

There are No Reasons for Emotions or Desires

Barry Maguire
The University of North Carolina at Chapel Hill

Abstract

A dogma of contemporary ethical theory maintains that the nature of normative support for emotions and desires is the very same as the nature of normative support for actions. The prevailing view is that normative reasons provide the support across the board. I argue that the nature of normative support for emotions and desires is importantly different from the nature of normative support for actions. Actions are indeed supported by reasons. Reasons are gradable, contributory, and they participate in quantitative explanations of all things considered normative facts. The support relations for emotions and desires have none of these properties. So-called right kinds of reasons for emotions or desires are instead facts that make those very attitudes fitting. Unlike reasons facts, fit-making facts for emotions or desires do not conflict with each other or contribute together in the explanation of further normative facts. More fit-making facts just make a more complex set of reactions fitting. I end with a speculative explanation for this difference. It is plausible that actions are a fundamental locus of normative assessment. This is also plausible for enduring states, such as caring about one's friends. But this is less plausible for emotions and desires. Rather, in many cases, our emotions and desires are made fitting by such enduring states.

There are no Reasons for Emotions or Desires

Section One: Introduction

There are no reasons for emotions or desires. There are no reasons for admiration, delight, fear, envy, gratitude, disgruntlement, contempt, Schadenfreude, relief, pity, shame, or amusement. There is no reason to want to listen to the Köln concert, or to want to avoid a grisly death. No reason to hope for a brighter day. There is no reason for your incredulous reaction to this announcement, but it might be *fitting* all the same. I do not deny that there are normative standards concerning these attitudes. Rather I claim that these two relations of normative support – being favoured by a reason or various reasons, and being made fitting – are importantly different in various ways. Contrary to a common assumption, token emotional responses are not supported by reasons but by facts that make them fitting, and likewise for desires. That is the central thesis of this paper.

I start out by clarifying the target notion of a normative reason, and distinguishing two classes of would-be reasons for emotions: the right kind and the wrong kind. I restrict my argument to the so-called *right kinds* of reasons for emotions or desires.¹ The heart of my argument is this. Putative right kinds of reasons for attitudes have a different nature from reasons. Reasons for action are *gradable* – they have weights – and they are *contributory* or *pro tanto* – they are incomplete parts of a specific kind of explanation of all things considered deontic facts, such as facts about what you ought to do. But putative right kinds of reasons for emotions and desires are neither gradable nor contributory. They do not have weights and they do not explain ought facts. They are themselves themselves, if you like, a kind of ought fact. Putative right kinds of reasons are facts that directly make different attitudes fitting.

Here's a case from my own life that got me thinking about these issues. My grandmother passed away recently. She was an amazing lady, the only woman driving a Jaguar in Glasgow in the 1950s. She had been suffering terribly with Alzheimer's for the last several years of her life, with no hope of improvement. She believed she would rejoin her beloved first husband in the afterlife. But she was our family matriarch, and she had great chat right up until the end. She used her last words to make fun of my dad's bald spot.

What is the fitting response to a death such as this? Here's one thought: the facts about her disease and her religious beliefs weigh against the facts about her being a matriarch and a good laugh. They cancel each other out. One should feel utterly unmoved. Clearly this is not right. Rather, in response to my grandmother's passing, both sadness and a certain kind of relief were fitting. The facts that made the sadness fitting didn't make the relief

¹ I am officially agnostic about whether there are wrong kinds of reasons for attitudes.

unfitting, nor did the facts that made the relief fitting make the sadness unfitting. These facts don't participate in any kind of weighing explanation. Each of them directly makes a specific response fitting.

So, at least, I'll argue. By way of conclusion, I'll offer an explanation for this divergence in kinds of normative support. It is plausible that *action* is a fundamental locus of normative support, and hence that there are normative reasons for and against alternative actions in a situation. So much, at least, is assumed by going theories of reasons and dominant ethical theories. But it is not similarly plausible that token emotional responses are – in general – a fundamental locus of normative support. It is more plausible in many cases that the fundamental locus in the case of emotions is an *enduring state* such as caring, and not a fitting manifestation of that state.

I hope that this discussion will be of as much interest to the theory of reasons as to the theory of the normativity of attitudes. The terminological question of what we should *call* a reason and what we shouldn't doesn't matter much, except that calling two different things by the same name invites confusion when the difference matters. The terminology I offer seems best, but it's the distinct natures of the relations of normative support that is important.

1.2 What is at stake?

There are several important theoretical issues at stake. Hopefully, like me, you'll find the business of trying to get clear about the nature of different normative properties and relations interesting for its own sake. In our ordinary normative practice, we use the term 'reason' profligately, and it seems reasonable to expect that this usage will occasionally conceal interesting differences. Indeed, it seems clear that even just in the case of reasons for action, the uniformity in our reasons talk conceals differences between, e.g. motivating reasons and normative reasons, and, more subtly, between considerations that metaphysically explain ought facts, and considerations that one ought to reason from (Wedgwood 2015). We should be open to the possibility that our talk about 'this or that reason for this or that attitude', while perfectly felicitous, might be best understood as picking out some different support relation from our talk about reasons for actions.

It is common nowadays for philosophers to unquestioningly assume various strong 'unity of reasons' theses: not just that moral and prudential reasons are reasons in the very same sense, but that there are reasons for emotions and reasons for beliefs in this very same sense. But this is a working hypothesis, not a datum. It has not been subjected to much scrutiny within the 'reasons' literature. To assess this latter hypothesis on its merits we need to attend carefully to putative *disunities*, attempting to give as clear an account of them as possible.

A number of ethical theories or arguments either appeal directly to reasons for attitudes or simply assume that the nature of normative support for actions and attitudes will be the same. The following theories presuppose

what I'm arguing against here, that there are reasons for emotions and desires (in whatever sense that there are reasons for actions). If what I will argue here is correct, some of the central arguments and theses need to be reconsidered.

- Some views explain the *fittingness* of an attitude in terms of the *balance* of the right kind of reasons for and against the attitude (see Schroeder, 2010:29). This view is a central target of the arguments to come.
- Other views explain *value* in terms of reasons for desires or for some other privileged class of attitudes. This is the so-called 'buck-passing' theory of value. Relatedly, some views define the aesthetic, or aesthetic value, in terms of the kinds of things we have reason to appreciate. Such views will need to be reconsidered. The prospects for a fittingness-first analysis of value is left on the table here.
- Other views explain *reasons for action* or what one *ought* to do in terms of reasons for attitudes. Douglas Portmore (2011:82) formulates consequentialism, to simplify, as the thesis that you ought to do what you have most reason to desire to do. He maintains: "Everyone seems to admit that in at least some instances an agent's reasons for performing a given action derive from her reasons for desiring its outcome. For instance, we think that, other things being equal, you have more reason to choose the more pleasure-inducing meal option to the less pleasure-inducing meal option if and only if and because you have more reason to desire the outcome in which you experience more pleasure than to desire the outcome in which you experience less pleasure."³ If there are no reasons for desires, this attempt to avoid the partiality challenge to traditional consequentialism will have to be reconsidered.
- The search for an account of the '*right/wrong reasons*' distinction assumes not just that right and wrong reasons are right and wrong *reasons*, but also that right and wrong reasons are right and wrong instances of the same category of thing. I'll suggest this is mistaken. So-called wrong kinds of reasons for attitudes are indeed *bona fide* reasons, either for the attitude itself, or perhaps for bringing it about that one has some attitude. (I'm officially agnostic about this.) But so-called right kinds of reasons are not reasons, they are fit-making facts.
- Some theorists have objected to the value-based theory of reasons and desire-based theory of reasons on the basis of a premise about the 'unity of reasons.' The challenge is based on the claim that these theories don't yield plausible accounts of reasons for emotions (Joseph Raz 2011, Jonathan Way 2013). But if there are no reasons for these attitudes, this important objection fails.
- There are also substantive implications. Once we distinguish these two distinct relations both commonly referred to as '*reasons*,' we see

³ I suspect Portmore's premise benefits illicitly from the plausibility of the thesis that you have a weightier reason to *eat* a meal that you have a stronger desire to eat. It doesn't follow that you have a weightier (or weighty) reason to *desire* to eat the meal.

that we cannot infer from the ‘normativity’ of reasons (see Raz 1999: 67) to the ‘normativity’ of fittingness in the same sense. It often does not go against *reason* – to borrow Raz’s phrase – to fail to have the attitude that is fitting. Envy, Schadenfreude, despair, and believing you are going to die miserable and alone, all have fittingness conditions. It is unlikely that you ought ever to have these attitudes. This raises an important issue. It is sometimes objectionable not to have a fitting attitude, for instance to be utterly unmoved by some tragedy that befalls a close friend. Once we distinguish fittingness and reasons, how exactly are they related? Is fittingness *genuinely* normative? I’ll hint at an answer towards the end.

1.3 Terminology

A few terminological stipulations will be helpful as we proceed. We have two relations: the *being a reason for* relation, and the *being a fitting response to* relation.⁴

I’ll call facts of the form ‘F is a reason to X’ *reasons facts*. We distinguish the reasons facts from the facts that are the reasons, *the reasons themselves*, namely the F. Then we have the following implicit definition: *Reasons facts* are facts about reasons *being* reasons to take options. (For simplicity I’ll assume that options are actions, not, e.g., activities or plans.)

I’ll call facts of the form ‘A is a fitting response to S’ *fittingness facts*. It will be convenient to have a name for the facts on the other end of the fittingness relation, the considerations that *fitting attitudes* are fitting responses to. I’ll call such facts *fit-making facts*.⁵ Then we have the following implicit definition: *Fittingness facts* are facts about certain *attitudes* (or other things) being a fitting response to *fit-making facts*. I will be arguing that putative right kinds of reasons for emotions or desires are, in most cases, fit-making facts.

I will use the word “supporting” when I want to remain non-committal about the kind of normative or even non-normative support, whether favoured by a reason or made fitting or partially grounded or something else.

A final word on the class of attitudes directly under consideration. I restrict the discussion to putative reasons for emotions and for desires. I’ll talk mostly about emotions. I address the nature of support for beliefs elsewhere, and ignore this question completely in this essay.⁶ I will be focusing on occurrent mental states rather than dispositions or traits of character (with

⁴ I assume that the *being a reason for* relation and the *being a reason against* relation are distinguished just by their valence.

⁵ Really, I should talk about propositions believed rather than facts, since justified false beliefs about outcomes can make responses fitting. But issues about evidence-sensitivity are largely orthogonal to the central points here (apart from briefly in subsection 4.1).

⁶ ‘Explaining Epistemic Normativity,’ co-authored with Jack Woods.

episodes of anger rather than with people being ornery) and rather than enduring states (such as loving Charlie or caring about poverty). Also I am focusing on object-directed attitudes rather than moods or feelings (with people getting angry about their low wages, not with their 'having gotten out of bed on the wrong side.')

The goal for today is to draw the advertised contrast as clearly as possible. I leave to another occasion the task of characterizing more precisely which things stand on the ends of these two relations.⁷

Section Two: Characterising Reasons

Let's start by characterising reasons, by which I will always mean, *pro tanto normative* reasons. I will work with a precise and restrictive characterisation. This suits my dialectical goal of making the contrast between the reasons relation and the fittingness relation as clear as possible. I will generalize the account to an alternative class of theories of reasons (including those of W.D. Ross and John F. Horty) later on.

Reasons theory is unpopular in some circles. But there is a simple argument for the importance of theorising about reasons, that is too infrequently made, and that helps to motivate the specific phenomenon that I have in mind.⁸

There are few, if any, informative principles relating finely-individuated act-types with the all things considered ought. Take any finely-grained action you like, no matter how gruesome, and a sufficiently imaginative ethicist will describe you a situation in which, alas, that is the thing you ought to do. But it doesn't follow that there is nothing systematic to be said about the explanation of ought facts. There are often various different conflicting considerations bearing on what you ought to do in any situation. This conflict is resolved by appealing to facts about the relative weights of these competing considerations. Anyone who might once have looked for exceptionless principles relating finely-grained action-types with oughts should look instead for an account of the nature and content of considerations that count in favour of actions to varying degrees. That is the simple argument for the importance of reasons.

Take Peter Singer's famous example:

'If you are walking past a shallow pond and see a child drowning in it, you ought to wade in and pull the child out. This will mean getting your clothes muddy, but this is insignificant, while the death of the child would presumably be a very bad thing.' (Singer 1971)

⁷ I don't necessarily deny that some actions are fitting (e.g. expressive actions in symbolic contexts) or that there are reasons for some attitudes (e.g. of the wrong kind, or for enduring states). Importantly, intentions may well fall on the 'action' side of the distinction. For the purposes of my argument I'll focus on paradigmatic cases.

⁸ This argument equally supports my 'weight-first' approach and such 'weightier-than first' views as Ross's and Horty's that I'll discuss later.

The fact that by wading into the pond you can save the life of a child is a reason to do so. The fact that by doing so you'll get your clothes muddy is a reason not to wade into the pond. However, compared to the importance of the saved life, getting your clothes muddy pales in significance.

On their own, reasons are normatively impotent. The mere fact that some fact is a reason to do something, even a weighty reason, doesn't yet make it the case that you ought to do that thing, or that you would be criticisable for not doing it. It is a mistake – albeit a common one – to say that *individual* reasons ever justify or require anything on their own. The fact that *this* is a reason for *that*, even a weighty reason, is simply silent about what other reasons there might be. Even a weighty reason, if significantly outweighed, would fail to justify the action it favours.⁹ It is implicit in Singer's example that you don't have any reasons that are weightier than your reason to save the child. However, in combination, facts about reasons explain what you ought to do. There are two properties here: being normatively impotent in isolation, and being such as to play a specific role explaining distinct normative facts. Let us call these the properties of being *individually impotent* and *contributory*. These properties can come apart: e.g. the fact that you are obligated to ϕ is not normatively impotent in isolation, but may play a similar explanatory role. I will be arguing that putative right kinds of reasons for emotions and desires are neither individually impotent nor contributory.

We can also say something rather specific about the nature of the contribution that reasons make to facts about what you ought to do, or to say the same in other words, about the kinds of explanations that reasons participate in.

Each reason has some *gradable* property that we can call its 'weight.' These weights are commensurable, by and large.¹⁰ The 'saving the child' reason has more weight than the 'getting your clothes muddy' reason. Moreover, these two properties work together. Reasons explain other normative facts in virtue of their weights.

Think about the last stages in the explanation of what you ought to do. The fact that you ought to x is immediately¹¹ explained by the fact that you have most reason to x . Or so it is plausible to suppose. This fact about most reason is in turn immediately explained by facts about the net weight of reason supporting each of the different options you have in the choice situation (together with some fact about these being all the options). These are purely

⁹ For discussion of the relations between such *pro tanto* reasons and 'sufficient reasons,' 'decisive reasons,' and 'conclusive reasons,' see Lord & Maguire 2016. I do not take a stand on whether reasons have a univocal normative force, or whether there are 'justifying' and 'requiring' reasons (on which see Joshua Gert 2003). Both types of Gert's reasons are gradable and contributory.

¹⁰ I don't need to get into questions about whether and how moral and prudential reasons are compared. For recent discussions see Harman 2016 and Schafer 2016.

¹¹ See Kit Fine 2012 for the immediate / mediate explanation distinction.

quantitative comparisons. At these stages of the explanation of what you ought to do, it does not matter at all *what* the reasons are.¹² All that matters to this explanation are facts about comparative weights. As we say, what matters is how *much reason* there is do one thing or another. Similarly, all that is relevant to determining the net weight of *reason* for any given option is the weight of the reasons for the option and the weight of the reasons against the option (together with the fact, for each option, that these are all the reasons for and against that option). More generally, I shall say that reasons explain all things considered deontic facts *quantitatively*.

An explanation as to stew used the largest number of types of ingredients would be quantitative in the same way. The fact that one stew had carrots rather than turnips, and the other rabbit rather than goat, is not relevant. All that matters is how many different types of ingredients there were in each. Compare an explanation of the taste of either stew in terms of the nature of its ingredients and the cooking method.¹³ Neither are reasons explanations the only explanations with this particular quantitative structure. An explanation of why this student won the prize, in terms of having the top marks, which in turn is explained in terms of having more questions right and fewer wrong than anyone else, will have the same structure. Explanations of the 'overall value' of a complex state of affairs will often have this structure as well (on which more later). The point is not that *only* reasons explanations have this structure, it is that *all* reasons explanations have this structure.¹⁴ Neither is the point that quantitative facts about reasons *exhaust* the grounds of facts about ought, since 'normative totality' facts are part of the explanation, and maybe other facts too. The point is that quantitative facts about reasons – facts about their weights – are *necessary* parts of this explanation. It is in the nature of reasons to play this explanatory role.¹⁵

To summarise, reasons participate in a two-stage normative explanation:

At the LOCAL LEVEL, the facts about individual reasons and their weights are explained. To give two prominent examples of how such an explanation might go, for a fact to be a reason is for it to be or explain a fact of the form [ϕ

¹² Of course, the facts about what the reasons are will matter for other things, e.g. for a theory of moral worth (Markovits 2010). The point is that in this part of the explanation of what you ought to do there may still be no mention of *what* the reasons are, what *kinds* of reasons there are, whether you *want* to take any of these options, or whether you are *obligated* to take any of these options, or what kinds of *values* are at stake. These facts may be part of other normative explanations, of course – for instance if you violate an obligation you will owe an apology. And they may be part of the explanation of the weights of reasons themselves – the LOCAL LEVEL of explanation. But they are not in themselves necessary parts of the ALL THINGS CONSIDERED level of explanation of ought facts.

¹³ To paraphrase Honderich (1982: 61), neither the orangeness nor the Irishness of the carrots is relevant to the number of ingredients in the stew.

¹⁴ At least, so I am assuming for the sake of argument. As previously advertised, I will generalize the argument to other ways of thinking about reasons in Section Six.

¹⁵ This thesis is stated too strongly to be accepted by a Scanlon-style Reasons Fundamentalist who maintains that nothing informative whatsoever can be said about the nature of reasons. Such a theorist would still presumably accept the weaker thesis that reasons *do* play this grounding role, even if this latter isn't a definitive feature of reasons.

promotes S], where S is valuable or where S increases the probability that some desire of yours will be satisfied.¹⁶ These facts about reasons may be sensitive to both conditions and modifiers. The fact that you promised to ϕ may be a reason to ϕ only on the condition that ϕ -ing is not evil. The weight of your reason to pay for a future benefit may be attenuated by its distance in the future.

At the ALL THINGS CONSIDERED LEVEL (for short, the ATC LEVEL), the fact that you ought to x is explained by the facts about the net weight of reason to take each of your available options. The fact about the net weight of reason (mass noun) for any given option is explained by some function from the facts about the weights of each reason for and against that option. This explanation includes a *normative totality fact*: some fact to the effect that these are all the reasons bearing on all the available options.

There are various similarities between reasons facts and fittingness facts at the LOCAL LEVEL. For instance, fittingness facts often admit of conditions and modifiers. There are also standards or conditions of fittingness, which themselves admit of explanations.¹⁷ A central contrast between reasons facts and fittingness facts consists in the fact that reasons facts participate in ATC LEVEL explanations, whereas fittingness facts do not. The normative explanations of attitudes are more direct, in this sense. Think again about the fitting reaction to the death of a beloved matriarch suffering from a terminal illness.

It is important, finally, to distinguish reasons in this restrictive metaphysical sense from a broader functional role that considerations can play as *representatives* in reasons-giving contexts.¹⁸ It is well-known in the theory of explanation that the provision of salient considerations can play different pragmatic roles in virtue of variations in our shared background knowledge and specific interests in reasons-giving contexts. Quite 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.¹⁹ “Make your conversational contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged” (Grice 1975).

In particular, it is permissible to *say* of considerations that they are reasons even if they are not reasons in my sense. For instance, it is permissible to say

¹⁶ For the desire-based theory of reasons, see Schroeder 2007. For the value-based theory of reasons, see Maguire (2016).

¹⁷ Of course, the explanation of the fittingness of an attitude will still have structure, for instance as I’ll suggest later the explanation may advert, inter alia, to facts about what the agent cares about; more familiarly, at least some fittingness conditions are explained by facts about the nature of the relevant attitude.

¹⁸ On this distinction see Raz (1975:22), Broome (2013:49), Wedgwood (2015), Maguire (2016), and especially Daniel Fogal (2016:86), from whom I borrow the term ‘representatives’. This is an instance of a more general distinction between metaphysical or causal explanations and considerations that one can give as an answer to a ‘why’ question in a specific context (see David Lewis (1986:215)).

¹⁹ Raz 1975, Schroeder 2007.

that your reason for ϕ -ing was that 'there was no reason not to ϕ ' or that 'a reliable informant advised ϕ -ing' or that 'there was most reason to ϕ ' or that 'Mildred didn't mind'... But none of these are reasons in the sense of considerations that play the metaphysical role I have in mind – the one motivated by the simple argument for reasons theory. This is a crucial point, because it is often permissible to say that some consideration or other is your reason for some attitude. It doesn't follow that these considerations play the same metaphysical role as reasons as I understand them.

Section Three: Supporting Attitudes

Let's turn now to considerations supporting attitudes.

3.1 Distinguishing Right from Wrong

I need to start by distinguishing the right from the wrong kinds of supporting considerations in order to focus attention on the so-called 'right kinds of reasons' for attitudes. This distinction is fairly well-known and difficult to analyze, so I will introduce this distinction by examples.²⁰

First exclude *incentive* cases. The fact that the gunman will shoot unless you hope that the Socialist Candidate wins the election seems to be a reason to hope that she wins, but the wrong kind of reason. Intuitively the right kind of reason to hope that the Socialist Candidate wins might be that she favours universal health care, or free child care, or hefty anti-discrimination provisions. Intuitively the gunman's threat provides the wrong kind of reason to have this attitude in these contexts.²¹

There is another class of so-called reasons for attitudes that are not quite the 'right kind'. When Joseph Raz says "A murderous rage...appears never to be appropriate" (2011:49), this is true, with 'appropriate' understood broadly, but false, if understood to mean that there is no such thing as the (so-called) right kind of reason to fly into a murderous rage. Consider the circumstances of The Bride in Tarantino's *Kill Bill* after the massacre of her friends and unborn baby on her wedding day, for example. Sometimes a joke can be

²⁰ For more on this distinction, see Rabinowicz & Rønnow-Rasmussen (2004) and especially D'Arms and Jacobson 2000.

²¹ Some incentive cases are subtler. The fact that some attitude will be pleasant/unpleasant is usually the wrong kind of reason to have/not have the attitude. Even the fact that manifesting an attitude in a situation will help you to develop the disposition to manifest that attitude in situations like this – 'fake it 'til you make it' reasons – are the wrong kinds of reasons. Only in unusual cases would this incentive affect whether some attitude is fitting. Suppose I tell you I will shoot you unless you are scared of me, or that I'll give you a million dollars if you are excited about an offer I make you. Are these the right kinds of reasons for the attitudes? I feel my intuitions pulled both ways here. I suspect we need to distinguish different objects: the fact that someone is threatening you with being shot, and the fact that you will be shot unless you are scared of the person making the threat. The first is definitely the right kind of reason to be scared (if there are any). I'm not sure about the second. What is important is that these are fringe cases. This shows we are cottoning onto the relevant distinction.

funny even though it wouldn't be appropriate to laugh – an abortion joke at a pro-life rally. Often emotions can be fitting but unpleasant, or unwise: too much anxiety about whether you will fall off the rickety bridge and you just might. With D'Arms and Jacobson (2000), let us classify instrumental, prudential, and moral considerations such as these as of the wrong kind.

3.2 Putative Right Kinds of Reasons Are Not Contributory

I will now argue that putative right kinds of reasons for emotions or desires are not reasons but fit-making facts. I will argue that fittingness facts are neither gradable nor contributory.²² In the rest of this section I address contributoriness. I will consider some of the ways in which reasons combine and conflict and contrast these with the way putative right kinds of reasons interact. The fact that some attitude is fitting is characteristically²³ not part of an explanation of any *further* normative fact.

Suppose we start with some very specific attitude, distinguished by kind, degree, and content: being very pleased at your friend's promotion. Then we try to find conflicting considerations bearing on the fittingness of this attitude. The simplest way for considerations to conflict is for one to make that attitude fitting and for another to make that very attitude unfitting.²⁴

Suppose you didn't get some promotion, but your good friend Angela did. You know that you would have deserved the promotion, but that it will be really good for your friend to get it. The fact that you didn't get the promotion makes it fitting to feel disappointed. The fact that your friend got the promotion makes it fitting to feel pleased. But the considerations supporting these attitudes do not conflict. The fact that you didn't get the promotion doesn't make it unfitting to feel pleased, and the fact that your friend got it doesn't make it unfitting to feel disappointed. It is fitting to feel pleased *and* disappointed in these different respects in this case.

Suppose that Angela didn't really deserve the promotion (there were some shenanigans) but that it will enable her to avoid getting kicked out of her apartment. The fact that your friend will be able to stay in her apartment makes it fitting for you to be pleased about the promotion, but the fact that the promotion wasn't really deserved seems to make it unfitting for you to be pleased about the promotion *as such*. This is all plausible, but notice that we have again distinguished two attitudes both of which are fitting and that

²² I will allow myself to talk loosely about whether fittingness itself is gradable, contributory, etc., or whether the fittingness fact is gradable or contributory, or whether putative right kinds of reasons are gradable or contributory. Strictly speaking the question is whether the fit-making fact makes the relevant attitude fitting *to some degree*.

²³ This qualifier leaves room for the discussion of 'favouring overall' in section 5.3.

²⁴ There is a subtlety here. Distinguish the case in which F1 makes A1 fitting and F2 makes A1 unfitting from the case in which F1 makes A1 fitting and F2 makes A2 fitting (A2 might be some distinct option, or the negation of A1), where A1 and A2 are *incompatible* in some sense. There are various ways in which attitudes might be incompatible, and correspondingly different kinds of conflict that might be at stake.

are perfectly compatible with each other: being pleased that she can stay in her apartment and not being pleased as such about the promotion. Just as in the case of my grandmother's death, you feel a certain amount of ambivalence, which is best explained by giving this more detailed characterisation of your response. There is still no conflict here.

Could it be fitting to be pleased that your friend Angela can stay in her apartment, and also fitting to be displeased that Angela can stay in her apartment? Yes, though we would need to offer a plausible reinterpretation. Suppose that moving would be inconvenient for Angela but she has a chance of finding somewhere better. Another friend, Margaret, would move into the apartment; it would be a dream move for her. It is fitting to be sad for one friend's fortunes while pleased for another's. The considerations supporting these attitudes each, one by one, make a different attitude fitting. But to say that two opposing attitudes are both fitting is not to say that a single attitude is both supported by one 'right kind of reason' and opposed by another. More fitting attitudes just make a more complex set of reactions fitting. There is no conflict here, yet.

I have just discussed the simplest case, in which considerations seemingly conflict with each other in their support for or against some specific attitude. Let's consider some other prominent ways that reasons interact together, in order to construct the most cases in which it is most plausible that considerations supporting attitudes are interacting in these ways.

In addition to conflicting with each other, reasons can also add weight together to some option. Let's look for an analogue of this relationship in the case of support for attitudes. It is fitting to cry over the spilt milk. Here's why. It was the last drop of milk from your favourite goat before she was eaten by the bear. You were going to make cheese with it. People would have eaten the cheese and thought well of you. These all sound like reasons for you to be sad about the spilling of the milk. But as before, I suggest that these are actually fit-making facts for three different attitudes: sadness about the dead goat, nostalgia for the non-existent cheese, perhaps anxiety about your lack of esteem.

In other cases, seemingly distinct considerations supporting some specific attitude are intuitively part of the *same* source of support for that attitude. Here's an example (thanks to an anonymous referee).

I am a knight about to face the terrible dragon. Each of his two hands has 8 razor-sharp claws on it, each of his two feet has 12 even sharper claws, his tail has 9 vicious spikes, he has 112 jagged teeth, and he breathes fire. I should fear the dragon; fear is the fitting attitude to take toward him. The advocate of reasons for emotions can say that each of these features (each claw, spike, tooth, etc.) provides a reason to fear the dragon...

Is it the case that each claw, spike, tooth, etc. provides a putative right kind of reason to fear the dragon? If so, and if, as I am arguing, putative right

kinds of reasons each make distinct attitudes fitting, it would follow that each claw makes fitting a distinct attitude directed at it. But it would be problematic to say that each claw, spike, tooth makes fitting some *distinct* attitude towards the dragon: i.e. to fear him for the first claw, to fear him for the second claw, and so on.

The central question here concerns the individuation of normatively significant supporting considerations. Is it really true that these different considerations would each constitute distinct *reasons* to be fearful? Claws are not reasons; at best it would be the fact that the dragon has this claw, and the fact that the dragon has that claw, that putatively constitute distinct reasons to be fearful of the dragon. But it seems more plausible to say that these are parts of the same putative reason to be fearful, namely that the dragon's sharp claws would do tremendous damage if it were successful in swiping at you. Plausibly you have a different reason to be fearful about the fact that it breathes fire. In some sense these putative reasons overlap, of course, since you are fearful of being killed or maimed by the dragon, and these facts are parts of the explanation of how much power the dragon has to kill or maim you. Everything we would say about reasons, here, we can say about fittingness. Facts about individual claws do not make distinct attitudes fitting. Rather some fact about all the claws and the strength of the arm they are attached to make it fitting for you to be fearful about what might happen if they were swiped at you (even in your nice new armour).

Another challenge. Suppose you don't know how many claws the dragon has. Your scout is peering through a telescope, and reporting as follows: "Three! No, four!, No, Five! Oh God, Six!" and so on. It is natural to say that each pronouncement is a reason to be even *more* fearful of the dragon. Surely with each pronouncement it becomes fitting to be more fearful of the dragon than before. So isn't each pronouncement a reason that makes a different attitude (or the same attitude to a different degree) fitting?

There is an issue here about timing. The fact that your scout reports five sharp claws on the dragon at one time makes it fitting to be fearful to some degree at that time. The fact that your scout reports six sharp claws on the dragon at a later time makes it fitting at that later time to be fearful to some higher degree. These considerations are not parts of the same explanation, and *a fortiori* they cannot be conflicting parts of the same explanation, as reasons are.

There is also a more general issue here concerning what our theory should say about modifiers of fitting attitudes. Examples are easy to come by. Take some case in which some tragedy would have been greater, event more shameful, or victory more stunning, but for some consideration *p*. It would be natural in such a context to say that *p* is a reason to feel less pride, shame, or sadness than we would otherwise have expected. In some cases this wouldn't be right: this additional consideration doesn't attenuate the degree of an attitude that would be fitting; rather it makes a separate attitude fitting. Contrast the fact that someone ran a fractionally poorer time in the 100 metre sprint with the fact that she gave an arrogant acceptance speech. The speech

doesn't make the victory less impressive; it makes the victor less admirable. But this leaves our question on the table. What about the fact that she ran a slightly poorer time? Is that fact a reason to admire her race less? Well, less than what? For sure, it would have been fitting to admire a slightly faster time a little more. But she didn't get a faster time. It is fitting to admire this time this much. No conflict here, either.

We can apply this to conditionality as well. Take the case of Robert C. Byrd, a one-time member of the Ku Klux Klan. He famously filibustered the Civil Rights Act of 1964 for fourteen hours and thirteen minutes. Here we have someone doing something intensely difficult, for what they took to be an important cause, but which is not at all fitting to admire.²⁵ In fact, it is fitting for you or me to be repulsed by this performance. It is fitting to admire strenuous effort in the service of a worthy goal but not an unworthy goal. This kind of conditionality is familiar. It is fitting to be proud when your friend but not a stranger does something impressive, and not when she does something unimpressive. I noted above that conditionality pertains to the explanation of reasons, and values, and fittingness facts, all at the LOCAL LEVEL. We should not confuse the obtaining of some condition for the fittingness of an attitude with the obtaining of a *reason* to have that attitude. In fact, for our purposes, we don't even need to rely on any deep distinctions between reasons and conditions. The central point is that the condition is only *part* of what provides one distinctive source of support for the relevant attitude.

To summarise this discussion: I have been sorting cases into (1) those in which different considerations do not constitute *normatively distinct* sources of support for any attitude, and consequently aren't in a position to conflict with each other (e.g. the dragon's claws or Robert Byrd's effort and his goal), and (2) those in which considerations do constitute distinct sources of normative support, but for distinct attitudes (e.g. the athlete's time and the arrogant speech). I deny that there are distinct sources of normative support or normative opposition for any one attitude. By contrast, you might have lots of distinct complete reasons for and against some specific action, say, accepting the job in Malaysia.

Section Four: Gradability and Normative Explanations

4.1 Putative Right Kinds of Reasons Are Not Gradable

Normative reasons are gradable – they favour *to a degree*. Putative right kinds of reasons are not gradable. Here it will help to switch directly to talk about fittingness facts. Strictly speaking, it is a mistake to talk about some attitude's being more fitting or less fitting. This is unobvious, I suggest, partly because of the following important distinction: between a gradable

²⁵ I'm open to fittingness being relative to an individual's values rather than value itself. It might be fitting for a racist to admire Byrd's filibuster. I discuss this in the final section.

normative fact about an attitude and a normative fact about a gradable attitude.

Suppose you learn that two individuals have died. One is near the end of her life, her most important projects behind her. The other is a twenty-one-year-old, talented, educated, motivated, and popular. Generally speaking, it is worse for someone to die in their twenties than their eighties. So let's just assume that the death of the youngster is more tragic than the death of the peaceful octogenarian. Now consider the following two claims:

- (a) The fact that the death of the youngster is very tragic makes it *fitting* to feel *more sad* about her death.
- (b) The fact that the death of the youngster is very tragic makes it *more fitting* to feel *sad* about her death.

It is very plausible that the first claim, (a), is true. It is plausible that the fact that the death of the youngster is more tragic makes it fitting to feel more sad about her death.

The greater the tragedy, the greater the amount of sadness that is fitting. This seems to apply to many other attitudes that are fittingness-apt. The more shameful some event, the greater the shame that is fitting. The more stunning some victory, the greater the degree of pride that is fitting. The worse the thing you did, the more guilty you should feel – perhaps you should buy your friend dinner rather than merely a drink. The more claws the dragon has, the more afraid it is fitting to be.

What about (b)? Does the fact that the death of the youngster is more tragic make it *more fitting* to feel sad about her death, or more *unfitting* to feel unmoved by her death?

I don't think so. Consider an analogy. Suppose you asked two children for the product of 7 and 5. One says 42. The other says 33. Which answer is more correct? I think there is an error in the question. Any answer other than 35 is incorrect. The two children's answers are both incorrect. But there is more we can say. We can say that the first child probably gave the product of 7 and 6, so she was one multiple off, as it were. She probably used a better method. We can say the other was numerically closer to the correct answer. In different contexts we might implicitly invoke these alternative standards of assessment, to substantiate the claim that one is more or less incorrect. These are ways of generating some *alternative* standard by which to assess their degree of failure. This alternative standard is gradable; this is the property we are tracking with our 'how incorrect' talk.

So it is with fittingness. An attitude is either fitting or unfitting; there are no degrees. It can be perfectly intelligible to talk about degrees of fittingness, but when we are doing so we are subtly shifting to assessment by some

other pertinent standard, usually a gradable notion.²⁶ Consider an earlier example, and distinguish the claim that it is *more* fitting to be impressed by the slightly faster time, from the claim that it is fitting to be *more* impressed by the slightly faster time. I'm inclined to think that the former is a slightly sloppy way of saying the latter. But if one wanted to insist upon comparing the fittingness of two attitudes we could say: it is *more* fitting to be moderately impressed by this performance than to be mildly impressed. That claim may well be felicitous, even if it is not *literally* true. I maintain, however, that it would be true in virtue of the fact that it would be fitting to be impressed together with the fact that being moderately pleased is closer – according to the pertinent standard – to being impressed than being mildly impressed is.

Sometimes what is fitting is a range of attitudes or strengths of an attitude, e.g. the fact that rush-hour traffic in Los Angeles is light today makes it fitting to feel more than mildly pleased but less than ecstatic.²⁷ The boundaries will be vague, to be sure. That will yield another property that seems to come in degrees, namely *clearly falling within the fitting range of attitudes*. Then we can say that one attitude is more fitting if it more clearly falls within the fitting range. But it doesn't follow from the fact that a property is vague that it is gradable. It might be vague whether it is impermissible to use a particular curse word in your lecture. It doesn't follow that impermissibility itself comes in degrees.

Consider an objection (again, thanks to an anonymous referee):

'Although some -ible/-able adjectives pertain to alethic possibility (such as 'visible', which means 'able to be seen', or 'inimitable', which means 'not able to be imitated'), there is a large class of -ible/-able adjectives that pertain to fittingness (such as 'admirable', which means 'fit to be admired', or 'deplorable', which means 'fit to be deplored'). If it is a mistake to talk about being more fitting or less fitting, then it would seem to follow that these fittingness adjectives should not have comparative or superlative forms. However, they *do* have comparative and superlative forms; we can say that one person is less admirable than another (where this means that the first person is more fit to be admired than the second), we can say that one candidate is the most deplorable of all the Republican candidates (where this means that the candidate in question is more fit to be deplored than all the others), and so on. So...rejecting the gradability of fittingness is to make a bold linguistic hypothesis. Or so, at least, the objection goes.'

²⁶ Julia Staffel 2015 attempts to work out just such an alternative standard for Bayesians. She writes: "The standard framework of subjective Bayesianism only allows us to distinguish between two kinds of credence functions—coherent ones that obey the probability axioms perfectly, and incoherent ones that don't. An attractive response to this problem is to extend the framework of subjective Bayesianism in such a way that we can measure differences between incoherent credence functions. This lets us explain how the Bayesian ideals can be approximated by humans."

²⁷ For another nice example see the discussion of gratitude in Roberts (2004, p. 64).

My reply is in two parts. Firstly, 'more admirable' and 'more deplorable' mean 'more admiration is fitting' or 'it is fitting to deplore more vociferously,' just like 'more fearsome' and 'less beautiful' mean 'fitting to be feared more' and 'fitting to appreciate somewhat less.' Secondly, I accept that "fittingness" is a gradable adjective. It is perfectly felicitous to *talk* about this or that being more or less fitting, a better or worse fit, completely fitting, and even completely unfitting. However, I deny that the *property* of fittingness is itself gradable. In giving a semantics for fittingness language, we need both the property of fittingness and some other gradable property. This other gradable property sets the relevant scale. The semantics for these claims cannot be read directly from the claims themselves. This is a less bold linguistic hypothesis. It may help to think about the truth predicate here. We often indulge in talk about claims being more or less true, closer to the truth, completely true. But many theorists want to insist that the property of being true is not gradable all the same. We can do so without revisionism by giving a semantics for these comparative and superlative forms in terms of truth and some other gradable property, e.g. accuracy or truthlikeness.²⁸

Consider another objection.²⁹ Suppose that you receive some evidence that a wonderful, or tragic, event has occurred. Perhaps it is evidence that your daughter has been in an accident, or perhaps evidence that she has been accepted to her top choice of university. However, the evidence is weak. More evidence keeps coming in slowly, piece by piece. Each piece of evidence raises the probability that the terrible thing, or the wonderful thing, has happened. Here is the objection: doesn't this evidence increasingly make it *more* fitting to be delighted, or more fitting to be devastated?

I don't think so. At each moment, the available evidence (including evidence about the likelihood of further evidence) makes some specific response fitting. It is uncontroversial that different attitudes are fitting towards a tiny chance of something bad happening, and towards certainty that the bad thing happened. Hope that the bad thing didn't happen is fitting in the one case and not the other. Grief does not become *more* fitting as you get more evidence about a tragedy, though it becomes more understandable. Grief is always unfitting absent sufficient evidence that the bad thing happened. This is part of the tragedy of missing person cases. Quite generally, different attitudes are fitting towards different evidence about similar outcomes. A structurally similar point applies to our attitudes to nearby possible worlds. There are different fitting attitudes towards chopping your finger off and towards *nearly* chopping your finger off. But these attitudes are very different: in the one case intense anxiety, in the other relief. We care a great deal about whether some story is truth or fiction.³¹ Different attitudes are fitting for past events (disappointment, anger, envy) and future events (hope, excitement, anxiety), and for different degrees of likelihood (right now you know that it is possible that something will fall out of the sky and

²⁸ Compare Oddie 2014.

²⁹ For which, thanks to Daniel Star.

³¹ For a fascinating discussion of the importance of whether a story is truth or fiction, see 'The Lifespan of a Fact' by John D'Agata and his one-time fact-checker, Jim Fingal.

crush the building you are in, but unlikely). To summarise, fitting responses are textured: they are sensitive to the modal and evidential statuses of evaluatively significant states of affairs.

This is a debunking argument against a presumption that fittingness is gradable. I have offered an interpretation of fittingness that explains various different relations to gradable properties while insisting that, like correctness, fittingness is not gradable. If fittingness facts are not gradable then they are incapable of participating in quantitative explanations of other normative facts, in the manner of reasons, and so they are not reasons.

However, these remarks are consistent with fittingness being gradable all the same. Suppose, then, that in spite of the foregoing argument fittingness is gradable. It would, of course, not follow that fittingness facts are reasons facts. We still have a perfectly good alternative argument against that conclusion – in the previous argument that fittingness facts are not contributory. Facts that aren't contributory cannot be reasons, even if they are gradable. Plenty of gradable properties aren't reasons facts. Thus, we have two independent arguments for the conclusion that fittingness facts are not reasons: one based on the claim that they are not contributory, and another based on the claim that they are not gradable.

4.2 The Grounds of Fit-Making Facts Are Not Reasons Either

Perhaps you are thinking: sure, fittingness facts themselves may not be gradable, but putative right kinds of reasons for attitudes are not the fittingness facts themselves, but the facts that explain them. Even if some putative right kinds of reasons for attitudes are themselves fittingness facts, these facts themselves have grounds. Perhaps, even if the fit-making facts themselves are not gradable, their grounds are?

But it doesn't seem plausible that the grounds of fittingness facts are gradable, or that they characteristically participate in quantitative explanations, either. In characteristic cases, the explanation of fittingness facts is qualitative, not quantitative. Consider one last pair of cases. It is fitting for you to admire Nina's piano playing. She was playing Bach's fugues at extremely high speeds. She is only six years old. She is blind. The piano was on fire. These sound a lot like reasons for you to admire the performance, not different things you admire about it. But they do not separately contribute weight to some single overall response. Your admiration would be different for an equally impressive finger-less deaf nonagenarian violinist playing Tchaikovsky's violin concerto in D major. The nature of your admiration is different. In fact, you would be better described as having a variety of different attitudes in the two cases. You are in awe, full of respect and admiration for the will power, admiration for her talent, admiration for her sensitivity in the difficult passages, etc.

The important point is that neither fit-making facts nor their grounds are characteristically gradable, nor do either of them characteristically

participate in weighing explanations. The facts that make it fitting for you to feel sad or pleased are remarkably varied. These explanations might be occasionally quantitative in form, but it is not in their nature to be quantitative. The explanation of fit-making facts is more like the explanation of existence and weight of reasons, or the existence and amount of value, rather than the explanation of all things considered deontic facts. There is no structural analogue here of the ATC LEVEL in reasons explanations. There is just the LOCAL LEVEL.

Section Five: 'Overall' Attitudes?

5.1 The Overall Attitudes Challenge and Direct Normative Support

But what about 'overall' attitudes, that is, attitudes that constitute distinctive responses to evaluatively complex states of affairs? It is common to ask questions, of oneself or others, such as the following. What do you *really* want? Are you upset? Is he enviable, desirable, admirable? Here we seem to be asking for an account of an *overall* attitude towards something. We might even ask explicitly contrastive questions – whether you want S *or not*, whether you are excited about R *or not*, whether you fear T *or not* – and explicitly comparative questions – whether you want S or R (given that you can't have both), whether you are more excited about R or T, whether you are more afraid of S or T. In many such cases there appear to be considerations supporting both sides. Aren't these reasons for and against these attitudes?³² Aren't thoughts like these support for the thesis that there is an ATC LEVEL as well as a LOCAL LEVEL in the normativity of attitudes?

No. Suppose someone asks whether you *want* your grandmother to pass away peacefully in her sleep. Here we are deliberately asking for an overall verdict about your emotional state. To answer this question, it may seem as though your more specific attitudes have a contributory character: as though you weigh your desire for her to stop suffering against your desire not to lose her. But notice that if the conversation continued much longer, you would go into further detail. You would say that you really do not want her to pass away, but... This suggests that really your attitude (which we assume is fitting) is complex. The putative grounds of the putative fittingness facts for putative overall attitudes are really just fit-making facts for more specific attitudes.

Take another example. Someone asks you whether you like the summer in New York City. You proceed to explain that you dislike the heat and the humidity, but the wind blowing off the water helps, and it is easy to spend the heat of the day indoors. And the city comes alive in the heat: people are more boisterous and enthusiastic and the streets are full of bustling life. There are concerts in the park, although tickets are increasingly restricted to the wealthy... In short, it is complicated. What is the answer to the original question? Yes. You do like the summer. But you also like grapefruit juice and

³² Compare Schroeder (2010) and the brief reply in Way & McHugh (2015:35).

violent Korean films. The question is not probative. Insofar as your attitudes towards the summer in New York are fitting, you are sensitive to these many evaluatively significant features of the experience.³³

It is worth reemphasizing here that our ways of talking about our attitudes are often misleading. A stranger asks you how you feel about spending the summer in New York. You might not feel inclined to say more than “rather pleased,” even though your attitude is more complex. Often your interlocutor won’t have much patience for a longer answer than this anyway.³⁴

5.2 Scepticism and Disagreement About Overall Attitudes

Consider the sceptical hypothesis that there are *no* overall attitudes – that talk about overall attitudes is an abstraction or an oversimplification, just like contour lines on a map are an abstraction or an oversimplification of the slope on a hill. You do not have some overall attitude towards staying in New York for the summer, or towards whether your grandmother dies peacefully tomorrow, or towards the promotion of one friend, or the loss of the apartment. You just have the many more specific attitudes. If there are no overall attitudes, there are no standards pertaining to them. Hence there could be no contributory considerations bearing on any standards for overall attitudes.

The central objection to scepticism about overall attitudes is that it cannot make sense of *disagreement*.³⁵ But I think reflection on this objection actually supports my overall case here. Consider the following dialogue, upon seeing a photograph in a newspaper of a supporter of Donald Trump’s candidacy for president earnestly giving a Nazi salute. First person: “Ugh, that’s awful.” Second person: “No, this is great. The more photographs like this come out, the poorer Trump’s chances.” Importantly, they seem to be *disagreeing* about something – this is signified by the “No,...”. This disagreement seems to be about what constitutes a fitting response. Presumably the second person agrees with some kind of “Ugh,” i.e. that it is disturbing and disappointing to see an earnest Nazi salute at a presidential primary. But still, they seem to be disagreeing, and it is natural to assume

³³ How precisely? How finely grained? There is an interesting question here about whether ‘fittingness implies can.’ I omit discussion for reasons of space.

³⁴ On the other hand, it may be that your actual attitude is quite simple, but only because you failed to respond fittingly to the complexity of the situation. You may be, simply, delighted at the passing of the bill, even though it involved the rolling back of funding for important social services – because you never took the time to find that out or because you failed to update.

³⁵ Here is another objection: what about attitudes towards individuals rather than states of affairs? Aren’t these overall attitudes? It would be highly revisionary to deny that we can have attitudes towards individuals. But notice we are not talking about enduring states such as loving Mildred or despising Alfred. We are talking about token attitudes such as being proud of Mildred for her pirouette or hoping Alfred will miss the penalty. It is not implausible that token attitudes are always directed towards or based on states of affairs involving individuals.

they are disagreeing about an *overall* fact about the fitting response to this picture, and that this overall fact is explained by some weighing of the expected instrumental value that this picture will have in galvanizing support against Trump, against the intrinsic awfulness of it.

This is a nice challenge, but my interpretation is somewhat different. Let me start by saying that it is not clear there are any conflicting attitudes here. It is fitting to be disgusted at the salute while also being pleased at the reduction in Trump's chances at the presidency. (Really, even this latter claim is itself too coarse. What is fitting is a set of attitudes towards the probabilities of this photograph becoming well-known, towards Trump's losing because people rejected his fascistic tendencies, etc.) I don't think the two speakers would disagree, in the end, about the relative significance of the intrinsic horror and the instrumental value.

More likely this disagreement concerns the conjecture that *expressing* just the attitude of disgust is evidence that the first speaker may have overlooked the upside. The "No,..." indicates that the original "Ugh!" leaves something out of the full fitting response. Furthermore, the 'No...' might be intended to signify that *if one were to express just one response* to the photo, it might as well be rueful satisfaction as disgust. This further thought might indeed be based on some comparison of the values at stake. The thought might be this: "well, sure, this is disgusting, but the disvalue of the disgust is significantly less than the expected value of this sort of publicity." (Assume this is true.) There may well be a disagreement here, but this is not a disagreement about which responses to which states of affairs would be fitting. It is a disagreement about the values of the relevant states of affairs.

So I don't think the phenomenon of disagreement about the expression of attitudes provides support either for the existence of overall attitudes, or for the thesis that there are reasons for and against attitudes.

5.3 Overall Attitudes with Derivative Fittingness Conditions?

However, I do not want to rest much of my case on a speculative thesis in the philosophy of mind about the individuation of attitudes. So, for the sake of argument, let's *stipulatively define* an overall attitude with overall fittingness conditions. Let *favouring overall* be an attitude such that it is fitting to favour overall a state of affairs if and only if that state of affairs is valuable overall. Why not say that the facts about the particular values and disvalues instantiated in this complex state of affairs – the various respects in which the complex state of affairs is valuable or disvaluable – are reasons for and reasons against favouring the state of affairs overall?

Consider a specific case. Suppose that both welfare and equality of welfare are valuable for their own sakes. Compare two states of affairs: one in which everyone has 10 units of welfare apart from one person who has 20 and another distribution in which everyone including this person has 10. This second state of affairs is worse in one respect and better in another. It is

worse in terms of aggregate welfare but better in terms of equality. Our question is this: Why not say, in this case, that the fact that the state of affairs is worse in terms of aggregate welfare is a reason against favouring it overall, and the fact that the state of affairs is better in terms of equality is a reason to favour it overall, and that the weights of these reasons are proportional to the values of these states of affairs, and hence, insofar as the value of equality is absolutely greater than the disvalue of the loss of welfare, you have more reason to favour overall this state of affairs, and that since there are no other significantly weighty reasons, you have most reason to favour it and that therefore it is fitting to favour it? These facts meet the conditions: they are gradable, contributory, and they participate in a quantitative explanation of an all things considered deontic fact, in this case, whether it is fitting to favour overall this state of affairs. Aren't these reasons for and against favouring overall the state of affairs? We have analogues of the LOCAL LEVEL and ATC LEVEL of explanation here. We have facts about the values of this state of affairs in respect to its well-being and equality, corresponding to the LOCAL LEVEL of explanation. Then comparative facts about these values explain the fittingness of favouring overall. This would correspond to the OVERALL LEVEL of explanation.³⁷

Before proceeding, let me set aside an important class of values that feel a bit like 'overall values,' but aren't really. *Organic* values are values instantiated by some complex state of affairs as such, and not just a function of the values of the parts of the complex.³⁸ For instance, the state of affairs consisting in your playing the cello part for Pachelbel's Canon in D is pretty uninspiring on its own, but quite wonderful when your friend joins on the violin. This complex state of affairs might well directly make appreciation fitting. A distinctive value emerges at the level of this whole, and plausibly a distinctive kind of appreciation is fitting in response. To take another example, perhaps a particular kind of admiration is fitting when someone is both extraordinarily gifted and genuinely humble (perhaps surprise would also be fitting). The fitting attitude to this complex is, again, to be added to the fitting attitudes to their talent and humility. In such organic cases, we don't have conflict between considerations for fitting attitudes. There is no weighing explanation here and no tension with my thesis. On the contrary, we have even more complexity. The fitting attitude towards the organic value is *added alongside* the fitting attitudes towards the various parts.

Let's simply stipulate that there are no organic values in our comparison between the distributions of welfare equality. Consider the question whether the state of affairs involving less welfare and more inequality is better overall than the state of affairs involving more welfare and less inequality. Is one of these better *overall* than the other? Presumably there is some fact of the matter about this (assuming conventional views about the objectivity and commensurability of value). However, and this is the key point: this is not obviously an *evaluatively significant* fact. This fact does not describe any new

³⁷ Notice the difference between favouring overall the state of affairs, which is our stipulative attitude, and favouring the state of affairs overall. It is plausible that the latter picks out some *aggregative* fact about your more specific favouring attitudes.

³⁸ On which, see Moore 1903 and Hurka 1998.

evaluatively significant features of the situation. The fact about overall value is instead a quasi-mathematical fact about the aggregation of other facts about value.³⁹ There is a question about whether such a quasi-mathematical fact about other values, that one is already responding fittingly to, constitutes a normatively distinct source of support for any attitude at all. For instance, if fitting attitudes correspond to evaluatively significant states of affairs (as many buck-passers and value-firsters would allow), and overall values are not evaluatively significant, then there will not be any fitting attitudes corresponding to overall values, simply as such.

Bear in mind that all the more locally fitting attitudes remain fitting. Consequently, another option at this point is to introduce a distinction between derivative and non-derivative fittingness conditions. Some attitude *A* is *derivatively* fitting just in case the fact that *A* is fitting obtains in virtue of the facts that attitudes *B*, *C*, *D*, ... are fitting and that *A* stands in some relation (to be specified) to *B*, *C*, *D*,... There are two ways to develop this distinction, depending on whether the derivative fittingness of an attitude is a function of the *objects* of other fitting attitudes (e.g. their values), or a function of the more specific fitting attitudes themselves (e.g. their intensities).⁴⁰ This opens up the possibility that an overall attitude might be fitting in response to the overall value of some state of affairs, but only *derivatively* so. This fittingness fact obtains in virtue of various facts about the fittingness of responses to various distinctively evaluatively significant states of affairs, together with some fact about the relations between these attitudes or their objects. Insofar as there are overall attitudes, their fittingness conditions are derivative: they obtain in virtue of the fittingness conditions for the various specific attitudes and some non-normative relationship between the specific attitudes and the overall attitude.

So even if there are overall attitudes with fittingness conditions, this doesn't yet threaten my central thesis. Still, at this point it will help to distinguish strong and weak versions of the thesis I am arguing for here. The strong version maintains that fittingness facts never admit of anything like an ATC LEVEL explanation. This strong version of the thesis would be in tension with the existence of non-derivative fittingness conditions for the attitude of overall favouring. To rule out the existence of such fittingness conditions one would need to defend some principled restrictions on what sorts of non-derivative fittingness conditions there can be. Perhaps this would involve some thesis restricting non-derivative fittingness conditions to non-derivative values. I do not relish this task and will not undertake it here. A weaker version of my central thesis maintains that fittingness facts *characteristically* fail to admit of explanations analogous to the ATC LEVEL explanations that reasons for action participate in. This weaker thesis, it seems to me, is almost impossible to deny (consider again the fingerless violinist).

³⁹ Interestingly, if we accept the value-based theory of reasons, then we don't need to assume that *value* is additive. We never need to countenance 'overall value' at all. For if 'overall value' is only ever *practically* important, and values correlate with reasons, then we can do all the aggregating we need just with the reasons instead of the values.

⁴⁰ This strategy is briefly sketched in Way & McHugh 2015:35

So even if sometimes, in cases like favouring overall a positively valuable complex state of affairs, the fittingness facts appear to admit of a quantitative explanation structurally analogous to the explanation of facts about what action you ought to perform in a situation: even if that is so, I deny is that this is in any way characteristic, far less necessary. Unlike the explanations of facts about the all things considered ought, there are many different kinds of explanations of fitting attitudes. The explanation of fittingness facts is a LOCAL LEVEL explanation, not an ATC LEVEL explanation. This is true, even if sometimes LOCAL LEVEL explanations share structural affinities with ATC LEVEL explanations. In this way, as in so many others, the explanation of fittingness facts is analogous to the explanations of finally valuable states of affairs: compare the explanation of the disvalue of an inequalitarian distribution of opportunity for welfare (which will essentially involve quantitative explanans) with the explanation of the value of drinking negronis by the sea with your beloved on a sunny afternoon.

5.4 The Practical Importance of Overall Judgements

Consider a diagnosis for our inclination to assume not just that there are 'overall' attitudes but that they are centrally important to the normativity of attitudes. We often need something like overall attitudes when choosing between actions.

For sure, sometimes we *need* to make an all-or-nothing judgement about some coarse-grained question, such as whether a particular judge is admirable. Perhaps you are Chair of the Town Admiration Committee and the judge is one of the favourites for the Annual Prize. Then you do the best you can. But this is a question about reasons for action. The question is not whether it is fitting to admire the judge. The question is whether you ought to give her the prize. Lots of reasons will bear on this question: how people would react either way, whatever precedents there might be, what the Town Admiration Committee rule-book says.

It is the same with desires, for example. Suppose you want to spend some time in the mountains, and you want to spend some time on the beach, and you'd also like to spend some time on your book. What do you want overall? Well, all of those things. We can introduce a specific choice situation, and ask what you want to do most. For instance, we say that the mountains are awfully far away and you need a break from your book and what do you want to do this weekend *overall*? But this is clearly a practical matter. What we do here is precisely to move from considerations that are fitting, to reasons for and against different options.

The main point is that there is no real *need* for overall attitudes or derivative fittingness conditions. Perhaps we can persist in asking questions about which overall attitude is fitting, and even give felicitous answers. But it would be a mistake to infer that these facts play a central role in the normativity of attitudes, in the way that facts about which options are

supported by the overall weight of reason is central in the normativity of actions. Even if there were some fitting overall attitude in a situation, this would not *displace* the fitting local attitudes. It would still be fitting to be minutely sensitive to all the various different evaluatively significant features of the situation, and their various probabilities. Even if someone is decidedly not admirable, all things considered, that person may yet be admirably punctual.

Here is the punchline for this discussion of overall attitudes: When it comes to actions, the all things considered normative property is all important. It is what really matters. Indeed, this is one of the reasons why 'reasons theory' is so late on the scene: Consequentialists and Kantians have been working away with all things considered deontic properties this whole time. By contrast, in the case of attitudes, all things considered facts hardly matter at all. It is contentious whether there are overall attitudes at all, and if so what their fittingness conditions are, and whether they obtain in virtue of other facts about fittingness or directly in terms of overall values. But whatever you think about this, we have a clear contrast. With action, the all things considered normative property is all important. With attitudes, it hardly matters at all.

To summarise this discussion, there are five claims it will be helpful to repeat:

1. Fitting responses correspond to the specific modal and evidential status of distinct evaluatively significant states of affairs.
2. Even if we do have some 'overall' attitude, it is likely just an abstraction from a more sophisticated attitude towards *all* the various supporting considerations.
3. Even if there is a 'overall' attitude, and even if it has derivative or even non-derivative fittingness conditions, all the various non-overall attitudes *remain* fitting.
4. Overall attitudes are most often considered in response to a practical need. In such circumstances we are primarily concerned with reasons for action.
5. It is also worth reiterating that the way we *talk* about supporting considerations is often misleading. It is sometimes permissible to say that some fact is a reason for an attitude, even if this is strictly speaking false, even when we are talking about reasons for actions.

Section Six: What About Other Characterisations of Reasons?

If we do not place restrictions on what constitutes a reason for an action – perhaps just saying that a reason is any fact that plays a role in explaining what one ought to do – then it will be more difficult to establish the relevant contrast between reasons and fittingness. But such an unrestricted view of reasons is implausible. It would receive no support from the simple argument for reasons, since it fails to pick out distinct contributory considerations. The most prominent defender of the view that reasons can be understood in terms of ought and explanation – John Broome – spends a

significant amount of energy drawing further distinctions in precisely trying to characterize this contributory role. His attempt faces serious objections.⁴¹ Indeed, Broome might appeal to the structure of the theory of reasons developed here to defend himself from some of these objections.

Still, there are theories of reasons distinct from the characterisation we are working here, that have enough structure to receive support from the simple argument.

We can classify theories of reasons as 'weight-first' or 'weightier-than-first' views. The characterization I have been working with is 'weight-first': each individual reason has some gradable property distinguishable as its weight independently of facts about the weights of other reasons. But the contrast in normative support for actions and emotions or desires also applies to 'weightier-than-first' characterisations of reasons, according to which the weight of a reason is explained in terms of facts about what other reasons it would outweigh.

W.D. Ross is widely held to have introduced contributory notions to contemporary ethical theory. But on a plausible interpretation of Ross's account of *prima facie* or *conditional duties*, this is a paradigmatic 'weightier-than' first account.⁴² Here is an illustrative passage:

...normally promise-keeping, for example, should come before benevolence, but that when and only when the good to be produced by the benevolent act is very great and the promise comparatively trivial, the act of benevolence becomes our duty. (Ross 1930, p. 19)

Here is a simple gloss: There is some small number of act-types associated with duties. If there is at least one token of one of these act-types, and no tokens of any other, in a situation, then your *prima facie* duty to perform tokens of that act-type becomes your duty *sans phrase*. If tokens of more than one of these act-types are available in a situation, then which of these becomes your duty *sans phrase* is determined by some higher-order rule, or by 'moral perception.' The duty of non-maleficence is more important than the duty of beneficence (*op. cit.* 21). The duties of fidelity, reparation, and gratitude in general trump the duty of beneficence (e.g. *op. cit.* 19, 30, 41-42; also see fn. 1 on page 42).

Our question is whether the structure of the explanation of your duty *sans phrase* is more similar more to the explanation of fittingness facts than was the explanation of the all things considered ought in terms of reasons. This would restrict the scope of my conclusion by showing that the relevant contrast depends on the adoption of a member of a specific class of theories

⁴¹ Broome 2013; for objections, see Kearns & Star 2008 and Brunero 2013. Notice that this challenge lends additional support to my objection against the 'grounds of fit-making facts' hypothesis in subsection 4.2.

⁴² For more contemporary versions of 'weightier-than first' theories of reasons, see Horty 2012 and Scanlon 2014. For critical discussion, see Lord & Maguire (2016).

of reasons. Perhaps this would even support an argument against theories in this class. I will now argue that my conclusion applies to both classes of theories of reasons.

The case in which you can only fulfil one duty in the situation is the most similar to the explanation of fittingness facts, since in that case your *prima facie* duty is your duty *sans phrase*. But notice that even in that case, the fact that ϕ -ing is your *prima facie* duty does not fully ground the fact that ϕ -ing is your duty *sans phrase*. You also need some fact about whether you have any other *prima facie* duties in that situation. You need a normative totality fact: that there is no alternative to ϕ that would satisfy a *prima facie* duty. We have a LOCAL LEVEL of explanation corresponding to the fact that you have some particular *prima facie* duty to ϕ , and an ATC LEVEL of explanation corresponding to the fact that given the absence of any other *prima facie* duties in this situation this particular *prima facie* duty is your duty *sans phrase*. This explanation is not quantitative, but it is distinguishable by this totality fact about other *prima facie* duties – analogous to the fact about the net weight of reason favouring different options – and this is precisely where we locate our distinction between reasons and fittingness.

This difference in levels of explanation is clearer in the more common case in which you have a *prima facie* duty to ϕ and a *prima facie* duty to φ (both established at the LOCAL LEVEL) and then, at the ATC LEVEL, according to some higher-order rule⁴³ the duty to ϕ trumps the duty to φ and hence it is your duty *sans phrase* to ϕ . Here we clearly have a LOCAL LEVEL at which facts about *prima facie* duties to do various things are explained, and then an ATC LEVEL of explanation taking these facts as inputs and generating an all things considered normative fact. This is analogous to the explanation of what you ought to do in terms of facts about the net weight of reason favouring each alternative option. Both of these are distinguishable from the case of fitting attitudes, which are made fitting directly, one by one, by evaluatively significant states of affairs.

Section Seven: A Speculative Conclusion

I have been arguing that actions and attitudes characteristically enjoy different kinds of normative support. If I'm right, the obvious next question is, why? One natural answer is that reasons are tied to *choice* and/or *selection* and/or what is under one's *control*. This isn't too implausible, but there is lots of *parti pris* on this issue, and I don't have anything particular original to add to the dispute. Let me instead end with a conjecture. It is fairly plausible that actions are a *fundamental* locus of normative support. By this I mean that actions are assessed directly in terms of their own properties. Act-centeredness in ethics is a pretty familiar idea. All direct consequentialisms, and both Rossian and some broadly Kantian deontologies, are act-centered

⁴³ Even in the case of imperfect duties where Ross maintains that some act of moral perception is called for to decide the case, this still constitutes an ALL THINGS CONSIDERED level of explanation. The virtuous judgment of the agent takes the place of the higher-order rule in this explanation.

(compare traditional virtue theory). It is sufficient for current purposes to point out that the reasons approach with which I have been contrasting the normative explanation of attitudes is explicitly act-centered.

By contrast it is not plausible for many particular emotions and desires to be the fundamental locus of normative support in this sense. The most fundamental assessment of such attitudes does not take place at the point of manifestation or triggering of a particular attitude, for instance some desire that p or fear about q or excitement about r . "It is important to understand emotions not as isolated mental states, but rather in terms of their rational connections to other mental states, including other emotions" (Bennett Helm 2001, p.67). Here is Agnieszka Jaworska on a similar theme:

Joy at the successes of the object of care rationally requires sadness at the object's failures; if things go well for the object, fear or hope are rationally required to turn into an emotion such as relief; and so forth. Thanks to their rational and referential interconnection, the individual caring emotions are intelligible as mental states of one agent. But, more important, by virtue of his steadfast emotional attunement to the ebb and flow of the fortunes of the object, the caring subject imbues the object with importance. A subject cognitively sophisticated enough to employ, at least implicitly, the concept of importance would *comprehend* the object's importance, and this can inspire further cognitive activity, for example, further inquisitiveness about the object, or the formation of stable intentions, plans, and policies concerning the object. This cognitive sophistication is necessary for genuine caring. (Jaworska 2007, p. 484)

Here is a simplified version of the idea. I have been distinguishing *enduring states*, such as caring about a friend, from *particular attitudes*, such as sadness about the friend's pain. Plausibly, in many cases, enduring attitudes provide the enabling conditions for the fittingness of some particular attitude. You care a great deal about the Nikkei index, or Liberace décor, or unusual styles of waltz. These enduring states explain the fittingness of your being upset at a sudden change in stock prices, or excited by a new exhibition of bespoke white pianos, or fearful of the rising dominance of West Coast Swing. It is fitting for the mother of the murderer to feel intense grief at his sentencing, but not for us.

In all these cases, it is plausible that the fundamental locus of normative support for attitudes is some temporally extended state such as care or love, rather than the token emotional responses.⁴⁴ Particular emotions or desires are often important because they are ways of manifesting these enduring states. Indeed, this change of focus from the token attitudes to the enduring states for which they constitute fitting modes of expression might help to

⁴⁴ This won't apply perfectly to all the cases, of course. Perhaps some attitudes are fitting for everyone, e.g. fear of one's own death. Perhaps some attitudes are fitting in virtue of commitments or roles that are not themselves enduring attitudes, e.g. being a priest.

explain some puzzling questions in value theory. If a theorist looks just at the intrinsic features of sadness in response to a friend's loss, it can be hard to explain why such an emotion is valuable. By looking instead at this emotion as a way of being 'attuned' to one's friend (to borrow Helm's expression), the value of this reaction seems less mysterious.

So particular emotions and desires are not a fundamental locus of normative support, and they are not governed by reasons. But this doesn't mean they are not important. On the contrary, they are extremely important. Fitting attitudes fit *us* together. They are manifestations of a particular way of relating directly to other things that really matter – our future selves, our families, friends, careers, our passions. But it is not the manifestation of fitting attitudes, on its own, that is important, but rather what they *signify* or help to *constitute* – which is this openness or attunement to others. Clearly it is good, in some sense, to be directly emotionally sensitive to significant events involving people and things we care about – to feel joy when they feel joy, and even to suffer when they suffer.

Quite separately, we can assess the value of these enduring states. Do we have reasons to change what we care about? Do we have reasons to change whatever it is that makes this attitude fitting in this situation? My long-suffering father has asked himself this question many times about his loyal support for the hapless Kilmarnock Football Club. What is important is that this is not a question about our reasons for *particular* emotional attitudes.

At this point it is helpful to remember the distinction that John Rawls articulated in his justly famous 'Two Concepts of Rules' (1955). The relationship between a particular attitude and an enduring state is a bit like the relationship between a move in chess and the playing of the game of chess. It might be fitting in the context of a game to take someone's rook, or even to sacrifice your own rook. The normativity governing these moves is quite different from your reasons to be playing the game at all. So it is, I suggest, with suffering with one's friends, and caring for them. If reasons are applicable to anything in the ethics of psychology, it will be to issues concerning differentially valuable enduring states and other commitments, and perhaps the value of certain conventions for manifesting care and concern. These values may, in turn, explain the conditions under which fittingness facts are *genuinely* normative.⁴⁵

⁴⁵ Thanks to (many people).

Works Cited

- Broome, J. 2013. *Rationality Through Reasoning*. London: Wiley-Blackwell.
- Brunero, J. 2013. 'Reasons as Explanations,' *Philosophical Studies* 165: 805-24.
- D'Arms, J. and Jacobson, D. 2000. 'The Moralistic Fallacy: on the Appropriateness of Emotions.' *Philosophy & Phenomenological Research*, volume 61, number 1.
- Fine, K. 2012. 'A Guide to Ground.' In *Metaphysical Grounding*, Fabrice Correia & Benjamin Schnieder (eds.), Cambridge: Cambridge University Press.
- Fogal, D. 2016. 'Reasons, Reason, and Context,' in *Weighing Reasons*, Lord & Maguire (eds.). Oxford: OUP.
- Gert, J. 2003. 'Requiring and Justifying: Two Dimensions of Normative Strength' *Erkenntnis* Vol. 59, No. 1, pp. 5-36.
- Grice, P. 1975, 'Logic and conversation.'. In Cole, P.; Morgan, J. *Syntax and semantics*. 3: Speech acts. New York: Academic Press.
- Harman, E. 2016. 'Morally Permissible Moral Mistakes.' In *Ethics*, volume 126, number 2.
- Helm, B. 2001. *Emotional Reason: Deliberation, Motivation, and the Nature of Value*, Cambridge: Cambridge University Press.
- Honderich, T., 1982, "The Argument for Anomalous Monism", *Analysis* 42, 59-64.
- Horty, J. 2012. *Reasons as Defaults*. Oxford: OUP.
- Hurka, T. 1998. 'Two kinds of Organic Unity.' *The Journal of Ethics*.
- Jaworksa, A. 2007. 'Caring and Full Moral Standing.' *Ethics*.
- Kearns, S. & Star, D. 2008. 'Reasons: Explanations or Evidence?' *Ethics* 119; 31-56.
- Lewis, D. 1986. 'Causal Explanation.' In *Philosophical Papers*, volume 2, 214-240. Oxford: Oxford University Press
- Lord, E. & Maguire, B. 2016. 'An Opinionated Guide to the Weight of Reasons.' In *Weighing Reasons*, Lord & Maguire (eds). Oxford: OUP.
- Maguire, B. (2016). 'The Value-Based Theory of Reasons,' *Ergo*.

- Markovits, J. 2010. 'Acting for the Right Reasons.' *Philosophical Review*, Vol. 119, No. 2.
- Moore, G. E. 1903. *Principia Ethica*. Cambridge: Cambridge University Press.
- Oddie, G. 2014. 'Truthlikeness.' In the *Stanford Encyclopedia of Philosophy*. Available here: <http://plato.stanford.edu/entries/truthlikeness/> Accessed 16 March 2016.
- Parfit, D. 2011. *On What Matters*. Oxford: OUP.
- Portmore, D. 2011. *Commonsense Consequentialism*. Oxford: OUP.
- Rabinowicz, W. & Rønnow-Rasmussen, T. 2004. 'The Strike of the Demon: On Fitting Pro-Attitudes and Value.' 114 (3):391-423.
- Rawls, J. 1955. 'Two Concepts of Rules.' *The Philosophical Review*.
- Raz, J. 1975. *Practical Reason and Norms*. Oxford: OUP.
- Raz, J. 1999. *Engaging Reason*. Oxford: OUP.
- Raz, J. 2011. *From Normativity to Responsibility*. Oxford: OUP.
- Roberts, R.C. 2004. 'The blessings of gratitude: A conceptual analysis.' In Emmons & McCullough (eds.) *The Psychology of Gratitude*, Oxford: OUP.
- Scanlon, T.M. 2014. *Being Realistic About Reasons*. Oxford University Press.
- Schafer, K. 2016. 'The Modesty of the Moral Point of View' in *Weighing Reasons*, Lord & Maguire (eds), Oxford: OUP.
- Schroeder, M. 2007. *Slaves of the Passions*. Oxford University Press.
- Schroeder, M. 2010. 'Value and the Right Kind of Reason.' In *Oxford Studies in Metaethics*.
- Singer, P. 1971. 'Famine, Affluence, and Morality.' *Philosophy & Public Affairs*, I (3): 229-243.
- Staffel, J. 2015. 'Measuring the Overall Incoherence of Credence Functions.' *Synthese*, Volume 192, Issue 5, pp 1467-1493.
- Way, J. 2012. 'Transmission and the Wrong Kinds of Reason.' *Ethics*.
- Way, J. 2013. 'Value and Reasons to Favour', in Russ Shafer-Landau (ed), *Oxford Studies in Metaethics*, volume 8.
- Way, J. & McHugh, C. 2015. 'Fittingness First.' *Ethics*

Wedgwood, R. 2015. 'The Pitfalls of 'Reasons'.' *Philosophical Issues*, 25(1): 123-143.