

## The Weight of Reasons<sup>1</sup>

Daniel Fogal (NYU) and Olle Risberg (Uppsala)

*Penultimate version - please cite published version where possible*

### 1. Introduction

In contemporary (meta)normative theorizing, normative reasons—henceforth ‘reason(s)’—receive a lot of attention. They are crucial to the Reasons-First project, according to which *all* normative facts can be identified with or explained in terms of facts about reasons—roughly, facts of the form [*r* is a reason for *S* to do *A*].<sup>2</sup> But one need not be a Reasons-Firster to find it plausible that what you ought to do (and believe, feel, etc.) often depends on the relative *weight*, *strength*, or *significance* of the considerations that speak for and against each available alternative. This idea underlies what Mark Schroeder calls the ‘classical’ argument for Reasons-First:

In every case, what you ought to do (all-things-considered, as Ross put it) is a matter of the *competition* between different factors [...] Ross called these competing moral forces *prima facie* duties, and later theorists have called them *reasons* (2021: 3112; italics in original).

If you’re considering whether to donate money to charity, for instance, the fact that doing so would alleviate suffering speaks in its favor, while the fact that it would worsen your children’s lives speaks against it. What you ultimately ought to do, the thought goes, depends on which consideration is stronger, weightier, or otherwise more significant.

As this suggests, one of the most important features of reasons is that they vary in strength or weight, and their strength or weight in turn helps determine what we ought to do. In view of this, one would expect reasons-enthusiasts in general, and Reasons-Firsters in particular, to have a good account of how the weight of reasons should be understood. But they don’t. Reasons-Firsters have largely neglected this question, and what little they have said is highly problematic. And they’re not alone—nearly all existing analyses of reasons fail to provide plausible accounts of their weight.

---

<sup>1</sup> For helpful comments, we would like to thank Krister Bykvist, Nicholas Makins, Chris Tucker, an anonymous referee for *Philosophical Studies*, and audiences at the 2022 Chapel Hill Normativity Workshop, the Spring 2022 Virtual Metaethics Colloquium, and the 2022 Stockholm June Workshop in Philosophy. The work on this paper was supported by Grant 2020-01955 from the Swedish Research Council.

<sup>2</sup> Adherents include Scanlon (1998, 2014), Skorupski (2010), Schroeder (2007, 2021a, 2021b), Lord (2018), and Rowland (2019). Opponents include Kearns & Star (2008), Broome (2013), McHugh & Way (2016), Howard (2019), Wodak (2020), and Fogal & Risberg (forthcoming).

Our aim in this paper is two-fold: to provide an analysis of the weight of reasons, and to argue that it's more plausible than alternative accounts. Our analysis of weight is a natural corollary of an independently-motivated analysis of reasons developed elsewhere (Fogal & Risberg forthcoming), according to which reasons are sources of normative support. On this view, the weight of a reason is straightforward: it is simply a matter of how much support it provides. This account directly conflicts with Reasons-First, as it implies that the notion of a normative reason is analyzable in terms of normative support, and hence isn't normatively fundamental.<sup>3</sup>

Some philosophers, including some reasons-enthusiasts, find the idea that reasons have weight to be mysterious (e.g., Schroeder 2007: 124–125; Lord 2018: 194; McHugh & Way 2022: ch. 5). In some cases (to be discussed below), this leads them to endorse complicated, indirect accounts of various weight-related phenomena, such as what it is for a reason to outweigh another reason. However, we don't see how the idea that reasons have weight is any more mysterious than facts about reasons themselves. As we'll suggest, this selective concern about avoiding putatively mysterious notions has led past theorizing astray, since paying attention to the weight of reasons reveals important things about their nature.

The plan is as follows. After some preliminary remarks (§2), we introduce the view of reasons as *sources of support* and explain the account of the weight of reasons it provides (§3). We then criticize competing views of weight: some that have been suggested by Reasons-Firsters (§4) and some that are related to competing analyses of reasons (§5), including other analyses of reasons in terms of explanation (§5.1), the 'Reasons as Evidence' view (§5.2), and the view of reasons as premises in good reasoning (§5.3).

## 2. Preliminaries

Our main question in this paper is: What is it for a reason to have a certain strength or weight, and for one reason to be 'stronger' or 'weightier' than another? (Terminological note: we use 'strength' and 'weight' interchangeably.) We take it to be an adequacy constraint on accounts of reasons that they are compatible with—and ideally themselves provide—plausible answers to these questions.

It's common to distinguish between *basic* and *derivative* reasons. Basic reasons are the considerations that most fundamentally speak for or against various alternatives, whereas derivative reasons are "only reasons in virtue of their relationships to the [basic] ones" (Maguire & Snedegar 2021: 366). Instrumental reasons, or reasons to take the necessary means to ends that we have independent reason to pursue, are paradigmatic examples of derivative reasons. We focus on the weight of basic reasons.

We are interested in (basic) reasons of all kinds: moral, epistemic, prudential, and so on. Some authors provide analyses of only one kind of reason, such as moral reasons. However, even an analysis that only concerns a particular kind of reason should provide

---

<sup>3</sup> We use 'notion' as a fudge term to cover both concepts and properties and relations.

(or at least be compatible with) a plausible account of the weight of that kind of reason. And ideally that account of weight should be plausible, *mutatis mutandis*, for other kinds of reasons as well. After all, in many or all normative domains, it is natural to think about the weight of reasons in much the same way; for instance, just as we have stronger epistemic reasons for some beliefs than for others, we also have stronger moral and prudential reasons for some actions than for others.

By an ‘account’ of weight, we mean a theory about how facts of the form [ $r$  is a reason of weight  $d$  to  $\varphi$ ], [ $r$  is a weightier reason to  $\varphi$  than  $r^*$ ], and so on, are to be analyzed, or otherwise understood. (The view that some such facts are unanalyzable, and thus not to be analyzed at all, counts as an account in this sense.) The relevant type of analysis, we assume, is broadly metaphysical: it says *what it is* for a reason to have a certain weight, and doesn’t speak directly to, say, the *concept* of reason’s weight. That said, a similar debate could also be had on the conceptual level, and our main points could be modified accordingly.

Lastly, we will speak freely about ‘weight(s)’, ‘amount(s)’, and other magnitudes, though we remain neutral on whether such talk is ultimately best understood as invoking genuine, ‘reified’ entities (such as properties or abstract particulars) or best understood in more nominalist-friendly terms. An analogous issue arises with respect to physical weight and other empirical magnitudes—it is controversial, for example, whether we should countenance genuine entities such as *700 kilos*, or only particular concrete objects (such as cows and statues) that *weigh* 700 kilos.<sup>4</sup> The various accounts of reasons we discuss are neutral on such issues. Hence, the idea that reasons have weight does not itself require the existence of a distinctive class of entities: *weights*. We also leave open questions about how these magnitudes can be measured, such as whether they can be perspicuously represented by real numbers, whether they are completely ordered, and which ‘units’ one should use in such representations.<sup>5</sup> In particular, we don’t exclude the possibility of incommensurability, indeterminacy, or vagueness with respect to whether a particular

---

<sup>4</sup> For recent discussion, see, e.g., Kim (2016); on corresponding questions about value magnitudes, see Bykvist (2021).

<sup>5</sup> The question of what unit to choose is part of what leads McHugh and Way (2022: ch. 5) to deny that reasons have weights. Ultimately, however, this may simply be a conventional matter. Krister Bykvist (2021: 6335) makes the analogous point about value magnitudes, or units for measuring value, nicely:

“It is true that value magnitudes differ from empirical ones, such as length, in that we do not have conventional names for them. But this does not seem to be a significant difference. That we can so easily think of measure phrases for empirical magnitudes has to do with the fact that we have widely accepted conventions about the standards for these magnitudes. For example, an object (according to the current definition, a certain path) is stipulated to be 1 meter. But nothing prevents us, in principle, from doing something similar for value magnitudes. For example, we could use the value of a certain token pleasure, stipulated to be 1 hedonic value unit, as the standard for value and meaningfully say things like ‘ $x$  is 1 hedonic value unit good’, which means that  $x$  has the same value magnitude as the token pleasure. Of course, this is just one possibility. There might be better ways to select a standard for value.”

reason is ‘weightier’ than another particular reason. On the contrary, some imprecision is to be expected, since not even facts about physical magnitudes are always fully precise.<sup>6</sup>

### 3. The S.O.S. View: Reasons and Support

#### 3.1 *Reasons as Sources of Support*

To set the stage for our analysis of weight, we’ll first introduce the analysis of reasons defended in Fogal and Risberg (forthcoming). On what we’ll call the ‘Sources of Support’ (S.O.S) view, to be a reason is to *provide*—or be a *source of*—*normative support*. More formally:

**S.O.S.<sub>R</sub>** For  $r$  to be a reason for  $S$  to  $\varphi$  is for  $r$  to provide a certain amount of normative support for  $S$ ’s  $\varphi$ -ing.

Two key notions should be clarified: *provides* and *support*. To say that  $r$  *provides*—or *generates* or *is a source of*—a certain amount of normative support is just to say that  $r$  *makes it the case* that there is that amount of normative support, or to say that the latter obtains *in virtue of* the former. We take this to be an instance of a non-causal, non-pragmatic kind of explanation that is often called ‘metaphysical’ or ‘grounding’ explanation.<sup>7</sup> For example, just like the fact that a barn is crimson might make it the case that it is red, so the fact that it seems to Anandi that  $p$  might make it the case that there is a certain amount of support for her believing that  $p$ . If so, on the S.O.S. view, that fact is a reason for her to believe that  $p$ .<sup>8</sup>

Although we use the mass noun ‘support’ as a technical term, what it denotes is familiar. In ordinary thought and talk, the relevant notion of normative support is commonly expressed by the mass noun ‘reason’.<sup>9</sup> Thus, to say that there is normative support for  $\varphi$ -ing is essentially to say that there is reason to  $\varphi$ . More colloquially, it’s just to say that there is something (though not necessarily some *thing*) to be said in favor of  $\varphi$ -ing. The relevant notion of support is thus typically *defeasible*—i.e., it can be opposed and

---

<sup>6</sup> For instance, the height of the Eiffel Tower, measured in nanometers, is presumably imprecise. A simple way to represent imprecise weights is by sets of real numbers rather than individual ones, just like imprecise credences can be represented by sets of precise credences. However, see Nebel (2023) for an approach to the measurement and aggregation of well-being that does not assume numerical representability at all. A similar approach may be viable for measuring support and reasons. Moreover, note that although standard accounts of extensive measurement presuppose that the quantities to be measured and aggregated are completely ordered, it’s arguable that extensive measurement can be reconciled with incommensurability (see Carlson 2008; Nolt 2022).

<sup>7</sup> The relevant notion is *direct, full* grounding; see Fogal & Risberg (forthcoming: §4.1.1).

<sup>8</sup> Throughout the paper, we set aside the role that general normative principles might (and we think do) play in normative explanations—see Fogal & Risberg (2020).

<sup>9</sup> Cf. Fogal (2016).

undermined to varying degrees<sup>10</sup>—and *pro tanto*, in that there can be support for actions and attitudes that ultimately ought not to be done or had.

Given the close relationship between the mass nouns ‘support’ and ‘reason’, the S.O.S. view of reasons could be rephrased colloquially as follows: to be a reason (count noun) to  $\varphi$  is to be a source of reason (mass noun) to  $\varphi$ . The relationship between count and mass uses of ‘reason’ thus parallels that between count and mass uses of ‘light’ and ‘pleasure’, for example, with the (denotation of the) mass noun being explanatorily prior to the (denotation of the) count noun. Just as lights are sources of light, and pleasures are sources of pleasure, so reasons are sources of reason.

In addition to its intuitive plausibility, the S.O.S. view enjoys two further attractive features: its *generality* and its (meta)normative *neutrality*. It is general insofar as it applies to reasons of all kinds—just as we can distinguish between (e.g.) moral, prudential, and epistemic reasons, we can also distinguish between moral, prudential, and epistemic support, and in each case use the latter to analyze the former. And it is (meta)normatively neutral insofar as it doesn’t take a stand on a variety of substantive issues, including:

- (a) The *sources* of support, or which particular facts are reasons (e.g., whether the fact that it appears that  $p$  always provides support for, and hence is always a reason for, believing that  $p$ );
- (b) The *varieties* of support, including both how many different normative *domains* there are (e.g., moral support, prudential support, epistemic support, etc.) and how many different *perspectival* notions of support there are within each (e.g., what there’s support for doing relative to the facts, or relative to one’s evidence, or relative to one’s attitudes, etc.);
- (c) The *ultimate basis* of support, both within and across domains (e.g., whether all facts about moral support, and hence moral reasons, are ultimately grounded in facts about well-being, and whether support-facts in all domains have the same ultimate basis);
- (d) The *objects* of support, or what there can be support, of one kind or another, *for* (e.g., whether there can be normative support, and hence reasons, for affective attitudes, or whether there can be moral support, and hence reasons, for belief);
- (e) The *imperialistic reach* of support, or whether any normative facts besides reasons-facts can be analyzed in terms of support;

---

<sup>10</sup> To illustrate, suppose that Andreas, a trusted friend, tells us that  $p$ , thereby providing some support for believing that  $p$ . If Mira, another trusted friend, testifies that  $p$  is false, then that provides support that *opposes* our support for believing  $p$ . Were she instead (or in addition) to testify that Andreas is a liar, then that would *undermine* the support for believing  $p$  that his testimony provides (at least to some degree).

- (f) The *reduction* of support, or whether support-facts can be analyzed in other (normative or non-normative) terms.

All of these issues should be settled by further theorizing. The S.O.S. view of reasons is thus itself a rather conservative metanormative view, in at least three ways: it analyzes *one* normative notion (as opposed to more than one) in terms of *another normative* notion (as opposed to a non-normative one) that is *intuitively closely related* (as opposed to not intuitively closely related) to it.

### 3.2 *Weights as Amounts of Support*

A further advantage of the S.O.S. view is that it, unlike many other analyses of reasons, gives a straightforward answer to the main question of this paper: the weight of a reason is simply a matter of how much support it provides. Thus:

**S.O.S.<sub>w</sub>** For  $r$  to be a reason of weight  $d$  for  $S$  to  $\varphi$  is for  $r$  to provide a corresponding amount of normative support,  $d$ , for  $S$ 's  $\varphi$ -ing.

The weight of a reason, on this view, is treated as a 'scalar' quantity, whereas support is akin to a 'vector' quantity. Weight is similar to scalar quantities like mass and speed insofar as it is completely described by its magnitude (*the amount of support provided*), while support is similar to vector quantities like force and velocity insofar as it has both magnitude (*amount of support*) and direction (*support for*).<sup>11</sup>

The above analysis of the weight of reasons is effectively built into our analysis of reasons themselves, and thus inherits many of its attractive features, including its generality, neutrality, and intuitive plausibility. As a result – and again unlike many other

---

<sup>11</sup> The scalar notion of weight should be distinguished from the vector notion. In science, for example, weight is commonly defined as the force exerted on a body by gravity, though a scalar notion of weight as the magnitude of the gravitational force is also used. Weight is sometimes also treated as a vector quantity in ethics, such as in the work of Chris Tucker (for the basic idea, see Tucker 2022: 371, 378, and note 14; for an overview, see Tucker ms). There are a variety of similarities as well as differences between Tucker's framework and our own, and we plan to engage with his work at proper length elsewhere. For now we'll settle for a few quick remarks. First, much of what Tucker cares about most, both normatively and metanormatively, centers around the distinction between two kinds of weight: *justifying* weight and *requiring* weight. Although the support-based view isn't committed to such a distinction, it's compatible with such: justifying weight for  $A$  can be understood as support for  $A$  (i.e. reason to  $A$ ) while requiring weight for  $A$  can be understood as opposition to  $\sim A$  (i.e. reason not to  $\sim A$ ), where opposition is the negatively-valenced counterpart of support. Second, there are non-trivial differences between the weight metaphor and the support (or, more generally, force) metaphor. For example, while the weight of a reason is essentially "attached" – it's the weight of a reason – facts about support are not – it doesn't automatically follow from the fact that there is support for  $A$  that there is a source of that support (and hence that there is a reason for  $A$ ). This is why, in order to allow for the possibility of some actions or attitudes having a 'default' positive normative status, Tucker is forced to posit default reasons – i.e. facts that have a certain weight by default – even though it's hard to identify plausible candidates for such. Positing default support, in contrast, doesn't force one to find a fact that serves as its source.

analyses – it also gives a straightforward account of what it is for a reason to be *weightier than* another, by vindicating the following attractive principles:

**Simple Weightier-Than** For a reason  $r$  to be weightier than a reason  $r^*$  is for  $r$  to have greater weight than  $r^*$ .<sup>12</sup>

**Simple Outweighing** For a reason  $r$  to  $\varphi$  to outweigh a reason  $r^*$  to  $\Psi$  is for  $r$  to have greater weight than  $r^*$ , where  $\varphi$  and  $\Psi$  are conflicting alternatives.<sup>13</sup>

Simple Weightier-Than and Simple Outweighing seem platitudinous. To illustrate, suppose again the fact that donating money would alleviate suffering is a reason to do it, while the fact that it would worsen your children’s lives is a reason to keep the money. On S.O.S<sub>W</sub>, the reason that is weightiest (and the reason that outweighs the other) is the one that provides the most support for the relevant alternative, and hence has the greatest weight. How much support (if any) a given fact provides is, in turn, also a question for first-order normative theorizing; for instance, deontologists of different stripes may disagree about how much moral support the fact that one has promised to  $\varphi$  provides for  $\varphi$ -ing, and hence disagree about the strength of the corresponding moral reason. Our view does not answer such questions, nor should it.

We can thus expand the list of substantive issues that the S.O.S. view is neutral on:

- (g) The *amount(s)* of support, or how much support particular facts provide, and hence how strong of reasons they are.

We can also add:

- (h) The *aggregation* of support, or whether (and if so, how) different kinds and amounts of support can be aggregated, both within and across domains (e.g., whether morality is “overriding” in the sense that if there’s most moral support for  $\varphi$ -ing, then there’s most support for  $\varphi$ -ing all things considered);<sup>14, 15</sup>

---

<sup>12</sup> McHugh and Way (*forthcoming*: ch. 5) formulate a similar principle, and note that it “might seem highly, even overwhelmingly, natural”, but ultimately propose a different view. We return to their view of weight below (§5.3).

<sup>13</sup> We assume that a reason for a given alternative can only outweigh a reason for a conflicting – i.e., non-compossible – alternative, since reasons to  $\varphi$  don’t ‘compete’ with other reasons to  $\varphi$  and so can’t outweigh them.

<sup>14</sup> Questions about aggregation are intimately connected to questions about how to measure support (cf. footnotes 4 and 5).

<sup>15</sup> Schroeder (2007: 125–126) argues that the fact [there is a reason to  $\varphi$ ] can itself be a reason to  $\varphi$ , and notes that its weight cannot plausibly be aggregated with other reasons to  $\varphi$ . However, this issue can be set aside for present purposes, since we are concerned with the weight of basic normative reasons, and the fact [there is a reason to  $\varphi$ ] is at best a non-basic reason to  $\varphi$ .

- (i) The *modification* of support, or how the weight of reasons can be modified.

Regarding modification: it is common, for example, to distinguish between reasons and various background conditions, such as ‘enablers’ and ‘disablers’ (Dancy, 2004; Bader, 2016). To illustrate, many philosophers deny that the fact that you can  $\varphi$  is itself a reason, or part of a reason, for you to  $\varphi$ ; instead, it merely *enables* the fact that you ought to  $\varphi$ . S.O.S.<sub>w</sub> can accommodate these nuances: enablers and disablers can play their characteristic roles and hence be part of the full explanation of why there is a certain amount of normative support for  $\varphi$ -ing, but insofar as they aren’t themselves sources or providers of that support, they won’t count as normative reasons to  $\varphi$ . The same goes for so-called ‘intensifiers’ and ‘attenuators’, if such there be: these will be facts that increase (intensifiers) or decrease (attenuators) the amount of normative support generated by other facts without themselves being sources of support – they modify existing amounts of support (and hence the weight of reasons) without generating it on their own (or themselves being reasons). Whether there really is a need to draw such fine-grained distinctions, and if so how we determine which facts play which roles, are important questions, but ones that we think an analysis of reasons should leave open.

Simple Weightier-Than and Simple Outweighing illustrate one advantage of taking the (monadic) notion of weight seriously, as S.O.S.<sub>w</sub> does, rather than just focusing on relations (such as outweighing relations) among reasons. A related advantage is that by focusing on the monadic notion of weight, S.O.S.<sub>w</sub> naturally allows for a distinction between the overall *balance* of support and the total *amount* of support at issue in a given case.<sup>16</sup> The overall balance of support is a matter of which action or attitude has the most net support, and hence is best supported, all things considered. Notably, the balance of support can remain the same even though the total amount of support present varies. Suppose, for example, that you are handed a coin and are told it is fair, and so have credence .5 that it will land heads if flipped. Suppose that you then proceed to flip the coin a thousand times, resulting in an equal number of heads and tails. What should your credence be that it will land heads if flipped? The same as before. In the second scenario, however, you have significantly more evidence, and hence more support, for believing the coin is fair than you did originally. So while the overall *balance* of support doesn’t change, as reflected in the epistemically appropriate credence (namely, .5 in both cases), the total *amount* of support does – it increases. This increase is reflected in the increased *stability* of one’s credence – after the tosses, one’s .5 credence is much more (rationally) resilient, or less likely to (rationally) change, in the face of additional evidence than it was at the outset.<sup>17</sup> Analogous cases arise in the practical realm as well.<sup>18</sup>

A final advantage of S.O.S.<sub>w</sub> is that, unlike some competing views to be discussed below, it straightforwardly allows for *interpersonal comparisons* of the weight of reasons (at

---

<sup>16</sup> We thank Nicholas Makins (personal communication) for this point.

<sup>17</sup> Cf., e.g., Keynes (1921) and Joyce (2005), though the exact terminology used to make the point varies.

<sup>18</sup> See especially Makins (ms).

least of the same type). Anna's reasons to rescue a drowning child might be stronger than Jens's reasons to comfort a scared child, for instance, and Jonas's reasons to believe that it's raining might be stronger than Krister's reasons to believe in string theory. On S.O.S.<sub>w</sub>, this is because Anna's reasons provide more support for helping the drowning child than Jens's reasons provide for comforting the scared child, and Jonas's reasons provide more support for believing that it's raining than Krister's reasons provide for believing in string theory. Interpersonal comparisons of this type plausibly underlie many further normative judgments; e.g., that Anna would be more blameworthy for not helping the drowning child than Jens would be for not comforting the scared child, and that Jonas's failing to believe that it's raining would be a more serious rational shortcoming than Krister's failing to believe in string theory. Again, as we'll see shortly, interpersonal comparisons like these pose a challenge for many other views.

#### 4. Reasons-First accounts of weight

We have seen how the S.O.S. analysis of reasons provides a straightforward analysis of the weight of reasons—one that analyzes the gradable notion of weight in terms of the gradable notion of support. As this approach is incompatible with Reasons-First, the question arises: how can Reasons-Firsters account for weight?

Perhaps surprisingly, many Reasons-Firsters assume that the reasons-relation only involves argument places for *non*-gradable things—e.g., that  $r$  is a reason for  $S$  to  $\varphi$ , where  $r$  is a fact,  $S$  is a subject, and  $\varphi$  is an action or attitude.<sup>19</sup> Since neither  $r$  nor  $S$  nor  $\varphi$  are such that there can be 'more or less' of them, this makes it hard to see how the gradable notion of a reason's weight should be understood. Few Reasons-Firsters have tried to address this issue, and the proposals that exist are problematic. In this section we critically consider the most prominent suggestions.

##### 4.1 Schroeder's recursive view

Perhaps the most well-known Reasons-First account of weight-related phenomena is due to Schroeder (2007). Schroeder says that it's a "mistake" (2007: 124) to think that reasons have weight, and thus rejects Simple Weightier-Than and Simple Outweighing. Instead he proposes an account of when one reason is weightier than another that focuses on which reason it is correct to 'place more weight on' in deliberation.<sup>20</sup> Simplified slightly:

---

<sup>19</sup> We ignore the possibility that the reasons-relation might involve additional argument places that are unrelated to their weight, such as one for the agent's circumstances (Scanlon 2014).

<sup>20</sup> The notion of a reason's weight also plays a crucial role in the defense of Reasons-First provided by Lord (2018). Despite that, Lord says little about how weight should be understood. He claims to find "mysterious" the idea of what he calls 'atomic weight', which is "the weight that a reason has independently of how it interacts with the other reasons in a particular situation" (2018: 193–194), and invokes a view of intensifiers and attenuators of weight that resembles Schroeder's view of weight, on which they are understood as reasons to place more or less weight on certain reasons (2018: 33).

**Recursive Weightier-Than** For  $S$ 's reason  $r$  to be weightier than  $S$ 's reason  $r^*$  is for it to be correct for  $S$  to place greater weight on  $r$  than on  $r^*$  in deliberation.<sup>21</sup>

The view is recursive because Schroeder understands correct deliberation in terms of how one's reasons support deliberating, which in turn depends on the weight of *those* reasons. He seeks to avoid a vicious regress by appealing to a "base clause" that does not mention weight (2007: §7.3). We set this issue aside – for critical discussion, see Enoch (2011) – as the view faces more straightforward problems. To begin with, once psychological and normative senses of '(placing) weight' are distinguished, the view is counterintuitive as it stands. It is also subject to apparent counterexamples, since there may be reasons which it is not correct to place weight on in deliberation. Indeed, as one of us notes elsewhere (Risberg 2016), this seems to follow from Schroeder's own view of reasons: the fact that there's a surprise party waiting for you next door might be a reason to enter the room, for instance, but since placing weight on it would entail awareness of it – which in turn would render the relevant fact a non-reason – it would not be correct to do so in deliberation. Third, since the view only focuses on how it is correct for an agent to place weight on *her* reasons in *her* deliberation, it fails to account for interpersonal comparisons of the weight of reasons.

Finally, it is often unclear what the reasons to place weight on reasons in deliberation *are* (Enoch 2011; Gregory 2014; McHugh & Way 2022). The reason to place weight on  $r$  cannot be that  $r$  has a certain weight, since that would be both circular and inconsistent with Schroeder's view that reasons lack weights. Moreover, the few examples of reasons to place weight on reasons that Schroeder provides fail to generalize. One example is that the relevant reason is agent-neutral (Schroeder 2007: 142). As McHugh and Way (2022: ch. 5) note, however, this suggestion fails to explain how some agent-neutral reasons can be weightier than other agent-neutral reasons, and how some agent-relative reasons can be weightier than other agent-relative reasons.

All of this points to a fundamental problem with Schroeder's view: that "[p]retheoretically [...] it seems hard to deny that we think that the correctness of placing weight on reasons in deliberation is derivative of their weight, not the other way around" (Enoch 2011: 435). This intuition is easily accommodated by the S.O.S view.

#### 4.2 Scanlon's counterfactual view

In another influential defense of Reasons-First, Scanlon (2014: ch. 5) similarly denies that reasons have weight (or 'strength(s)', in his terminology), and thus denies Simple Weightier-Than and Simple Outweighing. His main argument turns on the claim that there is no unique property – such as the extent to which  $\phi$ -ing would promote value, or

---

<sup>21</sup> Although we formulate this view as an analysis, Schroeder himself seems inclined to accept it only as an extensional thesis (2007: 139–140). Understood as an extensional thesis, however, the view says nothing about what it *is* for a reason to be weightier than another.

pleasure, etc. – that could determine the weight of reasons in all cases (2014: 109–111). Even if Scanlon is right about this, however, it’s hard to see how it motivates denying that reasons have weight – the weight of reasons might depend on different things in different cases, and Scanlon says little to rule this out. Indeed, this possibility is supported by the fact that reasons for action are not the only kind of reasons that have weight; so do reasons for belief and other attitudes. Thus, while the weight of a reason to perform some action may depend on (e.g.) the extent to which the action promotes value, the weight of a reason to believe that  $p$  may depend on something quite different, such as the extent to which the reason indicates that  $p$  is true. These, again, are questions for first-order theorizing. So Scanlon’s conclusion is at best premature.

Rather than appealing to facts about the weight of reasons, Scanlon opts for a counterfactual approach. In addition to the relation of *being a reason for*, he posits the relation of *being a sufficient reason for*, and uses both to provide the following analysis of outweighing (2014: 108):

**Counterfactual Outweighing** For  $S$ ’s reason  $r$  to  $\varphi$  to outweigh  $S$ ’s reason  $r^*$  to  $\Psi$  is for it to be the case that:

- (i)  $r$  is a sufficient reason for  $S$  to  $\varphi$ ,
- (ii)  $r^*$  is a reason, but not a sufficient one, for  $S$  to  $\Psi$ ,
- and
- (iii) if  $r$  had not obtained, then  $r^*$  would have been a sufficient reason for  $S$  to  $\Psi$ .<sup>22</sup>

The basic idea is that  $r$  outweighs  $r^*$  just in case  $r^*$  would have been a sufficient reason in the absence of  $r$ , but fails to be a sufficient reason in the presence of  $r$ , as  $r$  is a sufficient reason to do something else instead. Suppose, for example, that the fact that soda is tasty is a reason for Nils to drink soda, and that this reason is outweighed by the fact that soda is unhealthy. On Scanlon’s view, this holds because (i) the fact that soda is unhealthy is a sufficient reason not to drink soda, (ii) the fact that soda is tasty is a reason, but not a sufficient one, to drink soda, and (iii) if soda hadn’t been unhealthy, then the fact that it’s tasty would have been a sufficient reason to drink it.

Counterfactual analyses are notoriously problematic, and Scanlon provides no reason to think that his account avoids the usual pitfalls for such views. Nonetheless, the fundamental problem for Scanlon’s approach lies elsewhere. It resembles the fundamental problem for Schroeder’s view: the relevant counterfactuals intuitively depend on the weight of the relevant reasons, rather than the other way around. This in turn gives rise to several further difficulties:

---

<sup>22</sup> Our formulation is somewhat simplified. For instance, we ignore the complication that Scanlon requires the relevant counterfactual case to be “as normatively similar as possible” (2014: 108) to the actual case (except that  $r$  does not obtain in it), without explaining how this idea of normative similarity is to be understood.

1. As Scanlon's account only concerns outweighing, it fails to explain how a reason to  $\varphi$  can be weightier than another reason to  $\varphi$ . Scanlon notices this problem and suggests that "such cases of relative strength might be captured by saying that  $r$  is stronger than  $r^*$  as a reason for doing  $\varphi$ ... if there is some consideration  $r^{**}$ , which is a reason for doing some  $\Psi$  that is incompatible with doing  $\varphi$ , such that  $r^{**}$  outweighs  $r^*$  but does not outweigh  $r$ " (2014: 108, notation adjusted). But this suggestion is insufficiently general; for example, there could be two reasons to  $\varphi$  that differ in strength, but no reason, or only weaker reasons, to do anything else.
2. Like Schroeder's view, Scanlon's view fails to explain how interpersonal comparisons of the weight of reasons are possible, since it focuses exclusively on a given agent's reasons and their counterfactual behavior.
3. Because of condition (i), the analysis entails that only sufficient reasons can outweigh other reasons. This is questionable, however, as two insufficient reasons can differ in weight, with one outweighing the other. Suppose, for instance, that Tobias has three options: (a) to donate lots of money to charity, (b) to donate only a little money to charity, or (c) to steal money from the charity. Even if (a) is what Tobias ought to do, his reasons to do (b) can clearly be stronger than (and hence outweigh) his reasons to do (c). But since neither his reasons for (b) nor his reasons for (c) are sufficient, this verdict is incompatible with the counterfactual analysis.

### 4.3 The brute-force view

The final way of accounting for weight within the Reasons-First framework that we will consider is to simply add an extra argument place for weight to the reasons-relation. Skorupski (2010: ch. 2), for example, takes the reasons-relation to include an argument place for the reason's 'degree of strength'. This strikes us as the best of the three approaches. On this view, which we call the *brute-force* view, reasons-facts have the form [ $r$  is a reason of weight  $d$  to  $\varphi$ ], or [ $r$  is a reason that supports  $\varphi$ -ing to degree  $d$ ]. To vindicate Reasons-First, these facts must in turn be taken to be normatively fundamental.

A problem for the brute-force strategy, however, is that facts of the form [ $r$  is a reason of weight  $d$  to  $\varphi$ ] and [ $r$  is a reason that supports  $\varphi$ -ing to degree  $d$ ] intuitively involve two distinct, separable ingredients: (i) the degree to which  $\varphi$ -ing is normatively supported (or: the amount of the normative 'weight' that the reason has), and (ii) the fact,  $r$ , that *provides* this support. This suggests that such facts are *not* normatively fundamental, but can instead be further analyzed in terms of these ingredients – just as the S.O.S. view predicts. For on the S.O.S. view, for  $r$  to be a reason that supports  $\varphi$ -ing to a certain degree *just is* for  $r$  to provide a corresponding degree of support for  $\varphi$ -ing. The amount of normative support provided by a fact is not a relatum in a normatively fundamental reasons-relation, but instead a more fundamental normative feature in terms of which the reasons-relation should be understood. (Consider an analogy: if  $X$  financially supports  $Y$  to degree  $d$ , we can distinguish the amount of financial support provided ( $d$ ) from the person that provides it ( $X$ ), and we analyze the complex fact in terms of these distinct ingredients.)

Thus, reflection on the brute-force view actually points away from Reasons-First, and in the direction of the S.O.S. view instead.

Another reason for preferring the S.O.S view over the brute force view is that unlike facts about the weight of a reason, there can be facts about kinds or amounts of normative support that make no essential reference to reasons, and it is these facts that (at least usually) determine what you ought to do, or what attitudes you ought to have. To see why, consider the widely accepted slogan that you ought to do what you have most reason to do. Given what we said above about the relationship between the mass nouns ‘reason’ and ‘support’ (§3.2), the near-truism can be rephrased as follows: you ought to do what there is most normative support for doing – not what there is most weight to do (whatever that means). Facts about support thus provide the most proximal explanation of (at least many) ought-facts. Although reasons also play a role – namely, as sources of support – it’s less proximal, and the metaphysical relationship between reasons and oughts is mediated by facts about support. (An analogy might help: lights – e.g. streetlights, flashlights, etc. – are sources of light. The most proximal explanation of facts about visual experiences involves facts about light, not lights – lights do play a role, but it’s less proximal, and the causal relationship between lights and visual experiences is mediated by facts about light.)

#### **4.4 Summary**

Schroeder’s, Scanlon’s, and Skorupski’s proposals illustrate the difficulty of accounting for the weight of reasons in the Reasons-First framework. Without such an account, it is difficult to accommodate various putatively platitudinous sources of support for Reasons-First, such as Schroeder’s ‘classical argument’ (cf. §1) and the slogan that ‘you ought to do what you have most reason to do’. These are considerations that the S.O.S. view accommodates with ease: the ‘competing forces’ that determine what we ought to do can be understood as competing amounts (and/or kinds) of normative support, and what you have most reason to do is what there’s most support for doing.

### **5. Competing analyses of reasons and their weight**

We now turn to a critical discussion of what other analyses of reasons say – or in many cases: do not say – about their weight.

#### **5.1 Other explanationist analyses**

The S.O.S. view is a member of the more general family of *explanationist* views, on which reasons are analyzed as facts that explain other normative facts. In this (sub-)section, we’ll consider the three most prominent forms of explanationism and how they can analyze the weight of reasons: ought-explanationism (§5.1.1), fittingness-explanationism (§5.1.2), and good-explanationism (§5.1.3).

##### **5.1.1 Ought-explanationism**

The simplest form of explanationism focuses on facts about *ought*:

**Ought-explanationism<sub>simple</sub>** For  $r$  to be a reason for  $S$  to  $\varphi$  is for  $r$  to make it the case that  $S$  ought to  $\varphi$ .

Simple ought-explanationism, however, is *too* simple. Given that the relevant *makes it the case that*-relation is factive, the view implies that there are no reasons for actions and attitudes that ought not to be done or had. And that's implausible – even if you ultimately ought to donate your money, for instance, you may well have *some* reason to keep it. A related problem is that the view has nothing to say about how reasons can be outweighed (as it does not even allow for the existence of outweighed reasons), or more generally, about how some reasons can be stronger than others.

John Broome (2013) defends a variant of ought-explanationism that is intended to allow for outweighed reasons. He appeals to a distinction between two kinds of reasons: *pro toto* and *pro tanto*. According to Broome, simple ought-explanationism is true of *pro toto* reasons, but not of *pro tanto* reasons: a *pro toto* reason for  $S$  to  $\varphi$  is a fact that explains why  $S$  ought to  $\varphi$ , whereas a *pro tanto* reason for  $S$  to  $\varphi$  is a fact that plays a certain role in a 'weighing explanation' of why  $S$  ought (or ought not) to  $\varphi$ . A weighing explanation of an ought-fact is one that appeals to the weight of reasons:  $S$  ought to  $\varphi$  if (and because) the reasons for  $\varphi$ -ing outweigh the reasons against  $\varphi$ -ing, and  $S$  ought not to  $\varphi$  if (and because) the reasons against  $\varphi$ -ing outweigh the reasons for  $\varphi$ -ing. (The view thus allows that there can be reasons for actions and attitudes that ought not to be done or had.) Not all explanations of ought-facts are weighing explanations, according to Broome, but many are, and *pro tanto* reasons are defined by the role they play in such: for  $r$  to be a *pro tanto* reason for  $S$  to  $\varphi$  is for  $r$  to play the 'for- $\varphi$ ' role in a weighing explanation, and for  $r$  to be a *pro tanto* reason for  $S$  not to  $\varphi$  is for  $r$  to play the 'against- $\varphi$ ' role in such an explanation. Both *pro toto* and *pro tanto* reasons help explain facts about what you ought to do, though they do so in different ways. Thus:

**Ought-explanationism<sub>Broome</sub>** For  $r$  to be a reason for  $S$  to  $\varphi$  is for  $r$  to either  
(i) explain why  $S$  ought to  $\varphi$ , or  
(ii) play the for- $\varphi$  role in a weighing explanation of why  $S$  ought (or ought not) to  $\varphi$ .

Broome's view has been extensively criticized (e.g., Kearns & Star 2008; Brunero 2013; Dancy 2015; Hawthorne & Magidor 2018; Fogal & Risberg forthcoming). In this context, the most important problem is that while Broome's account assumes that reasons have 'weights' that determine what an agent ought to do, Broome says nothing about these weights or how they are determined, other than that they play this characteristic role (cf. Kearns & Star, 2008; Dancy, 2015). The view is thus silent on our central question – namely, what *it is* for a reason to have a certain weight. And it's hard to see how the view could be modified so as to answer this question without violating the spirit of Broome's ought-centric approach. Here again we see the difficulty of analyzing the *gradable* notion of weight in terms of some *non-gradable* notion (in this case, that of ought).

### 5.1.2 Fittingness-explanationism

An alternative explanationist approach appeals to the notion of *fittingness* (Chappell 2012; Howard 2019). Most straightforwardly:

**Fittingness-explanationism<sub>simple</sub>** For  $r$  to be a reason for  $S$  to  $\varphi$  is for  $r$  to make it the case that  $S$ 's  $\varphi$ -ing is fitting.

This view faces similar problems as ought-explanationism since, as standardly conceived, fittingness (like ought) is not gradable. In particular, fittingness is often identified with *correctness*, which is clearly a binary, non-gradable notion.<sup>23</sup> As Maguire (2018) puts it:

Suppose you asked two children for the product of 7 and 5. One says 42. The other says 33. Which answer is more correct? I think there is an error in the question. Any answer other than 35 is incorrect. The two children's answers are both incorrect. (2018: 791)

The same, Maguire thinks, is true of fittingness: “[a]n attitude is either fitting or unfitting; there are no degrees, no comparisons [...] We can talk about degrees of fittingness, to be sure. But when we are doing so we are subtly shifting to assessment by some other pertinent standard, usually a gradable notion” (2018: 791; for similar claims, see McHugh and Way 2022, ch. 3). In the math case, for instance, one of the answers may be *closer* to the correct one, and hence in an extended, derivative sense count as ‘more correct’ than the other. But strictly speaking, both answers are incorrect, and hence unfitting.<sup>24</sup>

Given that fittingness is not gradable, fittingness-explanationism implies that there are no reasons for unfitting actions and attitudes. Again, this is implausible: even if the only fitting intention is to intend to take a break from work, for instance, one might still have some reason to intend to keep working (e.g., that one needs to make a deadline).

What's more, problems remain for fittingness-explanationism even if one invokes a non-standard, gradable notion of fittingness.<sup>25</sup> That's because on many views, ought and

---

<sup>23</sup> Following Brentano, Danielsson and Olson (2007: 516) write that “[c]orrectness is a technical philosophical notion and the point of it is that it is, as Brentano says, analogous to the notion of truth”. The idea is the relationship between truth and belief mirrors the one between correctness and conative attitudes (cf. McHugh 2014). Hence, given that truth is not gradable, correctness plausibly isn't either.

<sup>24</sup> Compare: one woman may be said to be ‘more pregnant’ than another in the sense that the former is further along in her pregnancy than the latter, but whether someone is pregnant is binary, all-or-nothing matter – at least modulo vagueness.

<sup>25</sup> Howard (2019: 233, n. 30) assumes that “a response can be fitting to some degree”, and says that fittingness-explanationism thus “accounts for the existence of outweighed reasons”, though without how explaining how this works or how it can be reconciled with the underlying motivations for fittingness-based views (e.g., that fittingness is correctness). Berker (2023) argues that “thin fittingness properties”, such as fittingness, appropriateness, and aptness, are gradable, on the grounds that this is required for fitting attitude-analyses of “thick fittingness properties”, such as admirability and desirability, to work. In particular, contrary to Maguire (2018) and McHugh and Way (2022), Berker thinks that *being more admirable than* should be analyzed as *being*

fittingness can come apart: sometimes you ought to  $\varphi$  even though  $\varphi$  is unfitting, and vice versa. This makes it hard for the view to account for the relationship between the weights of reasons and oughts. This problem arises especially clearly for Chris Howard’s (2019) disjunctive version of the view, which is the best-developed version of fittingness-explanationism to date:

**Fittingness-explanationism<sub>disjunctive</sub>** For  $r$  to be a reason for  $S$  to  $\varphi$  is for  $r$  to either

- (i) make it the case that  $S$ ’s  $\varphi$ -ing is fitting, or
- (ii) make it the case that  $S$ ’s *wanting to  $\varphi$*  is fitting.

This view is designed to accommodate the existence of genuine *pragmatic* (or what Howard calls *value-related*) reasons for attitudes – ones which are unrelated to whether the attitude fits its object. For instance, Howard claims that “the fact that a deplorable dictator will order your execution unless you admire him is a fact that counts in favor of your admiring the dictator, since it’s a fact that explains why your admiring him would be good” (2019: 218), but denies that the admiring the dictator would be fitting, since the dictator isn’t *worthy of admiration*. The disjunctive view accommodates this verdict because although admiring the dictator isn’t fitting, it’s fitting to *want to admire* him, and the fact that he will order your execution if you don’t admire him explains why that’s fitting.

According to Howard, it is “an attractive normative hypothesis” that we normally ought to have fitting attitudes. However, he also thinks that when having an unfitting attitude is required to bring about a great good “it seems that, all-things-considered, you ought to have that attitude, despite its being unfitting” (2019: 218–219). Presumably, this is because in such circumstances the reasons favoring the unfitting attitude are *stronger* than the reasons against it. The problem is that it’s difficult to see how the disjunctive view gets this result. Why is it that reasons to  $\varphi$  that explain the fittingness of *wanting to  $\varphi$*  sometimes, but not always, outweigh reasons against  $\varphi$ -ing that explain the unfittingness of  $\varphi$ -ing? Even granting a gradable notion of fittingness, a plausible answer doesn’t seem forthcoming. For the strength of a reason to  $\varphi$  can’t depend solely on the degree of fittingness (whether it be of  $\varphi$ -ing or of  $S$ ’s wanting to  $\varphi$ ) that is explained by  $r$ , and so what attitudes you ought to have can’t depend solely on the degree to which they (and/or wanting to have them) are (un)fitting. Despising the dictator, for instance, may well be maximally fitting and yet something you ought not do, given the terrible consequences it would have.<sup>26</sup>

---

*such that admiration is more fitting* rather than *being such that more admiration is fitting*. He does not present an argument for the view that thin fittingness properties are gradable that is independent of the fitting attitude analysis, however, and he recognizes that this view is not supported (and possibly even disconfirmed) by linguistic evidence.

<sup>26</sup> In another paper, Howard (2020) presents a model for how epistemic and pragmatic reasons for belief interact to determine what one ought to believe. The model is independent of fittingness-

Finally, both versions of fittingness-explanationism face problems concerning reasons for belief.<sup>27</sup> It's very plausible that we sometimes have reasons to believe that  $p$  even though  $p$  is false. But it's widely held, at least among fans of fittingness, that a belief is fitting only if it is true. Given this assumption, simple fittingness-explanationism leads to an extreme form of infallibilism, as it unacceptably implies that there are no reasons to believe any false proposition. And disjunctive fittingness-explanationism fares only slightly better: it has infallibilist implications in all cases except when it is fitting to want to believe the relevant false proposition. It seems, then, that there are many more reasons for false beliefs than what either version of fittingness-explanationism allows.

In sum, there are (at least) three features of fittingness that make it unsuitable in an explanationist analysis of reasons and their weight: as standardly understood, it is not gradable; it is not connected closely enough to what one ought to do; and it risks generating infallibilism about reasons for belief. The S.O.S. view, by contrast, faces none of these problems.

### 5.1.3 Good-explanationism

The last version of explanationism we'll consider focuses on facts about *goodness* (see especially Finlay 2014 and 2019, as well as Gardner & Macklem 2004 and Maguire 2016):

**Good-explanationism** For  $r$  to be a reason for  $S$  to  $\varphi$  is for  $r$  to explain why  $S$ 's  $\varphi$ -ing would be good (in some way, to some degree).

Since goodness is gradable, good-explanationism can provide a straightforward account of the weight of reasons: the weight of  $r$  is proportional to the degree of  $\varphi$ 's goodness that is explained by  $r$ . This account faces a problem, however: there doesn't seem to be a notion of goodness that can be 'plugged in' to good-explanationism to make the view plausible.

To begin with, it's implausible that  $r$  is a reason to  $\varphi$  if  $r$  merely explains why  $\varphi$ -ing would be good *in any sense of 'good' whatsoever*. This generates entirely too many reasons —  $\varphi$ -ing might be good for the purposes of doing something stupid, for instance, but a fact that explains why this is so isn't thereby a reason to  $\varphi$ . On the other hand, if we appeal to a more restricted notion of goodness, such as intrinsic value, the view instead risks generating too few reasons.<sup>28</sup> For example, we often have prudential reasons to perform actions (e.g., going to the store or tying our shoes) that are intrinsically neutral but extrinsically good for us. This conflicts with versions of good-explanationism that focuses

---

explanationism, however, and he doesn't discuss whether (and if so how) it can be combined with an analysis of weight that is congenial with fittingness-explanationism.

<sup>27</sup> For similar points, see Lee (2022).

<sup>28</sup> Maguire (2016) suggests an analysis of reasons for actions in terms of the promotion of intrinsic value. As he notes (2016: 239–240), this analysis is difficult to extend to reasons for attitudes, such as belief and intention. This makes it hard to see how it could provide an analysis of reasons as such. A similar dialectic arises concerning promotion-based views more generally: though they may have initial plausibility concerning reasons for action, they typically fail to generalize to reasons for attitudes, and especially reasons for belief (cf. Sharadin 2015: §3).

only on intrinsic value. Reasons for belief provide another illustration: we arguably have strong reasons to believe many highly uninteresting claims, concerning, e.g., the colors of things in our immediate surroundings, but our having those beliefs isn't correspondingly very intrinsically good.<sup>29</sup>

Thus, the challenge for good-explanationists is to find a notion of goodness that allows their view to make plausible predictions about what reasons we have, what their weights are, and how their weights can be modified. We don't have a conclusive argument that such a notion of goodness doesn't exist, but it seems unlikely.

## 5.2 *Reasons as evidence*

Some philosophers analyze reasons as evidence (RAE), and in particular evidence concerning what one ought to do (e.g. Kearns & Star 2008; Thomson 2008; for a related view, Whiting 2018):

**RAE** For  $r$  to be a normative reason for  $S$  to  $\varphi$  is for  $r$  to be evidence that  $S$  ought to  $\varphi$ . (Kearns & Star 2009: 216; notation adjusted)

RAE is a non-explanationist view since  $e$  can be evidence for  $p$  without even partially explaining why  $p$ . (In short, not all evidence explains.) On Kearns and Star's (2009) version of RAE, the weight of a reason  $r$  for  $S$  to  $\varphi$  is analyzed in terms of how strong evidence  $r$  is that  $S$  ought to  $\varphi$ , where strength of evidence is in turn understood probabilistically (i.e., how strong evidence  $e$  is for  $p$  is a matter of how much  $e$  raises the probability of  $p$ ). This allows the view to vindicate the platitudes about weight outlined in §2 – namely, Simple Weightier-Than and Simple Outweighing.

RAE has nonetheless been extensively criticized. For our purposes there are three main concerns. First, since not all evidence explains, RAE forfeits the Rossian insight, noted in §1, that what you ought to do is (at least normally) explained by reasons and their weight. Second, the probabilistic gloss of evidence and its strength is insensitive to the distinction between the *balance* and *amount* of support, or (alternatively put) the balance and 'weight' of reasons. As noted in §3.2, the evidential probability of  $p$  can be the same in two situations, even though in one situation there are greater amounts of evidence for and against  $p$ . But perhaps the most pressing problem for the view is that, as Worsnip (2016) argues, it conflates something first-order (what one has reasons to do) with something higher-order (what one evidence indicates that one ought to do). This is clearest in the epistemic case. Intuitively, one has evidence of a certain strength for  $p$  iff

---

<sup>29</sup> Finlay (2014: ch. 4.3) suggests, roughly, that reasons to believe that  $p$  can be analyzed as facts that explain why believing that  $p$  is good relative to the end of *believing that p if and only if p is true*. This account faces the challenge of explaining why this end is more normatively significant than other ends that may be contextually salient, such as that of *believing that p if and only if p is funny*. See further Fogal & Risberg (forthcoming: §3.2).

one has a corresponding amount of reason to believe that  $p$ .<sup>30</sup> Thus, on RAE, one has a certain amount of reason to believe that  $p$  iff one has evidence (of a corresponding strength) that one ought to believe that  $p$ . It follows that the strength of evidence for  $p$  must be the same as the strength of evidence that one ought to believe that  $p$ . And that seems wrong. At least, many views about higher-order evidence allow that one can have evidence (of some strength) that one ought to believe that  $p$  while also not having evidence (of that same strength) that  $p$ , and vice versa.<sup>31</sup>

This objection can be sharpened and generalized. As Worsnip (2016) notes, the following are plausible:

- (1) One ought to  $\varphi$  iff one has most reason to  $\varphi$ : there is no gap between these two possibilities.
- (2) It's not always the case that one ought to  $\varphi$  iff one's evidence (on-balance) supports believing that one ought to  $\varphi$ : there *is* a gap between these two possibilities. (For, one can have misleading evidence about whether one ought to  $\varphi$ .)

RAE entails that:

- (3) One has most reason to  $\varphi$  iff one's evidence (on-balance) supports (belief in) the proposition that one ought to  $\varphi$ .

(1)-(3) are jointly inconsistent. Proponents of RAE thus face a choice: they must either reject (1) or (2). If they reject (1), they are "opening a gap where there seems to be none", and if they reject (2), they are "closing a gap that appears genuine" (Worsnip 2016). Both are unhappy results.

This problem could be avoided if one could simply analyze a reason to believe that  $p$  in terms of evidence for  $p$  (rather than in terms of evidence that one ought to believe that  $p$ ). However, such an analysis doesn't generalize to reasons for action or for attitudes other than belief, and so would result in an undesirably disjunctive view of reasons-as-evidence-(in the case of belief)-or-evidence-that-one-ought-(in the case of action and other attitudes). Avoiding this outcome is why RAE is formulated in higher-order terms: reasons are evidence that one ought. The problem, as Worsnip emphasizes, is that it seems possible for reasons to  $\varphi$  and evidence that one ought to  $\varphi$  to come apart.

The support-based view suggests a natural diagnosis of what might make RAE seem initially plausible and also why it goes wrong. To begin with, note that 'evidence' has properties that straddle the traditional mass/count noun divide – it's a so-called "object" or "collective" mass noun, akin to 'jewelry', 'ammunition', and 'furniture'. Just as furniture includes individual tables and chairs, so evidence includes individual items

---

<sup>30</sup> We're here setting aside pragmatic reasons to believe, as well as reasons to believe claims with probability 1 (which on some conceptions there can't be evidence for). See Reisner (2009).

<sup>31</sup> See, e.g., Worsnip (2018).

(facts, states, objects, or whatever one thinks pieces of evidence are). We can thus use expressions like ‘piece of evidence’ to pick out, well, pieces of evidence, just as we use ‘piece of furniture’ to pick out pieces of furniture. Morphosyntactically, ‘evidence’ is a mass noun – it doesn’t take plural form (‘evidences’ is ungrammatical) nor combine with numeral modifiers (‘two evidence(s)’ is also ungrammatical) – but semantically it’s akin to a count noun since its denotation includes individual, countable things (i.e. pieces of evidence).

What makes RAE seem initially attractive, we think, is that ‘evidence’ – when used predicatively, as in RAE – and the count noun ‘reason(s)’ are closely related. Both are non-gradable: just like something can’t be more or less of a reason, something also can’t be more or less of a piece of evidence. Moreover, both ‘reason(s)’ and ‘evidence’ pick out “sources” of support, or (equivalently) things that explain why there’s (a certain amount of) support. In the case of ‘evidence’, the support is epistemic in nature, and the same is true of the epistemic use of ‘reason(s) to believe’. However, rather than either being defined in terms of the other, we think both are better functionally defined in terms of what they do: *provide epistemic support for beliefs*. The *strength* of evidence, just like the weight of reasons to believe, depends on how much support the relevant consideration provides. Support-facts then in turn explain ought-facts (which proponents of RAE tend to treat as primitive). On this view of reasons and evidence, unlike on RAE, there’s no risk of conflating what’s first-order with what’s higher-order.

### 5.3 *Reasons as premises of good reasoning*

Some philosophers understand reasons in terms of (good) reasoning. The exact relationship posited between reasons and reasoning varies, but for illustrative purposes we’ll focus on one prominent account – namely, the view of reasons as premises of good reasoning (PGR) defended by Conor McHugh and Jonathan Way (e.g. Way 2017; Way & McHugh 2022):

**PGR** For  $r$  to be a reason for  $S$  to  $\varphi$  is for there to be a good pattern of reasoning from the belief that  $r$ , perhaps together with other correct attitudes which  $S$  has, to  $\varphi$ -ing.<sup>32</sup> (Way 2017: 254; notation adjusted)

One way to understand reasoning-based views such as PGR is as capturing a distinct notion of a reason than the one we’ve been focusing on. Ralph Wedgwood (2015), for example, thinks we should distinguish between two notions of a normative reason: one associated with a “justificatory story” – roughly, “a story that explains the truth about which action or attitude one has, all things considered, most reason to do” – and one associated with “an ideal deliberative or motivational procedure” – roughly, “what

---

<sup>32</sup> Common examples of good patterns of reasoning include the modus ponens pattern (reasoning from a belief that  $p$  and a belief that *if  $p$  then  $q$*  to a belief that  $q$ ) and the means-end pattern (reasoning from an intention to  $\varphi$  and a belief that *in order to  $\varphi$ , you must  $\Psi$*  to an intention to  $\Psi$ ).

would be our motivating reasons if we were suitably well informed and rational” (127).<sup>33</sup> Wedgwood calls them *normative-explanation* reasons and *ideal-deliberation* reasons, and thinks each (in principle) earns its keep by playing distinct theoretical role(s). Accordingly, one could hold the S.O.S. view about normative-explanation reasons and an alternative view, such as PGR, for ideal-deliberation reasons.

Problems for reasoning-based views like PGR arise, however, if the view is put forward (as it often is) as something meant to capture all the relevant theoretical roles. Among other things, such a view will have to account for the defeasibility and weight of reasons, both of which are central to the normative-explanation notion of a reason. To date, most attention has been devoted to defeasibility – if patterns of good reasoning are only defeasibly good, then PGR may provide an account of the defeasibility of reasons.<sup>34</sup> But an account of defeasibility doesn’t yet provide an account of the gradable notion of weight, and thus fails to provide other things we should want from an account of (the weight of) reasons, including a story about interpersonal comparisons and an explanation of why some reasons are stronger or weightier than others.

Given PGR, one natural place to look for gradability is the goodness of different patterns of reasoning. Perhaps some patterns of good reasoning are better than others? The main proponents of PGR cannot avail themselves of this strategy, however, since they tend to operate with a binary notion of good reasoning as *fittingness-preserving* reasoning (see, e.g., McHugh & Way 2022: ch. 2.3). Given the common assumption that fittingness is not gradable (cf. §5.1.2), this means that the goodness of patterns of reasoning is not gradable either.

Conor McHugh and Jonathan Way (2022) – henceforth ‘McWay’ – recognize this problem. Rather than giving an account of the monadic property of a reason’s weight – a property they say is “obscure” (forthcoming: ch. 5) – they provide an account of the binary relation of *outweighing* in terms of the notion of good reasoning, and tentatively consider a generalization of that idea to the relation of *weightier-than* among reasons.<sup>35</sup> Here’s a simplified version of that generalized view:

**Weight as Defeat** For reason  $r$  to be weightier than reason  $r^*$  is for it to be the case that:

- (i)  $r$  is a reason in virtue of a good pattern of reasoning  $A$ ,
- (ii)  $r^*$  is a reason in virtue of a good pattern of reasoning  $B$ ,
- and
- (iii)  $A$  defeats  $B$ , in the sense that the pattern of reasoning from the premise-responses of both  $A$  and  $B$ , together with beliefs which make their conclusion-responses conflict, to the

---

<sup>33</sup> Schmidt (2021) makes a similar suggestion, and outlines various challenges for PGR.

<sup>34</sup> Cf. Way (2017), among others.

<sup>35</sup> Way (2017, §5) also outlines a related PGR-based analysis of outweighing, though without generalizing it to the *weightier than*-relation.

conclusion-response of *A*, is good, but the pattern of reasoning from these premise-responses to the conclusion-response of *B* is not good (where a pattern of reasoning is good if and only if it is fittingness-preserving).

Given space constraints, we're unable to properly evaluate McWay's proposal. Instead, we'll restrict ourselves to two points. The first is that, as McWay acknowledge, Weight as Defeat is incompatible with apparent platitudes such as Simple Weightier-Than and Simple Outweighing. We take this to speak against the former, given the plausibility of the latter. The second is that Weight as Defeat illustrates the difficulty of trying to capture something fundamentally gradable (namely, weight) in wholly non-gradable terms – the explanation and defense of Weight as Defeat is complicated and, as McWay recognize, it faces a variety of challenges (cf. McHugh & Way 2022: ch. 5).

## 6. Concluding Remarks

We have defended an analysis of the weight of reasons that directly falls out of the support-based analysis of reasons. On this view, to be a reason is to provide normative support, and the weight of a reason simply corresponds to the amount of support it provides. We have also argued against several competing views about reasons and their weight – both ones on which reasons are 'first' and ones one which they are not.

Even for those who remain unconvinced of the support-based view, however, we at least hope to have shown how important the issue of weight is when theorizing about reasons. Though there are notable exceptions, this issue has received significantly less attention than it deserves: many metanormative theorists either ignore it completely or else bring it up only 'last minute', once their position has already been developed and motivated without taking it into consideration. This has generated several accounts of reasons that have nothing plausible to say about their weight. Thus, to make progress in this area, we propose that accounting for the weight of reasons should be treated as an adequacy constraint on any analysis of reasons.

## References

- Bader, R. 2016. 'Conditions, Modifiers and Holism.' In E. Lord & B. Maguire (Eds.), *Weighing Reasons* (27–55). Oxford University Press.
- Berker, S. Forthcoming. 'The Deontic, the Evaluative, and the Fitting'. In C. Howard & R. A. Rowland (Eds.), *Fittingness*. Oxford University Press.
- Broome, J. 2013. *Rationality Through Reasoning*. Wiley-Blackwell.
- Brunero, J. 2013. 'Reasons as explanations.' *Philosophical Studies* 165: 805–824.
- Bykvist, K. 2021. 'Taking values seriously.' *Synthese* 199: 6331–6356.
- Carlson, E. 2008. 'Extensive Measurement with Incomparability.' *Journal of Mathematical Psychology* 52: 250–259.
- Chappell, R. 2012. 'Fittingness: the sole normative primitive.' *Philosophical Quarterly* 62, 684–704.
- Dancy, J. 2004. *Ethics Without Principles*. Oxford: Oxford University Press.
- 2015. 'Reasons for Broome.' In I. Hirose & A. Reisner (Eds.), *Weighing and Reasoning* (177–188). Oxford: Oxford University Press.
- Danielsson, S., and Olson, J. 2007. 'Brentano and the Buck-Passers.' *Mind* 116: 511–522.
- Enoch, D. 2011. 'On Mark Schroeder's Hypotheticalism: A Critical Notice of *Slaves of the Passions*.' *The Philosophical Review* 120: 423–446.
- Finlay, S. 2014. *Confusion of Tongues*. Oxford University Press.
- 2019. 'A "good" explanation of five puzzles about reasons.' *Philosophical Perspectives* 33: 62–104.
- Fogal, D. 2016. 'Reasons, reason, and context.' In E. Lord & B. Maguire (Eds.), *Weighing Reasons*: 74–103. <https://doi.org/10.1093/acprof:oso/9780199315192.001.0001>
- Fogal, D., and Risberg, O. 2020. 'The Metaphysics of Moral Explanations.' In Russ Shafer-Landau (ed.), *Oxford Studies in Metaethics, Volume 15*: 170–194. Oxford University Press.
- Forthcoming. 'Explaining Normative Reasons.' *Noûs*.
- Gardner, J. and Macklem, T. 2004. 'Reasons.' In J. Coleman, K. Himma, & S. Shapiro (Eds.), *The Oxford Handbook of Jurisprudence and Philosophy of Law* (440–475). Oxford University Press.
- Gregory, A. 2014. 'A Very Good Reason To Reject The Buck Passing Account.' *Australasian Journal of Philosophy* 92: 287–303.
- Howard, C. 2019. 'The Fundamentality of Fit.' In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics, Vol. 14*. Oxford: Oxford University Press.
- 2020. 'Weighing Epistemic and Practical Reasons for Belief.' *Philosophical Studies* 177: 2227–43.
- Hawthorne, J., and Magidor, O. 2018. 'Reflections on Reasons.' In D. Star (ed.), *The Oxford Handbook on Reasons and Normativity*.
- Joyce, J. 2005. 'How Probabilities Reflect Evidence.' *Philosophical Perspectives* 15: 153–178.
- Kearns, S., and Star, D. 2008. 'Reasons: Explanations or Evidence?' *Ethics* 119: 31–56.
- 2009. 'Reasons as Evidence.' In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics, Vol. 4*. Oxford: Oxford University Press.
- Keynes, J. M. 1921. *A Treatise on Probability*. London? Macmillan.
- Kim, J. 2016. 'What are quantities?' *Australasian Journal of Philosophy* 94: 792–807.

- Lee, W. 2022. 'Fittingness first? Reasons to withhold belief.' *Philosophical Studies* 179: 3565-3581.
- Lord, Errol. 2018. *The Importance of Being Rational*. New York: Oxford University Press.
- Makins, N. manuscript. 'The Balance and Weight of Reasons.'
- Maguire, B. 2016. 'The Value-Based Theory of Reasons.' *Ergo* 3(9).
- 2018. 'There Are No Reasons for Affective Attitudes.' *Mind* 127: 779-805.
- Maguire, B. and Snedegar, J. 2021. 'Normative metaphysics for accountants.' *Philosophical Studies* 178: 363-384.
- McHugh, C. 2014. 'Fitting Belief.' *Proceedings of the Aristotelian Society* 114: 167-187.
- McHugh, C., and Way, J. 2016. 'Fittingness First.' *Ethics* 126: 575-606.
- 2022. *Getting Things Right*. Oxford University Press.
- Nebel, J. Forthcoming. 'Ethics without Numbers.' *Philosophy and Phenomenological Research*.
- Nolt, J. 2022. *Incomparable Values*. Routledge.
- Reisner, A. 2009. 'The possibility of pragmatic reasons for belief and the wrong kind of reasons problem.' *Philosophical Studies* 145: 257-272.
- Risberg, O. 2016. 'Weighting Surprise Parties: Some Problems for Schroeder.' *Utilitas* 28: 101-107.
- Rowland, R. 2019. *The Normative and the Evaluative: The Buck-Passing Account of Value*. Oxford: Oxford University Press.
- Scanlon, T. M. 1998. *What We Owe to Each Other*. Cambridge: Harvard University Press.
- 2014. *Being Realistic about Reasons*. Oxford: Oxford University Press.
- Schmidt, E. 2021. 'Where Reasons and Reasoning Come Apart.' *Noûs* 55: 762-781.
- Schroeder, M. 2007. *Slaves of the Passions*. Oxford: Oxford University Press.
- 2021a. 'The Fundamental Reason for Reasons Fundamentalism.' *Philosophical Studies* 178: 3107-3127.
- 2021b. *Reasons First*. New York: Oxford University Press.
- Sharadin, N. 2015. 'Reasons and Promotion.' *Philosophical Issues* 25: 98-122.
- Skorupski, J. 2010. *The Domain of Reasons*. Oxford: Oxford University Press.
- Thomson, J. J. 2008. *Normativity*. Open Court.
- Tucker, C.. 2022. 'The Dual Scale Model of Weighing Reasons.' *Nous* 56: 366-92.
- Manuscript. *The Weight of Reasons*.
- Way, J. 2017. 'Reasons as Premises of Good Reasoning.' *Pacific Philosophical Quarterly* 98: 251-270.
- Wedgwood, R. 2015. 'The Pitfalls of Reasons.' *Philosophical Issues* 25: 123-143.
- Whiting, D. 2018. 'Right in some respects: reasons as evidence.' *Philosophical Studies* 175: 2191-2208.
- Wodak, D. 2020. 'Who's On First?' In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics*, Vol. 15. New York: Oxford University Press.
- Worsnip, A. 2016. 'Review of *Knowing Better: Virtue, Deliberation, and Normative Ethics* by Daniel Star.' *Notre Dame Philosophical Reviews*.
- 2018. 'The Conflict of Evidence and Coherence.' *Philosophy and Phenomenological Research* 96: 3-44.