

*The Explanation View***3.1 Varieties of the Explanation View**

A reductionist approach to normative reasons can be characterised as a proposal to reduce to and explain normative reasons in terms of some other, presumably more robustly graspable, properties. Of course, not all sorts of reduction will be deemed plausible. So how are we to pick out plausible candidates for such a reduction? One common, if often only implicit, strategy among reductionist views in general is to reduce/explain an x by focusing on what appear to be central functions of the ordinary concept of x , the idea being to work out individually necessary and jointly sufficient conditions (within a substantive analysis) for something to be an x in terms of what appears to refer to the central functions of the ordinary concept of x .

This seemed to be the strategy to which proponents of the Reasoning view were attracted. One way of interpreting the Reasoning view is to start with a plausible idea that being able to play a role in good/correct reasoning is an important aspect of our concept of *reasons to act/hold an attitude*. Thus, it appears only natural to try to work out a full theory of normative reasons by an appeal to this aspect of our concept of *reasons to act/hold an attitude*. However, we also observed that a full-blown theory that attempts to define normative reasons only in terms of a role they can play in good reasoning faced some serious problems and is not fully satisfactory. But does it mean that all reductionist proposals about normative reasons are deemed to fail? Not necessarily, of course. But if not good/fitting reasoning, what else could plausibly be appealed to in order to explain normative reasons? Where do we even start looking for alternative reductionist accounts? How are we supposed to proceed?

One natural thing to do at this stage is to look once more into functions of our concept of *reasons to act/hold an attitude* (e.g. the concept that

corresponds to normative reasons). Surely, in our common-sense understanding, reasons should play a role in reasoning; in a sense, reasoning is made of reasons (see Grice 2001 and Chapter 2). But there is something else, it seems. Remember our initial normative versus motivating/explanatory reasons distinction. The reason why I ate chips was that I was hungry. It is part of an explanation of why I ate chips. It provides a descriptive explanation of why I did what I did. One thought is that it is natural to think that explaining is also what normative reasons do, except that the explanation here is not about descriptive but normative facts. Of course, the business of normative reasons is not descriptive explanation, but it might nevertheless be part of another sort of explanation. Indeed, considerations that we identify as normative reasons seem to be part of normative explanation in a sense to be seen. Think, for instance, of facts about oughts. It is natural to think that in many cases, if not in all of them (depending on whether you think that the principle of sufficient reason is true, for instance), when it is the case that S ought to F, it is meaningful to ask why S ought to F. A natural understanding of such a question is to see it as a question about what makes it right that S F-s. Think, for instance, of Karl who ought to call his grandmother. We can meaningfully ask: why is it the case that Karl ought to call his grandmother? A meaningful answer to this sort of question will typically consist in our locutor pointing to the relevant considerations, or a list of considerations that, taken together, make it the case that Karl ought to call his grandmother. Perhaps it will contain something like, that she is really sick, that he hasn't called her for a while, that it would really make his grandmother happy that he calls, and so on. The same kind of reasoning also seems to apply if you replace 'ought' in Karl's case with 'good in some respect'. It would be good that Karl calls his grandma, and the relevant considerations will be used to explain that. But what are these considerations? Well, it appears only natural to think of them as normative reasons. Indeed, when asked what makes it right or good and so on that Karl calls his grandma, the balance of normative reasons comes naturally to mind. It is because there are all these normative reasons that it is right/good that Karl calls his grandma, indeed, that Karl ought to call her. If this sort of reasoning is intuitive, and it seems it is, a theorist of reasons can start building an alternative reductionist account of normative reasons from here. That is, one may take seriously this natural role of normative reasons in normative explanations and define normative reasons in terms of it. And this is exactly what various recently increasingly popular Explanation views of normative reasons propose to do.

There is a variety of explanationist accounts of normative reasons¹ and this chapter will consider the most prominent ones in some detail. The common element of these views is, as we just saw, an appeal to a sort of explanation. Unsurprisingly then, the differences among these views will consist in how exactly the details of ‘a sort of’ are worked out. Most of the existing versions of the explanation views can be classified as being either about explanations of (aspects) of why someone ought to F (which is, arguably, equivalent to why some F-ing is (in some respect) right/correct/fitting) or about explanation of (aspects of) why some F-ing is good. We can call the first group the ‘Deontic Explanation views’, given their focus on the deontic facts (oughts), and the second the ‘Axiological Explanation’ views, given their focus on values. In what follows, we first explore the Deontic Explanation views (first arguments in favour and then objections), and then we turn to the Axiological Explanation views (again, starting with positive arguments before exploring the most serious objections).

3.2 The Deontic Explanation View: Ought, Explanation, and Weighing Explanation

The general approach. According to one prominent version of the Explanation approach, roughly understood, normative reasons can be defined by appeal to explanation and ought. This section focuses on variants of this approach.

Reasons as (parts of) explanation of oughts. Probably the best-known proponent of this sort of approach within contemporary debates is John Broome (2004, 2013, 2015, 2018). According to Broome, while ‘ought’ in the relevant (central) sense cannot be defined, normative reasons can be defined (Broome only specifies the sense of ‘ought’ that he is interested in as ‘normative, owned, unqualified and prospective’; cf. Broome 2013: 45). He undertakes to provide a functional definition of normative reasons – that is, a definition that is supposed to appeal to the central role of reasons (cf. Broome 2013: 53–54). The role in question is to figure in an

¹ I will use the terms ‘such and such Explanation view (of reasons)’ and ‘such and such explanationism’ interchangeably in what follows. Of course, the views we discuss here should be clearly distinguished from more specific accounts about the nature of evidential (or indeed justificatory) support in contemporary epistemology (see, for example, McCain 2013, 2018, and further references therein). The term ‘explanationism’ in these debates is often reserved for a view that opposes probabilistic accounts of evidential support and insists instead on the role of evidence/justifiers in inferences to best explanation. While there might be, and arguably are, connections between the views we discuss here and explanationism in contemporary debates about the nature of evidential support, the focus here is more general, so to say, on normative reasons in general.

explanation of deontic facts, where a deontic fact just is a fact about an ought. More specifically, a deontic fact is a fact that, for example, S ought to F. For instance, the fact that Agnes ought to feed her cats is a deontic fact involving Agnes and her cats.

Now, Broome provides two specific ways in which something can figure in an explanation of a deontic fact: something can explain why one ought to F, and something can be part of a 'weighing explanation' of why one ought to F. To see the difference between the two, we need to look into the details of what Broome means by 'weighing explanation' (and, in particular, why it is not a (mere) explanation of why S ought to F). Before that, note first that Broome doesn't commit himself to the view that these are the only two ways of how something can figure in an explanation of a deontic fact, and second, that Broome doesn't commit himself to the idea that all normative reasons have to correspond to one or another of these ways of playing a role in an explanation of a deontic fact. According to Broome, there are other possible ways in which something can figure in an explanation of a deontic fact, and it is in principle possible that some sort of normative reasons are associated with other ways in which something can figure in an explanation of a deontic fact (see Broome 2013: 61–62). That being said, however, Broome doesn't really provide any other example of a potential way in which a normative reason could figure in an explanation of a deontic fact. In a more recent article, Broome (2018: 12) does mention the conception of reasons that is present in John Harty's *Reasons as Defaults* book (Harty 2012) as another sort of subsidiary reasons, where subsidiary reasons can be seen as one general class of normative reasons, including reasons that play a role in weighing explanations and potentially other reasons, but he doesn't really elaborate on details of why Harty's 'reasons' are distinct from weighing-explanation reasons. He does provide an example of a plausible candidate for another way in which something could figure in an explanation of a deontic fact without being itself a normative reason. Namely, he mentions the cancelling role of some considerations. An example of this is when the fact that one releases a friend from an obligation cancels what would have otherwise been a reason for the friend to keep the promise (see Broome 2013: 62). Thus, strictly speaking, what we find in Broome's work is a theory of two sorts of normative reasons that are associated with two sorts of explanations of deontic facts, but not a full-blown theory of all possible sorts of normative reasons or explanations of deontic facts. Let us examine both of these in due order.

First, Broome observes that there are cases where a fact, *r*, explains why S ought to F. Broome assumes a simple and somewhat plausible view

about explanation according to which the explanation relation just is the reverse of the because relation. When *a* explains *b*, we assume that *b* is the case because *a* and vice versa. Broome also endorses the view that both *explanans* and *explanandum* in any genuine explanation are factive. That is, if *r* explains *q*, then both *r* is the case and *q* is the case. Roughly put, only facts can explain and be explained. Finally, without providing us with a full-blown theory of explanation, Broome notices that many alternative explanations may hold of the same thing, and that this is not really problematic. For example, it may well be the case that the room is dark because it's night and also be the case that the room is dark because the light bulb has burned out, and finally it may also be the case that the room is dark because I forgot to change the light bulb. There are in this case three alternatives and individually satisfying explanations of why the room is dark. But these three are not really in conflict. Which of these three explanations will get to be called 'the' explanation of why the room is dark may depend on purely contextual aspects. Note, nevertheless, that Broome does think that there is one complete explanation that takes all these relevant facts into account.

Now turning back to reasons, the first sort of normative reasons that Broome proposes to define is *pro toto* reasons, the normative reasons that gain their normative aspect because they are *explanations* of why one ought to *F*. The reason relation in itself is not normative, according to Broome; it is merely a relation of explanation, and there is nothing normative in the relation of explanation. The normative aspect comes from it being an explanation of a deontic fact. He writes: '[I]n "X is the reason why you ought to *F*", the "reason why" is so closely attached to the normative "ought" that the two tend to slide into each other. [. . .] The "reason why" (meaning explanation) bumps into the normative "ought", yielding a normative sense of "a reason" that combines the meaning of both' (Broome 2013: 50). Thus, one sort of normative reasons are merely explanations of some deontic facts. Their normative aspect is directly inherited from the normativity of deontic facts (e.g. oughts). (Note also that because *pro toto* reasons are mere explanations, they also inherit, in Broome's view, the relevant aspects of explanations *tout court*. For instance, there might be two *pro toto* reasons for the same deontic fact without there being any conflict between those two.) Here is Broome's definition of *pro toto* reasons:

(Pro toto normative reasons): 'A *pro toto* reason for *N* to *F* is an explanation of why *N* ought to *F*.' (Broome 2013: 50)

Of course, Broome is well aware that accounting for *pro toto* reasons only is not satisfactory for a theory of reasons (even for a non-exhaustive theory). For there are normative reasons for one to *F* when it is not the case that one ought to *F*. We saw earlier that Broome assumes the factivity of explanation. If all normative reasons were always explanations of deontic facts, there wouldn't be any outweighed reasons. The looming deadline is a reason for me to stay at home and continue to work; the fact that my kids are waiting for me to pick them up in a few minutes is a reason for me to run out; the second consideration clearly outweighs the first, but the first still continues to count as a reason to stay and work – it is not thereby cancelled. If reasons were always mere explanations of oughts, then such reasons would not be possible (assuming the factivity of explanation). However, ignoring outweighed reasons would be a serious cost for any theory of reasons. Thus, Broome introduces a second and somewhat more sophisticated sort of normative reasons.

The second sort of normative reasons, according to Broome, are what he calls *pro tanto* normative reasons. In Broome's view, (some) deontic facts have 'normative weighing explanation'. Note that Broome considers the question of whether all deontic facts have a corresponding weighing explanation, and leaves open the possibility that some deontic facts may not have a weighing explanation – for example, the fact that one ought not to have contradictory beliefs, conditional on evidentialism about belief (the view that only evidence determines what one should believe) being correct (cf. Broome 2013: 58–60). A normative weighing explanation is analogue, according to Broome, to a mechanical weighing explanation, in which that a balance tips to one side or another is explained by the total weight that there is on both the right-hand pan and the left-hand pan (we are imagining a two-pan traditional balance here). More specifically, in a case where the balance tips right, we explain it by the fact that the total combination of weight on the right-hand pan exceeds the total weight on the left-hand pan. Broome applies this analogy to normative explanation. Normative weighing explanations then, according to Broome, are explanations where in a case where *S* ought to *F*, this deontic fact is explained by the fact that the total 'weight' of in-favour-of-*F* considerations exceeds the total 'weight' of against-*F* considerations. A *pro tanto* reason, according to Broome, then just is whatever plays the relevant in-favour or against *F* role in a normative weighing explanation of this general sort. In Broome's terms: 'A reason for you to *F* is analogous to an object in the left-hand pan, and a reason for you not to *F* is analogous to an object in the right-hand pan' (Broome 2013: 52). Somewhat more formally, Broome proposes the following definition of *pro tanto* reasons:

(Pro tanto reasons): ‘A *pro tanto* reason for *N* to *F* is something that plays the for-*F* role in a weighing explanation of why *N* ought to *F*, or in a weighing explanation of why *N* ought not to *F*, or in a weighing explanation of why it is not the case that *N* ought to *F* and not the case that *N* ought not to *F*.’ (Broome 2013: 53)

Consider an example where I ought to return to the school after having picked up my kids from school. There is a weighing explanation of this ought – namely, the fact that we took another kid’s scarf inadvertently and it is winter outside outweighs the fact that we will be late for dinner and that we will have lost some time. The fact that the scarf we took is not ours and it is cold outside speaks in favour of returning to the school. The fact that we will be late for the dinner and lose some time (and energy) speaks against returning. On balance, the combination of things speaking in favour of returning have more ‘weight’ than the total combination of things counting against it. This is one example of Broome’s weighing explanation of why I ought to return to the school. The fact that the scarf is not ours and it is cold outside, which brings in the risk that the kid whose scarf we took may get cold, plays a for-returning role. These considerations are on this account *pro tanto* reasons for returning. Other considerations about losing time and energy and being late for dinner play the against-returning role and hence are *pro tanto* reasons not to return (compare this to a case from Broome 2013: 55–56).

Arguments in favour. There are several considerations that seem to speak in favour of Broome’s Deontic Explanation view. Let us consider briefly what appear to be the eight most important among them.

First, contrary to what some critics seem to think (cf. Wodak 2020: 2), this approach is in a position to provide a theoretically well-motivated account of why mere enabling conditions are not normative reasons. See Section 2.5 on how a major competitor of the Explanation view, the Reasoning view, doesn’t have an easy and well-motivated way to distinguish mere enablers from normative reasons. The account here relies on the well-known distinction with respect to explanations in general: it is common to distinguish *explanans* – that is, things that explain – from mere background conditions. This general distinction can be applied to normative reasons, given that normative reasons are, according to the Deontic Explanation view, entities that play a role in explanations of deontic facts (are either *explanans* or elements of weighing explanation). To take Dancy’s examples once more, the fact that a promise was not made under duress and the fact that I am able to keep my promise are, presumably, not normative reasons for me to do the promised thing (cf. Dancy 2004: 38).

Yet they seem to matter normatively nevertheless. A proponent of Broome's Deontic Explanation view is perfectly suited to explain why and how this is so. Ability and absence of duress are, on this view, mere background (enabling) conditions, and not [parts of] an explanation of the relevant deontic facts. If one accepts the general and common distinction between explanation and background conditions, one is in a position to apply it in the normative case as well.

Second, and relatedly to the previous point, some considerations are not strictly speaking reasons to F, they are not [parts of] an explanation of why S ought to F, but nevertheless entail that S ought to F. Presumably, cases of entailing evidence that S ought to F are of this sort. Again, proponents of the Reasoning view have some trouble explaining why this is so, given that considerations that entail p can function perfectly in good arguments/patterns of reasoning towards p (given some assumptions). The Deontic Explanation view explains how this is possible. In general, considerations that entail p are not [typically] parts of a good explanation of why p . This aspect of explanations in general presumably also applies to deontic explanations. Thus, given this general aspect of explanations, the Deontic Explanation view seems to have an important advantage over some of its main rivals at least. Broome himself appears to hint towards this line of argument in favour of the Deontic Explanation view in his brief remarks on what is missing in Raz's view (and on Dancy's objection to Raz's view). He writes:

Dancy points out that, by Raz's criterion [i.e. a version of the Reasoning view], conclusive evidence that you ought to perform an action would be itself a reason to perform the action. That is not necessarily so. Facts that merely entail that an agent ought to perform the action are not necessarily reasons for her to perform it; to be reasons they must explain why she ought to perform it.[. . .] If a newspaper publishes an article saying that a minister ought to resign, that is evidence that she ought to resign. If the newspaper is extremely reliable it may be conclusive evidence. But it is not a reason for the minister to resign. (Broome 2013: 51)

The Deontic Explanation view can easily explain why considerations that entail that one ought to F are not normative reasons to F, in cases where these considerations are not normative reasons to F.

Third, the Deontic Explanation view can also easily deal with the Moore-paradoxical (and self-undermining) cases that we introduced in the preceding chapter. Remember that in Moore-paradoxical cases, we have considerations like 'the building is on fire but I don't know that the building is on fire'. Intuitively, in situations where these considerations are

true, they can be normative reasons in favour of, say, checking, considering, reconsidering whether the building is on fire, and perhaps for running away. Some views – for example, the Reasoning view, that tie normative reasons too closely to one being able to reason from them towards the relevant (intention to) F-ing in a valid/good or fitting way – predict that such considerations cannot be normative reasons, even though they appear naturally to *speak in favour* of some or other reaction. The Deontic Explanation view seems to be able to account for these cases while respecting the pre-theoretical intuition that these considerations are some sort of normative reasons. Suppose that I really ought to check out whether the building where I am now is not on fire. Presumably, in such a case, there is a weighing explanation of such an ought. The weighing explanation of such an ought will appeal to elements from the relevant Moore-paradoxical considerations. To take the example from the Chapter 2, the fact that there is a fire of which I am not aware counts strongly in favour of checking/investigating whether the house is not on fire. It is also a part of the weighing explanation of why I ought to check/investigate/run away from the house on fire. If this is right, then proponents of the Deontic Explanation view are in a position to explain why the Moore-paradoxical considerations appear to be normative reasons. I ought to check the house because there is a fire of which I am not aware. Presumably, a similar account also applies to self-undermining considerations and, in general, about cases where it appears that a subject is not able to F on the basis of the relevant reason to F. Being in a position to explain these cases puts the Deontic Explanation view at an advantage compared to some of its main rivals.

Fourth, the Deontic Explanation view is designed to account for graded normativity. It is often recognised that a theory that has only ‘strict’ (or categorical) normative notions as opposed to a theory that can account for strict and for graded notions will be somewhat lacking (see Lord and Maguire 2016). It seems that we need both strict notions like ‘ought’ and ‘obligation’ and graded or weighed notions that can be appealed to in order to account for apparent normative conflicts (considerations pulling in different directions). This theoretic need for two sorts of normative notions is well known at least since W. D. Ross, who introduced the notion of *prima facie* oughts. As we observed earlier (cf. Chapter 1), normative reasons are supposed to fit the bill. They are supposed to come in different degrees. By its very design – for example, its appeal to ‘weights’ of reasons and the weighing explanation of deontic facts – the Deontic Explanation view is poised to account for the graded, non-categorical/non-binary aspect of normative reasons.

Fifth, the Deontic Explanation view is general enough to be easily applicable to action (intention) as well as to the case of belief and other attitudes. Having this level of generality certainly speaks in its favour as long as we expect that a serious contender to the status of the correct theory of normative reasons is able to explain all pre-theoretically plausible cases of normative reasons (and not only to be applicable to some cherry-picked cases). As long as there are oughts that apply to attitudes and as long as there are weighing explanations of (some) such oughts, the Deontic Explanation view predicts that there will be normative reasons for and against attitudes to which such oughts apply.

Sixth, the view also preserves a good level of theoretical generality with respect to some vexed questions in contemporary meta-ethics. For instance, it is, in its simple form, neither committed to nor incompatible with so-called buck-passing accounts of value, where buck-passing accounts suggest that values can be reduced to reasons (see Scanlon 1998, for instance; cf. Chapter 1). Given that buck-passing (in this general form) is not a claim about the relation between oughts and reasons, a proponent of the Deontic Explanation view can accept that values reduce to reasons even though, ultimately, reasons reduce to oughts and explanations. On the other hand, the Deontic Explanation view is also compatible with values ultimately not being reducible to normative reasons; the only thing that the Deontic Explanation view is committed to is that normative reasons do reduce to a combination of oughts and explanations. Thus, it is uncommitted to this much-debated issue in meta-ethics and is compatible with both sides of the debate.

Seventh, and somewhat connectedly but also distinctly, the view can also easily explain why there are no normative reasons for values (instantiations of values). Suppose that there is a value in you being physically fit. It is a valuable state of you being fit and healthy. But as we have already seen in Chapter 2, there are no normative reasons for you to be healthy. That is, there are no reasons for you to be healthy as opposed to there being reasons for you to undertake some training or adopt a good diet, as actions that could lead to you being healthy. Of course, there are reasons for you to go to the gym and eat healthily, but it is not the same thing as being healthy. The Deontic Explanation view is, again, perfectly able to explain why this is so. There are no oughts for states of value instantiation, there is no ought for you to *be* healthy, as opposed to undertaking sport and adopting a good diet. Given that there are no oughts, there are no weighing explanations, and hence the view predicts that there are no normative reasons for you to be healthy. There are, on the other hand,

reasons for you to go to the gym, given that you ought to go to the gym and there are weighing explanations of such an ought. Thus, the Deontic Explanation view seems to be able to predict correctly the difference between values and deontic properties and is able to explain in a theoretically well-motivated manner why there are no reasons for values (instantiations of values).

Eight, and finally, the view is in a good position to explain our pre-theoretical judgment that there is some kind of connection between normative reasons and explanation. Requests for explanations of some apparent ought to F and requests for reasons for F-ing can often be interchanged without any loss of information. Suppose that I ought to go to the grocery shop and I tell you just that – that is, that I ought to go to the grocery shop. It seems that in such a case, your request for me to give you reasons that I take there to be for me to go to the grocery shop amounts to exactly the same as your request for me to explain why I ought to go to the grocery shop. That is, in such a situation, your reply ‘What reasons are there for you to go to the shop?’ can be replaced without any loss of information with ‘Why ought you to go to the shop?’. Now, it would be somewhat odd if this were a mere accident. The Deontic Explanation view explains this intuitive connection without any need to postulate such an accidental connection. There is an intuitive link between reasons and explanations of oughts, because reasons just are explanations or play a role in a weighing explanation of oughts, according to the Deontic Explanation view.

We have enumerated some positive arguments in favour of the Deontic Explanation view of normative reasons. We have briefly reviewed eight of them. Unfortunately, the view also faces some pretty serious objections. Let us now turn to some of the main worries about the view.

Worries. The first worry appears with respect to Broome’s account of *pro toto* reasons. Roughly, the worry is that not all genuine explanations of why one ought to F appear naturally to be reasons for one to F. This worry, or at any rate a worry similar enough to this one, is due to John Brunero; see Brunero (2018) and also Brunero (2013) for a related worry for earlier (and presumably more problematic) versions of Broome’s account. Consider the fact that I ought to save a drowning toddler. Is there an explanation of why I ought to save the toddler? If there is an or the explanation of why I ought to save the toddler, then, according to Broome’s account, that explanation (*explanans*) is a *pro toto* reason for me to save the toddler. Let’s assume that there is an/the explanation why I ought to save the toddler. Given this assumption, Broome’s view predicts

that this explanation is a *pro toto* reason to save the toddler. But what could possibly be an explanation of why I ought to save the toddler? Presumably, it has to have as part of it the consideration that the toddler will die if I don't save her. But, presumably, it is not only that. It is also that I am not risking anything of a comparative value to the toddler's death by saving her. For instance, I am not risking certain death myself, or the death of my loved ones, or the death of a hundred other innocent babies, and so on (see Singer 1972 for the original example and a theory of altruism based on it; see Timmerman 2015 and Logins 2016a for a recent further discussion). But, crucially, the fact that I am risking neither my life nor the lives of my loved ones or numerous innocent babies are not (parts of) reasons why I ought to save the drowning toddler. These are considerations that correspond to the absence of 'undercutting defeaters' (see Brunero 2018), but are hardly seen as parts of reasons to save the toddler. These are certainly parts of an explanation of why I ought to save the toddler, but they don't appear to be themselves reasons to save the toddler. That the toddler is drowning is a reason to save her and not that that I am not risking my life by jumping into the water to save her. Broome's view seems to predict otherwise.

One possible reply to this worry is to insist on the distinction that the Deontic Explanation view postulates – namely, the distinction among genuine explanation and background conditions, as well as contextual and pragmatic considerations. As in the case of non-normative explanations, what will count as a/the explanation in given circumstances depends partially on background conditions and contextual/pragmatic considerations (e.g. what is of interest for the subject; see Broome 2013: 49). Thus, one might reply that the absence of undercutting defeaters may well be part of the bigger, complete explanation of why one ought to F, but will not be part of some *pars pro toto* explanations of why one ought to F (where a *pars pro toto* explanation is a partial explanation that stands in for the one big explanation, and in calling a *pro toto* explanation the explanation we are using a synecdoche, namely the so-called *pars pro toto* figure of speech, see Broome 2013: 49). But such a reply is unsatisfying, for it presupposes that the complete, big explanation *is* the explanation of why one ought to F, even though there are some *pars pro toto* explanations that don't include the mention of all the background conditions, the absence of undercutting defeaters, and interests. But a complete, big explanation of why one ought to F still counts as an explanation of why one ought to F and hence, on Broome's account, has to be a/the reason why one ought to F. (If this complete explanation is not a reason, it is then

difficult to see why partial explanations would be allowed on Broome's view to be *pro toto* reasons.)² But this is precisely what we dispute. The complete explanation is not the/a reason why one ought to F. For parts of the complete, bigger explanation are not parts of a/the reason why one ought to F. The toddler example is just about this. Certainly, the absence of undercutting defeaters has to be part of the complete, bigger explanation of why I ought to save the toddler. But the absence of such undercutting defeaters doesn't appear to be part of the/a reason why I ought to save the toddler.³

The second worry concerns the relation between *pro toto* and *pro tanto* reasons on the Deontic Explanation view. Let us take Broome's own example of a *pro toto* reason, the case where you ought to visit Mr. Reed (cf. Broome 2013: 50–51). Broome (2013: 50) suggests that

The explanation [of why you ought to visit Mr. Reed] might be that he is the best dentist around. [...] In a different context, the explanation might be that you ought to visit the best dentist around. [...] A fuller explanation would be the conjunctive fact that you ought to visit the best dentist around and Mr. Reed is the best dentist around.

Now, it is quite plausible that in this case, there is also a weighing explanation of why you ought to visit Mr. Reed: the total 'weight' of all for-Mr.-Reed-visiting considerations 'outweigh' the total 'weight' of

² That the complete, bigger explanation counts as an explanation of why S ought to F seems to be implied by Broome himself: 'So long as something explains why you ought to F, it is a *pro toto* reason for you to F' (Broome 2013: 50). The complete explanation does explain why one ought to F. Hence, on a natural reading of these remarks, it is a/the reason why one ought to F. He also writes: 'A *pro toto* reason therefore need not be a unique canonical reason' (Broome 2013: 50). It is natural to take 'x need not be Q' to imply that x nevertheless may be Q. Thus, it is natural to take it that a *pro toto* reason may be the canonical, complete reason.

³ This worry may be further developed as a problem for the very idea of there being any *pars pro toto* or perfect explanations of why one ought to F. Broome himself struggles in providing an obvious example of a perfect (*pars pro toto*) rather than a weighing explanation of why one ought to F. Strict deontic rules and beliefs in contradictions are two tentative examples. But they do depend on other, debatable assumptions. See Kearns and Star (2008: 36) for a line of worry based on this sort of observation:

Given Broome's own doubts about these examples [e.g. the examples of rigid deontic rules and there being reasons to never believe contradictions], it is unclear that this third feature of his account [e.g. that the view accommodates the possibility of reasons that are not weighed against each other, cf. Kearns and Star 2008: 36] is really a positive feature. Perhaps it should be discounted.

Without clear cases of *pars pro toto* explanations of why one ought to F, we are left only with the category of alleged weighing explanations of why one ought to F. However, in what sense 'weighing explanation' is an explanation and provides ground for a reductive analysis of reasons remains to be seen. More objections against the account of *pro tanto* reasons in terms of weighing explanation are discussed ahead.

against-Mr.-Reed-visiting considerations. The same line of reasoning that Broome proposes in the cases of typical weighing explanation cases seem to apply here too. There definitely are costs for you in visiting the best dentist around. If anything, this will certainly cost you some money, energy, or time. It will certainly cost you more than visiting some less good but, say, more nearby cheaper dentist. But if we assume that in this case there are both the explanation of why you ought to visit Mr. Reed and a weighing explanation of why you ought to visit Mr. Reed, then the Deontic Explanation account predicts that there are both a *pro toto* reason for you to visit and some *pro tanto* reasons for you to visit. Crucially, it seems that the same considerations will count as *pro toto* and *pro tanto* reasons. Presumably, the fact that Mr. Reed is the best dentist around will come out on this account as both a *pro toto* and a *pro tanto* reason for you to visit Mr. Reed. But if so, aren't we then double-counting a given consideration as a reason to visit Mr. Reed? Aren't we over-generating normative reasons? Does it mean that you have two reasons, two independent elements that count in favour of visiting Mr. Reed on this account? If so, then it certainly counts against the view (see Brunero 2013 for an objection against the Deontic Explanation view along similar lines).

It seems that whether there is a genuine problem here will depend on how exactly proponents of deontic explanationism conceive of the relation between the explanation of a deontic fact and a weighing explanation of a deontic fact. Yet it is not really clear how this relation is conceived. The mere fact that it is not clear how the relation between the two is conceived is in itself somewhat problematic, since it is incumbent on the proponents of the Deontic Explanation view to explain why their view, contrary to what one not completely far-fetched interpretation of their view might imply, is not committed to the problematic consequences of double-counting reasons.

Finally, one might also think that the very fact that Broome's view entails the existence of two radically different sorts of reasons counts somewhat against the view. As Kearns and Star observe:

[On Broome's view] [s]omething is a normative reason if it is an explanation of a normative fact or it weighs in favor of a certain action. These properties are so different that it is tempting to interpret Broome as claiming that 'normative reason' is ambiguous between them. Whether we understand Broome to be claiming that 'normative reason' is ambiguous or simply that there are two very different ways of being a normative reason, this is an unattractive feature of his account. (Kearns and Star 2008: 45)

The unattractiveness according to Kearns and Star comes from the consequence that such a disjunctive view is less simple and elegant than its non-disjunctive competitors. It doesn't seem to fit well with our ordinary way of talking about reasons. If so, it would imply that the concept of reasons doesn't pick out a non-artificial property (see Kearns and Star 2008).

The third concern is more specifically about *pro tanto* reasons. Broome has made it clear that he thinks of the explanation relation as the inverse of the because relation. According to Broome, that X explains Y just means that Y 'is so because' X 'is so' (cf. Broome 2008: 100, 2013). In other terms, X explains Y just when X makes Y the case (cf. Broome 2008: 100). But if this is so, this should also apply to weighing explanations. This, however, may appear implausible in the case of outweighed reasons (see Brunero 2013, 2018, for this line of objection). Take the case where you ought not go on vacation to the seaside. Let's assume that the proponents of the Deontic Explanation view accept that there is a weighing explanation of why you ought not to go on vacation to the seaside. Part of this weighing explanation will be that you broke your leg three weeks ago and it is still not advisable for you to travel far, and the seaside is a long way away. But if this is really a weighing explanation, then according to the Deontic Explanation view, against-not-going-to-the-sea considerations also play a role in explaining why you ought not to go to the seaside. Let us say that a friend of yours has proposed that you stay at their parents' house just next to the sea during the vacation for free, and that the weather is really nice at this time of the year; these are considerations that play the against-not-going role in the relevant weighing explanation. The free-of-charge accommodation and the good weather are considerations speaking in favour of going. According to the Deontic Explanation view then, the free accommodation next to the sea and the good weather then partly makes it the case that you ought not to go on the seaside vacation. But this result appears counterintuitive. If anything, these considerations could make it the case that you ought to go. How can these considerations make it (partially) the case that you ought not to go? Is it really the case that you ought not to go partially because you can have free accommodation next to the sea and the weather is especially clement there at this time of the year? To make the oddness even more vivid, imagine a discussion with a friend. 'You: I ought not to go to the seaside on vacation. Friend: Why? You: for one thing the weather there is really nice now. For another, I have this free house just for us for a whole week.' To my ear, such a conversation sounds

really like a joke. However, on the Deontic Explanation view, of course, this conversation shouldn't come out as a joke, but only as a partial explanation of why you ought not to go to the seaside. This seems to be a strongly counterintuitive implication of the view and hence it speaks against the Deontic Explanation view's account of *pro tanto* reasons (see Brunero 2013, 2018, for more details on this objection).

This verdict seems to point to an even more substantial issue with Broome's account of *pro tanto* reasons in terms of weighing explanation. The issue is that it is simply unclear what weighing explanations really are and, in particular, how to understand the idea that some considerations play a 'for-F' (or 'against-F') role in the alleged weighing explanation of an ought fact. For Broome (2013: 54), '[t]he for-F role can be identified from the structure of the explanation itself. How exactly? 'In a weighing explanation of why you ought to *F*, the for-F role is the winning one, and that is how it can be identified' (Broome 2013: 55). So, fundamentally we are supposed to grasp which considerations play the for-F role in a weighing explanation and thus be able to grasp which considerations are reasons by grasping considerations that are *the winning ones* in a weighing explanation. However, as Kearns and Star have rightly observed, this suggestion is problematic. In particular, they rightly wonder 'how the winning considerations are meant to be identified *as* winning, if not by weighing up different considerations in order to see which considerations together most strongly *count in favor* of particular actions' (Kearns and Star 2008: 43, original emphasis). The problem is that Broome doesn't actually seem to provide a substantive analysis of (*pro tanto*) normative reasons, since at the end of the day and contrary to what he suggests, we do need to grasp what counting in favour amounts to in order to understand reasons. To grasp what winning in a weighing explanation amounts to just is to be able to grasp various degrees of weight/strength of counting in favour. Following Kearns and Star (2008: 43), we may ask: 'Why should we think we could grasp what it is for certain considerations to be *winning* without a prior understanding that facts can *count in favor* of actions?' Thus, it is not even clear that Broome is offering a substantial or, to use Kearns and Star's terminology, full-blooded, account of *pro tanto* reasons in terms of explanation (of a deontic fact). At the end of the day, the account of *pro tanto* reasons as considerations that play a for-F role in a weighing explanation seems to be parasitic on our prior grasp of reasons counting in favour of F-ing. Perhaps even the very construct of weighing explanation of deontic facts appears intelligible only insofar as we have a prior grasp of reasons

speaking/counting in favour of or against F-ing and being comparable in this respect, and not the other way round.⁴

Fourth, and perhaps even more fundamentally, one may worry whether the whole analogy of weight is an appropriate way to capture what normative reasons are. Yet the weight analogy and the very idea of there being 'weighing explanations' is central to Broome's account. Thus, it is not clear whether the Deontic Explanation view tells us anything substantial about reasons once we take away considerations about weight of reasons and weighing explanation (see Hawthorne and Magidor 2018 for this line of objection).⁵ It is really dubious that the relative importance of normative reasons can be correctly captured by appeal to 'weights' of reasons. For one thing, weight respects additivity; that is, roughly, the principle according to which adding an object with the weight x to an object with the weight y results in there being something with the total weight of $x + y$ (leaving it open whether it is a new object or a sum of two objects). So, for instance, if you add 2 kg of oranges to 1 kg of kiwis (and nothing more), you have a total of 3 kg of fruit. This simple principle applies universally to anything that has a weight. However, it is violated in the case of normative reasons. Adding a normative reason of the degree of importance x to F to a normative reason of the degree of importance y to F doesn't always add up to there being two/combinatory normative reasons with the total degree of importance of $x + y$. Some reasons just don't add up. To take an example inspired by Hawthorne and Magidor (2018): that this piece of jewellery contains a red diamond is a reason of, say, degree x for you to buy it and that this piece of jewellery contains a ruby diamond is a reason of degree y for you to buy it. But it may still be the case that there are not two reasons/a new combinatory reason with the total degree of $x + y$ for you to buy the jewellery. That the piece of jewellery contains a specifically ruby diamond doesn't add anything to the total degree of reasons that there are for you to buy the jewellery. Or,

⁴ See also Brunero (2013: 812–816) for a similar objection. Brunero concludes that depending on how we interpret Broome's definition of *pro tanto* reasons as playing the for-F role in a weighing explanation, we get either a false or an uninformative (not a substantial) account of reasons. He writes: '[I]t seems as though Broome's account of reasons, depending on how it is interpreted, is either uninformative or false' (Brunero 2013: 816).

⁵ See also Kearns and Star (2008: 44) for a related worry:

A related worry is that Broome's purported analysis of reasons does not give us an analysis of the strength of reasons. He seems to believe it is a brute fact that reasons simply have certain weights of some kind. What these weights are seems to be a mystery. [...] [T]his is another respect in which Broome fails to provide a substantive analysis of reasons.

to take a somewhat different but equally relevant case, discussed more by Schroeder (2007) and Fogal (2016), that you like dancing is a reason for you to go to the party and that there will be dancing at the party is a reason for you to go to the party. But it doesn't follow from these two facts that you have two reasons of some sum of total weight to go to the party. It seems that you cannot generate more reasons or a higher degree of the importance of reasons just by putting these two considerations together.

Moreover, adding two reasons to *F* sometimes results in having less reason to *F* than having only one of these (see also our discussion in Section 2.5, which contains further references relevant for this issue). So, for instance, suppose that you like pizza and you like Nutella. That the item on the menu is a pizza seems to be a reason for you to order it. Similarly, that the item on the menu contains Nutella is also a reason for you to order it. However, that the item in question is a Nutella pizza is not a weightier reason for you to order it. Actually, it is not a reason at all for you to order it (say, you hate Nutella pizza). It seems that adding two reasons to *F* doesn't always add up to having a 'weightier' total of reasons to *F*. Sometimes adding different reasons to *F* cancels them all mutually out. See Nair (2016) for similar cases that are adapted from Horty (2012: 61), involving running in wet heat versus running in wet weather and running in heat.

Another aspect of reasons that is incompatible with the simple principles that apply to weighing is that whereas subtracting an element with the weight x from a total (stuff with the) weight of $x + y$ will result in there remaining an object with the weight y . Nothing similar can be universally applicable to normative reasons. It is not always the case that subtracting one normative reason to *F* from the total of considerations that speak in favour of *F*-ing will leave you with a smaller amount of considerations speaking in favour of *F*-ing. Let's take the Nutella pizza example again, but let's modify it a bit. Suppose that that the item on the menu is a Nutella pizza is a tiny reason for ordering it. Say, you are really hungry and you have a reason to eat anything minimally edible. However, that the item on the menu is a Nutella pizza is not a more important reason for you to order it than that the thing is pizza alone, or that the thing is made of Nutella alone. Indeed, any of these two has a much higher degree reason-wise for you than the two of them together. Thus, it seems that normative reasons don't respect the simplest of the principles that apply to weight and weighing. Without these, however, it is not clear in what sense the analogy with weight is still theoretically useful. Broome recognises that the analogy is not to be taken to be perfect, but as Hawthorne and Magidor (2018)

note, it is not just about the analogy being perfect; it seems that the analogy breaks down completely, and hence it is not clear in what sense the talk of 'weighing' may still be insightful at all. Without the analogy with weight, it doesn't seem that there is much of substance left in Broome's version of the Deontic Explanation view, given how crucial weighing explanation is for this account.

The fifth consideration that speaks against the Deontic Explanation view, and also points towards what's wrong with it more fundamentally, is that it focuses on one central aspect of reasons as we commonly understand them at the expense of another central aspect that we commonly attribute to reasons. This will be also relevant later, when we discuss our positive account. Remember that a fundamental worry with the Reasoning view mentioned earlier was that it focused exclusively on reasoning (and motivating reasons) and acting/having an attitude on the basis of normative reasons at the expense of other important aspects of reasons. It was not well suited to account for all normative reasons. Some considerations seem to count in favour of F-ing without being premises of good patterns of reasoning (cf. Moore-paradoxical considerations). Of course, the explanation views of reasons can easily account for these cases, as we saw earlier. In Moore-paradoxical cases, for instance, the relevant considerations may still count as parts of an [weighing] explanation of why one ought to F in the relevant ways. However, now we have the reverse problem, it seems. For the Deontic Explanation view is going to the other extreme, so to say. It is neglecting the reasoning aspect in accounting for normative reasons. As Kearns and Star (among others) have rightly observed, being practically relevant is a central feature that we commonly attribute to reasons: '[i]ndeed, the philosophical importance of reasons is due in large part to their importance in everyday life' (Kearns and Star 2008: 39). They suggest that this importance comes from reasons being the kind of thing that can help us work out what we ought to do. They write: 'That is, reasons are our guide to what we should do' (Kearns and Star 2008: 39). And they are so, '*in virtue of being reasons*' (Kearns and Star 2008: 39), not just because the facts that happen to be reasons are also facts that can guide us towards finding out what our deliberative oughts are. In short, we are back to the importance of reasons in deliberation.

One problem here is that even if Broome's account somehow manages to be compatible with the view that facts that are reasons for S to F are effectively facts that S can use in working out what S ought to do/what attitude to hold, this result would come out at best as an accidental by-product of the view, not as a central feature of reasons, which one might

think is already odd, given the centrality of guidance for one common way of understanding reasons.

More worrisome still is another problem – namely, on Broome’s account of reasons, there will be plenty of cases where one figures out what one ought to do on the basis of some other considerations and not on the basis of her reasons to F. That is, one might come to know that one ought to F, on Broome’s view, on the basis of some considerations without even knowing what her reasons to F are. The now famous cabbage example from Kearns and Star (2008: 40) is a case that illustrates the point: ‘the fact that a clearly reliable book says that you ought to eat cabbage may reasonably convince you that you ought to eat cabbage, even though this fact doesn’t explain why you ought to eat cabbage.’ Clearly one can come to know that one ought to eat cabbage on the basis of the reliable book saying so. However, that the mere fact that the book says it is not an explanation (nor part of a weighing explanation) of why one ought to eat cabbage. Thus, a consideration that guides one towards knowledge of one’s deliberative ought – that is, the testimony from the book – cannot be, on Broome’s view, a *pro toto* or *pro tanto* reason to eat cabbage. This seems excessive. Fundamentally, the problem is that in such cases Broome’s reasons play no practical role at all, no role in a deliberation of what one ought to do. Thus, we can agree with Kearns and Star (2008: 41) that ‘even if we think that such an action-guiding role is not essential to reasons, it is clear that reasons are generally practically important. However, if agents are able to work out what they ought to do without knowing what explains what they ought to do, then reasons, as Broome conceives of them, are not a vital part of practical life.’ Kearns and Star also provide an insightful analogy that sheds light on what might be the root of the problem here. They suggest that deontic explanationism is ‘backward looking’, whereas their own (and I think we can also say the Reasoning accounts) are ‘forward looking’.⁶ That is, whereas the Explanation view

⁶ Consider:

Broome’s view is backward looking. Typically, reasoning that concerns explanation proceeds as follows: A person knows a fact and wishes to explain it. She then infers some other fact by inference to the best explanation. By applying this to the case of normative facts, we get the following picture. We know certain normative facts which we wish to explain. We then infer other facts by inference to the best explanation. These facts are, according to Broome, the reasons we have to act. (Kearns and Star 2008: 41)

This is contrasted with an alternative account: ‘[o]ur view is a forward-looking view of reasons. That is, reasons are those things that are used to figure out what ought to be done. One is first in possession of reasons to act in certain ways and then uses these to determine how one should act’ (Kearns and Star 2008: 41).

starts with an ought, the Evidence and Reasoning views start with considerations that help us to find out what we ought to do. And the Explanation view gets the phenomenology of deliberation wrong. The idea is that in deliberating about what to do, what we *should do*, we don't start with the ought fact, but, typically, we engage in deliberation and reasoning to find out the relevant ought fact. Given that we don't know what we ought to do in such contexts, we are not looking for explanations of the ought facts. Yet reasons are the things we are looking for in deliberation. Thus, normative reasons cannot all be just explanations (or parts of weighing explanations) of some ought facts (cf. Kearns and Star 2008: 41). Exploring differences and similarities between reasoning and explanation will be the key and, indeed, the fundamental element of our own positive account to be developed later. Let us not anticipate that discussion yet.

Now, someone sympathetic to the Deontic Explanation view may want to try to amend it in ways that would fix the aforementioned problems. We conclude this section by considering briefly an alternative view that also combines an appeal to both oughts and something close to an explanation, in order to characterise reasons.

Reasons as right makers. The aforementioned five worries appear rather problematic for the Deontic Explanation view. Yet, as we observed earlier, there was also a *prima facie* case in favour of it. What should we think, then, if it's mistaken? What about the eight arguments that we have listed in its favour? Were we merely confused? Later, I will argue that such apparent confusion – a situation where some considerations draw us towards something like the Explanation view, but other considerations draw us towards a different understanding of reasons – is only to be expected, given what reasons are. However, let us now consider another option for those who are sympathetic to a broad deontic explanationism. It is based on a line of thought that would keep the spirit of the Deontic Explanation view and focus on reasons as constitutive with respect to oughts, while relaxing the theoretically ambitious aims of the approach. The approach proposed by Laura and Francois Schroeter about [practical] reasons as right makers seems to be one way of doing just this (Schroeter and Schroeter 2009). On their view, we should not aim to provide a reductionist account of normative reasons. According to them, some substantial constraints on what are normative reasons will always depend on common-sense intuitions that cannot be captured by purely formal accounts. Their approach is a way of characterising normative reasons, which are to be understood in functional terms, by appeal to their role in

constituting rightness (more on this in a moment) rather than to define or reduce reasons to a given role in explaining deontic facts.

The approach might appear to be closer to the reasons-first approach than to an explanationist approach since they refuse to reduce reasons to something else. However, their characterisation of reasons is very much in the spirit of the Deontic Explanation approach. For the positive things that the approach does say about reasons appeal to making conditions and rightness, which can be understood as corresponding to oughts. Despite being of the same broad family of views (i.e. the focus here is on the reverse of 'because' and ought), the right-making approach differs importantly from the Deontic Explanation approach. As we saw already, they don't aim to provide a reductive analysis of reasons in other terms. But there is yet another difference, namely it focuses on the constitution of rightness rather than on mere explanation of deontic facts. The relevant difference between constitution and explanation appears to amount to a difference between two species of the same genus. Contrary to what some passages from Broome might seem to suggest, not everything that can count as (part of) an explanation will count as (part of) the constitution of the relevant thing, which seems to give the constitution approach some advantage over Broome's deontic explanationism. For one thing, the right-maker view doesn't entail that any explanation of why one ought to F has to be a reason to F (cf. the first worry for Broome's explanationism). That I don't risk my life by jumping into the pond can be part of an explanation of why I ought to jump, without being itself a reason for me to jump (that a toddler is drowning is a reason). One might think that such a consideration points only towards the absence of undercutting defeaters for the claim that I ought to jump, and that the mere absence of such a possible undercutting defeater is not, strictly speaking, part of what makes it right for me to jump into the pond.⁷ Similarly, the correct meta-ethical theory will provide an explanation of why something is a reason to F while not being itself a normative reason to F. The focus on constitution rather than explanation explains how this can be the case. The correct meta-ethical view doesn't constitute the rightness of F-ing.

⁷ Yet note that whether this line of thought is well grounded will depend on how we think about constitution. If one thinks that constitution, like explanation, can be of a more or less complete sort, then, arguably, one might also insist that a more complete story of what makes my jump into the pond right – that is, what constitutes its rightness in a more complete sense – has to take the absence of possible undercutting defeaters into account. If so, then the right-maker view doesn't have an advantage over Broome's deontic explanationism here after all.

Another advantage that Schroeter and Schroeter are claiming to have over the Deontic Explanation approach is that, given their relaxed theoretical ambitions – that is, not aiming to provide a reduction of normative reasons to something else – they can also easily deal with cases where some reasons-wise irrelevant facts may play a role in an explanation of why S ought to F, without being counted as genuine reasons to F for S. The job of picking out the relevant facts as reasons is done not by any formal constraints according to this view but by concrete common-sense considerations on a case-by-case basis (cf. Schroeter and Schroeter 2009). Perhaps this line of thought might equally constitute first steps for dealing with the worry of double-counting reasons (cf. the second worry mentioned earlier).

Even though the right-making account might avoid some of the problems of Broome's deontic explanationism, it still has important and similar pitfalls. For one thing, it is not clear how the constitution of rightness works in the outweighed reasons case. Schroeter and Schroeter propose to think of constitution here using an analogy to how bricks constitute a house. It is bricks arranged in some way that constitutes a house. Reasons are supposed to constitute rightness in a somewhat similar manner. Yet in the case of outweighed reasons, it is not clear at all how the fact that a friend proposes that I stay in a house by the seaside can constitute (even if only partially) the rightness of me staying home for my holidays or how is it possible on this view that my promise to meet my friend for a coffee is a reason for me to go to the meeting place in a situation where I witness a traffic accident and am the only one able to help. It doesn't seem that my promise constitutes the rightness of me leaving the injured and going to the place where I promised to meet my friend instead. It is clearly wrong to leave the injured and meet my friend in such a situation. So, my promise cannot constitute the rightness of going to meet my friend here. But it doesn't seem to play any role in constituting the rightness of staying and helping the injured either. The promise to meet my friend is irrelevant with respect to the rightness of helping the injured. Also, it is not clear that the right-maker approach captures the gradable aspect of reasons. In particular, one might think that exactly as in the case of weighing explanation, the right-maker view fails to do justice to the failure of additivity of reasons. For constitution seems to satisfy additivity. If a set A of bricks constitutes a house and a set B of bricks constitutes a house, then, taken together (and without destroying either A or B), bricks A and B constitute either a larger house or two distinct houses that taken together are larger than the A house or the B house taken individually. More fundamentally, and exactly as in the case of deontic explanationism, the right-maker

approach appears to minimise too radically the role of reasons in deliberation. Yet, in a sense, our common-sense understanding of reasons is also as of things that help us to figure out what we ought to do or which attitudes to hold. Hence, it is not clear that the right-makers approach is really an improvement on the Deontic Explanation account, at least on these problematic points.⁸

3.3 Axiological Explanation View

Another possible way of capturing the pre-theoretical insight that normative reasons have to be tied to explanation of something normative, broadly understood, is to focus on explanation of values. It is common to see values as distinct from obligations, and more generally axiological properties (e.g. values, good, bad) from deontic properties (including, oughts, permissions, requirements, and so on). We have referred quite vaguely to the common category of these as broad normativity. At that level of abstraction, the Explanation approach of reasons is characterised as the view that a normative reason to F is an (partial) explanation (of an aspect) of a normative feature/fact with respect to F-ing. The preceding section focused on one prominent way of making this abstract idea more precise, namely defining reasons as explanations of aspects of deontic facts. The present section looks at the axiological alternative, an alternative that can be endorsed with or without endorsing a more ambitious project of defining/analysing other normative properties/notions, including oughts/obligations and fittingness in terms of values and thus vindicating the idea of value being the most fundamental and explanatorily prime normative property/notion (e.g. the value-first approach). Very roughly then, according to Axiological Explanation views (be they value-first accounts or not), a normative reason to F is (a part of) an explanation of why F-ing would be good (has value). More specifically, we will focus on two recent and arguably the most promising versions of the Axiological Explanation approach. The first is

⁸ Another option within the broad deontic explanationism family that has recently been suggested by Nebel (2019) is to deny the factivity of reason-why constructions but maintain that normative reasons are reasons why one ought to F. We will not go into details of this suggestion here, in order to keep our discussion manageable. But note, as Wodak (2020) has recently observed, that while, strictly speaking, normative reasons – that is, reasons-why one ought to F – on Nebel's account are not (parts of) explanations of why one ought to F, it is not really clear what exactly they are. If reasons-why are not explanations, then what are they? Nebel doesn't provide many positive details about these, and it is not clear that we have a clear pre-theoretic grasp of 'reasons-why' that are not understood as explanations.

defended by Stephen Finlay, the second by Barry Maguire. After presenting each of these, we will also look at what appear to be their pitfalls.

The central idea of Stephen Finlay's comprehensive treatment of the explanation-plus-value-based view of reasons, elaborated in a number of recent publications (Finlay 2006, 2012, 2014: 85–115, 2020), is that a consideration is a normative reason for one to F just in case the consideration explains why it would be somewhat good for one to F (cf. Finlay 2020: 1). On this view, for relevant F-ings, there is a degree of goodness of a kind in a subject's F-ing. Normative reasons are considerations that explain why this is so. Again, as is common, the approach has to be understood in functionalist terms. Reasons are considerations or, more specifically, facts that play some specific function (or stand in a specific relation, the reasons relation). That function, in this case, is the function of explaining why some F-ing would be good in a sense and to a degree (see Finlay 2014: 85). Thus, on this view, to say that the consideration that my kids are hungry is a reason for me to prepare them some food is to say that that my kids are hungry explains why it would be good that I prepare them some food. It would be good to prepare them some food because they are hungry. Note also that strictly speaking Finlay's account is about 'reasons' statements – that is, a theory of meaning of 'reasons' statements. In what follows, however, for reasons of manageability of discussion, we avoid the repetition of this aspect of his view and will talk about it directly as a theory of reasons.

The approach has a number of considerations in its favour. Before rehearsing some of them, let us, however, unpack slightly the specifics of Finlay's view. Let's see what his view entails beyond its central idea. Exploring these further aspects of Finlay's view will also enable us to sketch later how Finlay's version of axiological explanationism can respond to standard objections to value-based accounts of reasons. The fact that in its more elaborated form Finlay's explanationism has the resources to counter (at least some of) the objections that other, more rudimentary value-based views of reasons cannot respond to successfully will count as another point in its favour.

Presumably, the first and foremost clarification that we have to make about Finlay's account is that he endorses an end-relational view of *goodness*. Again, for the sake of brevity, we will sometimes talk as if Finlay directly provides a view about goodness, while bearing in mind that strictly speaking he is providing a theory of the meaning of 'goodness' and is interested in the related concept of *goodness*. According to Finlay, there is no such thing as goodness *simpliciter*. Things (features, actions,

attitudes, states, what have you) are always good with respect to some end, on this view. So, to take Finlay's own example, the sentence 'Umbrellas are good' can be analysed as having the following underlying syntactic form: 'It is good for S_1 's *doing* A_1 , for S_2 to do A_2 ' (Finlay 2020: 6), as, for example, in 'It is good for people's staying dry in the rain, for them to use *umbrellas*' (Finlay 2020: 6). The idea here is that "good" fundamentally expresses a relation between two propositional arguments' (Finlay 2020: 7). To take the same example, then, 'people use umbrellas' expresses, according to Finlay (2020: 7), 'the *object* proposition, p ' and 'they stay dry in the rain' expresses 'the *end* proposition, e '.

It should be also noted that ultimately something x being good for some end e is understood here in probabilistic terms. Namely, x 's obtaining promotes or, more specifically, raises the probability of the end e obtaining (cf. Finlay 2020). So, on this view, to say that it is good for people's staying dry in the rain for them to use umbrellas just is to say that that people use umbrellas increases the probability of them staying dry in the rain (i.e. that they stay dry in the rain).

Another aspect of Finlay's end-relational account of goodness that is important for our purposes is that ends are context-sensitive on this account. So whether something is good will also depend on which ends are salient in a given context. Or as Finlay (2020: 7) puts it, '[o]n this contextualist view, something is correctly said to be "good" (sans phrase) in a particular context only if it is good relative to an end that is in some way salient or privileged in that context'. So, according to this approach, what exactly 'umbrellas are good' means in a context (if meaningful at all) depends on which end or outcome connected to someone using umbrellas is salient or privileged in the context. If the relevant end in a context is that of people staying dry in the rain, then to say that umbrellas are good amounts to saying that people using umbrellas is good for (i.e. it increases the probability of) people staying dry in the rain, in that context.

Finlay's approach is highly flexible, given its many moving parts. Before we reproduce the whole official complete account, let us precis one last aspect of his account. Goodness-for-an-end on this account is understood in probabilistic terms, namely as an increase in probability of the relevant end given the relevant 'goodness' considerations (that the end is more probable given the relevant consideration than it is without it). But to make sense of an increase in probability (of a hypothesis h , given some consideration e), we need to know what probability we are talking about and, in particular, what is the relevant background information b , or in Finlay's terms 'information-base' (cf. Finlay 2020: 10). The relevant

background, or information-base, is also highly flexible on this account. It can be constituted by shared knowledge in a given context, but it can also be relativised to what one or another subject believes in a context (interpreted in intensional terms; cf. Finlay 2020: 10); it can also be relativised to ‘an information-base defined in some objective way (e.g. *the world’s state at time t*)’ (Finlay 2020: 10). Note also that despite appealing to a sort of (substantially revised) Deductive Nomological understanding of explanation, Finlay (2020: 10) sees explanation also as context-sensitive and for simplicity assumes ‘by default that in statements about normative reasons, explanation and probability/goodness are relativized to the same background information’. With all these remarks and specifications in mind, we are now in a position to reproduce Finlay’s complete analysis of normative reasons (i.e. of normative reasons statements):

[Reasons as Explanations of Goodness + End-Relative Theory of Goodness, REG+ERT] To say in a world w that R is a ‘reason’ for S to do A is to say, of some end e and information-base b , that R is an explanation in w why *given b* it would be good/probability-raising for e , if S does A. (Finlay 2020: 10)

Given the specifics of the view, I suggest referring to this Reasons as Explanations of Goodness + End-Relative Theory of Goodness view as the Axiologic Contextualist Explanation view (or Axiologic Contextualist Explanationism, for short).

So, to recapitulate, why is the consideration that my kids are hungry a normative reason for me to prepare some food for them? Well, this is so because (or it just means that) in a world w , there is some end e and an information base b and that my kids are hungry is an explanation in w why given b it would be good/probability-raising for e , if I make my kids some food. Presumably, the end here is that they don’t starve (or perhaps, less dramatically, that they just have a ‘healthy’ dinner), and the information base is that what I know and/or believe at this time in w , including that preparing food for dinner normally helps reduce hunger (if I make an effort) or something similar.

Arguments in favour. A number of considerations speak in favour of axiologic contextualist explanationism. First, and perhaps foremost, it appears to have a straightforward account of a central constraint for a theory of normative reasons, namely it respects the ‘gradable’ aspect of normative reasons. Indeed, it can explain where the ‘weight’ of normative reasons comes from. Thus, the comparability, combination, outweighing, and cases of additivity/subtraction of normative reasons (when applicable)

can be easily accounted for within this account. This is done within this account by reference to the gradable property of goodness. Goodness comes in degrees. This much seems uncontroversial. The ‘weight’ of a reason to F within axiological explanationism is tied to the degree or ‘weight’ of (relative) goodness of F-ing. Thus, that there is an injured person in front of me is a normative reason for me to stay and help. And this reason outweighs my reason to leave, which is grounded in my promise to meet a friend for a coffee, because the degree of goodness of helping is higher than the degree of goodness of meeting a friend for a coffee in this context. Presumably, the further explanation of why this holds would appeal to the difference in the ‘weight’ that is a function of the importance of the end of helping and the conditional probability of achieving that aim by staying, and the ‘weight’ that is a function of the importance of the end of keeping my promise and the conditional probability of achieving that end given my going to the coffee shop instead of staying. At any rate, axiological explanationists seem to have at their disposal means to the gradable aspect of normative reasons. By definition, it will be tied, one way or another, to the degree of goodness of the relevant F-ings. Note that this may be one of the most important advantages that axiological explanationism has over deontic explanationism. As we saw in the previous section, the existing deontic explanationism proposals seem to have trouble with the gradability aspect of normative reasons. (Recall that contrary to what, for example, John Broome suggests, the appeal to the special sort of explanations, the ‘weighing explanations’, doesn’t really fit our pre-theoretical judgments about how outweighed reasons work.)⁹

Second, like deontic explanationism, the Axiological Explanation view is also perfectly situated to explain the intuitive connection between explanation and reasons. In general reasons, be they normative or not, seem to be associated in a sense with explanations. It is natural for us to say that global warming is both an explanation and a reason why patterns of bird

⁹ Finlay also makes a more specific point that axiological explanationism fits perfectly with the observation that apparently at least some normative reasons don’t respect the additivity constraint on combining the ‘weight’ of reasons. It seems that at least sometimes having r_1 and having r_2 as distinct reasons to F doesn’t add up to having more ‘weight’ for F-ing than having only r_1 or r_2 as a reason to F. Axiological explanationism accounts for this, given that (i) it accounts for the weights of reasons by the degree and kind of goodness they explain and (ii) we can have non-competing correct explanations of the same thing’ (Finlay 2020: 3). I take this feature as being a specific instance of the more general point made in the text, namely that, given the central role of the goodness in the axiological explanationist view of reasons, and goodness being paradigmatically gradable and flexible in the relevant ways, the view is well suited to account for all aspects of the gradability of normative reasons. Non-additivity seems to be just one specific feature of the sort of gradability involved in reasons that the flexibility of axiological explanationism can take care of without any difficulty.

migration are changing. That is, that global warming is occurring can be referred to *interchangeably* in this context as the reason or explanation for the fact that bird migration patterns are changing. It seems there is no difference in the meaning of (a1) 'that global warming is occurring is the explanation why patterns of bird migration are changing' and (a2) 'that global warming is occurring is the reason why patterns of bird migration are changing'. 'Reason why' and 'explanation why' can be interchanged in this context without any change in meaning. And it would seem that a similar observation can also be made in the context of normative reasons. It seems that 'reason why' and 'explanation why' in the following two sentences can be interchanged without affecting the meaning of either of the two: (b1) 'that my kids are hungry is a reason why it would be good (for me) to prepare them a dinner' and (b2) 'that my kids are hungry is an explanation why it would be good (for me) to prepare them a dinner'. Crucially, it seems that (b1) can also be paraphrased by (c) 'that my kids are hungry is a reason for me to prepare them a dinner'. Arguably, (c) just corresponds to a standard way of expressing normative reasons. Thus, the argument goes, accounts of normative reasons that don't appeal to explanation in defining/characterising them have an extra burden of explaining why it can be natural to use 'reasons' and 'explanation' interchangeably in both non-normative *and* normative contexts. The proposal here is not that there cannot be an independently plausible story why this is happening that non-explanationist accounts could offer but rather that such accounts will be necessarily more complex than explanationist accounts on which reasons just are explanations in *all* contexts. Non-explanationist accounts will probably involve appeals to the ambiguity of 'reasons' (perhaps by suggesting that only in non-normative contexts do 'reasons' express explanations). No such assumptions are needed on explanationist accounts, which are simpler in this aspect (cf. Finlay 2020: 2).

Third, again, as in the case of deontic explanationism, axiological explanationism is also well suited to explain why mere enabling considerations are not normative reasons (see also Finlay 2020: 2–3 for a suggestion along sufficiently similar lines). Again, this is an advantage of any explanationist view over, say, alternative views that define reasons in terms of reasoning or evidence alone (see Chapters 2 and 4). Being (more or less) able to prepare my kids a dinner seems to matter with respect to why it would be good for me to prepare the dinner. But, again, it certainly doesn't matter as a normative reason, since that I am able to prepare the dinner is not a reason to prepare the dinner. That is, if my ability matters at all with respect to the goodness of *me* preparing the dinner for my kids, it should

matter otherwise than it constituting a normative reason for me to prepare the dinner. It should matter as a mere enabling condition. Explanationists have an easy way of explaining this: exactly as in the case of non-normative explanations, we can also distinguish in the normative case between *explanans* and mere background conditions (e.g. that there was oxygen is an enabling condition, that the match was lit was an *explanans* of the *explanandum* that there was fire). My ability to cook (even if rudimentary) is a mere enabling condition, but that my kids are hungry is a reason for me to prepare the dinner. Similar considerations apply for, say, not F-ing under duress as a mere enabling condition rather than a normative reason to F. It should be noted, though, that I am not suggesting here that for all normative reasons there have to be some enabling conditions, or that ability to F is universally an enabling condition for some F-ing having a normative aspect. The point is only that if there are cases where some conditions seem to matter normatively but appear to be more like enabling conditions rather than reasons, then axiological explanationism has an easy account for that, since with respect to all sorts of explanations, it makes sense to distinguish between enabling conditions and the *explanans*.

A fourth general advantage of axiological explanationism is that it also has the means to explain why entailing conditions are not normative reasons. A fact that entails that F-ing would be good (say, in virtue of being conclusive evidence that F-ing would be good) needs not be a normative reason for one to F. That a highly reliable newspaper reports that it would be good (for people living in cities) to go to the mountains this weekend need not be a normative reason for me, a city-dweller, to go to the mountains this weekend. Now, axiological explanationism can explain why this is so. Typically, *p* cannot be an explanation of why *p*. This is just another aspect of explanations in general that helps to account for why mere entailing conditions are not [always] normative reasons. By the same token, this is a problem for non-explanationist views, like the Reasoning view of reasons (see Chapter 2) and the Evidence view (see Chapter 4). Hence, this constitutes an advantage for axiological explanationism (and the generalised point, for all explanationist views) over these rival accounts of normative reasons.

Fifth, axiological explanationism might have another advantage over some alternative views – for example, the Reasoning view of normative reasons. I say ‘might’ because, as will become clear, this depends on which further assumptions axiological explanationism takes on board. Once more, similarly to deontic explanationism, it might account for the possibility of cases where Moore-paradoxical considerations (and

self-undermining considerations more generally constitute normative reasons for one to F). On axiological explanationism, that the building is on fire but I am unaware of the fire is a reason for me to check the state of the building given that that there is a fire of which I am unaware is an explanation why it would be good (presumably for the aim of self-preservation) for me to check the state of the building. Axiological explanationism might explain such a possibility if it is not also committed to the view that one needs to be able to reason in a fittingness-preserving way from reasons as contents of true beliefs (and other states) to the relevant F-ing. That is, axiological explanationism has an advantage here if it doesn't require that all explanation has to be understood as a valid/good pattern of reasoning that the subject has to be able to follow.

So, this raises some questions about the specifics of Finlay's axiological explanationism, and in particular about the nature of his commitment to a version of the DN model (i.e. the Deductive-Nomological model of scientific explanation; cf. Hempel and Oppenheim 1948; Hempel 1965) as applied to explanations in general. On a standard understanding of the DN model of *scientific* explanation,¹⁰ scientific explanation ultimately just is an argument of a specific sort. It has, in particular, among its premises a reference to universal laws (universal generalisations) and concrete conditions from which the relevant *explanandum* can be deduced. One of the well-known worries with it (cf. Salmon 1971) is that by definition, arguments can be valid and even sound and yet (explanatory) irrelevant; whereas good explanations cannot be irrelevant (see Chapter 5 for more on the fundamental differences between explanations and arguments). For instance, adding some (law-like) necessary truths to a valid argument will not undermine the validity of the argument. But, of course, adding some random necessary truth to an explanation would typically undermine the goodness of the explanation.¹¹ Finlay (2020: 5) tells us that on his account: '*p* is a *complete explanation* of *q* in case *p* is a set of true propositions that

¹⁰ DN models as elaborated and defended by logical positivists have been proposed as models of scientific explanation.

¹¹ Consider Salmon's (1971: 34) well-known example that illustrates the problem of irrelevant truths for the DN model of scientific explanation (a case that is different from the necessary aforementioned truth case in our main text): 'John Jones avoided becoming pregnant during the past year, for he has taken his wife's birth control pills regularly, and every man who regularly takes birth control pills avoids pregnancy.' The example satisfies the constraints imposed by the DN model on scientific explanation (the explanandum is entailed by the universal generalisation and specific conditions), but of course no one would accept that the universal generalisation together with the specific conditions in this example constitute a genuine, good explanation of why John Jones is not pregnant.

logically entails q but doesn't include q .' He then adds that '[l]ogical entailment may seem too strong here, but p should be understood as including any relevant conceptual truths or essential definitions, occupying the role played by scientific laws in Hempel's theory of scientific explanation' (Finlay 2020: 5). Of course, to avoid the objection from irrelevant truths, Finlay needs not only to say that complete explanations not only include 'any conceptual truths or essential definitions' (Finlay 2020: 5) but also that it should *not* include any irrelevant conceptual or other truths. And Finlay does seem to make a remark along these lines. See the relevant passage in Endnote 20, where he claims that the standard objections against the DN model will not apply to his account: 'Although this DN theory is admittedly too broad, the problems of irrelevancies and asymmetries don't arise because the explanantia I'll propose will be clearly both relevant and metaphysically prior to their explananda, given my reductive account of goodness' (Finlay 2020: 33). There is a reading of this remark on which it just amounts to an ad hoc move in the light of the problem of explanatory irrelevancies. To see that reading, compare two sets of considerations. The first, R_1 , contains merely the claim that I promised to call a friend of mine. Clearly R_1 entails that with respect to some background conditions b and an end e (say, the end of respecting one's promises), the probability that e is achieved given that I call my friend and given b is higher than the probability of e without me calling my friend (given b). Now, consider R_2 , which contains the claim that I promised to call my friend later today and the claim that I am 1.89 m tall. R_2 also clearly entails that with respect to b and e (e.g. respecting one's promises and background conditions), the probability of e being achieved given that I call my friend and b is higher than the probability of e without that I call my friend (given b). Thus, on the face of it, both R_1 and R_2 seem to satisfy the DN model of explanation. This would, of course, be problematic for Finlay's proposal, since it doesn't seem that R_2 is a genuine, good explanation of why the probability of respecting the promise by calling my friend is higher than the probability of respecting the promise without me calling my friend. It contains irrelevant information that undermines the goodness of the explanation. Compare R_2 to R_1 , which doesn't contain such irrelevant information. R_1 is a good candidate for explaining the relevant explanandum, not R_2 . On the DN account, however, both should be understood as equally good explanations of the explanandum (e.g. that the probability of respecting a promise given that I call my friend is higher than the probability of respecting the promise without me calling my friend). Now, one might think that a proper treatment of this problem

is not just to add a general remark that the version of the DN account that we have here is such that it doesn't allow there to be irrelevant information among its explanata. Of course, strictly speaking, Finlay is not making such a move. However, what he suggests might sound like it comes dangerously close to such a move. He suggests that the *specifics* of explanata that we have in the case of normative reasons are such that the possibility of cases of irrelevant information in explanata doesn't even arise. It is not clear why this should be so, however. One might want to see *why* exactly this is so. My aforementioned case about promises and promises conjoined with some irrelevant information seem to be just such a case that respect the entailment condition: both R₁ and R₂ entail the relevant explanandum (the goodness of an end given the F-ing of promise keeping). But R₁ seems to be able to genuinely explain the explanandum, since R₂ contains some irrelevant information. It is not clear *on what theoretical grounds exactly* R₂ is ruled out on Finlay's account from counting as an explanata.

However, let us assume here for the sake of the argument that there probably is a way to fix this and to give an independently plausible story of why our R₂ cannot be an explanata on Finlay's account. Returning now to the point about whether Moore-paradoxical cases will count as another line of argument in favour of axiological contextualist explanationism, or instead will constitute counterexamples to it, will ultimately depend on what is meant by 'entailment' in Finlay's version of the DN theory of explanation (recall: '*p* is a *complete explanation* of *q* in case *p* is a set of true propositions that logically entails *q* but doesn't include *q*'; Finlay 2020: 5). For instance, if it is required that the subject (for whom there is a reason) is able to reason [properly] following the relevant pattern of 'entailment' in a truth-preserving way, then Finlay's general account might actually collapse into a version of the Reasoning view seen earlier (Chapter 2). If so, Moore-paradoxical cases will count as counterexamples for this view as well, for general reasons explored in Chapter 2. If, on the contrary, the 'entailment' here is not understood as corresponding to a pattern of reasoning/argument that the subject should be able to instantiate in her reasoning, but be rather of subject-independent, metaphysical sort, then Moore-paradoxical cases may well be accounted for within this approach. Note, however, that this latter option might appear somewhat alien to the initial neo-positivist motivation for the DN model. One might reasonably wonder in what sense the view here is really a 'version' of the DN theory after all. In what follows, I leave this exegetical worry aside. On such an understanding, the Moore-paradoxical cases (e.g. (a) 'the building is on fire but I don't know

that the building is on fire’) can be categorised as cases of normative reasons to F (e.g. checking the state of the building), since Moore-paradoxical considerations (e.g. (a)) explain or ground/constitute why the probability of achieving some end *e* (e.g. self-preservation), given the F-ing (e.g. checking the building), is higher than the probability of achieving *e* without F-ing. Now, if this interpretation is on the right track, then axiological contextualist explanationism doesn’t predict that contrary to our pre-theoretical judgments, Moore-paradoxical considerations can never be normative reasons for one to F. Respecting the pre-theoretical judgments that they can be reasons to F is then another *prima facie* consideration in favour of axiological contextualist explanationism (given that the aforementioned assumptions are in place).

Worries. Finlay (2020) considers five objections (‘puzzles’) to the explanation of goodness-based accounts of reasons (axiological explanationism in general). He suggests that given the specifics of his more elaborated axiological contextualist explanationism, all these objections can be dealt with. That is, even if these objections undermine the plausibility of some more rudimentary versions of axiological explanationism, his contextualist version of it can be maintained. If this is right, then, of course, this would also constitute an additional argument in favour of axiological contextualist explanationism – namely, it has an advantage over other axiological explanationist accounts.

Among the five objections that Finlay considers are (i) the problem of the ‘right kind’ of reasons (in short: how are the ‘right kind’ of reasons for attitudes even possible, that is, reasons connected to the fittingness of an attitude rather than to, say, some benefit of holding it, on a value-based account of reasons?); (ii) the objection from normative facts being themselves reasons (in short, if that it would be good to F is a reason to F, then axiological explanationism would predict that some considerations are auto-explanatory, which is problematic); (iii) the objection from evidence being sometimes a reason to F (in short, normative testimony/evidence that F-ing would be good may sometimes constitute normative reasons to F without explaining why F-ing would be good); (iv) the objections from subjective; and (v) motivating reasons (in short, the existence of problem cases where an agent seems to have a reason to F but unbeknownst to the agent F-ing would actually not be good, and the existence of problem cases where F-ing would be good but unbeknownst to the agent the relevant reasons apparently constituting consideration R are false and thus cannot explain why F-ing would be good). We will not go through the details of *all* these ‘puzzles’ and Finlay’s treatment of them here. Instead, let me

merely sketch the general recipe of Finlay's responses to these. The guiding idea in Finlay's responses to all these is his appeal to contextualism about either the relevant ends in a situation or the contextualism about the relevant base-information or the contextualism about both of these. Among the elements that can 'move' in axiological contextualist explanationism are, for instance, the focus on the agent's desired/preferred outcome as the relevant end, the focus on a [potential] adviser's preferred outcome as the relevant end, the focus on some of the agent's subagential aims (as is suggested with respect to the problem of the 'right kind' of reasons, for instance; see ahead for discussion), the focus on the agent's knowledge/beliefs as the relevant information-base, the focus on the knowledge of a better-informed [potential] adviser as the relevant information-base, and so on. In short, according to Finlay, selecting the right focus, the right interpretation of the relevant end and information-base can account for all the apparent 'puzzles' in (i)–(v) without giving up value-based explanationism. The high flexibility of the account is fundamentally what does the job in dealing with these objections, according to Finlay.

However, it is also this high flexibility that one might find somewhat problematic with this account. The first general worry for axiological contextualist explanationism is a worry that any contextualist approach with respect to any normative concept has to tackle. Namely, one needs to explain how we are supposed to make sense of a view on which normative concepts are always relativised to some further aspect and even the relevant normative concept *sans phrase* has to be understood as relative to some aspect, say, a common ground. This is a very general worry and may just conceal differences in fundamental theoretical commitments that meta-normative theoreticians might have. Someone who is not attracted to contextualism in the normative domain in general may also find Finlay's fine-grained contextualism about normative reasons unattractive.¹² The worry here is the apparent lack of independent theoretical motivation for such an extreme contextualism given that our pre-theoretical judgments about reasons to do something or to have an attitude don't seem on the face of it to be so radically context-dependent. At least on the surface, they don't always appear to be relative to some given end or some given information-base. At any rate, a contextualist about reasons, axiological or not, may need to provide some error theory to explain why we are

¹² Thanks to Jacques Vollet for the suggestion that this may indeed be a serious problem with axiological contextualist explanationism.

wrong in pre-theoretically thinking that reasons are not so *radically* context-sensitive. Granted Finlay's contextualism can account for both 'the intuitions of Humean internalists about reasons as the result of privileging the *agent's* preferences for outcomes (e.g. Hitler may well have had "no reasons" to refrain from genocide, assuming that none of his preferred ends were threatened by it)' (Finlay 2020: 7) and 'the intuitions of externalists, and the categorical nature of moral claims, as the result of privileging the *speaker's* or *audience's* preferences for outcomes (*of course* Hitler had reasons to refrain from genocide, given the harm it inflicted on innocent people!)' (Finlay 2020: 7). But I doubt that either Humean internalists or externalists would agree with Finlay's proposal. Humean internalists would not accept 'reasons' in the externalist sense as genuine reasons at all, nor would externalists accept 'reasons' in the internalist sense as genuine reasons.¹³ Contextualism has the burden of explaining not only why both are right in a sense but also why strictly speaking both are wrong. The worry here is not that it cannot be done but that it is an extra challenge for a theory of reasons that has to be met.

Our second and more specific worry concerns the account's treatment of reasons for attitudes and beliefs in particular. To see this worry, I would like first to return briefly to Finlay's response to the 'puzzle' of the 'right kind' of reasons (cf. aforementioned (i)). First of all, the observation that simple value-based explanationist accounts of reasons cannot explain why there would be any right kind of reason for attitudes at all seems right.¹⁴ Finlay captures the worry precisely when he writes:

Something is a reason of the 'right kind' for an attitude if and only if it makes the attitude fitting, but raising the probability of an end desired by the speaker, audience, or agent is neither necessary nor sufficient for fittingness. (Finlay 2020: 8)

Recall that that someone's threatening me to admire them as contrasted to them being admirable is a 'wrong kind' of reason to admire them (cf. Section 1.4); believing *p* being practically advantageous is also often seen as being the 'wrong kind' of reason to believe that *p*. According to evidentialists about reasons to believe, considerations of this sort are not normative reasons at all to believe that *p*; pragmatists reject this and see practical considerations as possible genuine normative reasons to believe (see

¹³ And this seems to be exactly what we find in the literature. See, for instance, Williams (1989) for internalism and Parfit (2011: 58–110) for arguments against internalism.

¹⁴ And a little further, we will see briefly a value-based view of reasons that bites the bullet and endorses this conclusion.

Chapter 6 for more on the evidentialist–pragmatist debate about reasons to believe). Now, even if we put aside for a moment the debate over whether the ‘wrong kind’ of reasons are genuine reasons, the fact that simple value-based explanationist accounts of normative reasons entail that there are no ‘right kind’ of reasons for attitudes (e.g. belief) is already dramatic. For it would be a dramatic cost to a view if it were to entail that that a consideration makes *p* more probable or entails its truth cannot be by itself a normative reason to believe *p* and that a consideration that entails that someone is admirable cannot be by itself a normative reason to admire them.

How exactly is Finlay’s contextualism supposed to help axiological explanationism to respond to this problem – that is, to make axiological explanationism compatible with the existence of the ‘right kind’ of reasons? The move is to treat what has been known in the literature as constitutive norms of attitudes as possible salient ends in a situation, thereby making the ‘right kind’ of reasons fall within the definition provided earlier (REG +ERT). Thus, Finlay writes:

[w]hen talking about reasons for *attitudes* there is a competing source of salience for ends, in the reference to the attitudes themselves. These are attitudes that are commonly said to have subagential, ‘constitutive’ ends of their own, which need not be ends desired or intended by anybody. It is commonly said, for example, that the constitutive end of belief is *truth* (or knowledge). How to precisify this idea of ‘constitutive ends’ is controversial, but for our purposes this doesn’t matter.[. . .] All we need is that talking about ‘belief’ is sufficient, in normal contexts, to make salient an end like truth. (Finlay 2020: 8)

This then brings us to the specifics of our worry. The worry is that it is not clear in what sense, if at all, constitutive aims can count as one’s salient ends and be of genuine use for a value-based account of reasons. Finlay talks about ‘constitutive ends’ of attitudes, but it is clear from the context (and from the references he appeals to) that this talk should be understood in the sense of the ‘constitutive aim’ of attitudes as it has been commonly discussed in recent debates. But in these debates, the ‘aim’ talk is understood at best as a metaphor. An attitude is not the kind of thing that can aim at anything at all. Let us focus specifically on the ‘aim’ of belief.¹⁵ As it is commonly understood, to say that the belief aims at truth is to say that truth sets the standard of correctness or fittingness for belief, namely a belief is correct or fitting just in case it is true. It is assumed that a standard

¹⁵ See Fassio (2015) for a comprehensive overview of the aim of belief debate.

of correctness of an attitude defines or is essential for the sort of attitude in question. Thus, belief is an attitude that is defined as an attitude that is correct just in case it is true. Its standard of correctness allows us to distinguish it from other attitudes that have distinct standards of correctness (e.g. guessing is not an attitude that is defined as an attitude that one can correctly have if and only if one is guessing correctly). A common way to understand the standard of correctness for belief is to appeal to a fundamental norm of belief, the norm that defines what belief is. Various proposals exist (with respect to the content and with respect to the form of the norm), but let us assume, just for the purpose of illustration, that the relevant norm is something like: (TN) one ought to [believe that p only if p] (see Velleman 2000; Wedgwood 2002, 2013; Boghossian 2003; Shah 2003; Engel 2004, 2013; Shah and Velleman 2005; among many others). Thus, the idea is that only attitudes that are subject to (TN) are beliefs; (TN) defines what sort of attitude belief is. An attitude has to be subject to the fundamental (TN) if it is to count as belief at all, according to a standard approach with respect to the aim of belief. The suggestion is that the nature, the very essence of belief, is to conform to this norm. This is, then, the sense in which beliefs are said to aim at truth (the same applies for the alternative of the knowledge norm of belief). But if we understand the aim of belief in this sense, it is not clear how it could be one of the salient ends that can be plugged into the axiological contextualist explanationist definition of reasons. Here are two more specific considerations that suggest that the aim of belief cannot be a 'competing source of salience for attitudes' (see the aforementioned quotation) for an axiological explanationist and thus cannot help Finlay's attempt to block the objection from the 'right kind' of reasons.

The first reason why the appeal to constitutive aims of attitudes as presented earlier cannot help an axiological explanationist is simply that allowing for standards of fittingness (e.g. fundamental truth norm of belief) to play an indispensable and irreducible role in an analysis of reasons is giving up on the purely 'axiological' aspect of axiological explanationism. In other terms, appealing in this way to standards of fittingness makes it unclear in what sense the view would still count as a value-based view of reasons. Of course, it is not a problem per se. Maybe fittingness is indeed fundamental and should be appealed to in explaining other normative notions. It is only that this move doesn't seem available to someone who takes values as more basic and aims to reduce reasons to goodness. At best, the resulting view would count as a hybrid view combining both value and fittingness in explaining reasons. At this point, one might think

that existing alternative non-hybrid views may be preferable to this account on the basis of simplicity and parsimony. Be that as it may, the first point is that explanationists that are also value purists – that is, who want to appeal only to goodness in explaining reasons – cannot really appeal crucially to irreducible standards of fittingness in their explanatory analysis of normative reasons.¹⁶

The second reason why the appeal to constitutive aims of attitudes won't help axiological explanationists to deal with the worry from the existence of the 'right kind' of reasons is that it is not clear in what sense, if any, the aim of belief could be on a par with the ends that we pursue. Finlay talks about 'constitutive ends' of attitudes constituting 'a *competing* source of salience for ends' (see the aforementioned quotation, emphasis added). So the idea would be that standards of fittingness are competing and thus are on a par with one's ends or desired/preferred outcomes. But how should we understand such a suggestion? It is not clear that the appeal to constitutive aims being 'subagential' (cf. Finlay 2020) really helps here. The truth norm of belief, for instance, is not a subagential end one is having. The truth norm of belief is an *abstracta*, and as such, it is not, on the face of it, something one can properly have as an end. One could, of course, set oneself an end of believing only truths and avoiding falsehoods as an end towards which one is striving. But that's not how the constitutive aim of belief is understood in the literature. The constitutive aim of belief, as a standard of correctness, defines what belief is. It is a norm that defines the attitude of belief (belief is just the sort of attitude that is subject to the truth norm). Definitions, in terms of norms or otherwise, are not the kind of thing that one can have. There is no useful sense for analysis of reasons, it seems, in which a definition in terms of constitutiveness can be said to compete with one's desired outcomes. They are not the sort of thing that can compete, let alone be meaningfully compared. Definitions, contrary to one's ends (subagential or not), don't obtain and are not realised. Thus, our second problem here is that it doesn't seem that one can plausibly

¹⁶ Note also, in passing, that Finlay's proposal that REG+ERT, together with the constitutive understanding of ends of attitudes, provides an explanatory analysis of fittingness seems overenthusiastic. Finlay (2020: 9) claims that '[a]n added bonus [of the view] is that it also offers an explanatory analysis of fittingness, which some philosophers claim to be an unanalyzable primitive:[...] for an attitude to "fit" its object is for that attitude to realize or promote its constitutive end when directed at that object'. The problem, however, is that given the standard understanding of constitutive aims (see the main text earlier), such an analysis would be problematically circular. For it would amount to the claim that for an attitude to fit its object is for that attitude to satisfy its fittingness conditions (= realize its constitutive end), which is uninformative and thus unhelpful as a putative explanatory analysis of fittingness.

supplement the REG+ERT account of reasons with the thought that an attitude's constitutive aim is an end whose increase in probability of obtaining, given one having the attitude, would be explained by a consideration that would count as a reason to have the attitude in question. In short, it just doesn't seem possible to plug constitutive aims (correctness conditions) into the REG+ERT analysis without making it somewhat incoherent. It would seem, then, that while running away from the 'right kind' of reasons problem, Finlay runs into the 'wrong kind' of ends problem.¹⁷

Now, it might still be the case, as observed earlier, that an agent has an end of believing only truths and avoiding falsehoods. Such an end could, of course, plausibly function in the REG+ERT account as a salient end in a situation. But it should be noted that this sort of end plausibly can be reduced to believing truth and avoiding falsehood being the agent's preferred/desired outcome. Such an outcome, of course, can obtain or fail to obtain, be promoted, realised, or made more probable. But it is crucial to note that focusing on this sort of end as the salient end with respect to beliefs (and, *modulo* necessary amendments, to other attitudes as well) would consist in falling back to a simpler axiological explanationism that cannot explain *why* the 'right kind' of reasons are normative reasons at all for attitudes, given that, as Finlay (2020: 8) notes, 'raising the probability of an end desired by the speaker, audience, or agent is neither necessary nor sufficient for fittingness'. The 'right kind' of reasons for attitudes are essentially connected to fittingness of the relevant attitudes. And despite it being much more sophisticated and indeed it being an ingenious move,

¹⁷ See also Way (2013: 19ff) for a somewhat related objection to an arguably less sophisticated version of axiological explanationism on which reasons are defined in terms of partial explanation of goodness of F-ing for its own sake or instrumentally (Way's focus is specifically on pro-attitudes). The objection there is that such a view seems to be committed to an implausible assumption that 'pro-attitudes towards outcomes which are good in some respect are good for their own sake' (Way 2013: 22). An example concerning egalitarians illustrates the following point:

Consider an outcome in which wealth is distributed equally, but everyone is extremely poor. If egalitarians are right, this outcome is good in some respect. But even egalitarians can admit that this outcome is not good overall. And it does not seem plausible that favouring this outcome is good for its own sake. It does not seem to be good for its own sake to hope for everyone to be extremely, if equally, badly off, or to be glad if this outcome comes about. (Way 2013: 23)

The parallel with our objection here is that on the view reconstructed and rejected by Way, an axiological explanationist aims to specify the goodness in a way somewhat intrinsic to the attitude. Finlay's account appeals to constitutive aims, whereas Way's constructed axiological explanationists focus on overall goodness of attitude for its own sake. Both versions of axiological explanationism that focus on F-ing's goodness are problematic as is shown in main text and by Way's counterexample.

the contextualist version of axiological explanationism still has the same problem of explaining the existence of the ‘right kind’ of reasons. I conclude tentatively that this alone is a strong argument against this approach. The rest of the present section is devoted to an exploration of an alternative recent axiological explanationism that proposes a different move with respect to this problem.

Alternative axiological explanationism: in virtue of promoting valuable states of affairs. In the light of problems met by the contextualist version of axiological explanationism, someone who is sympathetic to the idea that reasons should be understood in terms of value (e.g. a value-first proponent) might want to explore alternative options for defining normative reasons in terms of explanation and goodness. One such alternative is elaborated in a recent and promising value-based account of reasons by Barry Maguire (2016). Let us rehearse rapidly the main tenets of his proposal.

The official version of the view is as follows:

Value-based theory of reasons: Some fact of the form [ϕ would promote S] is a reason to ϕ if and only if and if so in virtue of the facts that ϕ would promote S and that S is valuable. (Maguire 2016: 237)

There are a few things to note before comparing this view to Finlay’s account. First, without going into too much exegetical detail, it may nonetheless be useful to unpack the proposal slightly. On Maguire’s account, reasons are for options.¹⁸ Options can be understood in a variety of ways, according to him, so ‘our options are the things we choose between or rationally control – whether these are actions, omissions, activities, or plans’ (Maguire 2016: 239). That F-ing would promote S is on Maguire’s account a ‘canonical fact’.¹⁹ F-ing is the option that one may take (choose/rationally control). S is a placeholder for ‘state of affairs’. Not all canonical facts promote valuable states of affairs, and not all canonical facts are reasons (cf. Maguire 2016: 236). Value on this account is ‘a gradable monadic property of states of affairs’ (Maguire 2016: 237). Value is understood in a neutral way and not as ‘value-relative-to-me-or-you’ – for example, ‘the disvalue of the drowning of the small child in Peter Singer’s pond has nothing to do with you or me, the passers-by, and *a*

¹⁸ He writes: ‘The central claim is that to be a reason for an option is to be a fact about that option’s promoting some state of affairs, on the condition that that state of affairs is valuable’ (Maguire 2016: 237).

¹⁹ More precisely: a canonical fact is ‘a fact of the form [ϕ would promote S], where ϕ is some agent’s taking some option in some situation and S is a state of affairs’ (Maguire 2016: 236).

fortiori not to any reasons we might have to care or want to help. It has entirely to do with the child, his suffering, his loss of life prospects, and so on' (Maguire 2016: 238–239). Note equally that value here is supposed to be final, not instrumental. Promotion is also characterised in neutral terms. Maguire wants his view to be compatible with a number of possible interpretations of promotion: 'an option *promotes* a state of affairs by instantiating it, causing or partially causing it, constituting or partially constituting it, preventing the preventing of it, or non-superfluously probabilifying it' (Maguire 2016: 238) (here it is important that '[t]he relevant valuable state of affairs may be instantiated by the action itself'; cf. Maguire 2016: 239).

Second, the view qualifies as a version of explanationism given how its details are worked out. An important thing to note in this respect is that Maguire's proposal is a revisionary one in the sense that he doesn't take the surface structure exhibited by our ordinary statements or reasons ('reasons talk') to correspond exactly to metaphysics of reasons. Strictly speaking, on his account, only canonical facts can be reasons; canonical facts have the form of 'F-ing would promote state of affairs S'. Our ordinary talk of reasons doesn't reflect this form, at least not on its surface. We commonly say 'that it is raining is a reason to take an umbrella'. But 'that it is raining' doesn't have (on its surface) the form 'taking an umbrella would promote state of affairs S'. This is no problem, according to Maguire (2016: 254), since '[q]uite generally, the considerations we actually offer as reasons are those that are saliently needed, in the specific conversational context, to pick out a larger explanatory structure'. Maguire (2016: 254, fn 37) also notes, referring to a point from David Lewis, that '[t]his is an instance of a more general distinction between metaphysical explanations and considerations that one can give as an answer to a "why" question in a specific context'. So, we can imply here that canonical facts that are reasons are these larger explanatory structures or, in other terms, genuine metaphysical explanations. Thus, for Maguire, our common reasons' statements are '*good representatives* [...] either for the basic reasons themselves or for other significant chunks of this overall structure' (Maguire 2016: 255; the term 'good representatives' is attributed to Fogal 2016). Two examples illustrate Maguire's point. First,

[s]uppose that Frank enjoys flowers because they remind him fondly of his grandmother, or because he likes to marvel at the fragile beauty of nature, or because he likes their smell. Any one of these facts – the fact that flowers remind him of his grandmother, for instance – may be said to be a reason to give him flowers. Each is part of an explanation of why giving him flowers

would make him happy with some intensity, for some duration – which is the basic reason. (Maguire 2016: 255)

And, second, '[t]ake the fact that Teresa's tyre is flat. Since Teresa drives to work, the fact that her tyre is flat partially explains the fact that giving her a lift will realise the state of affairs of her being in work on time. That's a reason to give her a lift to work' (Maguire 2016: 255, fn 39). Thus, given how the details of the account are worked out (and, in particular, given the specifics of his proposed metaphysics of reasons, as opposed to our common and somewhat loose talk of reasons),²⁰ Maguire's proposal qualifies as a version of axiological explanationism. *Metaphysical* explanation is a key element in his analysis of normative reasons.

Turning now more specifically to a comparison of Maguire's account to Finlay's, we can note that a major difference between these two is that where the latter appeals to R explaining the goodness of *F-ing* for some end/outcome *e* (i.e. *F-ing* increasing the probability of *e* obtaining), the former appeals to the value of *an end/outcome* (e.g. state of affairs) that would be promoted by *F-ing* (cf. Maguire 2016: 234, fn2) in defining reasons. Adapting a distinction from Way (2013), we can say that the former sort of value-based account of reasons focuses on *F-ing*-based goodness (in Way's original terminology: attitude-based); whereas the latter's focus is on object-based goodness/value (the goodness of the relevant end).

The focus on neutral value in the case of Maguire versus a fine-grained contextualism of end-relative goodness/value in Finlay's account is another notable difference between the two. Thus, someone having the sort of worry we alluded to the aforementioned about Finlay's contextualism being, well, too relativised (to ends and information-bases) need not have a parallel worry for Maguire's object-based value-neutral approach. Maguire's proposal doesn't involve that level of context-sensitivity (given that *F-ing* would promote a neutrally valuable state of affairs, the canonical fact of the form [*F-ing* would promote *S*] just is a reason *sans phrase* to *F*, and this need not be relativised to some further parameter on Maguire's account).

²⁰ See also: 'Strictly speaking, only the basic reasons play a grounding role in determining the total weight of reason favouring each option, and hence (by way of further principles not defended here) in turn, in determining what you ought to do' (Maguire 2016: 254). For matters of brevity we haven't defined what's meant by 'basic reasons' here, but it has to do with the value of distinguishable states of affairs (cf. Maguire 2016: 252). For our purposes this may be left unspecified here.

Yet another, even more important difference for our purposes between the two views is that Maguire, contrary to Finlay, baldly rejects the very existence of genuine ‘right kind’ of normative reasons for attitudes. Finlay attempts to provide an explanation of these within the framework of axiological contextualist explanationism, while Maguire embraces the consequence that his account may rule out the ‘right kind’ of reasons for attitudes, where the ‘right kind’ of reasons are understood as ‘facts that make attitudes *fitting*’ (Maguire 2016: 240). Maguire suggests that this consequence need not be a problem, for it is not clear that there is a unified reasons relation anyway. The suggestion is that the focus of his account – namely, the ‘reason-for-an-option relation’ – may well not be the same relation as the “‘right-kind’ of reason for an attitude relation’ (cf. Maguire 2016: 240).

In a more recent publication, Maguire (2018) provides a further argument against the unity thesis (i.e. that a reason for an option and the ‘right kind’ of reason for an attitude are relations of the same sort). Before examining whether this more elaborated argument undermines the ‘unity thesis’, note the details of Maguire’s dialectic here with respect to the objection from the ‘right kind’ of reasons: ‘[h]owever, whether we should accept any such “unity” argument [e.g. argument against his value-based account relying on the premise that the “right kind” of reasons for attitudes are genuine normative reasons], or rather reject the relevant “unity” premise, will depend in part on the strength of the case for the Value-Based Theory of Reasons itself’ (Maguire 2016: 240). I tend to think that the qualification ‘in part’ is doing the heavy lifting here. Of course, if there is a strong *prima facie* case in favour of the value-based account of reasons, it will matter for the overall assessment of the view; however, at the end of the day, if the ‘right kind’ of reasons objection hasn’t received an independently plausible response *and* there are alternative initially plausible accounts of reasons that don’t have this problem (and are not subject to some further unanswerable challenges), then the conclusion should be clear: the value-based account cannot be rationally upheld. So, it is crucial for our overall assessment of the view to focus specifically on whether Maguire’s suggestion that the ‘unity thesis’ should not be accepted really works out. This is specifically what we will now focus on in the concluding part of this section. (The details of an alternative, new theory of reasons that is not subject to the objection from the ‘right kind’ of reasons and is independently plausible will be elaborated in Chapter 5.)

The main argument in Maguire (2018) is an argument against the possibility of the ‘right kind’ of reasons for affective attitudes. The

argument there is silent about the ‘right kind’ of reasons for beliefs (cf. ‘I take no stand on whether my arguments can be generalized to conative attitudes (the action-oriented attitudes such as preference and intention) or epistemic attitudes’; Maguire 2018: 782). However, a reference is made to a distinct manuscript where a ‘related account of normative support for epistemic states’ (cf. Maguire 2018: 782, fn 5) is defended, which is, presumably, Maguire and Woods (2020). However, note, crucially, that the proposal defended in that paper is highly contentious; it is basically a rather radical version of pragmatism about reasons to believe. A very popular view in contemporary epistemology – namely, evidentialism – according to which normative reasons to believe are evidential (i.e. truth-conducive) – is clearly incompatible with the account defended in that paper. Thus, one should be advised that taking on board the value-based explanationist response to the ‘right kind’ of reasons objection in general might have, at the end of the day, a rather high theoretical cost. It is far from being neutral with respect to substantive theoretical options. However, for the sake of the argument, let us put the ‘right kind’ of reasons for beliefs question aside for the time being. We will return to the debate about the nature of reasons to believe and the pragmatism–evidentialism opposition in Chapter 6.

The overall structure of the argument in Maguire (2018) is really simple and straightforward. Indeed it may appear to be rather appealing at first sight. It can, I think, be summed up as follows:

1. Normative reasons are non-strict, essentially contributory, and essentially gradable in a sense to be explicated.
2. For all x , such that x is a ‘right kind of reason’ consideration that normatively supports an affective attitude A , x is a fit-making consideration for A .
3. Necessarily, no fit-making consideration is non-strict, contributory, and gradable in the relevant sense.
4. Necessarily, no ‘right kind of reason’ consideration for an affective attitude A is a normative reason for A .

In other terms, Maguire identifies what he thinks are essential features of normative reasons; he, then, distinguishes normative reasons from fit-making considerations and argues that fit-making considerations have none of the essential features of reasons (an alternative that seems to be evoked in the last section of the article is that fit-making considerations don’t have *all* of the essential features of reasons). Crucially, the considerations that qualify as the ‘right kind of reasons’ for affective attitudes are all

fit-making considerations. It is concluded on this basis that no consideration that qualifies as a ‘right kind of reason’ for an affective attitude is a genuine normative reason for the affective attitude in question.

What exactly are these allegedly essential features of normative reasons? According to Maguire, these are being non-strict, being contributory, and being gradable. Being a non-strict normative consideration is basically to satisfy the constraint of holism (cf. Dancy 2004: chapter 5) – that is, a constraint according to which considerations that are reasons to F never ‘justify or require anything on their own’ (cf. Maguire 2018: 784). The idea here is that reasons on their own don’t imply anything about what normative considerations – for example, reasons against F-ing, there might be. To illustrate the idea, consider the following: ‘[t]he fact that the child will drown, for instance, would fail to justify wading into the pond if the alternative were defusing a bomb on dry land, set to destroy Chicago’ (Maguire 2018: 784). Reasons are not assumed by Maguire to be essentially non-strict (see Maguire 2018: 801); however, fit-making considerations are assumed to be essentially strict. Maguire (2018: 785) claims that ‘facts about what you overall ought to do are the paradigmatic strict facts’ and ‘[so called reasons of the right kind for affective attitudes] are themselves, if you like, a kind of ought fact’ (Maguire 2018: 780).

The gradability constraint on reasons, according to Maguire (2018: 783), just is that a reason ‘has some gradable property that is usually called its weight’. That is, ‘reasons are essentially *gradable*’ (Maguire 2018: 785). Again, Singer’s drowning child example is supposed to illustrate the gradability constraint: ‘the “saving the child” reason has more weight than the “getting your clothes muddy” reason’ (Maguire 2018: 783).

The contributory constraint on reasons is the idea that it is essential for reasons to contribute – that is, to interact in determining the overall normative status of F-ing. According to Maguire (2018: 785), ‘[c]-ontributoriness is the property of playing a specific role in a “weighing explanation” of an overall normative fact.’ How exactly does this work? The idea is that reasons have ‘weights’, and that the net weight of reason supporting an option is somehow explained by the weights of all reasons for and against that option and the fact that ‘these are all the reasons bearing on the given option’ (Maguire 2018: 784). It is assumed that ‘[t]hese facts about the net weight of reason supporting each option (together with some fact about what options there are) in turn explain the fact that there is *most* reason in support of some option’ (Maguire 2018: 784). And this, then, is supposed to explain that one *ought* to take the relevant option (it is assumed that one ought to do what one has most

reason to do). The ought fact – that one ought to take the relevant option – is on Maguire’s view an overall normative fact. ‘It obtains in virtue of all the contributory normative facts together with a “normative totality fact” to the effect that these are all the relevant normative facts’ (Maguire 2018: 784). Finally, Maguire (2018: 785) specifies that typically reasons fulfil this function of contributoriness ‘by combining or by competing with each other’.

Now, what about Maguire’s argument then? It does appear to be valid, the conclusion seems to follow if the premises are true, and there doesn’t seem to be any equivocation going on – the terms used in the premises do seem to have the same meaning as in the conclusion. How about soundness? I would like to submit that there are reasonable and distinct considerations speaking against each of Maguire’s premises. Let me elaborate on these a bit.

The first, and arguably the most problematic, consideration concerns premise 1. To see the problem, recall the overall dialectic we are facing at this point. The fundamental question that we are addressing in this part of the present section is whether any axiological version of explanationism is plausible overall and, in particular, does better than deontic explanationism, which we have already put aside due to a number of problems it faced. Now, we observed that a major problem for any axiological explanationism is the so-called right kind of reasons problem – it is not easy to see how any value-based explanationist proposal can account for normative reasons for attitudes, beliefs, and emotions in particular. Maguire’s proposal on this matter has a radical boldness. He seems to suggest that the problem can be avoided, since according to him there are no genuinely normative reasons for attitudes. But on the face of it, attitudes are paradigmatic examples of F-ings that can be supported by reasons. So his argument for that specific conclusion has to be convincing. However, note that his argument for that conclusion – that is, that there are no reasons for attitudes – relies on the assumption that (arguably, broad) explanationism about reasons is true, namely it is part of his premise 1. He assumes that reasons are essentially contributory, and as we saw earlier, this property just is an instance of a broad explanationist approach – that reasons are things that play a ‘role in a “weighing explanation” of an overall normative fact’ (Maguire 2018: 785). But to presuppose the truth of explanationism at this point would amount to begging the question in the context of the present discussion. The supposed truth of explanationism *is* what’s under question in our discussion. One could, of course, reply that that’s not the context of Maguire’s discussion. This might be correct and that’s why

I only claim that the move we are discussing *would* commit a proponent of explanationism to begging the question, not that Maguire is actually begging the question against us in the present context. However, note that the 'right kind' of reasons problem is a problem for any axiological explanationist account, including Maguire's proposal. And consequentially, presupposing the truth of general explanationism at any point within the context where the 'right kind' of reasons problem is discussed, is problematic, especially in the context where other, non-axiological versions of explanationism have already been set aside as not fully satisfactory.

One could reply to this by pointing towards the very general aspect of Maguire's contributoriness constraint on reasons. It does not assume *axiological* explanationism in an argument against an objection to axiological explanationism; it only assumes the truth of explanationism in general. One could insist that it is not a flaw to assume the truth of the general theory X in one's argument in favour of a theory XN and then argue in favour of the specific version XN. Surely, it is true that taking on board this sort of assumption need not *always* constitute a flaw in an argument. However, again, recall the present dialectic. We are investigating explanationism about reasons, in all its forms. Thus, within our present context, assuming that normative reasons essentially are things that play a role in a sort of explanation of overall normative facts would constitute begging the question. We have already put aside other versions of explanationism. Thus, we cannot at this moment in our discussion merely presuppose that a version of explanationism has to be true.

Moreover, some of the objections that we have explored against deontic versions of explanationism, and in particular against the very idea of there being a special sort of 'weighing explanation' of normative facts, can be brought up again in the present context. It is not really clear how the supposed weighing explanation works. If it is an explanation at all, it has to be radically different from any other explanations we know. Maguire doesn't propose an argument for the claim that the 'net weight of reason' supporting an option is explained by facts about all the reasons for and against that option, and thus that the facts about reasons for and against an option explain the facts about most reason for an option. It is not clear, in particular, in what sense the fact that a friend of mine offers me his beach house for free for the holidays (and we can add to this the fact that spending my holidays at a beach house would promote my happiness) *explains* or *partially explains* the fact that I have most reason and indeed ought not to go on holiday to that beach house. (Recall, I ought not to go on the holiday, since I just broke my leg.) It is not clear in what sense we

talk about explanation at all when we talk about the alleged ‘weighing explanations’ of normative facts (compare to the arguments from Brunero 2013, 2018). At this point, appeal to the metaphors of ‘weight’, ‘weighing’, and ‘balance’ only makes the discussion more obscure.

Another objection against what we have identified as Maguire’s premise (1) comes from a different angle. As David Faraci (2020) has recently observed in his critical discussion of Maguire’s article, the difference between reasons for action and ‘right kind’ of reasons considerations for attitudes might be explained by the difference between action and attitudes. More specifically, the contexts where there are reasons for action tend to be such that the subject cannot take several options at the same time. There is an actual limitation in terms of what the subject can do. Hence, it is not surprising to observe a competition among reason considerations. However, in contexts where we have ‘right kind’ of reasons considerations, we do often have, as Maguire himself observes, the possibility to have a number of distinct attitudes. Therefore, it is not surprising that we observe the absence of the competition aspect with respect to ‘right kind’ of reasons considerations.²¹ In short, we can take this to show that the aspects that Maguire takes to be proper aspects of normative reasons, indeed, as being essential characteristics of reasons, are merely side effects of his focus on reasons for action (and given certain background assumptions; see the previous paragraphs). But the fact that reasons for actions are always non-strict, gradable, and contributory (let’s grant this for the sake of this line of thought) doesn’t yet allow us to conclude that these features are features of normative reasons in general. All that follows is only that these might be features of normative reasons for actions. The fact that we are able to have various distinct attitudes with respect to some relevant objects may well explain why reasons for attitudes don’t have these same features that reasons for action have. In other terms, we don’t have sufficient grounds for holding that the characteristics that Maguire identifies are

²¹ See Faraci’s point:

What Maguire’s cases illustrate, I submit, is not that FMCs [that is, fit-making considerations] are not reasons but rather that the contexts in which we tend to think about reasons for action and reasons for attitudes differ, in part because actions and attitudes themselves differ. We tend to think about how we and others should act in contexts where the reasons for them compete, because actions are frequently in tension, and we need to pick one. By contrast, as Maguire’s cases show, we often consider how we and others should feel in contexts where there are no relevant tensions, because we can feel many different ways about different things, or about different aspects of the same thing. In such contexts, there is no reason to expect competition amongst the relevant FMCs, and therefore no reason to think they are not reasons. (Faraci 2020: 230)

essential to reasons as such, as opposed to being aspects of reasons for action, because of specifics of action as opposed to attitudes.

Second, and independently from our discussion of premise (1), one might also question the claim (i.e. premise (2)) that all 'right kind' of reason considerations that normatively support an [affective] attitude are fit-making facts, which are 'if you like, a kind of ought fact' (Maguire 2018: 780). If this verdict is on the right track, then we have a case where a consideration – say, the general, abstract consideration that it is raining – is a reason that speaks in favour of holding an attitude, namely feeling pessimistic about the prospects of a pleasant run, without also being a fit-making consideration for one to feel pessimistic about the prospects of a pleasant run. Given that it is also really hot outside, it is not a fitting attitude for the subject in this case (one can, of course, vary the case by changing the attitude of feeling pessimistic to being sad, being disheartened, unhappy, demoralised, and so on, without affecting our assessment of it). Of course, Maguire is free to deny that this sort of case is even possible. However, if one does this, then it is not clear how one could maintain the possibility of parallel cases involving action – for example, running/intending to run in the heat and the rain; see Nair (2016) and our discussion in Section 3.2, where we took these cases as showing that the additivity/accrual of reasons to F can fail. If this is on the right track, then reasons for attitudes might not always be fit-making facts, whatever exactly this implies.

The third and distinct complaint concerns premise (3) – that is, the claim that no fit-making consideration can be non-strict, gradable, and contributory (alternatively: at least have one of these characteristics). If we stick with a fairly common view within the rich literature of contemporary philosophy of emotions – with which oddly enough Maguire never engages, but really should have, given that the debates about fittingness of emotions and reasons for emotions have gone on in that field for decades and well before they made an appearance in general normative philosophy – an emotion is fitting when it respects its *formal object*. Here is a classic contemporary statement of this idea, applied to the example of fear: 'The formal object of fear – the norm defined by fear for its own appropriateness – is the Dangerous' (de Sousa 2002: 251; see also de Sousa 1987; Deonna and Teroni 2012: 1–12; Scarantino and de Sousa 2018 for an introduction of the relevant distinctions and further overview of the relevant literature). Thus, on a common view in philosophy of emotions, an emotion E is fitting if and only if 'its intentional object o [e.g. the thing *about* which the emotion is, as for example the dog in the case of fearing a

dog] exemplifies the formal object that E (re)presents *o* as having' (cf. Echeverri 2019: 543). Crucially for us and without going into too much detail, whether the emotion's intentional object *o* indeed exemplifies the formal object that the emotion is supposed to present *o* as having depends on a number of factors. Among these will figure background conditions, cultural influences, personal history, and so on. Take the case of the dog and fear. It will be only in virtue of some specific features of the dog, say, that the dog has big teeth and it moves erratically, that the dog exemplifies the formal object of fear – that is, the dangerous/the dangerousness (cf. Deonna and Teroni 2012: 1–12; Echeverri 2019). To put it in Maguire's terminology then, it is in virtue of some aspects that the fact about there being a dog makes the fear of the dog fitting (is a fit-making fact). However, that these features – for example, the teeth and erratic movements – exemplify the danger is not a 'strict' fact. It depends on the context, on background aspects and assumptions. For example, it may do so only in virtue of us having the background knowledge of there being some sort of correlation between dogs with big teeth moving erratically and one getting bitten by a dog. But consider, for instance, a community where no erratic dogs with big teeth have ever been observed biting people, and no reports about such incidents have ever been heard about from other places. Arguably, in such a community, that the dog in front of one has big teeth and is moving erratically *doesn't exemplify* the formal object (i.e. the dangerous/the dangerousness) that the emotion of fear represents its intentional object as having. Thus, the fact that the dog with big teeth in front of one moves erratically exemplifies the formal object of fear – the dangerousness that my fear represents the dog as having, holds *only* provided that I/we have the background knowledge of the correlation of dogs with big teeth moving erratically and one getting bitten by a dog. But this crucial aspect does sound very much like a 'non-strict' fact in Maguire's sense. That the dog has big teeth and moves erratically doesn't 'justify or require anything on their own' (cf. Maguire 2018: 784). We need the additional knowledge that in this community it is well known that dogs with big teeth who move erratically tend to bite. Only given this background assumption is the dangerousness present.

Now, with respect to the contributoriness, the abovementioned case of running can be taken as putting pressure on Maguire's idea that fit-making considerations are not contributory. And with respect to gradability, as contemporary discussions on gradable adjectives show, it is far from established that adjectives exhibiting a sort of crispness and absoluteness are not gradable (see for linguistics Kennedy 1999; Kennedy and McNally

2005; Kennedy 2007; Lassiter 2017; Logins 2020a, 2020b contains discussion of these points with respect to *correct/appropriate, confident, and supported*).²² Thus, overall we might well have sufficient grounds for doubting premise (3) as well.

Given the aforementioned discussion, I think we have good reasons for doubting the soundness of Maguire's argument that we schematised earlier – that is, (1)–(4). The claim that there are no genuine 'right kind' of reasons for affective attitudes doesn't withstand scrutiny. Given that the 'right kind' of reasons for attitudes are genuine normative reasons, Maguire's version of axiological explanationism is back at square one. That there are genuine 'right kind' of normative reasons for affective attitudes speaks against this version of axiological explanationism. And its proponents still don't have a satisfactory answer to this problem.

Two quick remarks before concluding are in order. First, it may be useful to note that contrary to what axiologists like Finlay and Maguire seem to think, their views do have problems not only with accounting for the 'right kind' of reasons for attitudes but *also* with accounting for the 'wrong kind' of reasons for attitudes. More precisely, their view seems to rule out the possibility of at least some sorts of normative reasons for attitudes that are not directly related to fittingness of attitudes ('fit-making') *and* are not directly related to promoting valuable states of affairs. The key case here is *epistemic reasons* for emotions – that is, reasons that contribute to making an affective attitude/emotion reasonable from an epistemic point of view. It is a commonly held view within philosophy of emotions – indeed, I would say it's more or less orthodoxy today – that emotions can be assessed from a number of perspectives. And in particular, it is a very popular view that, besides emotions, being fitting or not (i.e. meeting their formal objects), emotions can also be evaluated from a purely epistemic point of view (see Gordon 1987; Greenspan 1988; Mulligan 1998; Goldie 2002, 2004; Deonna and Teroni 2012; Pelser 2014; Epley 2018; Meylan 2018; Scarantino and de Sousa 2018; Drucker 2019; Echeverri 2019; Na'aman 2021; for a non-exhaustive list

²² Faraci also puts forward similar considerations against Maguire's claim that fit-making considerations lack non-strictness, contributoriness, and gradability. See in particular his suggestion that fit-making considerations for emotions might well combine: 'And it seems perfectly possible that while neither the sharp claws nor the murderous look alone is sufficient to outweigh the dragon's claim to be only interested in friendship, in combination they are; you ought to run. Most importantly, this is no less plausible if we take these to be reasons for and against fearing her' (Faraci 2020: 231). And with respect to non-strictness: 'arguably it is always fitting to feel disappointed when someone gets a promotion you deserve, whereas it is not always fitting to fear something with sharp claws' (Faraci 2020: 231, fn 11).

of philosophers who appear to endorse this idea).²³ Robert Gordon puts the idea about the epistemic status of emotions being distinct from their fittingness states explicitly in terms of different sorts of reasons one can have for an emotion. He writes: '[i] Tom is worried that his wife was on the two o'clock flight, because that's the one that was hijacked (attitudinal). [ii] Tom is worried that his wife was on the two o'clock flight, because she said she'd be arriving early in the evening (epistemic)' (Gordon 1987: 35). The example in (i) is a case of an attitudinal reason for an emotion (e.g. connected to the formal object of emotions, fit-making), whereas example (ii) illustrates an epistemic reason for an emotion (e.g. connected to the purely epistemic status of an emotion). In a sense, then, epistemic reasons for emotions are of the 'wrong kind', understood in the specific sense that only fit-making considerations (e.g. related to an attitude's formal object) can be the 'right kind' of reasons. But, given that epistemic reasons are connected to attitudes' epistemic status, and an attitude's epistemic status is not dependent on promoting some valuable state of affairs, the value-based accounts have trouble explaining why these 'wrong kind' of reasons – that is, epistemic reasons for emotions – are genuine normative reasons. The value-based accounts have trouble explaining not only the 'right kind' of reasons for beliefs and emotions but also in accounting for some (i.e. epistemic) 'wrong kind' of reasons for emotions.

Second, and in fact connectedly to the first point, this further weakness of the account seems to elicit a more fundamental problem with axiological explanationist accounts and, arguably, explanationist accounts *tout court*. While they do capture one important aspect of our common understanding of reasons – that is, that there is, in a sense, a connection between reasons and explanation of normative facts – they have hard time explaining another equally fundamental aspect of our common understanding of reasons, namely that there has to be some connection between reasons and reasoning broadly understood. In a sense, reasons are understood to be right foundations on which our attitudes can be based; that is, reasons are often perceived as considerations that can lead us to hold fitting responses, e.g. attitudes in a given situation. Somehow, views that limit their accounts of normative reasons only to their explanatory functions fail to appreciate this important aspect of our common view.

²³ If a quotation is needed, here is one paradigmatic statement of this idea: 'standards of correctness so conceived should be distinguished from epistemological standards by which we assess the justification of emotions' (Deonna and Teroni 2012: 6).

3.4 Concluding Remarks

This chapter has focused on views that analyse/define reasons in terms of their explanatory role for some normative facts. The two major approaches within this broad explanationist camp comprise those who appeal to the role of reasons in explaining deontic facts (e.g. ought facts), and those who appeal to the role of reasons in explaining axiological facts (e.g. value facts – either F-ing-based or object-based goodness). We looked at the details of what appear to be the most promising versions of both general approaches. We considered the positive arguments in their favour. A major advantage of these approaches is that they do indeed seem to capture an important insight, namely reasons have to be connected in a sense to explanation of normative (deontic or axiological) facts. However, we also saw that none of the more specific versions of the general explanationist approach withstands scrutiny. Both deontic and axiological versions of explanationism run into fundamental difficulties. For one thing it is not clear how exactly outweighed reasons can function in explanation of ought facts (deontic explanationism); for another, it is not clear how the ‘right kind’ of reasons for attitudes are even possible if reasons are supposed to explain value promotion or the value of an end. More specific worries were elaborated along the way. Thus, the conclusion we reached is that while explanationist views do seem to point towards an important insight in understanding reasons, they also seem to leave substantial aspects of reasons unexplained. Reasons matter for us not only in virtue of their explanatory roles but, sometimes, also for their role in reasoning/argument towards a fitting response.