of course there are further ways in which mass/count noun pairs can be related. Here are four of the main ones:

---

<sup>8</sup> There are actually two readings of (4b) with *much*, one of which means 'There are reasons to lie, but not much reasons' and the other of which means 'There are reasons to lie, but not to lie much'. It's the former that's relevant. Similar remarks apply to (3b).

<sup>9</sup> Cf. Payne and Huddleston (2002) and Gillon (2012).

<sup>10</sup> Notice that *cheese* also belongs to **U**. The categories are not meant to be mutually exclusive—it's quite common for a common noun to belong to multiple categories.

	Mass noun	Denotation of count noun
<b>K</b>	<i>cheese, beer, virtue, activity</i> “I ate a lot of cheese.”	<b>kinds</b> ( <i>cheeses, beers, virtues, activities</i> ) “I’ve tried many different cheeses.”
<b>U</b>	<i>beer, coffee, cake, pizza</i> “I drank too much beer.”	<b>units</b> ( <i>beers, coffees, cakes, pizzas</i> ) “I bought too many beers.”
<b>I</b>	<i>detail, thought, action, error</i> “It contains much detail.”	<b>instances</b> <sup>11</sup> ( <i>details, thoughts, actions, errors</i> ) “It contains many details.”
<b>S</b>	<i>sorrow, pleasure, light, anxiety</i> “It filled me with sorrow.”	<b>sources</b> ( <i>sorrows, pleasures, lights, anxieties</i> ) “Life is full of sorrows.”

So what about *reason*? There’s a good case to be made that it patterns like the nouns belonging to **S**, all of which license (what I’ll call) “generational” readings of verbs like *give* and *produce*. To see what I mean, consider the following pairs:

(5a) Taking warm baths is a simple pleasure of mine.

(5b) Taking warm baths gives me great pleasure.

(6a) The candle was the brightest light in the room.

(6b) The candle produced lots of light.

(7a) That we won fairly was a reason to celebrate.

(7b) That we won fairly gave us reason to celebrate.

Notice that in the first sentence of each pair above, the relevant count noun (*pleasure(s)*, *light(s)*, *reason(s)*) is being used to describe some (type of) thing—whether it be an activity, event, fact, or something else—and in the second sentence that same thing is being described as explanatorily responsible for (or a “source” of) the stuff denoted by the mass noun (*pleasure*, *light*, *reason*). This pattern unifies the family of mass/count pairs belonging to **S**.

This isn’t to say that members of **S** are *only* members of **S**—both ‘sorrow’ and ‘light’, for instance, belong to **S** but are also members of **K** and **I**, and hence can be used as count nouns to denote kinds as well as something like instances of sorrow and light, respectively. And of course there are many other members of **S**, including normative notions besides ‘reason’ (e.g. ‘merit’). As far as I can tell, though, ‘reason’ (in

---

<sup>11</sup> I realize this isn’t happy terminology, since ‘instance’ is too broad—e.g. conventional units of *X* are also instances of *X*. So I’m using it quasi-stipulatively in a way that is suggestive rather than definitive of the relevant family. And it really is like a family, as opposed to a well-defined category. The difference between (e.g.) a **S**- and **I**- interpretation of a given expression is sometimes unclear. But just because there is overlap between families in some cases doesn’t mean there is overlap between families in all cases.

its normative sense) only belongs to **S**. Also, notice that there is little metaphysical unity among the kinds of “stuff” picked out by the mass nouns belonging to **S**—it includes stuff that is physical (‘light’), experiential (‘pleasure’), emotional (‘sorrow’), and normative (‘reason’). Similar diversity arises in other families of mass/count pairs, such as **K** and **I**. And as with mass nouns in general, it’s a further question what, if anything, composes each kind of stuff. For although in the case of physical stuff (e.g. water, light) it’s reasonable to expect an answer to the question of what it’s composed of (e.g. H<sub>2</sub>O, photons), it’s far less clear we should expect an answer in cases of nonphysical or abstract “stuff” (e.g. sorrow, reason). In answering such questions we have no choice but to proceed on a case-by-case basis.

Importantly, however, there’s an explanatory asymmetry involved: with respect to each mass/count pair belonging to **S**, the count noun—i.e. what it applies to—is properly understood in terms of the mass noun—i.e. what is given rise to, or otherwise explained—rather than vice versa. To illustrate: imagine meeting someone blind from birth who asks you to explain what light is, as well as what lights are. There are various ways you might reply, but one thing you won’t (or at least shouldn’t) do is offer an explanation of light in terms of lights—to say, for example, that light is that which is emitted by lights. In contrast, you’d be hard-pressed to offer an explanation of what lights are without making reference to light—that is, to avoid saying that lights are things which emits light.<sup>12</sup> Or imagine meeting an insensate who has never experienced pleasure. He asks you to explain what pleasure is, and what pleasures are. This is undoubtedly a difficult task, but one thing it makes little sense to do is offer an account of pleasure in terms of pleasures—although it might be *true* to say that pleasure is a state that is generated by pleasures, it wouldn’t be *explanatory* in the relevant sense. In contrast, it makes perfect sense to characterize pleasures in terms of pleasure—to say, for example, that pleasures are things which give rise to pleasure.<sup>13</sup>

This provides the basis for a family resemblance argument in favor of understanding reasons in terms of reason, rather than vice versa. For just as sorrows, pleasures, and lights are naturally understood to be things (= facts, states, events, objects, whatever) which, in certain circumstances, are explanatorily responsible for there being sorrow, pleasure, and light, so we should understand reasons (for *S*) to  $\varphi$  to be things which, in certain circumstances, are explanatorily responsible for there being reason to  $\varphi$ . (Note: I’ll oftentimes omit reference to the agent (‘for *S*’) in reason(s)-claims out of convenience.) And just as it would be a mistake to characterize sorrow, pleasure, and light in terms of sorrows, pleasures, and lights, so it would be a mistake to char-

---

<sup>12</sup> Of course, this wouldn’t help much unless you gave the person a rudimentary explanation of light—but that’s as it should be. Also, I’m not denying that when you’re trying to provide an explanation of what it is to be a light it might sometimes be *adventitious* to first describe them as things that have a certain shape, give off warmth, etc., before characterizing them in terms of what they emit—namely, light. The order in which we explain things needn’t always follow the order of explanation itself.

<sup>13</sup> Once again this passes the explanatory buck to the mass noun—as it should be.

acterize reason in terms of reasons—to say, for instance, that reason to  $\varphi$  is that which is generated by reasons to  $\varphi$ . That would be getting things backwards. Instead, reasons are to be understood in terms of reason, together with the (admittedly vexed) notion of explanation: reasons to  $\varphi$  are things which are (at least partly) explanatory responsible for there being reason to  $\varphi$ .

The family resemblance argument isn't knock-down, of course. For as with any family there are going to be differences as well as similarities among members of **S**, and perhaps the relative priority of reasons and reason is such a difference. At the very least, though, the foregoing considerations shift the burden of proof onto those who wish to understand reason in terms of reasons, since they'll need to offer an alternative (and non-ad hoc) story about the relationship between our use of *reason* and *reason(s)* that vindicates the priority of the latter. Without such a story, however, we should be open to the possibility of following the argument where it leads. In the next section I'll consider some consequences of doing so.

### 3 The Priority of Reason: Some Consequences

#### 3.1 A unified account of reasons-talk

One immediate benefit of understanding reasons in terms of reason is that it allows us to provide a pleasingly unified account of both normative *and* non-normative uses of the count noun *reason(s)*—one according to which *all* reasons are explanatory reasons, and essentially so. To be a reason (of any kind) *just is* to be something that helps explain something else. Different “kinds” of reasons—normative, motivating, explanatory, and so on—merely differ in terms of what's being explained (the explanandum), how it's being explained (the kind or level of explanation), or what's doing the explaining (the explanans). For instance, motivating reasons differ from other explanatory reasons primarily with respect to the level of explanation, together with the nature of the explanans—to explain why an agent  $\varphi$ -s in terms of (the contents of) the agent's beliefs, desires, and the like is to provide an explanation at the so-called “personal” level that is distinct from, but compatible with, explanations at the sub-personal level (e.g. in terms of lower-level cognitive processes) or non-personal level (e.g. in terms of neural or neurophysiological activity). Normative reasons, on the other hand, are distinguished primarily in terms of the explanandum. A normative reason is something that helps explain why there is (some amount of) reason to  $\varphi$ . Since what normative reasons help explain is the presence of something mass-y (i.e. reason) rather than something count-y, this proposal also naturally accounts for the absence of mass-y counterparts of non-normative uses of the count noun *reason(s)*—the nature of the explanandum differs.

Among those who seek to provide a unified account of our use of *reason(s)*, the view on offer is most similar to those of John Broome (2004; 2013) and Stephen Finlay (2014). They both take normative reasons to be facts which (at least help) explain other normative facts; they just disagree over what the other normative facts are.



Whereas Broome takes normative reasons to be facts which play a certain role in explaining why some agent *ought* to  $\varphi$ , Finlay takes reasons to be facts that explain why it would be *good* (in some way, to some degree) if some agent  $\varphi$ s. Although these are important alternatives, I won't be able to engage with them here. Suffice it to say, both err in prioritizing reasons instead of reason.<sup>14</sup>

### 3.2 The possibility of “groundless” normative facts

Another consequence of understanding reasons in terms of reason is that it allows for the possibility of “groundless” normative facts. For although there will always be (at least some) reason to  $\varphi$  whenever there is a reason to  $\varphi$ , nothing I've said guarantees that there is a reason to  $\varphi$  whenever there is reason to  $\varphi$ . The fact that there is reason to  $\varphi$  might instead be *groundless*; it might not hold in virtue of any particular fact identifiable as a reason. (Compare the possibility of there being light that is not produced by any lights.)

I mention this possibility not to endorse it, but merely to point out that it's an open question at this point whether the following biconditional fails in the left-to-right direction:

(RR)      There is reason to  $\varphi \leftrightarrow$  there is a reason to  $\varphi$ .

I should note that I've been following Schroeder (2007) and others in assuming that claims of the form *there is a reason to  $\varphi$*  involve bona fide existential quantification, and hence that the following holds:

(RR+)     There is a reason to  $\varphi$  iff there is some  $r$  such that  $r$  is a reason to  $\varphi$ .

Enoch (2011) notes the importance of (RR+) for Schroeder's view, but raises a worry:

What if some actions (or some other things) enjoy a basic, ultimate kind of positive normative status? [T]hey are actions that are to be performed and there is nothing in virtue of which they are to be performed. (443)

The present point is that *even if* (RR+) is true, the possibility Enoch considers isn't ruled out.

I should also note that some hear claims about what there is reason to do as being stronger than claims about what there is a reason to do, and so balk at the claim that there is reason to  $\varphi$  whenever there is a reason to  $\varphi$ . One might therefore worry about

---

<sup>14</sup> Broome (2013) does briefly consider the possibility of understanding reasons in terms of reason, though he ends up dismissing it. Finlay, in contrast, wrongly takes constructions involving the mass noun to be idiomatic. For criticism of Broome's positive account of reasons, see Brunero (2013) and Dancy (forthcoming). For criticism of Finlay, see my (forthcoming) review.

the right-to-left direction of (RR) as well. But the difference in strength is arguably just pragmatic, and traceable to the threshold-y yet unspecific nature of bare (i.e. unquantified) reason-claims. To say there is reason to  $\varphi$  is to say that  $\varphi$ -ing enjoys a non-zero amount of normative support, but (semantically speaking) it's to remain silent concerning how much. So it would normally be pointless or misleading to make such claims unless the amount of support was significant. (Compare what's normally communicated versus strictly entailed by unquantified claims like 'There is water on the floor'.) Adding modifiers like (*at least*) *some*, *little/not much*, *lots of*, *sufficient*, *decisive*, and so on makes the strength of the intended claim explicit, and I for one have no trouble making sense of claims like 'There is reason to  $\varphi$ , but not very much'.

Indeed, the same sort of pragmatic story that Schroeder (2007) offers to account for the oddness of a claim like 'You have a reason to eat your bicycle' despite its literal truth can be used to explain the apparent infelicity of 'You have reason to eat your bicycle'. What Schroeder calls the "negative existential fallacy" is just an instance of a broader "fallacy" involving both mass-y and count-y quantificational claims—e.g. claims like 'There's no water left' are typically acceptable even when a couple drops remain. And this is itself an instance of a broader "fallacy" involving categorization in general—e.g. claims like 'I'm not angry' are often acceptable even if the speaker is angry, but not to a contextually salient degree.

### 3.3 Against the "reasons first" program

So far I've played a bit fast and loose with various linguistic/conceptual/metaphysical distinctions, resulting in them being blurred. Such blurring often arises in philosophy—most obviously with questions concerning what it is to "count as" a thing of a certain kind. Talk of "counting as" sounds linguistic/conceptual but talk of being "a thing of a certain kind" sounds metaphysical. Questions concerning what it is to count as a thing of a certain kind are perhaps best understood as an exercise in something like what P.F. Strawson (1959) called "descriptive metaphysics" or—better—what Emmon Bach (1986) calls "natural language metaphysics". The extent to which the broadly linguistic and conceptual argument I've given can be parlayed into a purely metaphysical (and not merely a "descriptive" or "natural language") one isn't clear, since it depends on thorny issues having to do with the relationship between language, thought, and reality—issues that I'm not able to address at any length here.

Nonetheless, at least in practice many philosophers who are fond of reasons proceed as if there's not much distance between our ordinary normative judgments involving the count noun *reason(s)* and the underlying normative facts, taking it for granted that claims of the form *r is a reason for S to  $\varphi$*  are true just in case it's a fact that *r* is a reason for *S* to  $\varphi$ —or in more theory-laden terms, just in case *r* has the relational property of being a reason for *S* to  $\varphi$ . Some are explicit about the presumed correspondence, but most are at least implicitly committed to it in virtue of their unabashed appeal to ordinary judgments about reasons when engaging in substantive normative (and metanormative) inquiry. If one takes our normative judgments involv-

ing the count noun *reason(s)* seriously, however, it's hard to see why one wouldn't also take our judgments involving the mass noun *reason* seriously—indeed, theoretical and methodological consistency would seem to require it. This is the dialectical space within which the family resemblance argument in favor of understanding reasons in terms of reason gains the most theoretical traction.

Notice that if we do end up applying the considerations above concerning the way we deploy *reason* and *reason(s)* in ordinary thought and talk to our normative metaphysics, we'll be committed to rejecting both of the following:

- Reasons Primitivism:** The property (and concept) of being a normative reason is unanalyzable, and hence primitive.<sup>15</sup>
- Reasons Fundamentalism:** Facts about reasons—i.e. facts expressed by (true) claims of the form *r is a reason (for S) to φ*—are normatively fundamental.<sup>16</sup>

Note that facts *about* reasons are to be distinguished from the facts *which are* reasons—i.e. the facts corresponding to *r* in claims of the form '*r is a reason (for S) to φ*'. Whereas facts about reasons are normative facts, the facts which are reasons needn't be—and usually aren't.<sup>17</sup>

Although Reasons Primitivism and Reasons Fundamentalism are not often explicitly distinguished, it's worth keeping them apart. The former concerns analyzability, after all, while the latter concerns normative fundamentality, and neither entails the other. Schroeder (2007), for instance, accepts Reasons Fundamentalism but rejects Reasons Primitivism, offering a reductive analysis of facts about reasons. Normatively fundamental facts needn't be fundamental *full stop*. Nonetheless, many philosophers who are attracted to one are also attracted to the other, and together they lie at the heart of the increasingly popular “reasons first” program in metaethics (or rather, metanormativity). Not everyone agrees on the exact contours of the program, of course, and there are various ways in which the core claims might be precisified. For example, we might distinguish “weak” and “strong” versions of Reasons Primitivism, where weak versions merely claim that the property (or concept) of being a normative reason is primitive while strong versions claim that it's the *only* normative property (or concept) that is primitive. Similarly, we might distinguish “weak” and “strong” versions of Reasons Fundamentalism, where weak versions claim that the fundamental normative facts *include* facts about which things are reasons to do what in which circumstances, while strong versions claim that the fundamental normative facts are *exhausted by* such facts. I should note that there is also disagreement concerning the

---

<sup>15</sup> See Scanlon (1998: 17), Skorupski (2010: 2), and Parfit (2011: 31) for claims to the effect that our concept of a normative reason is primitive. Presumably, however, they would say the same thing about the corresponding property.

<sup>16</sup> See, for instance, Raz (2000: 67), Schroeder (2007: 81), and Scanlon (2014: 2).

<sup>17</sup> Some have neglected this distinction; see Olson (2009) for critical discussion.

arity of the reasons-relation. Scanlon (2014), for instance, takes it to be a four-place relation holding between a fact, agent, action, and set of circumstances, whereas Schroeder (2015) argues that at least one of its argument places is redundant. It's perhaps most common to treat it as a three-place relation between a fact, agent, and action, and for concreteness I'll be following suit. But, importantly, these sorts of in-house disputes won't matter so long as it's agreed that Reasons Primitivism and Reasons Fundamentalism concern the notion of a normative reason operative in ordinary thought and talk, since that's what the broadly linguistic considerations I've adduced (and will adduce) speak to.

It should be clear by now why those who take the linguistic data and family resemblance argument above seriously should reject Reasons Primitivism. For such theorists will be in a position to answer Scanlon's (2014: 44) question of "whether there is something further to be said about what it is to be a reason...something further that might be said to identify the property signified by that [relational] concept". The answer will be 'yes': a given fact  $r$  has the property of being a reason (for  $S$ ) to  $\varphi$  just in case—and *because*— $r$  plays a certain role in explaining why there is reason to  $\varphi$ . The property of being a reason (and corresponding relation)—as well as the rest of properties corresponding to the count nouns in category **S** ("sources") above, such as the property of being a light, being a pleasure, and so on—will thus be analyzable in broadly functional terms, in much the same way that properties such as being a cook and being a gene are. To be a cook, it seems, is to play a certain (at least intended) role in the production of food, and to be a gene is to play a certain role in encoding and transmitting genetic information.<sup>18</sup> In general, a functional property is a property something has in virtue of the role it plays, whether that role be causal, teleological, computational, metaphysical, normative, or whatever. The details of the relevant roles are of course complicated, and functional analyses are rarely, if ever, expressible as a set of individually necessary and jointly sufficient conditions. But that's true of nearly all concepts and properties, and it's compatible with the claim that some concepts and properties can be analyzed partly in terms of others. If it wasn't, then nearly all concepts and properties would be on a par, and the corresponding forms of primitivism and fundamentalism would come cheap.

It should also be clear why such a theorist should reject Reasons Fundamentalism. For if the property of being a reason is a functional property defined partly in normative terms, then facts expressed by true claims of the form *r is a reason (for S) to  $\varphi$*  will obtain partly in virtue of other normative facts—in particular, (mass-y) facts about reason—and so not be normatively fundamental. Facts about reason will be normatively prior to facts about reasons, contra Reasons Fundamentalism.

### 3.4 An ecumenical middle ground?

By taking the mass/count distinction seriously, we've arrived at two preliminary con-

---

<sup>18</sup> Cf. Kim (1998: 25).

clusions: first, normative reasons (like all reasons) count as reasons in virtue of the explanatory role they play; second, what they help explain is some further (though closely related) normative fact. The first point suggests it's a mistake to take the property or concept of a normative reason as primitive, and the second suggests that it's a mistake to take facts about reasons to be normatively fundamental—they obtain partly in virtue of other normative facts.

Neither of these points may seem all that significant, however. And in many ways they're not. After all, to say there is reason to  $\varphi$  is just to say that there is a non-zero amount of normative support for  $\varphi$ -ing, and hence that  $\varphi$ -ing is to at least some degree worth doing or having. More colloquially, it's just to say that there is *something* (though not necessarily some *thing*) to be said in favor of  $\varphi$ -ing. This is little more than a paraphrase, though, not an analysis. And if that's all that can be said about what it is for there to be reason to  $\varphi$ , then it will turn out that Reasons Primitivists are at least partly right. Although they're wrong to take the property or concept of a normative reason to be primitive, they're right to insist that there's a positive pro tanto normative status in the vicinity that isn't amenable to further analysis. Replacing Reasons Primitivism with Reason Primitivism can be seen as simply shifting the focus to where it should have been (and to some extent was) all along.

A similar shift may appear plausible when it comes to Reasons Fundamentalism. Even if facts about reasons are not normatively fundamental, much of what Reasons Fundamentalists care about can be preserved by switching to Reason Fundamentalism. Facts about reason, after all, are facts about a normative status that is pro tanto and non-strict—it's a status that can be "opposed" or "defeated" in various ways, and one doesn't automatically do anything wrong by failing to do what one has reason to do, since one may have more (or equal) reason to do something else. Unlike reasons, however, reason is a quantitative notion—it's something that comes in degrees, something you can have more or less of. So whenever someone makes a claim about there being reason to do something, there is always a further question concerning *how much* reason there is. This is what allows for comparative claims, whether between people ('Joe has more reason to lie than Sally, but I have the most') or between actions/attitudes ('Joe has more reason to lie than to hide'). The quantitative nature of reason also helps account for our talk of the "strength" or "weight" of reasons—the intuitive strength of a reason can be understood in terms of how much reason that fact is explanatorily responsible for, with possible amounts ranging from minute to decisive.<sup>19</sup>

Various distinctions between different "kinds" of normative reasons—practical, moral, epistemic, etc.—can be similarly preserved by understanding them in terms of the different kinds of reason—practical, moral, epistemic, etc.—they help explain. The same goes for the intuitive distinction between the reasons there are to  $\varphi$  and the reasons an agent has to  $\varphi$ —it can be understood as the distinction between the facts that

---

<sup>19</sup> It would therefore be a mistake to treat the difference between, say, a "decisive" reason and a merely "sufficient" reason as a difference in kind rather than degree.

help explain why there is a certain amount (and kind) of reason to  $\varphi$  and the facts that help explain why an agent has a certain amount (and kind) of reason to  $\varphi$ .

It's in answering such questions—questions concerning kinds and quantities of reason—that our judgments about reasons are most relevant. After all, it's in virtue of helping explain the existence of reason that a fact counts as a (normative) reason at all. But we might go further and take reasons to play a *distinctive* role in explaining facts about reason: as suggested by the discussion in §2.2, we might take reasons to be what *generate* (and hence serve as “sources” of) reason. So even if facts *about* reasons are not normatively fundamental, the facts which *are* reasons might nonetheless play a privileged explanatory role—and one which would justify much of the attention that has been lavished on them in recent years.

Unfortunately, with respect to the last point, things aren't so simple. For as I'll argue in the following sections, the facts we cite as reasons are rarely (if ever) themselves sources of reason. At best they're *partial* sources—on their own they're normatively irrelevant, and it's only in combination with other facts that any amount of normative support is generated. Again, that's at best. For it's arguable that some of the facts we unproblematically cite as reasons aren't even partial sources, but instead play a different explanatory role. (Examples of such will undoubtedly be contentious; I provide one in §5.3.)

More specifically, in the following section I'll present new data concerning the contextual variability of our judgments about reasons that reveals a tension between the theoretically substantive role that reasons-claims are supposed to play and the more pragmatic role that reasons-claims in fact play. For it is widely regarded as a truism that reasons are those things which “count in favor of” performing certain actions or having certain attitudes—they are “sources” of normative support (aka reason). What this amounts to is the assumption that reasons are things which *themselves* explain—in a broadly *metaphysical, non-pragmatic* sense—why there is reason to  $\varphi$ . But, as I'll argue, the things we actually cite as reasons only *partially* explain—in a broadly *communicative, pragmatic* sense—why there is reason to  $\varphi$ . As I'll put it, they're merely “representatives” of sources.

This contrast between roles rests on two cross-cutting distinctions: first, the distinction between full and partial explanations, and second, the distinction between pragmatic and non-pragmatic senses of *explanation*. Both distinctions figure prominently in the literature on causal and scientific explanation. But since a detailed discussion of either is beyond the scope of this paper, a few suggestive remarks will have to suffice.

The first distinction—between full and partial explanations—should be familiar enough. For almost all the explanations we actually provide are partial, not full. Suppose, for instance, that I order a triple espresso and you ask me why. I tell you it's because I'm exhausted. This is a perfectly sensible reply. Notice, however, that the fact that I'm exhausted does nothing, *by itself*, to explain why I order a triple espresso. It's only in combination with the fact that, say, I want to perk up, as well as the fact that a

triple espresso will perk me up, and that I know (or at least believe) that a triple espresso will perk me up, that my reply comes close to fully answering your query. And even then there's significant room for further explanatory detail. The fact that I'm exhausted is thus only *part*—and in fact a rather small part—of the full explanation.

It's worth emphasizing that the provision of merely partial explanations is a general feature of our explanatory practice. It's true not just of action explanations, but also of all other explanations we provide as well—causal, scientific, historical, normative, metaphysical, you name it.

Talk of “providing” explanations brings us to the second distinction—the one between pragmatic and non-pragmatic senses of *explanation*. Very roughly, we can think of pragmatic explanations as answers to “why”-questions, understood as the type of questions that form the core of our explanatory practice. Such questions are standardly posed in a particular context, by particular people—people who in asking the question already possess, and therefore presuppose, a huge amount of the background information. A successful answer to such a why-question will be one that helps make sense of the phenomenon being asked about, or render it intelligible given the background information already possessed by the inquirer.<sup>20</sup> Slightly more fully, following James Woodward (2014) we can say that an explanation in the pragmatic sense is one that involves “irreducible reference to facts about the interests, beliefs or other features of the psychology of those providing or receiving the explanation and/or (ii) irreducible reference to the ‘context’ in which the explanation occurs”. A *non-pragmatic* explanation is thus one that *doesn't* involve irreducible reference to such things. Instead, it solely concerns the objective explanatory relations—whether they be causal, nomic, metaphysical, or something else—obtaining between facts ‘out there in the world’. A good pragmatic explanation will oftentimes be one that accurately captures (some of) the objective explanatory relations involved, but due to the vagaries of context, the limitation of our knowledge, and various situation-specific demands, the relationship between the two is likely to be complicated.

Although a lot more could be said about both distinctions, all that's needed for present purposes is an intuitive grasp of each. The goal in what remains will be to gain a better understanding of the actual role of reasons-talk, thereby making the tension between it and its (at least purported) theoretical role more apparent.

## 4 The Variability of Reasons

It has often been observed that what counts as a reason in one context may fail to count as a reason in another context, where the worldly (i.e. non-linguistic) facts are different. Reasons-claims thus exhibit *inter-contextual* variability. But what hasn't (yet) been recognized is that what counts as a reason in one context may fail to count as a reason *in that very same context*, where the worldly facts remain the same. Reasons-

---

<sup>20</sup> The notion of “making sense” is thus to be understood as implicitly relativized to particular inquirers or to the information they possess.

claims therefore also exhibit *intra*-contextual variability.<sup>21</sup> Notably, it's a form of contextual variability that also infects our thought and talk about causes, as Eric Swanson (2010) has argued. Indeed, much of my discussion in this section intentionally mirrors Swanson's discussion in order to make manifest the striking and systematic similarities between our thought and talk about reasons and our thought and talk about causes. I conclude the section by sketching an account of the role of reasons-talk that naturally accommodates the data.

## 4.1 Examples

Here's an example of the kind of *intra*-contextual variability I'm concerned with:

### DANCE<sup>22</sup>

Billy loves to dance and there is going to be dancing at the party tonight. Billy, however, hasn't yet been informed of the party. I tell you about Billy and the party. In deliberating about whether Billy should attend, you might say any of

(D1) The fact that there will be dancing at the party is a reason for Billy to attend.

(D2) The fact that Billy loves to dance is a reason for him to attend.

But it would be unacceptable to say any of

(D3) #The fact that there will be dancing at the party is a reason for Billy to attend, and so is the fact that he loves to dance.

(D4) #The fact that there will be dancing at the party and the fact that he loves to dance are both reasons for Billy to attend.

(D5) #Since there will be dancing at the party and Billy loves to dance, there are two reasons for Billy to attend.

Claims like (D3)-(D5) fail to appropriately describe the stipulated normative relations. The dominant reaction among those I've consulted is that such claims are guilty of something like "double-counting". I think that's exactly right. Summarily put, what's wrong with (D3)-(D5) is that they double-count a source of normative support for Billy's going to the party. (I'll expand on this below.)

It's worth noting that although DANCE concerns reasons *there are* (or "objective

---

<sup>21</sup> There are some technical notions of 'context' that individuate contexts extremely finely and so would count both types of variability as 'inter-contextual' variability. But it doesn't matter whether we decide for theoretical purposes to call the variability I'm concerned with 'intra-contextual' or not; what matters is just that it's different from the kind of variability usually appealed to, and more disconcerting for those who take reasons-claims seriously.

<sup>22</sup> This is inspired by one of Schroeder's (2007) favorite examples, though I'm using it for a very different purpose. For a similar example involving causation, see Swanson (2010).



reasons”), the same variability arises with reasons one *has* (or “subjective” reasons”). For instance:

### SMOKE

Sally opens the door to find the hallway filled with smoke. In explaining what grounds she has to believe there’s a fire you might say either of the following:

- (S1) That fact that she sees smoke is a reason for Sally to believe there’s a fire.
- (S2) The (known) fact that smoke is a sign of fire is a reason for Sally to believe there’s a fire.

But it would be unacceptable to say any of

- (S3) #The fact that she sees smoke is a reason for Sally to believe there’s a fire, and so is the fact that smoke is a sign of fire.
- (S4) #The fact that she sees smoke and the fact that smoke is a sign of fire are both reasons for Sally to believe there’s a fire.
- (S5) #Sally sees smoke and knows that smoke is a sign of fire, so she has two reasons to believe there’s a fire.

As before, reasons-claims like (S3)-(S5) fail to appropriately describe the stipulated normative relations. They too are guilty of double-counting a source of normative support, though in this case it’s a source of *epistemic* support rather than practical support. The variability I’m concerned with thus cross-cuts the objective/subjective divide as well as the epistemic/practical/moral/whatever else divide.

As I hope mundane cases like DANCE and SMOKE make clear, the intra-contextual variability of reasons-claims is extremely widespread. For those who take our ordinary judgments involving the count noun *reason(s)* at face value, however, such variability should be disconcerting. For it shows that ordinary reasons-claims are far more sensitive to conversational context than previously realized. In particular, it shows that what can legitimately count as a reason (or among the reasons) to do something in a particular context—just like what can legitimately count as “the” reason—depends in part on idiosyncratic facts about *us*, including what else has been cited as a reason.<sup>23</sup> But that’s not what we would expect if reasons-claims corresponded directly to the underlying normative facts, which presumably do *not* so depend. To put the point slightly differently: if (true) claims we make using sentences of the form *r is a reason (for S) to φ* expressed facts that were wholly objective (or “out there in the world”), then a given fact *r*’s putative status *as a reason* in a given context presumably wouldn’t depend on obviously irrelevant factors, such as those having to do with con-

---

<sup>23</sup> Again, this parallels Swanson’s (2010) point about the contextual variability of our judgments involving ‘a cause’, extending Lewis’s (1973) point with respect to ‘the cause’.

versational dynamics. Instead, we would expect the infelicity of any such claims to be *merely* pragmatic—to be pointless, irrelevant, misleading in virtue of generating a false implicature, or some such, despite being strictly speaking true. But the defectiveness of (D3)-(D5) and (S3)-(S5) is more than merely pragmatic, and isn't traceable to one of the usual suspects.<sup>24</sup> To insist that they're strictly speaking true but suffer from some familiar pragmatic malady doesn't do justice to the character of their defectiveness, which is stubbornly present in a way that familiar pragmatic maladies are not.

It's important to be clear that I'm *not* claiming it's impossible for (D3)-(D5) and (S3)-(S5) to be used felicitously—on the contrary, they can be. Suppose, for instance, that we discover that Billy gets paid handsomely for every dance party he attends. In that context (D3)-(D5) are fine. So the infelicity doesn't attach to the sentences themselves, but rather to their use in the context. For this reason, among others, the account I'll offer of ordinary reasons-claims won't be directly concerned with the *truth conditions* of such claims. Instead, the primary focus will be on their *acceptability conditions*—i.e. the circumstances in which it is acceptable for someone who knows all the relevant non-semantic facts to use sentences of the form *r is a reason for S to φ*—and how those conditions are sensitive to context.<sup>25</sup>

I'll provide an explanation of why (D3)-(D5) and (S3)-(S5) are acceptable in some contexts and not in others below, but I mention it now just to make clear the scope of my claim: I'm only claiming that (D3)-(D5) and (S3)-(S5) can't be used felicitously *given knowledge of the relevant non-semantic facts* in DANCE and SMOKE. This shows why the diagnosis that (D3)-(D5) and (S3)-(S5) are merely guilty of implicating something false is unsatisfactory—implicatures are characteristically cancelable (à la Grice (1989)), but I'm not aware of any continuations of (D3)-(D5) and (S3)-(S5) in the relevant contexts that eliminate the infelicity. As we'll see, a broadly semantic explanation—one that concerns the communicative role of reasons-talk—is more promising.

## 4.2 Reasons as representatives

The immediate moral to be drawn from the intra-contextual variability of reasons-talk

---

<sup>24</sup> The “usual suspects” qualifier is included because on the view I sketch below, the defectiveness can be seen as straddling the semantic/pragmatic boundary, and so might be considered at least quasi-pragmatic in nature. In this respect the defectiveness of (3)-(5) might be compared—but not assimilated—to the defectiveness involved in cases of presupposition failure, which similarly straddles the semantic/pragmatic line. What's important is just that the defectiveness is not pragmatic in any of the familiar or straightforward ways that we would expect were reasons-claims reliable guides to the underlying facts.

<sup>25</sup> Cf. Swanson (2010). If a truth-conditional account is desired, however, it can be provided. For example, one might relativize sentences of the form *r is a reason for S to φ* to conversational contexts, where such contexts are stipulated to include information about (e.g.) what else has been cited as a reason. See Strevens (2013) for an analogous account of the truth conditions of sentences of the form *c is a cause of e*, relativized to “explanatory frameworks”.

is that insofar as we're interested in carving normative reality at its joints, or at least as close to its joints as we can, we should be wary of putting much stock in our judgments concerning what and whether something is (i.e. *feliculously counts as*) a reason for something else. More generally, however, I think the lesson to be learned is that we shouldn't take singular explanations—i.e. explanations where the explanans is singular—of normative facts at face value, whether the normative facts to be explained are

**contributory** (e.g. *f is a reason for S to (not)  $\varphi$ ; f gives S reason to (not)  $\varphi$* <sup>26</sup>)

or

**threshold-y** (e.g. *f makes it good/bad/(un)reasonable for S to  $\varphi$ ; f is why S may/must/should (not)  $\varphi$* )

Given what we know about the nature of our explanatory practices in general, in which surprisingly few facts usually suffice (for communicative purposes, *not* metaphysically) to explain complicated facts or events, it should come as little surprise that when we make claims of the latter “threshold-y” form the cited fact *alone* will rarely be responsible for the relevant normative status of S's  $\varphi$ -ing, whether it be good, bad, reasonable, unreasonable, or whatever. Instead, other facts will almost always play an important role as well, with full responsibility for the normative status of S's  $\varphi$ -ing belonging to a cluster of facts that includes but is not limited to the fact cited. Call any cluster of facts that taken together—and only taken together—are normatively relevant (in some particular way, to some degree) to some action or attitude a “normative cluster”. For reasons having to do with communicative efficiency as well as our informational limitations, we rarely mention entire normative clusters. Instead, we single out one or two particularly relevant and/or accessible facts and ascribe responsibility to them. In essence, the cited facts function communicatively as representatives of, or proxies for, the normative clusters to which they belong, or are otherwise related.<sup>27</sup> On the resulting view, a fact playing such a role should not (by itself) be taken as indicative of anything metaphysically or explanatorily deep. In a different conversational context—one differing in terms of shared background information, preceding discourse, interests, etc.—some other fact belonging to the very same cluster that fully explains the relevant normative status of S's  $\varphi$ -ing might be more salient, or otherwise

---

<sup>26</sup> Judgments of the form *f gives S reason to  $\varphi$*  are subject to the same kind of contextual variability as those of the form *r is a reason for S to  $\varphi$* . They are both singular explanations of normative statuses—namely, that there is (a certain amount of) reason for S to  $\varphi$ —and so this is to be expected. Rarely do singular facts do all the relevant explanatory work on their own.

<sup>27</sup> There is an obvious analogy with Mackie's (1965) INUS account of causation (or causal explanation), though there are also some important differences. See Strevens (2013) for a recent defense of a (non-reductive) Mackie-style account of causal claims that is congenial to the view of reasons-claims that I develop.

relevant, and hence be referenced instead.<sup>28</sup>

What's more surprising—or at least what's been insufficiently remarked upon—is that precisely the same thing is true when it comes to “contributory” claims of the form *f is a reason for S to φ*, which (as I've argued) are essentially explanatory in nature. Accordingly, when making or evaluating such claims we shouldn't take a fact's status as a reason (or “reason-giver”) to be indicative of anything metaphysically or explanatorily deep. For as we've seen, whether a given fact can felicitously count as a reason in a particular context typically depends not only on there being other facts which taken together—and only taken together—fully explain the relevant normative status of an agent's action or attitude, but *also* on whether any of those other facts have already been cited as a reason. This helps explain why a dispute over whether we should accept (D1) or instead (D2) in the context of DANCE, or (S1) or instead (S2) in SMOKE, will (according to many people) seem shallow. Each fact is eligible to serve as a representative of the relevant normative cluster, and it's up to us whether we use one or the other (or their conjunction). The facts we cite as reasons merely function as representatives; they're not themselves the real thing.

Understanding reference to reasons in this way—i.e. as reference to particular facts that function as (something like) representatives of the normative clusters to which they belong, or are otherwise related—allows for a tidy explanation of the otherwise problematic contextual variability observed in cases like DANCE and SMOKE. For on this view it's natural to take our thought and talk about reasons to be implicitly governed by something like the following norm:<sup>29</sup>

#### **Use Good Representatives**

When you attribute some normative support for A to a normative cluster relevant to A, use good representatives of that cluster for the purposes at hand.

The relevant sense of ‘goodness’ is a broadly communicative one. The goodness of a representative is function of the amount of information it provides—either directly or indirectly, given the background information—about the normative cluster to which it belongs, as well as the economy with which it provides that information. *Ceteris paribus*, the better a representative satisfies this dual constraint, the better a representative it is. Typically the most apt representatives of a given normative cluster will be those that are the biggest “difference makers” in the sense that changes in them, holding other features of the situation fixed, would make the most difference to the normative status of the subject's attitude or action.

---

<sup>28</sup> See Broome (2013)—especially pp. 48-51—for similar remarks concerning our explanatory practice when it comes to (what I'm calling) threshold-y claims.

<sup>29</sup> This principle intentionally mirrors (in name and form) one that Swanson (2010) uses to explain the contextual variability of our causal-talk. Once again, much of my discussion in this section intentionally mirrors Swanson's discussion of causation in order to illustrate the similarities between our thought and talk about reasons and causes.

Unsurprisingly, the relative “goodness” of a potential representative will vary from context to context, and even within a context, depending on the interests of and background information shared by the conversational participants—including whether (and if so, which) other representatives have already been used. Consider again DANCE, where—let’s suppose—the normative cluster consists of just two facts (I’ll use ‘[*p*]’ as a noun phrase that refers to the fact that *p*):

(NC) { [Billy loves to dance]; [There will be dancing at the party] }

As our judgments concerning (D1)-(D2) demonstrate, the initial conversational context is one in which either fact belonging to (NC) is eligible to serve as a good representative:

(D1) The fact that there will be dancing at the party is a reason for Billy to attend.

(D2) The fact that Billy loves to dance is a reason for him to attend.

But although both facts start off on a communicative par, the act of citing either as a reason robs the other of its potential role as a representative. In this way, using the fact that there will be dancing to represent (NC) renders impotent—mid-utterance—the representative role that Billy’s love of dancing could have otherwise played. In the same way, using the fact that Billy loves to dance to represent (NC) renders impotent—mid-utterance—the representative role that the fact that there will be dancing could have otherwise played. To nonetheless cite both facts individually as reasons would be to add unnecessary (and indeed misleading) information.<sup>30</sup> Hence the infelicity of (D3):

(D3) #The fact that there will be dancing at the party is a reason for Billy to attend, and so is the fact that he loves to dance.

The same reasoning applies to (D4)-(D5), and an analogous explanation is available for (S3)-(S5) in SMOKE.

Use Good Representatives also helps explain why the default interpretations of (D3)-(D5) are ones where it’s assumed that Billy’s love of dancing and the fact that there will be dancing at the party are normatively independent of each other.<sup>31</sup> Since both facts are being used as representatives, it will be assumed that both are *needed*—

---

<sup>30</sup> Notice that despite being overkill it would be fine to cite both facts collectively:

(D1+2) The fact that there will be dancing at the party together with the fact that Billy loves to dance is a reason for him to attend.

<sup>31</sup> “Normatively independent” in the sense that the normative relevance of at least one of the facts doesn’t depend on the other for its belonging to a normative cluster relevant to Billy’s going to the party.

that both are playing an important representative role. But the only way for each of them to play such a role would be for them to represent distinct normative clusters. This is why (D3)-(D5) are fine in a context where (e.g.) Billy gets paid handsomely for attending dance parties. (The same point applies to (S3)-(S5) in SMOKE.)

#### 4.4 Two alternatives: narrowing and coarsening

What the foregoing suggests is that there's a significant amount of distance between our ordinary judgments involving the count noun *reason(s)* and the underlying normative facts. This conclusion isn't forced on us, however, and those who wish to maintain a closer connection between the two might try to provide an alternative account of the intra-contextual variability data. There are two main lines of response that I'll consider. To streamline the discussion, I'll focus on how each response handles (D1)-(D5) in DANCE, although everything I say applies, *mutatis mutandis*, to (S1)-(S5) in SMOKE, as well as countless other examples.

According to the first line of response, only one of (D1) or (D2) truly specifies a reason for Billy to attend. The other (at best) merely specifies a *background condition* for the reason—i.e. a condition that needs to obtain in order for the other fact to constitute a reason, but which is not itself part of the reason. On this view, then, only one sentence is true, and the acceptability of the other is to be explained on other grounds. Since this attempts to preserve a finely-individuated picture of reasons at the cost of reinterpreting or rejecting various ordinary reasons-claims, I'll call this the “narrowing” response.

The second line of response is to argue that neither (D1) nor (D2) identify the “real” or “full” reason for Billy to attend (at least not directly). Instead, they merely specify a *part* of the reason, with context providing the rest of the information needed to identify the reason itself—i.e. the conjunctive fact that Billy loves to dance and there will be dancing at the party. This nicely explains why (D3)-(D5) seem to involve double-counting, and why it would strike us as a mistake to list the two facts separately under the “pro” column while deliberating. Because it individuates reasons more coarsely than our ordinary reasons-claims do, I'll call this the “coarsening” response.

Neither response is fully satisfying. There's a lot that could be said, but I'll focus on just a couple points of dissatisfaction. Consider first the coarsening response. Assuming the proponent of this response purports to be using *reason(s)* in its ordinary normative sense, she's going to be committed to thinking that nearly all of our reasons-claims, including (D1) and (D2), are false. The fact that there will be dancing at the party, it turns out, is *not* a reason for Billy to attend, and neither is the fact that he likes to dance. Neither fact is a reason any more than half of a car is a car. Although the bitterness of an error theory might be sweetened by supplementing it with a plausible pragmatic story, it's still something that many would rather avoid. The most obvious way of avoiding an error theory is to say—as I do—that what counts as a reason is highly context-sensitive matter. But since the underlying normative facts are presumably *not* a highly context-sensitive matter (at least not in the same way), this is to

concede the point that our reasons-claims fail to directly correspond to the underlying normative facts.

Accordingly, the proponent of the coarsening response might shift tactics and try to distance herself from the ordinary sense of *reason(s)* by introducing a technical term—most creatively, she could use *Reason(s)*—to better capture what our ordinary reasons claims only capture in part.<sup>32</sup> The conjunctive fact, but neither of its components, would then be (e.g.) a Reason for Billy to attend, with (D1) and (D2) merely being used to direct our attention to its vicinity by identifying a part of it. There are various ways in which this strategy might be fleshed out, and the account of reasons-claims offered above can be understood as one way of doing so. The main point is just that such a response acknowledges the distance between ordinary reasons-claims and the underlying normative facts, and it's a substantive further issue just how close or far apart they are. Such claims would therefore need to be treated with caution and rarely, if ever, taken at face-value. Methodologically, it would be illegitimate to rely on ordinary reasons-claims in motivating (or attacking) a certain principle while formulating the official principle in other terms, without a careful story connecting the two.

What about the narrowing approach? Once again a lot could be said, but it's worth emphasizing just how widespread the intra-contextual variability of reasons-talk is. For if the narrowing response is adopted across the board, we'll end up having to privilege a large range of intuitively acceptable reasons-claims at the expense of others, and in each case we'll need to both motivate the purported asymmetry as well as explain why it's something we're insensitive to. In some cases the motivation and explanation might seem easy to provide. For instance, various philosophers have denied that desires are themselves reason-giving, and so would reject (D2) in DANCE.<sup>33</sup> Why do so many of us find (D2) acceptable, then? That'll depend. Perhaps we're just mistaken in thinking that desires are reason-giving, and so (D2)'s veneer of plausibility is simply the result of a widespread but nonetheless erroneous assumption.<sup>34</sup> Alternatively, perhaps desires are merely necessary background conditions for reasons and not (part of) reasons themselves.<sup>35</sup> In that case (D2)'s plausibility might be explained along more pragmatic lines—even though Billy's love of dancing isn't a rea-

---

<sup>32</sup> Raz (2000), for instance, introduces the term 'complete reason' for stand for "all the facts stated by the non-redundant premises of a sound, deductive argument entailing as its conclusion a proposition of the form 'There is a reason for *P* to *V*'". The utility of such a notion is questionable—see Dancy (2004) for criticism. Scanlon's (2014) introduction of the normatively fundamental four-place 'R' relation is a more promising move along such lines, though as with Raz there's a tension between the introduction of a technical term and the often rather casual reliance on ordinary reasons-claims. Indeed, the intra-contextual variability data strongly suggests that Scanlon's R-relation either fails to carve normative reality at its joints, or else is significantly more removed from ordinary reasons-claims than he realizes.

<sup>33</sup> Or at least provide a non-standard explanation of why it's true.

<sup>34</sup> Cf. Scanlon (1998) and Dancy (2004).

<sup>35</sup> Cf. Schroeder (2007).

son, it's closely related to something that is.

It's worth noting, however, that the explanatory debts don't end there. For it's not enough to explain why, say, (D2) in DANCE seems true even though it's not; one also has to explain why it *doesn't seem false*. That is, we have to explain why in the context of DANCE neither of these are acceptable:

- (D2\*) The fact that Billy likes to dance is not a reason for him to go.
- (D3\*) The fact that there will be dancing at the party is a reason for Billy to go, but the fact that he likes to dance is not a reason for him to go.

This additional explanatory burden applies to advocates of the coarsening response just as much as it does to advocates of the narrowing response. And it's one that the account of reasons as representatives is well-placed to discharge.

Let's suppose, though, that the narrowing approach is plausible in the case of DANCE. That does little to motivate it in general, since the intra-contextual variability of reasons-talk extends far beyond cases involving desires. Consider SMOKE, for instance. What grounds do we have for privileging one of (S1) or (S2) over the other? When it comes to believing there's a fire, is the sight of smoke epistemically relevant in a way that the known connection between smoke and fire is not? Here it seems even less plausible to insist that only one of (S1) or (S2) specifies the "real" reason, while the other merely specifies a background condition. The coarsening response has more plausibility here. The defender of a close relationship between ordinary reasons-claims and the underlying normative facts might therefore opt for a mixed strategy, treating each case on its own terms.

No matter which approach is adopted, however, we'll have to concede that our ordinary judgments about reasons do not reliably track the fundamental normative facts, and we'll have to explain why we're systematically inclined to cite as reasons facts that aren't in fact reasons. Although I'm not in a position to rule out the possibility that such a story can be told that applies to the full range of cases, I'm also not optimistic. And other things being equal, we should avoid an error-theoretic account.

## 5 The Variability of Reasons: Some Consequences

So far I've considered two different challenges facing those who place a lot of theoretical weight on our ordinary normative judgments involving the count noun *reason(s)*, the first arising from the relationship between the count noun and the mass noun and the second arising from the intra-contextual variability of reasons-claims. It's worth emphasizing, however, that the mere fact that a theory or principle is formulated primarily using the count noun *reason(s)*—in its ordinary normative sense—isn't always a problem. And that's because it will usually be possible to reformulate the theory or principle in such a way that it no longer makes use of the count noun but is in all important respects equivalent to the original. But other times it won't be possible to reformulate the theory or principle in this way, and it's in these cases (and only these



cases) that I think the theories or principles should be viewed with suspicion. The recent emphasis and burgeoning literature on reasons in (meta)ethics and epistemology furnishes us with a myriad of cautionary tales, but for the purposes of illustration I'll just consider two. In particular, I'll explain how the foregoing considerations threaten to undermine the main argument for moral particularism as well as Schroeder's recent defense of a Humean theory of practical reasons.

## 5.1 Against particularism (moral or otherwise)

Particularists—as I'll be understanding them—about a given normative domain  $D$  (e.g. moral, practical, epistemic) are concerned to deny that there are any substantive, finite, and exception-less principles concerning  $D$ .<sup>36</sup> As Berker (2007: 112) explains, in arguing for their view particularists have historically assumed

...a three-level framework according to which non-normative facts at the first level determine the facts about reasons at the second level, which in turn determine other normative facts such as the overall rightness and wrongness of actions at the third level, and they [claim] that there is variability of an inscrutably complex sort both in how the second level depends on the first and in how the third level depends on the second.

We can thus understand particularism as a two-fold thesis: it concerns both the way in which reasons arise out of a situation's non-normative features and the way in which the reasons that are present in a given situation combine to yield the overall normative status of a possible action or attitude (call it  $\varphi$ ).<sup>37</sup> More fully:

**holism about reasons:** for every nonnormative feature that is a reason for/against  $\varphi$ -ing in one possible context, there is another possible context in which that same feature either is a reason of opposite valence or else fails to be a reason of any kind.<sup>38</sup>

**noncombinatorialism about reasons:** the combinatorial function takes as input the valence and weight of all the reasons in a given possible situation and gives as

---

<sup>36</sup> A bewildering variety of theses have fallen under the label "particularism" over the past couple decades. I merely intend to capture one particularly prominent strain of particularist thought. A proper examination of the consequences for the many different particularist theses would require a much more exhaustive treatment.

<sup>37</sup> This, along with what immediately follows, is a generalization of Berker's characterization of moral particularists on pp. 119-122.

<sup>38</sup> There are weaker formulations of holism that are not committed to the universal generalization being true—as both Dancy (2004) and McKeever and Ridge (2006) note—though exceptions will be treated as "special".

output the normative status (e.g. rightness or wrongness) of each  $\varphi$  is not finitely expressible (and so not additive).

Although the argument for particularism based on holism has been challenged on a number of fronts (see especially Ridge and McKeever (2006)), whatever force it retains is undermined once we realize that reasons are merely representatives—that is, once we understand reference to reasons as reference to facts that serve as representatives of normative clusters. Since particular facts may belong to, and hence serve as representatives of, different normative clusters in different (non-conversational) contexts, it's only to be expected that, as Margaret Little (2001: 34) puts it, a “consideration that in one context counts [as a reason] for an action, can in another count against it or be irrelevant”. Such contextual variability is straightforwardly accommodated by the account of reasons as representatives—an account which is intentionally silent about the nature of the underlying normative facts—and so it doesn't, by itself, have any robust normative or metaphysical implications. In this way, the argument from holism (at least in its usual guise) is effectively neutered, since it wrongly takes the contextual variability of reasons-claims to reveal something significant about the normative realm when it's really just an artifact of our reasons-talk. Non-particularists can thus accommodate, and indeed predict, the relevant data without conceding the metaphysics. Support for particularism will have to come from elsewhere—the variability of our judgments about reasons is a non-starter.

Importantly, one might still deny that there are any substantive, finite, and exception-less principles in a given normative domain  $D$  by arguing that the normative forces (support, opposition, etc.) generated by  $D$ -clusters do not interact in a way that is finitely systematizable. (This concerns the link between the second and third levels above, rather than first and second.) But this would only vindicate a kind of normative *regionalism*, not holism, since the normative relevance of (at least some) normative clusters is arguably constant across contexts, even if the interactions between clusters that together determine all-things-considered statuses remain invariably complex.

## 5.2 Against Hypotheticalism

Schroeder (2007) ably and imaginatively defends a (neo-)Humean theory of reasons that he calls Hypotheticalism. At the heart of this view is the following claim:

**Reason\*** For all propositions  $r$ , agents  $x$ , and actions  $a$ , if  $r$  is a reason for  $x$  to do  $a$ , that is because there is some  $p$  such that  $x$  has a desire whose object is  $p$ , and the truth of  $r$  is part of what explains why  $x$ 's doing  $a$  promotes  $p$ .

If true, **Reason\*** would provide what many philosophers would love to have: a reduction of normative reasons (at least of the practical variety), where facts about reasons are explained in terms of, and ultimately grounded in, facts about the psychological states (in the case of Humean theories, desires) of agents. We already know that in

some cases some reasons are to be explained in terms of the agent's psychology, and it would be nice if we could explain all reasons in such terms.

Of course, Humean theories face a familiar battery of objections, and much of Hypotheticalism's attractiveness is due to its apparent ability to avoid or blunt many of them. In particular, central to Schroeder's defense of Hypotheticalism is the claim that there is a principled distinction to be drawn between *reasons* and *background conditions*, and that desires merely belong to the latter. Now, the distinction between a kind of thing and what explains why something is a thing of that kind is, in general, a perfectly good distinction. But Schroeder realizes there are certain cases in which such a distinction has "proved to be hard to maintain—e.g. in the literature on causation" (24). This is important because of the striking and systematic parallels between our talk of causes and our talk of reasons—parallels that I've only been able to hint at so far. In brief, however, *reason(s)* and *cause(s)* pattern linguistically in very similar ways, they both exhibit the same kind of intra-contextual variability explored above, and they are both plausibly governed by Use Good Representative-like norms.<sup>39</sup> (For a defense of the latter two claims in the case of causation, see Swanson (2010).) Such similarities become less surprising once it's realized that reasons and causes are both "because"—they are things (facts, states, events, whatever) that help explain things.

Given the systematic similarities between our reasons-talk and causal-talk, the challenge for Schroeder is straightforward: why think the needed distinction between reasons (in the ordinary normative sense) and background conditions is any more metaphysically significant than the distinction between causes and background conditions? Consider the standard example of a match being struck and catching on fire. Although some have argued that only the striking of the match is a (or "the") cause of the fire, with the presence of oxygen merely being a background condition, it's widely held that such discrimination between causally relevant factors is untenable and shouldn't be reflected in our metaphysics. For there are other contexts in which it's

---

<sup>39</sup> Notice that just as there are singular ascriptions, both specific and non-specific, of reasons, so there are singular ascriptions of causes:

- Singular**
- (R1) The fact that there will be dancing at the party is a reason for Billy to go.
  - (C1) Sally's sleep apnea is a cause of her fatigue.
  - (R2) There is a reason for Billy to go to the party.
  - (C2) There is a cause of Sally's fatigue.

Similarly, just as there are general ascriptions of reasons, so there are general ascriptions of causes:

- General**
- (R3) The fact that there will be dancing at a party is a reason to go.
  - (C3) Sleep apnea is a cause of fatigue.

Moreover, just as there are related "threshold-y" and "contributory" normative locutions, so there are "threshold-y" and "contributory" causal locutions:

- Threshold-y**
- (R4) Billy should go because there will be dancing at the party.
  - (C4) Sally is fatigued because of her sleep apnea.
- Contributory**
- (R5) The fact that there will be dancing at the party is relevant to whether Billy should go.
  - (C5) Sleep apnea contributes to Sally's fatigue.

clear that the underlying metaphysics remains the same and yet what we're inclined to count as a (or "the") cause changes. (Suppose, for example, the match is being repeatedly struck in a vacuum when oxygen suddenly rushes in.)

Why not think the same lesson applies in the case of reasons? That is, just as (e.g.) the striking of a match and the presence of oxygen are on a causal par—they're both part of the same causal process terminating in the presence of fire and the causal relevance of each depends on the presence of the other—why not treat the fact that (e.g.) Billy loves to dance and the fact that there is dancing at a party as being on a normative par? Why not think that they're both part of the same normative cluster bearing on whether Billy should go and the normative relevance of the one depends on the presence of the other (together belonging to what I've called a "normative cluster")? In the former case we have good reason to think the distinction between causes and background conditions isn't metaphysically robust, and that it tells us more about our interests and information than reality. Given the parallels between reasons and causes, the needed distinction between reasons and background conditions appears to fare no better.

### 5.3 The "flattening" worry

One important worry with the account I've explored so far is that it seems to ignore important distinctions between different explanatory roles that the facts belonging to a given normative cluster might play. Denying the distinction between reasons and background conditions makes this worry vivid. For recall that normative clusters are just clusters of facts that, taken together, are normatively relevant (in some particular way, to some degree) to some action or attitude. To take normative clusters as explanatorily basic—in the metaphysical sense—seems tantamount to "flattening" the explanatory landscape, treating all facts within them as being on a par. But this might seem problematic—we might want to distinguish between different (metaphysically robust) explanatory roles that facts belonging to a given normative cluster might play, such as (cf. Dancy (2004), Bader (this volume)):

- (a) that which *makes it the case that* there is (a certain amount of) normative support for  $\varphi$ -ing (the so-called "grounds" of support),
- (b) that which (merely) *enables* there to be such support ("enablers"), and
- (c) that which (merely) *modifies* the amount of support ("modifiers").

By ascribing ultimate explanatory responsibility to normative clusters rather than the facts cited as reasons, aren't I guilty of flattening the explanatory landscape?

No, I'm not. Everything I've said is compatible with the view that there are different metaphysically robust explanatory roles to be played by different (sub-clusters of) facts. Even if full explanatory responsibility resides with a given cluster, that's compatible with different facts within it playing different roles, whether on their own or in conjunction with others. Indeed, the distinctions above are just particular applications

of more general metaphysical distinctions that might, with equal or more justice, be applied to a wide range of normative notions, such as *value*, *evidence*, *justification*, *rights*, *obligations*, and, of course, *reason*.

In fact, similar distinctions can be (and are) drawn outside the normative realm altogether. For instance, we might want to distinguish the “grounds” of friendship—i.e. what it is in virtue of which someone is a friend—from the mere “enabling conditions” of friendship, as well as from various possible “modifiers” of one’s status or standing as a friend. Or we might not. Whether such distinctions mark genuine differences between explanatory roles is something to be decided on a case-by-case basis. And although the story I’ve told about ordinary reasons-claims and their relationship to normative clusters is largely silent on the extent to which the such distinctions apply in this particular case, the intra-contextual variability data does suggest that they normally fail to apply directly to the facts we cite as reasons. At best the distinctions apply to larger clusters of facts that include the ones we cite.

In fact, the promiscuity of reasons-claims suggests something stronger—namely, that the facts we regularly and unproblematically cite as reasons fail to play the same metaphysical (as opposed to communicative) role, even when taken together with other facts. For example, a lot of philosophers want to deny that the ability to  $\varphi$  is itself a reason—or part of a reason—to  $\varphi$ , but there are many contexts in which it seems perfectly acceptable to cite one’s ability to  $\varphi$  as a reason to  $\varphi$ . These include contexts in which  $\varphi$  is obviously choiceworthy—e.g. ‘The fact that you can help someone in need is a reason to do so’—as well as those in which one isn’t usually able to  $\varphi$ —e.g. ‘It’s rare to get the chance to ask the President a question. The fact that you’re able to right now is a reason to do so.’ We face a similar range of choices with respect to such claims as those in DANCE and SMOKE. One option is to bite the bullet and say that they’re strictly speaking false, offering some pragmatic story to explain (or explain away) their acceptability. Another—which I favor—is to accept them as perfectly good reasons-claims, since the facts cited are good representatives in the imagined contexts and that’s all that being (i.e. counting as) a reason really amounts to. Nothing would then follow about the specific metaphysical role the facts play within the cluster they represent—that would be a further, substantive issue, to be settled on broadly theoretical grounds.

Regardless of which option we choose, however, the point remains that ordinary reasons-claims tell us very little on their own. All they tell us is that the fact cited, together with other facts, plays some role in explaining—in the pragmatic, but not necessarily metaphysical, sense—why there is reason to perform the relevant action or adopt the relevant attitude. They don’t tell us *which* other facts those are, *what* the specific explanatory role is, *how much* reason there is, nor what *kind* of reason there is. All that a claim of the form *r is a reason (for S) to  $\varphi$*  entails is that there’s something (not: some *thing*) to be said in favor of  $\varphi$ -ing, and that *r* is something it makes sense to cite in an effort to explain why. We need to rely on background information and engage in further theorizing to fill in all the gaps.

## 6 Conclusion

I've explored various subtleties in how we ordinarily think and talk about reasons, revealing some of the limitations and liabilities of taking such talk seriously. Somewhat ironically, the challenges I've considered—those arising from the use of *reason* as both a count noun and as a mass noun as well as from the context-sensitivity of reasons-claims—are the result of taking the language of reasons *more* seriously than it has been taken by many, if not all, of its proponents. It turns out that our judgments about reasons are, at best, a partial and highly defeasible guide to what really matters, both normatively and metaphysically.<sup>40</sup>

---

<sup>40</sup> Thanks to Selim Berker, Paul Bloomfield, John Broome, Cian Dorr, Billy Dunaway, Ian Grubb, Liz Harman, Daniel Harris, Errol Lord, Conor McHugh, Matt Moss, Caleb Perl, Jim Pryor, Mark van Roojen, Stephen Schiffer, Justin Snedegar, Eric Swanson, Ang Tong, David Velleman, Jonathan Way, and audiences at University of Colorado Boulder, University of Southern California, University of Kent, New York University, University of California, Riverside, and the St. Louis Annual Conference on Reasons and Rationality for helpful discussion and/or comments. Special thanks go to Barry Maguire for multiple rounds of valuable feedback.

## References

- Bach, E. 1986. Natural Language Metaphysics. In *Logic, Methodology, and Philosophy of Science* VII: 573-595.
- Bader, R. (this volume). Conditions, Modifiers, and Holism.
- Berker, S. (2007). Particular Reasons. *Ethics*, 118(1): 109–139.
- Broome, J. (2004). Reasons. In *Reason and Value: Essays on the Moral Philosophy of Joseph Raz*. Oxford University Press.
- (2013). *Rationality through Reasoning*. Wiley Blackwell.
- Brunero, J. (2013). Reasons as Explanations. *Philosophical Studies*, 165(3): 805–24.
- Dancy, J. (2004). *Ethics Without Principles*. Oxford University Press.
- (2015). Reasons for Broome. In *Weighing and Reasoning*. Oxford University Press.
- Davidson, D. (1963). Actions, Reasons, and Causes. *Journal of Philosophy* 60(23): 685–700.
- Enoch, D. (2011). On Mark Schroeder’s Hypotheticalism. *Philosophical Review*, 120(3): 423-446.
- Finlay, S. (2014). *Confusion of Tongues: A Theory of Normative Language*. Oxford University Press.
- Fogal, D. (forthcoming). Review of Stephen Finlay’s *Confusion of Tongues: A Theory of Normative Language*. *Ethics*.
- (2015). *Bad Attitudes: Rationality and its Discontents*. Ph.D. Dissertation, NYU.
- Gillon, B. S. (2012). Mass Terms. *Philosophy Compass*, 7(10): 712–730.
- Grice, H. P. (1989). *Studies in the Way of Words*. Harvard University Press.
- Huddleston, R., and Pullum, G. K. (2002). *The Cambridge Grammar of the English Language*. Cambridge University Press.
- Kim, J. (1998). *Mind in a Physical World*. MIT Press.
- Lewis, D. K. (1973). Causation. *Journal of Philosophy*, 70: 556–567.
- Little, M. O. (2001). On Knowing the “Why”: Particularism and Moral Theory. *The Hastings Center Report*, 31(4): 32–40.
- Mackie, J. L. (1965). Causes and Conditions. *American Philosophical Quarterly*, 2(4).
- McKeever, S. D., and Ridge M. R. (2006). *Principled Ethics: Generalism as a Regulative Ideal*. Oxford University Press.
- Olson, J. (2009). Reasons and the New Non-Naturalism. In *Spheres of Reason: New Essays in the Philosophy of Normativity*. Oxford University Press.
- Parfit, D. (2011). *On What Matters*, Vol. I. Oxford University Press.
- Raz, J. (2000). *Engaging Reason*. Oxford University Press.
- Scanlon, T. M. (1998). *What We Owe to Each Other*. Harvard University Press.
- (2014). *Being Realistic about Reasons*. Oxford University Press.
- Schroeder, M. (2007). *Slaves of the passions*. Oxford University Press.
- (2015). Being Realistic About Reasons, by Scanlon, T.M. *Australasian Journal of Philosophy* 93(1): 195-198.
- Skorupski, J. (2010). *The Domain of Reasons*. Oxford University Press.
- Strawson, P. F. (1959). *Individuals: An Essay in Descriptive Metaphysics*. Routledge.
- Strevens, M. (2013). Causality Reunited. *Erkenntnis*: 1–22.
- Swanson, E. (2010). Lessons From The Context Sensitivity of Causal Talk. *The Journal of Philosophy*, 107(5), 221–242.
- (2012). The Language of Causation. In *The Routledge Companion to the Philosophy of Language*. Routledge.
- Woodward, J. (2014) Scientific Explanation. *The Stanford Encyclopedia of Philosophy*, <<http://plato.stanford.edu/archives/win2014/entries/scientific-explanation/>>.