

Explaining Normative Reasons

Daniel Fogal (NYU)

Olle Risberg (Uppsala University)

Penultimate version - please cite published version where possible

Abstract. In this paper, we present and defend a natural yet novel analysis of normative reasons. According to what we call *support-explanationism*, for a fact to be a normative reason to φ is for it to explain why there's normative support for φ -ing. We critically consider the two main rival forms of explanationism—*ought-explanationism*, on which reasons explain facts about ought, and *good-explanationism*, on which reasons explain facts about goodness—as well as the popular Reasons-First view, which takes the notion of a normative reason to be normatively fundamental. Support-explanationism, we argue, enjoys many of the virtues of these views while avoiding their drawbacks. We conclude by exploring several further important implications: among other things, we argue that the influential metaphor of 'weighing' reasons is inapt, and propose a better one; that, contrary to what Berker (2019) suggests, there's no reason for non-naturalists about normativity to accept the Reasons-First view; and that, contrary to what Wodak (2020b) suggests, explanationist views can successfully accommodate what he calls 'redundant reasons'.

1. Introduction

What do reasons do? Despite the high generality of this question, it has an answer that's both popular and plausible: reasons *explain*.

For many kinds of reasons, the idea that reasons explain seems close to undeniable. For example, when we say that the explosion is the reason (or reason why) the bridge collapsed, or that the pandemic is the reason the economy collapsed, this seems equivalent to saying that the explosion explains why the bridge collapsed, and that the pandemic explains why the economy collapsed. It's no wonder, then, that such reasons are standardly called *explanatory* reasons. Something similar is true of so-called *motivating* reasons—i.e., the kinds of things we invoke when we say that Talia's reason for becoming a doctor is to help people, or that Herman's reason for believing it rained is that the streets are wet. While the details are controversial, motivating reasons are standardly taken to be considerations which explain, in a distinctive way, why the agent in question did what she did or has the attitude that she has.

In addition to explanatory and motivating reasons, however, there are also *normative* reasons. These are the sorts of reasons we invoke when we say that the fact that the street is wet is a reason to believe it has rained, or the fact that donating to charity helps alleviate suffering is a reason to do it. Can the idea that reasons explain be extended to the normative domain? Though less obvious, such an extension is appealing. Besides the relatively superficial point that we use the count noun ‘reason(s)’ to express claims about both, the connection between normative reasons and explanation is motivated by reflecting on cases and, more generally, by the practice of normative ethics and other normative disciplines. To illustrate, suppose that:

- (1) Polluting the environment is morally wrong because it causes unnecessary harm.

Given (1), it is natural to accept both of the following:

- (2) The fact that polluting causes unnecessary harm is a (or: the) reason not to pollute.
- (3) The fact that polluting causes unnecessary harm is a (or: the) reason why pollution is morally wrong.

In this example, (1) says that a certain fact about polluting explains its moral status, (2) says that the same fact about polluting is a moral reason against it, and (3) says that the same fact about polluting is an explanatory reason in relation to its moral status. These claims are clearly closely related – so much so that normative theorists often go back and forth between them, seemingly without hesitation.

Our aim in this paper is to vindicate and illuminate the close connection between normative reasons and explanation—a task we take to be a constraint on any adequate account of normative reasons.¹ In particular, we defend a version of (what we call) *explanationism* about normative reasons. On explanationist views, facts about normative reasons are analyzed as facts about what explains what, and what makes normative reasons *normative* (as opposed to merely explanatory) is the normativity of what is explained. (Henceforth we’ll refer to normative reasons simply as ‘reasons’.) To date, two main forms of explanationism have been defended: *ought-explanationism*, where reasons explain facts about what we ought to do (Broome, 2013), and *good-*

¹ Hence, while there are views of normative reasons that sever the seemingly-intimate connection to explanation, we think such views are (other things being equal) at a disadvantage, and engaging with them is beyond the scope of this paper. One example is the view of reasons as premises of good reasoning (see e.g. Hieronymi, 2005; Setiya, 2014; Way, 2017). Another is the view of reasons as evidence (see e.g. Kearns & Star, 2008, 2009; Thomson, 2008; and, for a related view, Whiting, 2018).

explanationism, where reasons explain facts about what is good (Finlay, 2014).² While these views accommodate the idea that reasons explain, they also face important objections. We'll defend a novel alternative: *support-explanationism*, according to which reasons explain facts about normative support. This view, we'll argue, retains the strengths of its explanationist competitors while avoiding their weaknesses.

The distinction between reasons and support is subtle but substantial: it's an instance of the intuitively clear—and metaphysically real—distinction between a source or provider of something (in this case, being a reason) and that which it is a source or provider of (in this case, normative support). A recurring theme of this paper is that subtleties such as these matter, and paying attention to them pays off. In particular, doing so reveals the superiority of support-explanationism over another influential view of reasons: the one commonly called *Reasons-First* (see, e.g., Scanlon, 1998, 2014; Raz, 1999; Skorupski, 2002, 2010; Schroeder, 2007, 2021, forthcoming; Lord, 2018; Rowland, 2019). According to Reasons-First, reasons-facts—i.e., facts of the form [*r* is a reason for *S* to ϕ]*—are normatively fundamental, and all other normative facts can ultimately be explained in terms of them. As we'll argue, many considerations that have been taken to favor Reasons-First—such as the slogan that 'you ought to do what you have most reason to do'—are more straightforwardly and naturally accommodated by support-explanationism. Indeed, once we pay closer attention to the details concerning reasons, explanation, and the relation between them, it becomes clear that they actually speak against the idea that reasons-facts are normatively first, or otherwise fundamental. The plausibility of Reasons-First has thus been largely illusory, and is mainly a result of the view not being sufficiently distinguished from nearby but importantly different ones, such as ours.*

The paper is structured as follows. After some methodological preliminaries in §2, we argue against competing explanationist views and diagnose their central flaws in §3. This sets the stage for our articulation and defense of support-explanationism in §4, where we contrast it favorably with its explanationist competitors as well as Reasons-First. We then consider a variety of implications of our view: among other things, we'll argue that the influential metaphor of 'weighing' reasons is inapt, and propose a better one (§4.3); that, contrary to what Selim Berker (2019) suggests, there's

² A related view analyzes practical normative reasons in terms of explanation and *rightness* (cf. Schroeter & Schroeter, 2009). A more recent explanationist account seeks to analyze normative reasons in terms of explanation and *fittingness* (cf. Howard, 2019). We discuss the latter view, and a dilemma it faces, in footnote 9. Another family of views analyzes normative reasons in terms of explanation and *promotion*. The idea (roughly put) is that *r* is a reason to ϕ iff *r* explains why ϕ -ing would help promote some relevant objective (such as value or the satisfaction of the agent's preferences). While we lack the space to discuss these views here, we believe they're susceptible to problems we'll present for good-explanationism in §3.2. For further critical discussion of promotionalism, see Bedke (2008) and Behrens and DiPaolo (2011), among others.

no reason for non-naturalists about normativity to accept the Reasons-First program (§5.1); and that, contrary to what Daniel Wodak (2020b) suggests, explanationist views can successfully accommodate what he calls *redundant reasons* (§5.2).

2. Methodological Preliminaries: Reasons and Explanation

Before elaborating on support-explanationism and critically considering alternatives, we need to be clear about the target: normative reasons. There are various locutions we ordinarily use to talk about normative reasons, such as:

- (NR₁) *r* is a reason (for *S*/that *S* has) to φ .
'The fact that it's cold is a reason (for Sasha/ that Sasha has) to stay inside.'
- (NR₂) *S* has a reason to φ .
'Sasha has a reason to stay inside.'
- (NR₃) There is a reason (for *S*) to φ .
'There is a reason (for Sasha) to stay inside.'

Part of what distinguishes normative reasons is obvious: unlike other kinds of reasons, normative reasons are *normative*. We won't pretend to offer an informative analysis of what it is to be normative, though in §4 we will offer an informative analysis of what it is to be a normative *reason*.

Another way to distinguish normative reasons from other kinds of reasons is by focusing on the theoretical roles they play. Various heuristics for identifying normative reasons have been proposed (cf. Schroeder, 2007; Lord, 2018; Maguire & Snedegar, 2021). For example, normative reasons have been variously described as things which:

- (R1) *count in favor* of some agent's action or attitude φ .
- (R2) explain why an agent *ought* or is *permitted* to φ .
- (R3) are suitable as *premises in good reasoning* about whether to φ .
- (R4) could be offered as a *justification* by an agent for φ -ing.
- (R5) could be offered as *advice* to an agent in support of φ -ing.

We'll focus on (R1) and (R2), since we take them to be the most theoretically central roles normative reasons play.³

We can thus identify normative reasons by the locutions we use to talk about them as well as various theoretical roles that they are assumed to play. This is how metanormative debates often proceed. What's less often noted, however, is that there's a tension at the heart of this methodology. In particular, as has been argued

³ It's of course possible to use 'reason(s)' to refer to something that, say, only plays the (R5)-role, but that would be a largely stipulative use of the term, not a substantive disagreement.

elsewhere (see, e.g., Fogal 2016a, 2018; Maguire & Snedegar 2021), there's a tension between (a) assuming that normative reasons play the theoretical roles just mentioned and (b) assigning significant evidential weight to the ordinary normative judgments we're inclined to make using 'reason(s)'. To see why, suppose Candice looks through a window and sees a room filled with smoke. In a discussion of what reasons she has to believe there's a fire, you might say either of the following:

- (i) The fact that smoke is a sign of fire is a reason for Candice to believe there's a fire.
- (ii) That fact that she sees smoke is a reason for Candice to believe there's a fire.
- (iii) The fact that there's smoke in the room is a reason for Candice to believe there's a fire.

But it would be unacceptable to say any of them in combination, such as:

- (iv) The fact that she sees smoke is a reason for Candice to believe there's a fire, and so is the fact that there is smoke in the room.
- (v) Candice sees smoke, there is smoke in the room, and smoke is a sign of fire, so she has three reasons to believe there's a fire.

In such a context, sentences like (iv) and (v) are guilty of double- and triple-counting—the facts cited as reasons don't separately count in favor of the relevant belief.⁴ This suggests that ordinary 'reason(s)'-judgments often fail to *directly* track the things that play the relevant theoretical roles—like (R1)—that reasons are supposed to play. Such judgments also arguably fail to *reliably* track those roles. For example, many philosophers deny that the fact that the subject can φ is itself a reason—or part of a reason—for her to φ , and instead insist that such facts play some other role (such as an 'enabling' role). Nonetheless, there are many contexts in which it's perfectly acceptable to cite the fact that one can φ as a reason to φ . These include contexts in which φ 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 φ —e.g. 'The fact that you're able to ask the President a question right now is a reason to do so.'

The methodological tension just noted can be resolved by distinguishing between *ordinary* uses of 'reason(s)'—uses of the term in ordinary thought and talk—and more *theoretical* (and hence semi-technical) uses of 'reason(s)'—uses of the term

⁴ For other examples that illustrate this point, see Fogal (2016a; 2018) and Maguire and Snedegar (2021). It's common to distinguish between *basic* reasons and *derivative* reasons. Basic reasons are the primary sources of normative support, whereas derivative reasons are "only reasons in virtue of their relationships to the [basic] ones" (Maguire & Snedegar, 2021: 366). Thus, even if some or all of the facts cited in (iv) and (v) count as derivative reasons, they don't count as basic reasons. We focus on basic reasons in this paper.

by theorists meant to track facts that play various theoretical roles. Many philosophers want to deny that the fact that you can φ is *itself* something that counts in favor of φ -ing (recall (R1)) or makes it the case that you ought, or are permitted, to φ (recall (R2)). If they're right, then that fact isn't a reason in the relevant theoretical sense. Nonetheless, there are plenty of contexts in which that fact is perfectly fine to cite as a reason, and hence counts as such in the ordinary sense.

One way to distinguish between ordinary and more theoretical uses of 'reason(s)' is by considering the kind of explanation the relevant facts provide. As noted above, one of the central theoretical roles normative reasons play is explaining other normative facts, such as facts about what subjects ought or are permitted to do. Normative theorists aren't typically interested in just any kind of explanation of such facts, however. When asking why some subject S ought to φ , for example, normative theorists typically want to know what *makes it the case* that S ought to φ , or what it is *in virtue of which* S ought to φ . Like many others, we take this to be an instance of a non-causal, non-pragmatic kind of explanation that commonly goes under the heading of 'metaphysical' or 'grounding' explanation.⁵ This is different from the looser, everyday notion of explanation, which we'll call 'pragmatic' explanation. This is the sort of thing that we standardly ask for and provide concerning a variety of subject matters, and whose success depends on "facts about the interests, beliefs or other features of the psychology of those providing or receiving the explanation [or] the 'context' in which the explanation occurs" (Woodward & Ross, 2021: §6.1). In slogan form, we should distinguish between 'explains' in the sense of *makes it the case* (i.e., non-pragmatic explanation) and 'explains' in the sense of *makes sense of why it's the case* (i.e., pragmatic explanation). Although knowing what makes something the case will often help make sense of why it's the case, one shouldn't assume that pragmatic and non-pragmatic explanations are neatly aligned. Something might make something else the case without making sense of why it's the case—e.g. we might lack relevant background information—and something might make sense of why something else is the case without making it the case—e.g. we're often content with information that merely gestures at the real explanatory story.

The two distinctions just made—between ordinary and theoretical uses of 'reasons' and between pragmatic and non-pragmatic explanations—are closely related. Whereas reasons in the *ordinary* sense provide pragmatic explanations—they help *make sense of* why things are the way they are—reasons in the *theoretical* sense provide non-pragmatic explanations—they help *make it the case* that things are the way they are. The theoretical notion of a reason is thus related to, but nonetheless

⁵ This paragraph is indebted to Fogal and Risberg (2020). For more on the metaphysics of normative explanations, see Rosen (2017) and Berker (2018). For an explanation of why expressivists (who otherwise tend to eschew metaphysics) also need an account of normative explanations, see Berker (2020).

distinct from, the ordinary notion, and the former can be seen as a precisification of the latter.

We say all this to be clear: our focus in what follows is on normative reasons in the relevant theoretical sense. However, metanormative theorists are often less-than-transparent with respect to their methodology, and don't always draw the same distinctions we think should be drawn, which can make exegesis a fraught affair. It's nonetheless plain that metanormative theorists are typically interested in more-than-merely-pragmatic explanations.⁶ In particular, when it comes to (R2) above, the explanations that reasons provide are not supposed to depend on anyone's interests or information: normative reasons are supposed to help *make it the case that* someone ought or is permitted to do something, not merely help make sense of why it's the case. This is what we're calling a theoretical use of '(normative) reason(s)'.

Two last preliminaries are in order. First, in what follows we'll consider a number of different explanationist views concerning the nature of normative reasons, all of which deny that the reasons-relation is normatively fundamental. Each view takes the following canonical form: for r to be a reason for S to φ is for r to _____, where r is a fact or true proposition, S is a subject, and φ is an action, attitude, feeling, or whatever else there can be reasons to do or have. Each view appeals to a different normative notion in filling in the blank, specifying non-trivial necessary and sufficient conditions that together state what it is for something to be a normative reason, and thereby providing an informative account or analysis of the nature of reasons—one that entails the reasons-relation is not normatively fundamental. To deny that such an analysis can be provided is to accept *Reasons Fundamentalism*, which is one of two central tenets of the Reasons-First view.⁷ The other is *Reasons Imperialism*, which claims that every *other* normative property and relation can be analyzed in terms of the reasons-relation. The Reasons-First view is thus the conjunction of Reasons Fundamentalism and Reasons Imperialism.⁸ Lastly, although we will sometimes use the fudge term 'notion' to cover both concepts and properties, we will primarily be operating in a metaphysical mode throughout the paper. A similar (though not identical) debate could nevertheless be had at the conceptual level. Theorists who are afraid of metaphysics but comforted by concepts can translate accordingly.

⁶ Typically but not always. Finlay (2014), for example, is a prominent exception.

⁷ Note that Reasons Fundamentalism only requires that the reasons-relation be *normatively* fundamental; it's compatible with that relation being *metaphysically* non-fundamental (cf. Schroeder, 2007).

⁸ We're glossing over some further distinctions, since they won't be relevant in what follows. For example, one can distinguish between a version of Reasons-First formulated in terms of grounding and a version formulated in terms of analysis (cf. Schroeder, forthcoming), and between giving a metaphysical analysis of some property (cf. Dorr, 2016) and giving an account of the essence of that property (cf. Rosen, 2015).

3. Explanationism: Against Ought- and Good-based Views

3.1 Ought-Explanationism about Reasons

John Broome (2004, 2013) has defended a version of explanationism that focuses on facts about what we *ought* to do. The simplest version of this view is as follows:

Simple ought-explanationism: For r to be a reason for S to ϕ is for r to explain why S ought to ϕ .⁹

While this view makes straightforward sense of the connection between reasons and explanation (cf. §1), it also faces an obvious problem: we often have reasons to do things that we ought not do, since we have ‘stronger’ or ‘weightier’ reasons to do something else instead. For example, the fact that soda is tasty might be a reason for Nils to drink soda, but if there are other, weightier reasons not to do so – for example, that soda is unhealthy – then the fact that soda is tasty will be *outweighed*: it’s a reason to do something (i.e. to drink soda) that Nils, ultimately, ought not to do. Given that explanation is factive, an outweighed reason to ϕ can’t explain why we ought to ϕ , since it’s false that we ought to ϕ . The possibility of outweighed reasons is thus incompatible with simple ought-explanationism.

Broome recognizes this problem (Broome, 2013: ch. 4.3). His solution is to distinguish 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 ϕ is a fact that explains why S ought to ϕ , whereas a *pro tanto* reason for S to ϕ is a fact that plays a certain role in a ‘weighing explanation’ of why S ought (or ought not) to ϕ .¹⁰ A weighing explanation of an ought-fact is one that appeals to the weight of reasons: S ought to ϕ if (and because) the reasons for ϕ -ing outweigh the reasons against ϕ -ing, and S ought not to ϕ if (and because) the reasons against ϕ -ing outweigh the reasons for ϕ -ing. Not all explanations of ought-facts are

⁹ Simple ought-explanationism resembles some other explanationist views, including *fitting-explanationism*, on which reasons explain facts about *fittingness*. Fitting-explanationism is hinted at, but not developed, by Chappell (2012: 688), and a disjunctive variant of it is advanced by Howard (2019). It is difficult to assess such views, however, since important questions remain concerning how the relevant notion of fittingness should be understood – Howard (2018: 8), for example, concludes his survey by listing several important and so-far-unaddressed questions about how best to understand it). Nonetheless, depending on how the view is developed, it will likely face a dilemma. If the central notion of fittingness does not come in degrees – e.g. if fittingness is identified with correctness – then the corresponding form of explanationism will face problems similar to those facing ought-explanationism discussed in this section. On the other hand, if the relevant notion does come in degrees, then the view risks being a confusingly-formulated notational variant of the view that we’ll defend in §4, and thus not a genuine competitor to it.

¹⁰ Or in a weighing explanation of why it’s not the case that S ought to ϕ and not the case that S ought not to ϕ . We omit this for simplicity.

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 ϕ is for r to play the ‘for- ϕ ’ role in a weighing explanation, and for r to be a *pro tanto* reason for S *not* to ϕ is for r to play the ‘against- ϕ ’ 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:

Broomean ought-explanationism: For r to be a reason for S to ϕ is for r to either (i) explain why S ought to ϕ [*pro toto reasons*] or (ii) play the for- ϕ role in a weighing explanation of why S ought (or ought not) to ϕ [*pro tanto reasons*].

Broome’s analysis has been extensively criticized. We’ll focus on three criticisms of his analysis of *pro tanto* reasons specifically – namely, (a) it is uninformative, (b) the notion of normative weight remains unexplained, and (c) the weighing metaphor is inapt. All three will be useful points of contrast to the view we’ll develop in §4.¹¹

First, it has been argued that our understanding of what it is for facts to play the ‘for- ϕ ’ or ‘against- ϕ ’ role in a weighing explanation rests on a prior, more basic understanding of what it is for facts to count in favor of or against ϕ -ing, rather than vice versa, and that Broome therefore fails to provide an informative, noncircular analysis of reasons (cf. Kearns & Star, 2008; Brunero, 2013; Dancy, 2015).¹²

Second, Broome’s account of weighing explanations assumes that reasons have normative ‘weights’ that combine in some way to determine what an agent ought to do. However, Broome says nothing about the weight(s) of reasons or how they are determined, other than that they play this characteristic role (cf. Kearns & Star, 2008; Dancy, 2015). So even if (contrary to the first objection) Broome successfully provides an informative analysis of reasons in terms of ought, he doesn’t provide an analysis

¹¹ Additional criticisms we won’t consider concern Broome’s notion of a *pro toto* reason (cf. Brunero, 2013; Dancy, 2015; Nebel, 2019) as well as the disjunctive nature of Broome’s account. For instance, Nebel (2019) is sympathetic with Broome’s view, but thinks it comes “at the cost of an objectionably disjunctive account of normative reasons” (2019: 462; cf. Kearns & Star, 2008). Instead, Nebel defends the more simple view that “a (normative) reason for S to ϕ is just an (explanatory) reason why S ought to ϕ ” (2019: 462). This counts as a form of ought-explanationism, assuming that for something to be an explanatory reason why p is for it to help explain why p . Interestingly, however, Nebel denies that ‘reason why p ’ is factive with respect to p , and hence denies that explanatory reasons are just explanations (since ‘explains why p ’ is factive). While engaging with Nebel’s view is beyond the scope of this paper, it’s worth noting that Nebel is only concerned to defend the simple ought-based view from an objection, not to motivate it vis-à-vis its competitors, and the central data regarding the limited factivity of ‘reason why p ’ is compatible with competing analyses of reasons, including ours.

¹² Broome (2013: 54–55) addresses this worry. For convincing responses to Broome’s response, see Kearns and Star (2008: 43–44), who (rather than being oracles) were responding to a pre-publication draft of Broome’s book, and Dancy (2015: 181–182).

of the weight of reasons in terms of ought. Moreover, it's not clear that he *could* provide such an analysis, since when an ought-fact admits of a weighing explanation, the ought-fact will obtain *in virtue of* facts about the weights of reasons. Assuming the relevant kind of explanation is asymmetric, the facts about weight(s) can't then obtain in virtue of the ought-fact.

One reaction to this problem is to simply take the notion of normative weight as primitive. Doing so, however, is in tension with Broome's ought-centric picture, and it also suggests a much simpler account of reasons: for a fact to be a normative reason is for it to have normative weight. This points in the direction of the proposal we'll be defending in §4. At this stage, however, what matters is just that such a weight-based view is very different from ought-explanationism.

The third problem we'll consider besets both Broomean ought-explanationism and the simple weight-based view just mentioned, as it concerns the idea of normative weighing explanations. These explanations are supposed to be analogous to mechanical weighing explanations (cf. Broome, 2013: 52), where individual weights are placed on either side of a scale and the combined weight of each side determines whether the scale tips – and if so which way and by how much it tips. Similarly, in normative weighing explanations, there are reasons for and against an action, each of which has a normative 'weight', and the combined weight of each 'side' determines whether or not the action ought to be done. As John Hawthorne and Ofra Magidor (2018) note, however, the analogy is not obviously apt: for example, whereas "plac-ing weights on a balance is a monotonic process (in the sense that placing more weights on one pan can only increase the total weight on that pan), and it is additive (to deter-mine how much total weight there is on one pan of the balance, one simply adds the weights)", it's controversial whether the process of "balancing reasons" has these features (133). Broome himself recognizes these limitations of the weighing analogy, as well as others (2013: 52). However, as Hawthorne and Magidor also note, if the weights of reasons do not add up or interact in the way that the weights of objects add up or interact, then the nature of normative weighing explanations (which Broome characterizes via the weighing analogy) as well as of normative reasons (which Broome characterizes in terms of normative weighing explanations) becomes significantly less clear.¹³

While these problems for ought-explanationism may not be decisive on their own, they together motivate the search for an alternative that avoids them.

3.2 Good-Explanationism about Reasons

The second view we'll consider is good-explanationism. Here's a first pass:

¹³ For a defense of the weighing metaphor (but not of Broome's view), see Maguire and Snedegar (2021).

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

Good-explanationism improves over simple ought-explanationism because it straightforwardly accommodates outweighed reasons. That's because goodness is gradable—things can be more or less good, and some good things are better than others—and an action that ought not to be performed can still be good to some degree. In such a case, a fact that explains why the sub-optimal action is good to some degree counts, on this view, as a reason to do it. For example, the fact that soda is tasty might explain why Nils' drinking soda is good to some degree, even if, all things considered, Nils ought not to drink soda. Good-explanationism also avoids the problems for Broomean ought-explanationism outlined above, as it doesn't invoke the idea of a weighing explanation in its analysis of what it is to be a reason.

The most prominent advocate of good-explanationism is Stephen Finlay (e.g., Finlay, 2014, 2019; see also Gardner & Macklem, 2004; Maguire, 2016; Wedgwood, ms). Some care is called for, however, because Finlay's subject matter and methodology differ importantly from our own (cf. §2). What Finlay aims to do is to offer a semantics for normative terms, including (but not limited to) the count noun 'reason(s)', that captures their meaning and use in ordinary communication. He then uses his semantics as a guide in answering various conceptual and metaphysical questions about normativity. However, Finlay does not clearly distinguish between the ordinary notion of a reason and the theoretical one that, we think, is of primary interest in normative and metanormative theorizing. Indeed, in recent work he explicitly takes accounts of reasons as "answers to ('why') questions" to be roughly equivalent to accounts of reasons as "grounds/truth-makers" (2019: 62, fn. 1), whereas we want to sharply distinguish them. As far as we can see, there is no more reason to expect that a correct account of reasons-as-metaphysical-grounds can be arrived at through the study of ordinary normative language than there is to expect that a correct physical theory of mass and force can be arrived at through the study of the ordinary English terms 'mass' and 'force'. Finlay himself does little to motivate his language-first methodology, other than to say that the proof is "in the pudding" (2014: 14). Since Finlay's semantic project has been criticized elsewhere (e.g., Fogal 2016b; Dowell, 2020; Worsnip, 2020), we'll put Finlay's alternative methodology to the side. Instead, we'll treat good-explanationism as an account of the theoretical notion of a reason, rather than the ordinary language one.

Two general forms of good-explanationism can be distinguished, which differ on what they take the relevant kind of goodness to be. What we'll call *restricted* versions focus only on forms of goodness that are normatively significant in and of themselves, such as moral value, prudential value, and epistemic value. We'll call these 'robust' forms of goodness. On restricted versions of good-explanationism, a fact is a reason

if and only if it explains why the action in question is good in one of these robust ways. *Unrestricted* versions, by contrast, place no such restrictions. Instead, on such a view, pretty much any kind of goodness can play the relevant role, as long as it is expressible by '(is) good'.¹⁴ There are various standards relevant, for example, in determining whether a given car, knife, or candy bar is good – the goodness of a car depends (inter alia) on how well it drives, the goodness of a knife depends (inter alia) on how well it cuts, and the goodness of a candy bar depends (inter alia) on how it tastes. Similarly, whether an action is good for a certain purpose depends (inter alia) on whether it is a means for promoting that purpose. And on unrestricted versions of good-explanationism, all these forms of goodness (in addition to the robust ones) can be relevant to whether a fact is a reason to perform a certain action.

Both versions of good-explanationism face problems. The problem with restricted views is simple: when an agent's performing a certain action would be robustly good, this fact can *itself* be a reason to perform the action. This idea is not only intuitively plausible, but it also accords well with the various heuristics for identifying normative reasons considered earlier (cf. §2). In particular, the fact that performing an action would be robustly good seems capable both of counting in favor of doing it, thus fulfilling (R1), and of helping explain why we ought (or are permitted) to do it, thus fulfilling (R2). However, given that the relevant form of explanation is asymmetric, the fact that performing the action would be robustly good cannot explain why performing the action would be robustly good. Restricted goodness-explanationism is thus incompatible with facts about robust goodness themselves being reasons.¹⁵

Unrestricted versions of goodness-explanationism can avoid the problem just outlined. In addition to the normatively significant, or 'robust', ways in which an action can be good, there are also many other, not-clearly-normatively-significant ways in which it can be good. Hence, even if one grants that no fact explains itself, one can still maintain that when performing a certain action would be robustly good, this may explain why performing it would be good in some *other*, not-necessarily-

¹⁴ The unrestricted version is closer to Finlay's view, according to which (roughly) reasons explain why the relevant action would be good for some contextually salient end (see e.g. 2014: 145). After all, there seems to be no restrictions on what ends can (in some context or other) be contextually salient.

¹⁵ This kind of argument has been used to criticize so-called 'buck-passing' analyses of goodness (cf. Crisp, 2005). However, as Schroeder (2009) argues, the most plausible buck-passing analyses of goodness are compatible with thinking that goodness-facts can themselves be reasons. This is good news for buck-passers, since it is hard to think of independent reasons to deny that goodness-facts can indeed be reasons (independent, that is, from the idea that buck-passing analyses might entail such a denial).

robust way, and that this suffices for that fact to be a reason.¹⁶ One might claim, for instance, that when performing an action would be morally good, this fact explains why performing it is *good for the purposes of doing something morally good* (which, we may grant, is a form of goodness that is distinct from moral goodness). And on unrestricted versions of good-explanationism, when a fact explains that performing an action would be good in this latter way, that suffices for it being a reason to perform the action.

However, unrestricted views are *too* unrestricted, leading to undesirable implications. For example, just as performing an action can be good for the purposes of doing something morally good, so it can be good for the purposes of doing something morally *bad*. But a fact that explains why performing an action would be good for the purposes of doing something morally bad is not thereby a reason to perform it—it's not the sort of thing that actually counts in favor of performing the action, or that helps explain why it ought to be performed. Unrestricted good-explanationism is incompatible with this fact.

Good-explanationism thus faces a dilemma. On the one hand, if the forms of goodness appealed to are robust, then good-explanationism *undergenerates* reasons. This is because the relevant facts of the form [S's ϕ -ing would be (robustly) good] can plausibly be reasons themselves, which, given that nothing explains itself, is incompatible with the restricted form of good-explanationism. On the other hand, if any kind of goodness can play the relevant role, then good-explanationism *overgenerates* reasons—it implies, for example, that a fact that explains why ϕ -ing would be good for doing something morally bad is a normative reason to ϕ .¹⁷

To avoid this dilemma one must strike a balance between the two extremes. That is, one needs to formulate a version of good-explanationism that's more restricted than the most unrestricted ones and less restricted than the most restricted ones. As far as we can see, however, there's no reason to think that such a middle-ground view would avoid both horns of the dilemma—if anything, such a view would risk being impaled on both.

¹⁶ This strategy resembles the one endorsed by Finlay (2019: §4.1) who also seeks to accommodate the idea that goodness-facts can be reasons by taking them to explain some other kind of goodness-fact.

¹⁷ Note that this dilemma doesn't depend on assumptions about what value bearers are. If one thinks that states of affairs can be (robustly or non-robustly) good, for example, then the dilemma can be run in terms of facts of the form [the states of affairs in which *S* performs ϕ would be good], whereas if one thinks that actions can be good, then the dilemma can be run in terms of facts of them form [S's ϕ -ing would be good], and so on. Thanks to an anonymous reviewer for prompting us to clarify this.

3.3 Preliminary diagnosis

While the problems for ought- and good-explanationism may not be decisive, they do motivate the search for a better explanationist account. In this subsection, we'll formulate three desiderata a view should satisfy in order to avoid the problems that ought- and good-explanationism face. In the subsequent section, we'll show how our favored view – support-explanationism – satisfies all three.

As noted above, simple ought-explanationism fails to accommodate outweighed reasons. Good-explanationism, in contrast, faced no such difficulty. The reason was simple: goodness, unlike ought, is gradable. Goodness is also *quantity-like*, in that it makes sense to ask *how much* goodness (or value) something has. While appealing to a gradable and/or quantity-like notion isn't the only way to accommodate outweighed reasons (as Broome's view illustrates), we do take it to be the most natural one. This yields the following desideratum:

Desideratum #1: Appeal to a gradable and/or quantity-like normative notion in one's analysis of normative reasons.

Instead of appealing to a gradable and/or quantity-like normative notion, Broome accounts for outweighed reasons by appealing to the notion of a weighing explanation in his analysis of normative reasons. But as noted above, such an appeal is problematic. The second desideratum ensures that we avoid it:

Desideratum #2: Avoid appealing to the notion of a weighing explanation in one's analysis of normative reasons.

Good-explanationism satisfies both Desideratum #1 and Desideratum #2, and hence fares better than both forms of ought-explanationism. However, as we argued, it faces a dilemma: depending on the details, it either yields too few reasons, by denying that facts about robust goodness are reasons, or yields too many reasons, by implying that anything that explains any kind of goodness-fact is a reason. A view that satisfies desideratum #3 avoids this dilemma:

Desideratum #3: Allow facts about robust forms of goodness to be normative reasons without implying that all facts about non-robust forms of goodness are reasons.

Collectively, these desiderata motivate the search for a version of explanationism that appeals to a quantity-like normative notion that's distinct from goodness, thereby avoiding the need to appeal to weighing explanations while also allowing for more plausible verdicts about whether – and if so which – goodness-facts are normative reasons. This line of thought leads to the view we favor: *support-explanationism* about reasons.

4. Explanationism: The Support-Based View

The basic idea behind support-explanationism is simple: *to be a reason is to provide – or be a ‘source’ of – normative support*, where the ‘weight,’ ‘strength,’ or ‘significance’ of a reason is simply a matter of how much support it provides. In this section we’ll clarify support-explanationism in numerous respects (§4.1), compare and contrast it with Reasons-First (§4.2), and conclude by exploring some of its substantive implications (§§4.3-5.2).

4.1 Clarifying support-explanationism

Our view can be initially formulated as follows:

Support-explanationism: For r to be a normative reason for S to ϕ is for r to explain why there is a certain amount of normative support for S ’s ϕ -ing.

The property of being a normative reason is thus analyzed in terms of the notion of normative support together with the notion of explanation. This is compatible with there being distinct species of normative support: just as one might want to distinguish between moral reasons, epistemic reasons, prudential reasons, and so on, so one can distinguish between moral support, epistemic support, prudential support, and so on, and in each case analyze the former in terms of the explanation of the latter. Like other explanationist views, then, support-explanationism neatly secures the connection between reasons and explanation characterized earlier (§1).¹⁸

To illustrate the implications of support-explanationism, return to the possibility of outweighed reasons. In the soda case in §3.1, the fact that soda is tasty helps explain why there’s a certain amount of support for Nils’ drinking it, and the fact that soda is unhealthy helps explain why there’s a certain amount of support for Nils’ not drinking it.¹⁹ Both facts are reasons (one for and the other against), but the latter is ‘stronger’ or ‘weightier’ since the amount of support it provides is greater. This example also illustrates the fact that normative support is typically *pro tanto*, in the sense that there can be support for actions (such as Nils’s drinking soda) that, all things considered, ought not to be performed. We’ll mostly focus on *pro tanto*

¹⁸ Support-explanationism is also compatible with the distinction between ‘objective’ kinds of support and other, more ‘subjective’ or ‘perspectival’ kinds of support, akin to the distinction between objective (or fact-relative) reasons and other, more subjective (e.g. evidence-relative, belief-relative) reasons. For more on the nature and importance of such distinctions, see Fogal and Worsnip (forthcoming).

¹⁹ For simplicity, we’re assuming that reasons against ϕ should be understood as facts that provide support for not doing ϕ . However, another possibility (suggested to us by Selim Berker in personal communication) is to understand reasons against ϕ as facts that provide *dis-support* for doing ϕ . Support-explanationism is compatible with both these possibilities. For further discussion of reasons against, see Snedegar (2018; 2021) and Tucker (forthcoming).

support in what follows, leaving the ‘pro tanto’ qualifier implicit, though we leave open the possibility that some facts count as reasons in virtue of providing decisive (rather than merely pro tanto) support.²⁰

Given the distinctions between ordinary and theoretical uses of ‘reason(s)’ and between pragmatic and non-pragmatic notions of explanation (cf. §2), support-explanationism can be more carefully put as follows: reasons in the *ordinary* normative sense are considerations that help *make sense of why* (i.e. pragmatically explain why) there’s a certain amount of normative support, while reasons in the relevant *theoretical* sense – i.e. the one associated with roles (R1) and (R2) above – are facts that help *make it the case that* (i.e. non-pragmatically explain why) there is a certain amount of normative support. The theoretical notion of a normative reason is thus related to, but not identical to, the ordinary one. We can thus distinguish:

Support-explanationism_O: For r to be a normative reason (in the ordinary sense) for S to ϕ is for r to make sense of why there is a certain amount of (pro tanto) normative support for S ’s ϕ -ing.

Support-explanationism_T: For r to be a normative reason (in the theoretical sense) for S to ϕ is for r to provide a certain amount of (pro tanto) normative support for S ’s ϕ -ing.

As noted above, it is the theoretical notion, and hence support-explanationism_T, that primarily interests us here.

In the remainder of this section we’ll further clarify the key terms that figure in the analysis: ‘provide’ (§4.1.1), ‘a certain amount of’ (§4.1.2), ‘pro tanto’ (§4.1.3), and ‘support’ (§4.1.4). We’ll then make clear how support-explanationism satisfies the three desiderata outlined earlier (§4.1.5) and highlight three respects in which it is metanormatively neutral (§4.1.6).

4.1.1 Clarifying ‘provide’

On our view, reasons in the theoretical sense don’t just play any old role in helping make it the case that there is a certain amount of (pro tanto) normative support: as we variously put it, they *provide*, or *generate*, or *are sources of* such support. Each of these terms is meant to express the same thing: *direct full grounding*, where grounding is a form of metaphysical explanation (cf. §2). For the fact $[p]$ to *directly* (or *immediately*) ground $[q]$ is, roughly, for $[p]$ to ground $[q]$ without there being any grounds ‘between’ $[p]$ and $[q]$.²¹ (We use square brackets for facts, and sometimes for fact

²⁰ Our view is thus compatible with the possibility of there being indefeasible duties (cf. Reisner 2018), understood in terms of indefeasible support.

²¹ This is merely meant to be an intuitive gloss; we’re ignoring various possible complications (Fine 2012).

schemata—context will disambiguate.) This is to be contrasted with indirect (or mediate) grounding, which involves additional grounds between what's grounded and what grounds it. As Kit Fine (2012) notes, for instance, the disjunctive fact $[p \vee q]$ is directly grounded in $[p]$ and $[q]$ while facts about cities are (at best) indirectly grounded in facts about atoms, and it is the notion of direct or immediate ground that thus “provides us with our sense of a ground-theoretic hierarchy” (51). Further, for the fact $[p]$ to *fully* ground $[q]$ is, roughly, for $[p]$ to constitute a full—as opposed to merely partial—metaphysical explanation of $[q]$. $[p]$ fully grounds $[p \vee q]$, for example, but only partially grounds $[p \ \& \ q]$ —what fully grounds $[p \ \& \ q]$ is $[p]$ *together with* $[q]$.

Appealing to the general (and generally recognized) distinctions between direct and indirect grounding and full and partial grounding allows support-explanationism—as well as other forms of explanationism—to avoid certain counterintuitive consequences. Suppose, for instance, that it seems to Jens that p , and that this provides some support for him to believe that p . According to support-explanationism, [It seems to Jens that p] is thereby a reason for him to believe that p . By invoking direct grounding, however, the view avoids implausibly counting all further grounds of [It seems to Jens that p], such as facts about his neural state or the microphysical facts that ultimately ground that state, as reasons for him to believe p . The distinction between full and partial grounding is also important. Suppose, as seems plausible, that it is the fact that Candice sees smoke *together with* the fact that smoke is a sign of fire that provides support for her believing that there's a fire. By invoking full grounding, support-explanationism avoids counting [Candice sees smoke] and [Smoke is a sign of fire] as separate reasons, and hence avoids the sorts of 'double-counting' worries raised earlier (see §2).

However, understanding the provision of normative support in terms of full, direct metaphysical explanation might still be too coarse-grained. That's because there may be different roles for facts to play in the full, direct metaphysical explanation of other facts, similar to how there might be different roles for events to play in the full, direct causal explanation of other events. In the case of causal explanation, for example, one might want to distinguish causes (understood as causal 'producers') from mere background conditions, and distinguish both from the laws that govern causal relations between events. Yet all three might have distinctive roles to play in the full, direct causal explanation of an event. Similarly, we might want to distinguish grounds (understood as metaphysical 'producers') from mere background conditions, and distinguish both from the metaphysical laws that govern grounding relations between facts. All three might have distinctive roles to play in the full, direct metaphysical explanation of a given fact.²² In normative theorizing, for

²² On the role of laws in metaphysical explanations, see Glazier (2016), Schaffer (2017), and Fogal and Risberg (2020).

example, it's common to distinguish between reasons and various background conditions, including so-called 'enablers' and 'disablers' (cf. Dancy, 2004: ch. 3; Bader, 2016). To illustrate, as noted above (§2), many philosophers deny that the fact that the subject can φ is itself a reason – or part of a reason – for them to φ , and instead insist that such facts merely enable other facts to be reasons. Support-explanationism can accommodate these nuances: enablers and disablers can play their characteristic roles and hence be part of the full (non-pragmatic or metaphysical) explanation of why there is a certain amount of normative support for φ -ing, but insofar as they aren't themselves providers or sources of that support, they won't be reasons to φ in the relevant sense. 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 (i.e. reasons) without themselves being sources of support – they modify existing amounts of support without generating it on their own.

In general, however, we needn't take a stand on the exact nature or utility – or even the intelligibility – of the various (extra-)fine-grained distinctions introduced so far. We view them as intramural matters, and hence something our account is intended to be neutral with respect to.

4.1.2 Clarifying 'a certain amount of'

Turning to the locution 'a certain amount of' in the official formulation of support-explanationism, the basic idea is that for r to be a reason of a certain 'strength' or 'weight' is for r to provide a corresponding amount of normative support. Thus, strong reasons provide a lot of support, weak reasons provide little support, and so on. Generalizing, we can formulate a version of the support-explanationism that tells us what it is for a reason to have a certain strength:

Support-explanationism_T: For r to be a normative reason of strength d for S to φ is for r to provide d amount of (pro tanto) normative support for S 's φ -ing.

This analysis plausibly entails the more generic analysis of reasons above (i.e., support-explanationism_T), given that r is a reason if and only if r is a reason of some strength or other.²³

²³ There are of course further questions about how to understand amounts of support (e.g. whether they should be thought of as genuine abstract entities), corresponding to questions about how to understand lengths, heights, values, and any other magnitude. Support-explanationism is compatible with different views about these general issues. See Bykvist (forthcoming) for further discussion of value magnitudes in particular.

4.1.3 Clarifying ‘pro tanto’

The main point regarding the ‘pro tanto’ qualifier in support-explanationism is that the relevant notion of support is (at least in general) defeasible. There are two main forms of defeat, both of which admit of degrees: *opposition* (or rebutting) and *undermining* (or attenuation). 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).²⁴

If one accepts the (extra-)fine-grained distinctions introduced above (§4.1.1), we can further clarify the notion of pro tanto support by distinguishing the following four ‘levels’ in explanations of particular facts about what we ought or are permitted to do:²⁵

- (1) *the underlying level*: facts (often non-normative ones) that provide support for and against the relevant actions or attitudes available to a given agent in a given situation;
- (2) *the initial normative level*: facts about how much support the facts on the underlying level provide, prior to being modified by modifiers, such as intensifiers, underminers, and so on (if such there be);
- (3) *the modified normative level*: facts about how the support generated on the initial normative level is modified by modifiers (if such there be);
- (4) *the overall normative level*: facts about which actions/attitudes the relevant agent ought or may perform/have, and which obtain in virtue of the facts about modified support at the third level.²⁶

Given this quadripartite picture, the question arises how we should understand the various ‘pro tanto’-notions that have been invoked in the literature (pro tanto reasons,

²⁴ As Pryor (2013: 90) notes, although “[m]ost of us discern an intuitive kind” when considering such examples of epistemic undermining, “it’s difficult to say in a rigorous way what makes it distinctive”. See Pryor (2013: §2) for some possible complications.

²⁵ This picture is not meant to be exhaustive; among other things, it ignores the fact that general normative principles (or ‘laws’) concerning support may (and we ourselves think do) also play a role in fully explaining particular facts about support, similar to how causal laws arguably play a role in fully explaining the occurrence of particular events (see §4.1.1). For articulation and defense of the general principle-based account of normative explanation, see Fogal and Risberg (2020).

²⁶ We’re here drawing and expanding on Berker (2007: 116), who presents a three-level model for moral explanations (rather than for normative explanations more generally). Berker does not distinguish between what we call the initial and the modified normative levels; his model only includes the underlying level, the “contributory” level and the overall level.

pro tanto duties, etc.). For instance, does the amount of support generated by a pro tanto reason correspond to the amount of support provided at the initial normative level or the modified normative level? We think there's no need to choose: there are simply two useful notions of a pro tanto reason (and more generally, two useful groups of pro tanto notions), corresponding to the second and the third level, respectively. What we may call *initial* or *unmodified* support is the support that modifiers modify, and which is determined on the second (initial) level, whereas what we may call *modified* support is the support that results once the modifiers have modified the initial support, and which is thus determined on the third (modified) level. Given support-explanationism, we can thus say that r is an initial reason just in case r provides initial support, and r is a modified reason just in case r provides modified support.

In many cases, one and the same fact will be both an initial and a modified reason. However, one way that the initial and modified notions of a reason might come apart is in cases where a reason is attenuated or undermined entirely (reduced to zero, as it were).²⁷ Suppose again that Andreas testifies that p , thereby providing some initial support for us to believe that p . Suppose further, however, that we subsequently learn that Andreas is no more reliable than a coin flip. This clearly undermines the support provided by his testimony – so much so (let's suppose) that it is eliminated entirely.²⁸ Does the fact that Andreas testifies that p still count as a reason for us to believe that p ? Support-explanationism is compatible with different answers depending on the notion of support one chooses to invoke, but we think the best answer is: in one sense yes, in another sense no. It's a reason in the initial sense because it provides initial support, but not in the modified sense because it fails to provide modified support.

4.1.4 Clarifying 'support'

The fourth and final point of clarification concerns the locution 'support' that support-explanationism invokes. Here it's crucial to distinguish between the *verb* '(to) support' and the *mass noun* 'support' (each in their normative senses). The verb is very closely related to the count noun 'reason(s)' in its normative sense – indeed, we take the claim that r is a normative reason to ϕ to be essentially equivalent to the claim that r normatively supports ϕ -ing. All parties can, should, and often do accept this much. Indeed, even Reasons-Firsters – who reject explanationism about reasons –

²⁷ Thanks to an anonymous referee for urging us to address this kind of case. (A similar set of issues arises with respect to disabled reasons as well.)

²⁸ We ourselves actually doubt this is a case of complete elimination – like Schroeder (2007: §5.3) we think cases of complete elimination are at best rare.

often claim that being a reason to ϕ is basically equivalent to ‘favoring’ or ‘counting in favor’ of ϕ -ing, where the latter is basically equivalent to supporting ϕ -ing.²⁹

Can something informative be said, then, about what it is for r to normatively support (verb) ϕ -ing? According to Reasons-First, the answer is ‘no’. According to support-explanationism, however, the answer is ‘yes’: for r to normatively support ϕ -ing is for r to *provide*—or *be a source of*—*normative support* for ϕ -ing, where ‘support’ in its latter occurrence is a mass noun, not a verb. According to support-explanationism, then, what it is to normatively support (verb) something is analyzed in terms of normative support (mass noun) together with the relevant notion of metaphysical explanation, in much the same way that normative reasons are. (This is unsurprising, since, again, being a normative reason to ϕ is essentially equivalent to supporting ϕ .) So whereas a verb-based version of support-explanationism risks being a notational variant of Reasons-First, the mass noun-based view—which we endorse—is not. To insist otherwise would be to collapse the distinction, noted at the

²⁹ We focus on ‘support’ rather than ‘count(s) in favor’ or ‘favor(s)’—both of which are similarly intimately related to reasons but not similarly used as mass nouns—because ‘support’ is more expressively adequate in at least three ways. First, as a verb it expresses more-or-less the same thing as ‘count(s) in favor of’ and ‘favor(s)’ while avoiding some of the idiosyncrasies of both—like ‘favor(s)’ and unlike ‘count(s) in favor’, it naturally allows for passive constructions (compare ‘ ϕ -ing is supported/favored (by r)’ to ‘# ϕ -ing is counted in favor of (by r)’), and like ‘count(s) in favor’ but unlike ‘favor(s)’, it allows for a clear(er) non-contrastive reading—in a sentence such as ‘ r favors ϕ -ing’, it is natural to hear an implicit contrast: ‘ r favors ϕ -ing rather than ψ -ing’. Although reasons and support may ultimately be contrastive (though we ourselves doubt it), that’s a substantive issue, and so not something we should build into our terminology from the outset (see further Snedegar, 2017). Secondly, although all three terms can be used in ways that reflect the fundamentally gradable, quantity-like nature of the underlying normative phenomenon—e.g. we can ask, ‘How much does r support/favor/count in favor of ϕ -ing?’ and answer by saying things like, ‘ r strongly supports/favors/counts in favor of ϕ -ing’—‘support’, given its use as both a verb and mass noun, is the most flexible and systematic—for instance, we can answer the preceding question by quantifying or otherwise qualifying the *amount* of support (mass noun) rather than just qualifying the *relation* of support (verb). Lastly, although ‘reason(s)’, ‘favor(s)’, and ‘count(s) in favor’ are related in English, the latter two expressions appear to be somewhat idiosyncratic to English. In Swedish, for instance, there are no straightforward translations of ‘favor(s)’ or ‘count(s) in favor’; the expression closest to ‘count(s) in favor’ is ‘talar för’ (literally ‘speaks for’), and there is no expression that’s particularly close to ‘favor(s)’. Yet Swedish, and the Swedish term for ‘reason’ (‘skäl’) in particular, resembles English in many other relevant ways: (a) ‘skäl’ can be used both as a mass and a count noun; (b) ‘skäl’ is connected to a mass noun/verb-pair that translates to ‘support’/‘supports’ (‘stöd’/‘stödjer’); (c) ‘skäl’ and ‘reason’ are connected to explanation in the same way, with ‘skäl’ also having motivating and merely explanatory uses, causal as well as metaphysical ones, and so on; and (d) ‘skäl’ and ‘stöd’ both naturally combine with a locution corresponding to ‘provides’ (‘ger skäl’ and ‘ger stöd’; literally ‘gives reason’ and ‘gives support’). This provides some limited but suggestive cross-linguistic support for taking ‘support’ to be more reflective of the underlying phenomenon than ‘count(s) in favor’ and ‘favor(s)’. Thanks to an anonymous referee for prompting us to clarify this issue.

outset, between being a source or provider of something (in this case, being a reason) with that which it is a source or provider of (in this case, normative support).

Although ‘normative support’ is a theoretical term, what it denotes is familiar. In ordinary thought and talk, the notion of normative support is commonly expressed by the mass noun ‘reason’ (as well as, of course, ‘support’). To say that there is reason (mass noun) to ϕ is essentially the same as saying that there is normative support for ϕ -ing, and hence that ϕ -ing is, at least to 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 ϕ -ing. Support-explanationism could thus alternatively be called ‘reason-explanationism’, and be expressed in slogan form as follows: to be a reason is to provide reason (i.e., support). Indeed, this view parallels the account of the semantic relation between the mass noun ‘reason’ and count noun ‘reason(s)’ (in its normative sense) provided by Fogal (2016a), drawing on work in lexical semantics on the mass-count distinction. There are nonetheless important differences both in focus—Fogal focuses on ‘reason(s)’ in the ordinary sense whereas we focus on the theoretical sense—and in methodology—Fogal’s arguments are primarily linguistic in nature whereas ours are not. The present point is just that support-explanationism sits well with an independently-motivated linguistic view about the relationship between mass and count noun uses of ‘reason’ in its ordinary normative sense. That said, for the sake of clarity of contrast with reasons, we prefer the quasi-theoretical term ‘(normative) support’.

We should note, however, that the term ‘support’ is polysemous, and while some uses of it pattern like ‘(normative) support’, with the verb plausibly being analyzable in terms of the mass noun, others might not. Consider, for example, the financial sense of ‘support’. Here too the verbal form ‘to (financially) support’ is plausibly analyzed in terms of the mass noun ‘(financial) support’: for *A* to financially support *B* is for *A* to provide, or be a source of, financial support to *B*, and being a source of financial support (e.g. being a parent or investor) is clearly different from the support that is provided (e.g. money). This contrasts with the sense of ‘support’ invoked when talking about sports allegiances, such as when we say that Elton John supports Watford F.C. Here the verb ‘(to) support’ is not naturally analyzed in terms of the mass noun ‘support’: for *S* to support some team *T* is not intuitively a matter of *S* providing some kind of support to *T*. Indeed, talk of a fan ‘providing support’ to a sports team sounds odd, unless it means something like financial or emotional support, neither of which is usually the case. More generally, with notions of support where the verbal form is basic, talk of ‘providing’ support is usually stilted or otherwise unnatural. Since such talk is highly natural in the normative case, this

further supports taking the mass noun '(normative) support' to be prior to the verb 'to (normatively) support'.³⁰

Although the analysis of reasons in terms of support is meant to be informative, it is not meant to be revolutionary. Many other—and clearly correct—analyses are similar in this respect. What is it, for example, to be a beer (count noun)? The obvious answer is: to be a beer is to be a (conventional) unit or serving of beer (mass noun). What is it to financially support (verb) someone? It's to provide them with financial support (mass noun), such as money. What is it to be a cook? It's to be someone who prepares food for consumption. And so on. While such analyses might seem almost trivially true, that's not a reason to reject them—on the contrary, their obviousness indicates that they are *true*, not that they are false. Complaining that such analyses are uninformative is similarly misguided, since they clearly *are* informative. The analysis of being a beer (count noun) in terms of being a conventional unit or serving of beer (mass noun), for example, reveals the priority of the (denotation of the) mass noun over the (denotation of the) count noun. It also reveals a point of commonality with certain other count-mass pairs—a coffee, for example, is a (conventional) unit or

³⁰ This isn't to say that all uses of 'support' that pattern like '(normative) support' in *some* ways also pattern like it in *all* ways. There do tend to be other commonalities, however. For example, in cases where talk of 'providing support' is natural, talk of 'explaining why there is support' is also usually natural. If someone's parents provide them with financial support, then their parents (or facts about them) help explain why they have a corresponding amount of financial support, and when a group of senators provide support for a gun safety bill, they (or facts about them) help explain why there is a certain amount of political support for the bill. (We owe these examples to Selim Berker (personal communication), though they were originally offered as counterexamples.) But of course there are limits to the similarities between different uses of 'support'—they are different uses after all—and this is especially true of semi-technical uses of 'support', like ours. Political support is particularly notable, since it's a case where 'provision'-talk is apt even though the verb is not naturally analyzed in terms of the mass noun. Consider the sentence, 'There was some support for the gun safety bill in the Senate, but ultimately it was defeated by a vote of 55 to 45'. As Selim Berker points out (in personal communication), if we ask what explains why there was some support for the bill in the Senate, one answer is 'Because Senator #1 supported it (by voting for it)', another answer is 'Because Senator #2 supported it (by voting for it)', and so on. Here facts involving the relation denoted by the verb seem to explain facts about the quantity or amount denoted by the mass noun. While it might be thought that the same is true of normative support, we (of course) doubt it. That's because there are important disanalogies between normative support and political support in the sense at issue: whereas a fact can support ϕ -ing to a certain degree, or to a greater degree than another fact supports ϕ -ing, or to a greater degree than it supports ψ -ing, parallel claims aren't true of political support: either a senator supports a bill by voting for it or they don't. Similarly, if there's more political support for a bill in the Senate, then there are more supporters for the bill, but if there's more normative support for ϕ than for ψ , that doesn't entail that there are more reasons to ϕ (for instance, a strong reason to ϕ might outweigh two weak reasons to ψ). A modified version of the verb-based view of normative support might nevertheless be viable—one that includes an argument place for degrees. We criticize this possibility (formulated in the language of reasons) in §4.2.

serving of coffee. Such analyses also have informative implications: for instance, if someone has an alternative view of beers—perhaps they endorse a competing analysis, or take the notion of a beer to be unanalyzable (cf. ‘Beers-First’)—the analysis offered can reveal that they’re mistaken. Of course, the analysis of beers in terms of (conventional units or servings of) beer will only be informative to someone who already understands the analysandum—and in particular what beer is—but this is true of analyses in general.

All the points just made apply to support-explanationism as well. If the analysis of reasons in terms of the provision of normative support (or reason) seems to border on the trivial, this is a reason to *accept* the analysis, not to reject it. What’s more, the analysis is both informative in its own right—among other things, it reveals that normative reasons, just like motivating and explanatory reasons, can be understood partly in terms of explanation—and has informative implications—it reveals, among other things, that competing analyses of reasons should be rejected. (We’ll consider a number of further substantive implications of the analysis below.) Of course, if one has trouble grasping the notions that figure in the analysandum—whether it be normative support or the relevant notion of explanation or both—then one will not understand support-explanationism either. But again, this is true of any analysis whatsoever, including analyses of normative reasons such as good-explanationism and ought-explanationism.

Terminological care with regards to ‘support’, ‘reason’, ‘reason(s)’, ‘provides’, and related terms is called for not only because it helps make sense of their relationship to each other, but also because terminological carelessness can lead to confusion. Consider, for example, T. M. Scanlon’s influential characterization of a reason:

I will take the idea of a reason as primitive. Any attempt to explain what it is to be a reason for something seems to me to lead back to the same idea: a consideration that counts in favor of it. ‘Counts in favor how?’ one might ask. ‘By providing a reason for it’ seems to be the only answer. (Scanlon, 1998: 17)

Taken literally, this seems to entail that for something to *be* a reason is for it to *provide* a reason. To be a reason is thus to be something that provides something that provides something that... *ad infinitum*. That’s worse than uninformative.

Of course, the literal reading might be uncharitable. Maybe Scanlon just meant that the only answer to his question is ‘by *being* a reason for it’, or maybe he intentionally ignored such subtleties. Note, though, that there is another way to naturally adjust his formulation—instead of changing ‘providing’ to ‘being’, we might change ‘a reason’ (count noun) to ‘reason’ (mass noun). The answer to

Scanlon's question would then be 'by *providing reason* (i.e. support) for it', which is precisely the answer that support-explanationism predicts.

Similar issues arise with Derek Parfit's explanation of why the notion of a reason is "indefinable". Parfit writes:

[I]t is hard to explain the concept of a reason, or what the phrase 'a reason' means. Facts give us reasons, we might say, when they count in favour of our having some attitude, or our acting in some way. But 'counts in favour of' means roughly 'gives a reason for'. (Parfit, 2011: 31)

While this indicates why one might think that the notion of *giving* a reason is indefinable, it doesn't shed much light on what it is to *be* a reason (i.e. what it is that is 'given') or why it might be indefinable – unless one assumes that to be a reason is to give a reason, which again raises the spectre of regress.³¹

4.1.5 Satisfying the three desiderata

It should now be clear how support-explanationism satisfies the three desiderata presented in §3.3:

Desideratum #1: Appeal to a gradable and/or quantity-like normative notion in one's analysis of what it is to be a normative reason.

Desideratum #2: Avoid appealing the notion of a weighing explanation in one's analysis of what it is to be a normative reason.

Desideratum #3: Allow facts about robust forms of goodness to be normative reasons without implying that all facts about non-robust forms of goodness are reasons.

The notion of support is gradable and quantity-like (Desideratum #1), the view doesn't invoke the notion of a weighing explanation (Desideratum #2), and it allows that when an action is morally good, that fact provides some support for performing it, without also implying that when an action is non-robustly good (e.g. good for the purposes of doing something morally bad), that fact provides some support for performing it (Desideratum #3).

³¹ Interestingly, Broome (2021) presents an interpretation according to which Parfit's usage of the phrase 'gives a reason' reveals that Parfit's concept of a reason "differs from most philosophers'" (Broome, 2021: 299). This would be important (and surprising) if true, but as Broome himself acknowledges, there are good reasons to reject his interpretation (see Broome, 2021: 299, fn. 2).

4.1.6 Theoretical neutrality

A final point of clarification concerns the ecumenical nature of support-explanationism. There are at least three important respects in which it is metanormatively neutral.

First, unlike ought-explanationism and good-explanationism, support-explanationism is neutral with respect to the relative priority of reasons *vis-à-vis* both deontic and evaluative facts. Whereas ought-explanationism analyzes reasons in terms of the paradigmatic deontic notion—namely, ought—good-explanationism analyzes them in terms of the paradigmatic evaluative notion—namely, goodness. Each takes a substantive stand, then, on the relative priority of the deontic and/or evaluative over reasons. By contrast, support-explanationism avoids such a commitment: it analyzes reasons in terms of a notion—normative support—that clearly belongs to the same category of normative notions as that of a reason, and is neither uncontroversially evaluative nor uncontroversially deontic (cf. Berker, ms). Everyone should agree that reasons and support are in the same normative category, and that's all our view is committed to.³²

Second, support-explanationism is compatible with—but doesn't entail—pluralism concerning the grounds of support, according to which support-facts in different domains obtain in virtue of other (normative or non-normative) facts, which needn't be the same across domains, or even within a single domain. This is an issue to be settled by substantive first-order theorizing, which support-explanationism is neutral with respect to. Nonetheless, to illustrate, facts about support might obtain in virtue of facts about *value* in the objective practical realm (cf. Oddie & Menzies, 1992), *expected choiceworthiness* in the realm of substantive practical rationality (cf. Wedgwood, 2017; MacAskill & Ord, 2020), *epistemic justification* in the realm of substantive theoretical rationality (cf. Pryor, 2005), and *attitudinal pressure* in the realm of structural rationality (both practical and theoretical) (cf. Fogal, 2020). Each of these dimensions of evaluation can be said to correspond to a distinctive kind of pro tanto normative support, with few (if any) being reducible to the others, and in each case actions and attitudes can be ranked higher or lower, or better or worse, along that dimension, with the ranking being determined by how much normative support the relevant actions or attitudes enjoy.

Finally, support-explanationism is silent as to whether normative support can be analyzed in non-normative or naturalistic terms, and hence on whether a form of

³² The analysis of reasons in terms of support is thus neutral on whether support can be analyzed in terms of some other normative notion, such as value—it only rules out that support can be analyzed in terms of the notion of a reason (count noun), since that would be circular. Analyzing support in terms of value is nonetheless likely to face problems, including ones akin to those facing good-explanationism (cf. §3.2).

metanormative reductionism or naturalism about support (and so, given support-explanationism, also about normative reasons) is true. This, too, is a debate to be settled elsewhere.

4.2 Support-Explanationism vs. Reasons-First

Although we have already highlighted a number of differences between support-explanationism and Reasons-First, the reader might wonder what more might be said to motivate one over the other. Hence, in this subsection we'll highlight three further important differences and their consequences, before exploring three additional substantive implications of our view (§§4.3–5.2).

The first additional difference concerns the structure of the normatively fundamental facts. According to Reasons-First, those facts are what we earlier called 'reasons-facts—i.e., facts of the form [r is a reason for S to ϕ]. These facts involve at least three relata: (i) the fact, r , that is the reason, (ii), the subject, S , and (iii) the action, ϕ , that r is a reason to perform.³³ On support-explanationism, however, support-facts—i.e., facts of the form [There is (a certain amount of) support for S 's ϕ -ing], in which r is not a relatum—are more fundamental than reasons-facts. Although support-facts obtain (at least in part) in virtue of the fact that *is* the reason—i.e. r —they do not obtain in virtue of reasons-facts themselves. Given that disputes about relative fundamentality are substantive (and that facts with different constituents cannot be identical), the conflict between support-explanationism and Reasons-First is equally substantive.

All of this is relevant when evaluating one of the main contemporary sources of support for Reasons-First, which is that normative reasons are intuitively more fundamental than many other normative facts, such as facts about what we ought to do. This consideration is accommodated equally well, if not more so, by support-explanationism. The distinction between the fact, r , that *is* the reason, and the fact *that* r is a reason is crucial here. On support-explanationism, the facts that *are* normative reasons—i.e. the facts corresponding to r in reasons-facts of the form [r is a reason for S to ϕ], and that belong to the 'underlying level' discussed in §4.1.3—are simply the facts that help explain (in a distinctive way) why various actions or attitudes are supported to certain degrees, and thus help explain why certain actions ought to be performed or or attitudes ought to be had. However, those facts don't play the relevant explanatory role *because they are normative reasons*—instead, they are normative reasons because they play that explanatory role. For example, suppose there is (a certain amount of) support for giving your partner a gift because doing so would make them happy. It is the latter, non-normative fact about your partner's

³³ Further argument places can also be added; for instance, Scanlon (2014) includes an argument place for the agent's 'circumstances'.

happiness that explains the fact about normative support. The fact that there's support for giving your partner a gift is *not* explained by the fact that the fact that [*sic*] the action would make them happy is a reason to perform it. On the contrary, rather than playing an explanatory role, the reasons-fact—i.e. the fact *that the action would make your partner happy is a reason to perform it*—is itself explained: it obtains in virtue of the non-normative fact providing support for the relevant action. This point generalizes: reasons-facts obtain in virtue of facts about support, while particular facts about support obtain in virtue of the facts that *are* reasons. Although this vindicates the idea that reasons (that is: the facts that are reasons) are intuitively more fundamental than many normative facts, such as facts about what we ought to do, the relationship between reasons-facts and ought-facts is less straightforward. That's because neither reasons-facts nor ought-facts obtain in virtue of the other—instead, both obtain in virtue of support-facts. Support-facts are thus a 'common cause' (or, as it were, a 'common ground') of both reasons-facts and ought-facts.

A second notable difference between support-explanationism and Reason-First is that Reasons-First is like ought-explanationism (and unlike support-explanationism) in appealing to a non-gradable primitive: something can't be more of a reason than something else. A reason can be *stronger* or *weightier* than another, but that doesn't make it more of a reason—it just means that the normative support it provides is greater. This is notable since another feature of Reasons-First that has often been taken to be attractive is that it, unlike (e.g.) ought-explanationism, involves a normative primitive that is pro tanto and 'contributory'. For example, what Mark Schroeder (forthcoming) calls the 'classical argument' for Reasons-First focuses on the idea that "[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*" (6). We take it that reasons are more naturally thought of as things that *exert* 'moral force', not things that *are* forces. But there are two other things that we want to emphasize. First, the classical argument for Reasons-First isn't really an argument for taking the reasons-relation to be 'first', as opposed to some other contributory notion (or notions, for that matter). Support-explanationism, in particular, fits well with Schroeder's Rossian analogy, with the competing forces that determine what we ought to do being understood as competing amounts of normative support. Second, absent a plausible Reasons-First-friendly account of how a reason's 'strength' or 'weight' should be understood, the Rossian analogy actually speaks *against*, rather than in favor of, Reasons-First. For if the analogy is correct, and competing normative forces determine what we ought to do, then the relative strengths of those forces clearly play an important role in that determination.

The importance of giving an account of the weight of reasons is also illustrated by the popular slogan that 'you ought to do what you have most reason to do', which

has often been taken to support Reasons-First (e.g., Schroeder, 2018: 289–290). Given the close connection between support and reason (mass noun), support-explanationism makes perfect sense of this slogan: it simply expresses the idea that you ought to do what there is most support for doing. Without an account of the weight of reasons (and thus of what there is *most* reason to do), by contrast, Reasons-First actually struggles to make good sense of it. This is another putative source of support for Reasons-First that support-explanationism accommodates at least as well, if not more so.

Whether Reasons-First is plausible thus depends on whether the gradable notion of the strength or weight of a reason can be convincingly incorporated into the framework. We ourselves doubt it: not many Reasons-Firsters have tried to provide a theory of weight, and the proposals that exist are problematic. Though we won't consider the proposals in detail here, we'll briefly outline some concerns about them.

The first and most natural way to account for weight within the Reasons-First framework is simply to take the reasons-relation to include an argument-place for the weight of the reason, or the degree of support that it provides (cf. Skorupski, 2010: ch. 2). On this modified view, reasons-facts are of the form [*r* is a reason of weight *w* to ϕ], or [*r* is a reason that supports ϕ -ing to degree *d*]. However, as this modification makes clear, such facts involve two ingredients: (i) the degree to which ϕ -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 is evidence that that reasons-facts are *not* normatively fundamental, contrary to what this version of Reasons-First entails, and can instead be further analyzed in terms of these ingredients—just as support-explanationism predicts. For according to support-explanationism, for *r* to be a reason that supports ϕ -ing to a certain degree, *d*, *just is* for *r* to provide a corresponding amount of support, *d*, for ϕ -ing, and in that sense explain why it obtains. 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. Thus, the most natural way to account for weight within the Reasons-First view actually points away from Reasons-First, and towards support-explanationism instead.

It is notable that, perhaps because of this worry, several prominent Reasons-Firsters instead seek to analyze the weight of reasons in terms of yet more facts about reasons. The most well-known example of this strategy is due to Schroeder (2007), who analyzes the weight of reasons in terms of further reasons to 'place weight' on those reasons in deliberation. This has not been a popular view. For not only is it subject to apparent counter-examples (see e.g. Risberg, 2016), it is also highly counter-intuitive, at least once psychological and normative senses of '(placing) weight' are clearly distinguished. What Schroeder's strategy illustrates, however, is the pressure Reasons-Firsters face to try to account for the weight of reasons in terms of reasons

themselves—for otherwise it is difficult to maintain the idea that reasons are normatively ‘first’.³⁴ And until such an account has been provided, it is also difficult to accommodate various putatively platitudinous sources of support for Reasons-First, such as Schroeder’s Rossian analogy and the slogan that ‘you ought to do what you have most reason to do’. As noted above, these are considerations that support-explanationism accommodates with ease.

A third important difference between Reasons-First and support-explanationism is that the latter view does not entail an ambitious ‘Support-First’ approach, according to which every other normative notion—goodness, rightness, virtue, etc.—can ultimately be explained in terms of normative support. While such an approach would be more plausible than Reasons-First (in part for the reasons mentioned above), it would face versions of the many familiar problems for imperialistic reductionist programs.³⁵ Instead, support-explanationism only entails that the notion of normative support is *prior* to the notion of a normative reason, and as a result, that reasons-facts are not normatively first. This is compatible with thinking that there are other normative notions that are at least as fundamental as that of support.

4.3 Force vs. Weight

A further advantage of analyzing reasons in terms of support is that doing so avoids some of the limitations of the simple weight/weighting metaphor highlighted in §3.1. Our view instead suggests a more productive metaphor, one which appears in Schroeder’s presentation of the ‘classical argument’ above: that of *force*. Like forces in general, normative support has something akin to magnitude (which can vary) and direction (for/against), and we can usefully distinguish between component and resultant support on analogy with component and resultant forces. All of this helps make straightforward sense of various force-like notions employed by normative theorists, including not just support but also opposition and undermining, which are

³⁴ Another Reasons-First-friendly account of weight is proposed by Scanlon (2014), who analyzes weight in counterfactual terms. His view is threatened by the same problems that plague counterfactual analyses in general (cf. Shope, 1979), and he provides no reason to think that his view avoids them (for further criticism of Scanlon’s view, see Schroeder (2015)). The notion of a reason’s weight also plays a crucial role in Lord’s (2018) defense of Reasons-First, but he says little about how it should be understood. Lord says he finds “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). He also invokes a view of intensifiers and attenuators of weight that resembles Schroeder’s view, on which they are understood as reasons to place more or less weight on certain reasons (2018: 33).

³⁵ For example, an analysis of value in terms of supported pro-attitudes would face a version of the well-known ‘wrong kind of reason(s)’-problem (cf. Rabinowicz & Rønnow-Rasmussen, 2004). Moreover, as Wodak (2020a) argues, it’s not clear why we should expect any normative property or relation to be ‘first’ in the first place. For further criticism of X-first views, see Reisner (ms).

the two main ways in which support is defeasible (cf. §4.1.3). The force metaphor allows for different ways of understanding the nature of such interactions, as well as the possibility that they are non-monotonic and non-additive. In cases involving opposition, for instance, we can think of the initial component normative force (i.e. pro tanto support) as remaining constant even though the overall resultant force (i.e. net support) is altered by competing component forces, whereas in cases involving undermining we can think of the initial component force as itself being weakened (and, at the limit, eliminated), with the resultant force changing as a result. Hence, even if the resultant force (or net support) ends up being the same in both cases, the explanation of why each ends up that way differs, which allows one to straightforwardly distinguish opposition and undermining. Similar things can be said about other ways in which the strength or significance of a reason may be altered, such as by being ‘intensified’.

Support-explanationism thus provides a ready answer to a question posed by Hawthorne and Magidor, concerning whether “there are more subtle ways to expand the explanation[ist] view” in a way that goes beyond the simple model of weights and weighing (2018: 134). They don’t themselves provide a positive view, but they do highlight some important work that might serve as inspiration. In particular, they mention work in the philosophy of science on modeling causal forces (and their interactions), as well as work in epistemology that employs inference graphs (e.g., Pollock, 1995; Berker, 2015), where such graphs “encode directed patterns and strength of support between ‘prima facie’ reasons and various propositions as well as phenomena of defeat (nodes on inference graphs get marked as defeated or undefeated)” (Hawthorne & Magidor, 2018: 134). Both of these literatures hold promise when it comes to modeling relations of normative support—indeed, the epistemological literature making use of inference graphs in modeling relations of epistemic support can be seen as a special case of the more general project of modeling relations of normative support of all kinds.

5. Support-explanationism: Two applications

Before concluding, we’ll consider the implications of our view for two additional issues in metanormativity: the question of whether Reasons-First should be especially attractive to adherents of ‘non-naturalism’ about normativity (answer: no), and the question of whether explanationist views can accommodate what Daniel Wodak (2020b) calls ‘redundant reasons’ (answer: yes).

5.1 Support-explanationism, Reasons-First, and non-naturalism

Understanding reasons-facts as facts about explanations of normative support allows us to account for an observation made by Selim Berker (2019) about a way in which reasons-facts seem unusual. What Berker suggests is that particular normative

reasons-facts don't always obviously cry out for 'grounding explanations' by reference to natural (or descriptive or whatever) facts. If this is correct, reasons-facts differ in this regard from other particular normative facts, such as particular facts about what is right and wrong, which are widely assumed to always have such explanations. Berker puts the point as follows:

[W]hat about facts of the form [Fact r] is a reason for agent S to ϕ ? Must such facts always obtain in virtue of other facts? That they might not seems very much to be a live theoretical option. While it is mystery mongering to hold that the wrongness of what I did last week is a brute fact that cannot be explained, it is not necessarily mystery mongering to hold that nothing explains why what it's like to be in agony is a reason for me to avoid future agony. (Berker, 2019: 931; notation adjusted)

Based on this, Berker proposes that the Reasons-First view should be especially attractive to non-naturalists – roughly, those who take normative facts to be 'different in kind from', or 'irreducible to', natural facts. According to Berker, the best version of non-naturalism implies that at least some particular normative facts are not grounded in natural facts.³⁶ And reasons-facts, he thinks, are good candidates for being such facts – as he puts it, "the reason relation is an excellent candidate for a non-resultant normative relation" (2019: 931).

However, support-explanationism provides an alternative account of the possibility that reasons-facts are unusual in this way. If reasons-facts are *themselves* facts about what explains what, asking what explains why r is a reason to ϕ is akin to asking what explains why r explains why ϕ -ing enjoys (some amount of) normative support. It is thus akin to asking for an explanation of an explanation, which is quite different from asking for an explanation of a non-explanatory fact, whether normative or otherwise. (Compare the question 'Why did the bridge collapse?' to 'Why does the fact that there was an earthquake explain why the bridge collapsed?', or the question 'Why did you run?' to 'Why does the fact that you want to get in shape explain why you ran?') This, we submit, is why reasons-facts might appear to be exceptions to the general rule that particular normative facts are always at least partially explained by natural facts – the rule is not usually meant to include facts about normative explanations themselves. This isn't to say that facts about explanations can't themselves have explanations (cf. Bennett, 2011). Our point is just that it's unsurprising that facts like [The fact that donating to charity helps alleviate

³⁶ We disagree. Instead, we think non-naturalism is best understood as the view that one or both of the following are true: (i) that fundamental normative *principles* cannot be explained by natural facts (cf. Fogal and Risberg, 2020), and (ii) that *what it is* to be right, wrong, normatively supported, or to have any other normative property cannot be explained in fully naturalistic terms (cf. Leary, forthcoming).

suffering helps explain why you should donate to charity] are intuitively different from facts like [You should donate to charity], and the former don't always cry out for explanation in the same way, or to the same degree, as the latter.

Importantly, our diagnosis of the (at least purportedly) unusual status of reasons-facts is compatible with them being analyzed in terms of normative support, and with particular support-facts obtaining (at least partly) in virtue of particular natural facts. Thus, we submit, Berker's suggestion that some reasons-facts might lack grounding explanations does not provide reason for non-naturalists to be Reasons-Firsters, or otherwise take reasons-facts to be unanalyzable.

5.2 Redundant reasons

Support-explanationism also provides a straightforward account of what Wodak (2020b) calls *redundant reasons*—facts which count in favor of certain actions but cannot explain why we ought to do them. Wodak argues that the possibility of such reasons is incompatible with ought-explanationism. He provides the following example:

The Gods Reward the Rational. You face a choice between φ and φ^* . All things considered, you ought to do φ (for reasons r , r^* , and r^{**}). The Gods appear before you and tell you that, if you do whatever you ought to do, They will reward you with an eternity of pleasure in the afterlife. (Wodak, 2020b: 268; notation adjusted)

The problem, Wodak says, stems from two apparent truths: first, the fact that the Gods will reward you if you φ —call that fact [*Reward*]—is a reason for you to φ ; but, second, while it is true that you ought to φ —call that fact [*Ought*]—it is not the case that [*Reward*] explains [*Ought*]. Instead:

[*Reward*] is true only because of [*Ought*]: [*Ought*] cannot then be true because of [*Reward*]. In this sense, [*Reward*] is a 'redundant reason': it counts in favour of doing φ but cannot explain [*Ought*]. (Wodak, 2020b: 268; notation adjusted)

However, even if redundant reasons pose a potential problem for ought-explanationism, they are easily accommodated by support-explanationism. (Other explanationist views that analyze reasons in terms of a gradable normative notion can accommodate redundant reasons as well, but we'll focus on support-explanationism.) In *The Gods Reward the Rational*, facts r , r^* and r^{**} together provide *some* amount of support, d , for φ . Because φ^* is supported to some lesser degree, you ought to φ . When the Gods show up, their promise provides *additional* support for φ . In other words, [*Reward*] explains why there is (now) additional support for φ —i.e.,

to a higher degree than just d . Support-explanationism thus entails that [*Reward*] is a reason to ϕ , as desired.

Now, Wodak may be right that [*Reward*] is explained, at least in part, by the fact that ϕ is antecedently well-supported, and better supported than the alternatives. But it's unmysterious, and compatible with the asymmetry of explanation, that a fact about an amount of something can help explain a subsequent increase in the amount of that same thing. For example, if somebody has a lot of money, this can help explain future monetary gains – i.e., why her net wealth increases. This doesn't entail that her having a lot of money explains why she has a lot of money, but only that the fact that she started with a certain amount of money explains (together with other facts) why she tends up with a larger amount of money. She starts out rich and ends richer. Similarly, in *The Gods Reward the Rational*, ϕ starts out well-supported and ends with even more support.

A version of Wodak's objection that targets support-explanationism specifically thus fails. It would go something like this: suppose r provides a certain amount of (pro tanto) support, d , for ϕ -ing. The Gods tell you that if you do whatever there is d support for doing, they will reward you. As it turns out, the fact that they make this promise provides exactly the same amount of support, d , for ϕ -ing. Support-explanationism makes sense of this: initially, there is only one source of support for ϕ -ing (i.e., r), and when the Gods make their promise, that constitutes a second, equally strong source of support for ϕ -ing.

6. Concluding Remarks

A recurring theme in this paper has been that support-explanationism accommodates many considerations that have been taken to favor rival views while avoiding their flaws. Among other things, it captures the close connection between reasons and explanation without facing the objections that plague alternative explanationist views, such as the problem of outweighed reasons (see §3). Similarly, it accommodates the main motivations for the Reasons-First program – such as the slogan that 'you ought to do what you have most reason to do', and the insight that contributory normative notions should be seen as fundamental – while avoiding its problems (see §4). If this is correct, why has support-explanationism been overlooked? And why has Reasons-First instead become such an influential position in the metanormative debate? While we cannot do full justice to these questions here, we'll offer some partial and speculative answers to them in closing.

The basic idea behind the Reasons-First view is clear enough, and at first glance, this idea might seem highly attractive. Importantly, however, the central theses of this view are often expressed in ways that are, in crucial respects, less-than-fully precise. As noted, this is true even of the opening statement of Scanlon's 1998 book, *What We Owe to Each Other* – one of the most important works in the Reasons-First

tradition. To some extent imprecision in expression is understandable: after all, in ordinary language, terms like ‘reason’ and ‘reasons’ are arguably polysemous, context-sensitive, and combine with other expressions (such as ‘gives’, ‘is’ and ‘provides’) in ways that are both complex and confusing. As we hope to have shown, however, paying attention to such seemingly insignificant details can have a major impact on how one answers many of the most fundamental metanormative questions, including those that concern the nature of normative reasons themselves. Besides illustrating the importance of clarifying one’s terminology, methodology, and subject matter when formulating one’s metanormative views, it also raises the question of whether the apparent attraction of the Reasons-First view is really genuine. We suspect not – instead, as already suggested, it may be largely the result of Reasons-First not being clearly distinguished from views that are subtly but importantly different, such as support-explanationism (and the related but optional ‘Support-First’ program). We believe, however, that once their differences are made clear, support-explanationism emerges as the preferable alternative.³⁷

³⁷ For valuable feedback, we thank Joseph Cunningham, Matti Eklund, Chris Howard, Nathan Howard, Zoë Johnson King, Alex King, Christos Kyriacou, Nicholas Laskowski, Stephanie Leary, Errol Lord, Barry Maguire, Victor Moberger, Jonas Olson, Andrew Reisner, Knut Olav Skarsaune, Daniel Star, Folke Tersman, Daniel Wodak, Alex Worsnip, and Michael Young, as well as audiences at the Boston University Ethics Seminar, the Cyprus Metaethics Workshop, the Higher Seminar in Practical Philosophy at Uppsala University, the Normativity Reading Group at the University of Pennsylvania, the Slippery Slope Normativity Summit, and the Virtual Metaethics Colloquium. We owe an extra debt of gratitude to Selim Berker and an anonymous reviewer for this journal for multiple rounds of helpful feedback. Work on this paper was supported by Grant 2020-01955 from the Swedish Research Council.

References

- Bader, R. (2016). Conditions, modifiers and holism. In E. Lord & B. Maguire (Eds.), *Weighing Reasons* (pp. 27–55). New York: Oxford University Press.
- Bedke, M. (2008). A new theory of Humean reasons? A critical note on Schroeder's hypotheticalism. *Journal of Ethics and Social Philosophy*, 3, 1–5. doi: 10.26556/jesp.v2i3.119
- Behrens, J. & DiPaolo, J. (2011). Finlay and Schroeder on promoting a desire. *Journal of Ethics and Social Philosophy*, 6, 1–7. doi: 10.26556/jesp.v6i1.146
- Bennett, K. (2011). By our bootstraps. *Philosophical Perspectives*, 25, 27–41. doi: 10.1111/j.1520-8583.2011.00207.x
- Berker, S. (2007). Particular reasons. *Ethics*, 118, 109–139. doi: 10.1086/521586
- Berker, S. (2015). Coherentism via graphs. *Philosophical Issues*, 25, 322–352. doi: 10.1111/phis.12052
- Berker, S. (2018). The unity of grounding. *Mind*, 127, 729–77. doi: 10.1093/mind/fzw069
- Berker, S. (2019). The explanatory ambitions of moral principles. *Noûs*, 53, 904–936. doi: 10.1111/nous.12246
- Berker, S. (2020). Quasi-dependence. In R. Shafer-Landau (Ed.), *Oxford Studies in Metaethics, Vol. 15* (pp. 195–218). New York: Oxford University Press.
- Berker, S. (manuscript). The deontic, the evaluative, and the fitting.
- Broome, J. (2004). Reasons. In J. Wallace, M. Smith, S. Scheffler & P. Pettit (Eds.), *Reason and Value* (pp. 28–55). Oxford: Oxford University Press.
- Broome, J. (2013). *Rationality Through Reasoning*. Malden: Wiley-Blackwell.
- Broome, J. (2021). Giving reasons and given reasons. In J. McMahan, T. Campbell, J. Goodrich, and K. Ramakrishnan (Eds.), *Principles and Persons: The Legacy of Derek Parfit* (pp. 209–308). Oxford: Oxford University Press.
- Brunero, J. (2013). Reasons as explanations. *Philosophical Studies*, 165, 805–824. doi: 10.1007/s11098-012-9982-8
- Bykvist, K. (forthcoming). Taking values seriously. *Synthese*. doi: 10.1007/s11229-021-03071-1
- Chappell, R. (2012). Fittingness: the sole normative primitive. *Philosophical Quarterly*, 62, 684–704. doi: 10.1111/j.1467-9213.2012.00075.x
- Crisp, R. (2005). Value, reasons, and the structure of justification: how to avoid passing the buck. *Analysis*, 65, 80–85. doi: 10.1111/j.1467-8284.2005.00528.x
- Dancy, J. (2004). *Ethics Without Principles*. Oxford: Oxford University Press.
- Dancy, J. (2015). Reasons for Broome. In I. Hirose & A. Reisner (Eds.), *Weighing and Reasoning: Themes from the Philosophy of John Broome* (pp. 177–188). Oxford: Oxford University Press.
- Dorr, C. (2016). To be F is to be G. *Philosophical Perspectives*, 30, 39–134. doi: 10.1111/phpe.12079
- Dowell, J. L. (2020). Finlay's methodology: synthetic, not analytic. *Analysis*, 80, 102–110. doi: 10.1093/analys/anz081

- Fine, K. (2012). Guide to ground. In F. Correia & B. Schnieder (eds.), *Metaphysical Grounding* (pp. 37–80). Cambridge: Cambridge University Press.
- Finlay, S. (2014). *Confusion of Tongues*. New York: Oxford University Press.
- Finlay, S. (2019). A “good” explanation of five puzzles about reasons. *Philosophical Perspectives*, 33, 62–104. doi: 10.1111/phpe.12126
- Fogal, D. (2016a). Reasons, reason, and context. In E. Lord & B. Maguire (Eds.), *Weighing Reasons* (pp. 74–103). doi: 0.1093/acprof:oso/9780199315192.001.0001
- Fogal, D. (2016b). Review of Stephen Finlay’s *Confusion of Tongues: A Theory of Normative Language*. *Ethics*, 127, 281–288. doi: 10.1086/687342
- Fogal, D. (2018). Deflationary pluralism about motivating reasons. In V. Mitova (Ed.), *The Factive Turn in Epistemology* (pp. 193–218). Cambridge: Cambridge University Press. doi: 10.1017/9781316818992
- Fogal, D. (2020). Rational requirements and the primacy of pressure. *Mind*, 129, 1033–1070. doi: 10.1093/mind/fzz038
- Fogal, D. & Risberg, O. (2020). The metaphysics of moral explanations. In R. Shafer-Landau (Ed.), *Oxford Studies in Metaethics, Vol. 15* (pp. 170–194). New York: Oxford University Press.
- Fogal, D. & Worsnip, A. (forthcoming). Which reasons? Which rationality? *Ergo*.
- Gardner, J. & Macklem, T. (2004). Reasons. In J. Coleman, K. Himma, & S. Shapiro (Eds.), *The Oxford Handbook of Jurisprudence and Philosophy of Law* (pp. 440–475). New York: Oxford University Press.
- Glazier, M. (2016). Laws and the completeness of the fundamental. In M. Jago (Ed.), *Reality Making* (pp. 11–37). Oxford: Oxford University Press
- Hawthorne, J. & Magidor, O. (2018). Reflections on the ideology of reasons. In D. Star (Ed.), *The Oxford Handbook of Reasons and Normativity* (pp. 113–142). New York: Oxford University Press.
- Hieronymi, P. (2005). The wrong kind of reason. *Journal of Philosophy*, 102, 437–457. doi: 10.5840/jphil2005102933
- Howard, C. (2018). Fittingness. *Philosophy Compass*, 13. doi: 10.1111/phc3.12542
- Howard, C. (2019). The fundamentality of fit. In R. Shafer-Landau (Ed.), *Oxford Studies in Metaethics, Vol. 14* (pp. 216–236). Oxford: Oxford University Press.
- Kearns, S. & Star, D. (2008). Reasons: explanations or evidence? *Ethics*, 119, 31–56. doi: 10.1086/592587
- Kearns, S. & Star, D. (2009). Reasons as evidence. In R. Shafer-Landau (Ed.), *Oxford Studies in Metaethics, Vol. 4* (pp. 215–242). Oxford: Oxford University Press.
- Leary, S. (forthcoming). What is normative non-naturalism? *Ergo*.
- Lord, E. (2018). *The Importance of Being Rational*. New York: Oxford University Press.
- MacAskill, W. & Ord, T. (2020). Why maximize expected choice-worthiness? *Nous*, 54, 327–353. doi: 10.1111/nous.12264

- Maguire, B. (2016). The value-based theory of reasons. *Ergo*, 3. doi: 10.3998/ergo.12405314.0003.009
- Maguire, B. & Snedegar, J. (2021). Normative metaphysics for accountants. *Philosophical Studies*, 178, 363–384. doi: 10.1007/s11098-020-01435-w
- Nebel, J. (2019). Normative reasons as reasons why we ought. *Mind*, 128, 459–484. doi: 10.1093/mind/fzy013
- Oddie, G. & Menzies, P. (1992). An objectivist's guide to subjective value. *Ethics*, 102, 512–533. doi: 10.1086/293422
- Parfit, D. (2011). *On What Matters, Vol. 1*. Oxford: Oxford University Press.
- Pollock, J. (1995). *Cognitive Carpentry: A Blueprint for How to Build a Person*. Cambridge, Mass.: MIT Press.
- Pryor, J. (2005). There is immediate justification. In M. Steup & E. Sosa (Eds.), *Contemporary Debates in Epistemology* (pp. 181–202). West Sussex: Wiley-Blackwell.
- Rabinowicz, W. & Rønnow-Rasmussen, T. (2004). The strike of the demon: on fitting pro-attitudes. *Ethics*, 114, 391–423. doi: 10.1086/381694
- Raz, J. (1999). *Practical Reason and Norms*. Oxford: Oxford University Press.
- Reisner, A. (2018). Two theses about the distinctness of practical and theoretical normativity. In C. McHugh, J. Way, & D. Whiting (Eds.), *Normativity: Epistemic and Practical* (pp. 221–240). Oxford: Oxford University Press.
- Reisner, A. (manuscript). *The Pragmatic Foundations of Theoretical Reason*.
- Risberg, O. (2016). Weighting surprise parties: some problems for Schroeder. *Utilitas*, 28, 101–107. doi: 10.1017/s0953820815000394
- Rosen, G. (2015). Real definition. *Analytic Philosophy*, 56, 189–209. doi: 10.1111/phib.12067
- Rosen, G. (2017). What is a moral law? In R. Shafer-Landau (Ed.), *Oxford Studies in Metaethics, Vol. 12* (pp. 135–159). New York: Oxford University Press.
- 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.
- Scanlon, T. M. (2014). *Being Realistic about Reasons*. Oxford: Oxford University Press.
- Schaffer, J. (2017). Laws for metaphysical explanation. *Philosophical Issues*, 27, 302–321. doi: 10.1111/phis.12111
- Schroeder, M. (2007). *Slaves of the Passions*. Oxford: Oxford University Press.
- Schroeder, M. (2009). Buck-passers' negative thesis. *Philosophical Explorations*, 12, 341–347. doi: 10.1080/13869790903067725
- Schroeder, M. (2015). Being Realistic About Reasons, by T.M. Scanlon. *Australasian Journal of Philosophy*, 93, 195–198. doi: 10.1080/00048402.2014.910241
- Schroeder, M. (2018). Getting perspective on objective reasons. *Ethics*, 128, 289–319. doi: 10.1086/694270
- Schroeder, M. (2021). *Reasons First*. Oxford University Press.

- Schroeder, M. (forthcoming). The fundamental reason for Reasons Fundamentalism. *Philosophical Studies*. doi: 10.1007/s11098-020-01572-2
- Schroeter, L. & Schroeter, F. (2009). Reasons as right-makers. *Philosophical Explorations*, 12, 279–296. doi: 10.1080/13869790903067691
- Setiya, K. (2014). What is a reason to act? *Philosophical Studies*, 167, 221–235. doi: 10.1007/s11098-012-0086-2
- Shope, R. (1978). The conditional fallacy in contemporary philosophy. *Journal of Philosophy*, 75, 397–413. doi: 10.2307/2025564
- Skorupski, J. (2002). The ontology of reasons. *Topoi*, 21, 113–124. doi: 10.1023/A:1014801003579
- Skorupski, J. (2010). *The Domain of Reasons*. Oxford: Oxford University Press.
- Snedegar, J. (2017). *Contrastive Reasons*. Oxford: Oxford University Press.
- Snedegar, J. (2018). Reasons for and reasons against. *Philosophical Studies*, 175, 725–743. doi: 10.1007/s11098-017-0889-2
- Snedegar, J. (2021). Competing reasons. In J. Brown & M. Simion (Eds.), *Reasons, Justification, and Defeat* (pp. 247–268). Oxford: Oxford University Press.
- Thomson, J. J. (2008). *Normativity*. Peru: Open Court.
- Tucker, C. (forthcoming). Weighing reasons against. In R. Shafer-Landau (Ed.), *Oxford Studies in Metaethics*.
- Way, J. (2017). Reasons as premises of good reasoning. *Pacific Philosophical Quarterly*, 98, 251–270. doi: 10.1111/papq.12135
- Wedgwood, R. (2017). Must rational intentions maximize utility? *Philosophical Explorations*, 20, 73–92. doi: 10.1080/13869795.2017.1356352
- Wedgwood, R. (manuscript). Values First.
- Whiting, D. (2018). Right in some respects: reasons as evidence. *Philosophical Studies*, 175, 2191–2208. doi: 10.1007/s11098-017-0954-x
- Wodak, D. (2020a). Who's on first? In R. Shafer-Landau (Ed.), *Oxford Studies in Metaethics, Vol. 15* (pp. 49–71). New York: Oxford University Press.
- Wodak, D. (2020b). Redundant reasons. *Australasian Journal of Philosophy*, 98, 266–278. doi: 10.1080/00048402.2019.1620301
- Worsnip, A. (2020). Resisting relativistic contextualism: on Finlay's *Confusion of Tongues*. *Analysis*, 80, 122–131. doi: 10.1093/analys/anz082
- Woodward, J., and Ross, L. (2021). Scientific explanation. In E. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy (Summer 2021 Edition)*, URL = <<https://plato.stanford.edu/archives/sum2021/entries/scientific-explanation/>>.