15

Weighing Explanations

Stephen Kearns and Daniel Star

The primary goal of John Broome's new book, *Rationality Through Reasoning* (2013), is to outline and defend an account of reasoning that makes it clear how it is possible to actively become more rational by means of engaging in reasoning.[1,2] In the process Broome finds it necessary to also provide his own accounts of ought, reasons, and requirements. We will focus here on the account of reasons. This is not the first time we have done so. In an earlier paper (Kearns and Star 2008),[3] we contrasted Broome's account with our own favored account of reasons (*reasons as evidence*).[4] Although there are some differences between the views defended in the relevant chapters of Broome's book (chs. 3 and 4) and the draft manuscript and earlier papers that we used as the basis of our discussion in our earlier paper, these do not, for the most part, substantially affect our earlier arguments.

It is clear that in articulating an alternative account of reasons we were also heavily influenced by Broome, so we are particularly grateful to have this opportunity to contribute a piece to a *Festschrift* for him.[5] In a response to us and some other critics, Broome (2008) presented some challenges for our account of reasons, but did not address our criticisms of his own account (and we responded, in turn, to Broome's challenges in Kearns and Star 2013). Here we will first provide updated versions of our earlier concerns, since they mostly still seem pertinent. We will then turn to provide a fresh response to his account of reasons that focuses on the notion of a weighing explanation. On Broome's account, pro tanto reasons are facts cited in weighing explanations of what one ought to do; facts that have weights. It is not clear

[1] We would like to thank the editors of this volume for providing helpful feedback.

[2] All numbers in parentheses that are not clearly references to other works should be taken to be references to this book.

[3] From now on, we refer to this paper without mentioning authorship (in other words, any references to 2008 where an author is not mentioned are to this paper).

[4] The central claim of which is: Necessarily, a fact X is a reason for an agent N to F if and only if X is evidence that N ought to F.

[5] Judith Jarvis Thomson (2008) also has an account of reasons that is similar to ours, and she does not appear to have been influenced by Broome. We did not know about Thomson's views on this topic at the time of writing and publishing our early papers (2008, 2009).

to us what the idea that pro tanto reasons have weights really amounts to. While recognizing that a simple analogy with putative non-normative weighing explanations involving physical weights initially seems helpful, we argue that the notion of a weighing explanation, especially a normative weighing explanation, does not ultimately stand up to scrutiny.

1 Reasons as Normative Explanations, or Parts of Normative Weighing Explanations

According to Broome there are at least two types of normative reasons: pro toto reasons and pro tanto reasons. The first of these is defined as follows:

> A pro toto *reason* for *N* to *F* is an explanation of why *N* ought to *F*.
>
> (50)

Broome immediately adds "a *pro toto* reason for you to *F* is something that makes it the case that you ought to *F*," and he clarifies what he has in mind when referring to explanations by saying, "An explanation need not be full or complete, and what counts as an explanation may depend on the context" (50). He provides the following examples:

> Suppose you ought to visit Mr Reed. The explanation might be that he is the best dentist around. If that is indeed the explanation, it is a *pro toto* reason to visit him. In a different context, the explanation might be that you ought to visit the best dentist around. This would be a *pro toto* reason for you to visit Mr Reed, in that context. We now have two alternative *pro toto* reasons, but they do not compete with each other.
>
> (50)[6]

Like all of us working on reasons, Broome is well aware that there are clearly non-normative uses of the word "reason"; in particular there are general "explanatory reasons": "A 'reason' in this sense is more or less synonymous with 'an explanation of ' " (46). Given that this is a clear sense of "reason," should we not conclude that pro toto reasons are simply non-normative explanatory reasons (of normative facts), since their structure would seem to require this categorization? Broome responds to this question by saying "The 'reason why' (meaning explanation) bumps into the normative 'ought,' yielding a normative sense of 'a reason' that combines the meaning of both" (50). He immediately admits this is a somewhat "picturesque etymology," but we do not object to this element of his view. On *reasons as evidence*, it might also be objected that the reasons we take to be normative are not really normative (as distinct from oughts), but we claim that they "inherit" their normativity from oughts. This is less picturesque, but probably not much more illuminative, for all that.

[6] It is a little unclear whether "context" here refers to decision situations themselves, or contexts in which decision situations are described, or either. We take it he means either.

The second class of normative reasons that Broome is interested in analyzing are pro tanto reasons. These are reasons that appear to be more or less weighty, and might favor an action that we actually ought not do (pro toto reasons never do this). Broome offers the following "functional definition":

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

In order to understand this definition, we need to understand what a weighing explanation is, and what a for-*F* role is. Broome approaches normative weighing explanations by first providing an example of a non-normative weighing explanation that involves a mechanical scale with items of different weights placed on each side of it. When considering such an example, we can explain why the balance tips left (say) by referring to the combined weights of the object on the left and the combined weights of the object on the right and by noting that the combined weight on the left exceeds that on the right (52). In the case of normative weighing explanations, it is the weight of reasons we are considering, rather than the weight of physical objects, and we are looking for an explanation of what an agent ought to do. As for the "for-*F* role," its presence in Broome's definition is what makes the definition a functional one. This means that, in effect, "a *pro tanto* reason can be defined without any reference to the for-*F* role . . . it is something that has a weight, where the weights of reasons combine in some way to determine whether or not you ought to *F*" (54). Weights, then, are more fundamental than for-*F* roles.

In passing, we have three questions about Broome's account of reasons. These are not really objections, so much as things we would like to see clarified. The first question we have is whether every weighing explanation that you ought to *F* (used to locate pro tanto reasons) is also specifying a pro toto reason? Broome says, "The combined weights of the reasons for you to *F* exceed the combined weights of the reasons for you not to *F*. That is why you ought to *F*" (52).[7] The "that" here seems to refer back to the whole of the preceding sentence (or, more carefully, the fact that would be picked out in an accurate filling out of the schema provided in the preceding sentence), and we have been told that every *X* in a true statement of the form "*X* is why you ought to *F*" is a pro toto reason.

Second, it seems to us that, on Broome's analyses, every pro tanto reason that aligns with an ought (i.e. in every case where there is pro tanto reason to *F and* where the relevant agent ought to *F*) is also *ipso facto* a pro toto reason, and that this is true even when the pro tanto reason in question is quite weak and *would not be enough to*

[7] On the same page Broome also says, "We often say you ought to do something because the balance of reasons is in favour of your doing it. When an explanation takes this form, I shall call it . . . a 'normative weighing explanation'."

determine that one ought to F in the absence of other reasons that align with the ought in question. Is this right? This entailment would certainly not be the case if the notion of a pro toto reason were identical to the notion of a *complete* reason, but, as just noted, these notions are quite separate.

Why do we think Broome might be committed to thinking that even many very weak pro tanto reasons that happen to align with oughts must be pro toto reasons? This seems to follow from something he says when discussing facts about explanations that he takes to apply to his own account of pro toto reasons:

> Although I shall not try to describe the nature of the explaining relation, I do need to say something about the individuation of explanations. Suppose Joanne broke a slate a while ago, and as a result the roof leaks. It rained last night, and today the carpet is wet. When we enquire why the carpet is wet, you might say the explanation is that it rained last night. I might say that it is that Joanne broke a slate. Someone else might say that it is that the roof leaks. We respectively make these statements: "The explanation of why the carpet is wet is that it rained last night", "The explanation of why the carpet is wet is that Joanne broke a slate", and "The explanation of why the carpet is wet is that the roof leaks". Read literally, no two of our statements can be true together. Still, our explanations are not rivals . . . there is really one big explanation of why the carpet is wet. It is a complex fact that includes as parts all the separate facts the three of us described. Each of us is picking out a part to stand in for the whole. We call it *the* explanation because it is standing in for the one big explanation. . . . Which part we pick out will depend on our context: our background knowledge, our interests in the matter and so on.
>
> (48–9)

Notice that it is a feature of the example Broome uses that certain facts that count as explanations fail to be sufficient conditions for the obtaining of the fact that they explain (the missing slate will not make the carpet wet in the absence of bad weather). Consider the following example. Suppose Jake has promised to meet a friend in ten minutes, and is walking to meet him (we include this fact merely in order to ensure we are talking about a case involving pro tanto reasons). Suppose Jake comes across a person needing urgent medical assistance, which will require a trip to an emergency ward. Fortunately she is lying next to a Mercedes that has been left unlocked with keys still in the ignition. Jake judges that, despite the promise to his friend which he will not be able to fulfill if he drives to the hospital, he ought to drive her to the hospital. A bystander might correctly say, *The explanation of why he ought to drive her to the hospital is that she needs urgent medical assistance.* A different bystander might correctly say, *The explanation of why he ought to drive her to the hospital is that there is an unlocked Mercedes parked nearby.* If these claims are both appropriate (and we think they are, since, as Broome indicates, what is left out can depend on background knowledge),[8] then there seem to be two possibly problematic consequences: (1) the fact

[8] It might be thought that, "The explanation of why he ought to drive her to the hospital is that there is an unlocked Mercedes parked nearby," would not be a true sentence expressed in, or about, the imagined context. It is important to bear in mind Broome's point about background knowledge. Imagine the

that there is an unlocked Mercedes parked nearby gets to count as a (pro toto) *reason*, even though many would class it as a mere *enabler*;[9] and (2) the fact that there is an unlocked Mercedes parked nearby gets to count as a pro toto reason, even though it clearly is not a sufficient condition in this context for it being the case that one ought to drive to the hospital.

Now let us alter the example a little to make a third point. Suppose the person encountered is clearly merely somewhat unwell and probably should see a doctor soon, rather than in obvious need of urgent medical care, and let us suppose also that Jake has always wanted to drive a Mercedes, and let us also suppose that the law of the land happens to say it is legal to borrow unlocked cars in situations such as the one he is in, and that he will impress some friends standing nearby, and so on. This is intended to be a sketch of a case where none of the reasons that favor driving the car to the hospital are individually strong enough to justify doing so, but together do justify doing so (there presumably will be such cases, so we need not worry that our example has become a little outlandish). Then it will turn out that (3) *none* of these pro toto reasons is really sufficient to make it the case that Jake ought to drive the car.

One response Broome might have to these concerns is to say that we have misunderstood his proposal: the reason "*The* explanation why…" statements can be piled on top of each other without fear of contradiction is because each such true proposition expressed is really extensionally equivalent (for they purposely refer to the complete explanation, merely by explicitly highlighting one part of the explanation). The worries just expressed do not seem to get off the ground if this is true, but it means that for every ought that can be explained using a pro toto reason, there is, strictly speaking, only one pro toto reason, and that is the complete explanation. Admittedly, this is a reasonable way of reading the quotation we provided immediately above, but we are not convinced that this is the best way of explaining why explanatory statements of one and the same fact can often be piled on top of each other. More crucially, it does *not* fit well with the definition of a pro toto reason as "*an* explanation of why *N* ought to *F*" (50, our emphasis), and with the quotations we provided at the beginning of this section—they explicitly state that there is often more than one pro toto reason in play in relation to any one ought fact, and that pro toto reasons need not be complete explanations.

Our third question is related to our second. Broome claims that "the relation of explaining is the relation of making so" (48), but it seems to us that there is

following dialogue: *He ought to drive her to the hospital.—How could that be the case? He left his car at home.—Oh, but the fact that there is an unlocked Mercedes nearby is what explains why he ought to drive her to the hospital.* Of course, a complete explanation would include something about the fact that driving to hospital is the best means available, but that does not render the last sentence in the conversation false. In any case, we could use a different example to make our points here.

[9] Full disclosure: it seems we defenders of *reasons as evidence* also have to accept that enablers are reasons, since they are evidence concerning what one ought to do (see Fletcher 2013 and Setiya 2014). We do not think this is too bad, actually, but many would disagree with us about this.

something to be said for the view that none of the weak reasons in favor of going to the hospital in the second, altered case above individually *make it the case that* one ought to drive to the hospital, even though together they do (and even though they might be cited in true and informative "The explanation why he ought…" statements). It seems plausible to us that the relation of "making it the case that" should be reserved for a stronger metaphysical relation (perhaps a *grounding* relation) than the mere explanatory relation, which does seem somewhat more interest-relative, hence somewhat less metaphysically robust. The question is: how metaphysically robust are explanations meant to be on Broome's account? This is an important question, because, given Broome's account of reasons, it directly bears on the question of how metaphysically robust reasons are meant to be.

2 The Objections of "Reasons: Explanations or Evidence?"

Our earlier paper provided seven key objections to Broome's account of reasons. We take some of these to be stronger than others. Here is a brief restatement of our key objections. (In most cases, we have had to omit aspects of, or variations on, these objections; aspects and variations that tend to strengthen them.)

i. The deliberative role of reasons. We take it to be a truism about reasons that they play an essential role in actual deliberation with respect to answering the question, *What ought I do?* (2008: 40–2). Typically, one works out what one ought to do in the situation one finds oneself in *by* considering the reasons that apply to one in this situation; one generally does not first know what one ought to do, and only then ask oneself what reasons are applicable. On the other hand, explanations, or parts of explanations, of any facts are things we reason our way towards only after we have an explanandum in sight. The view that reasons are parts of explanations of ought facts, whilst not strictly inconsistent with the view that practical reasoning typically starts with a grasp of specific reasons in combination with ignorance about specific ought facts (explanans on Broome's account) and works towards conclusions about what specifically ought to be done (explanandum on Broome's account), certainly prima facie conflicts with this truism. At the very least, we are owed an account of what explains the truism (or, alternatively, convincing grounds for rejecting it). We will see this objection is also connected to another objection, (v) below.

ii. The "for-F role" in normative weighing explanations. It might be thought that Broome fails to provide a non-circular analysis of pro tanto reasons, since one might suspect it is impossible to understand what a "for-*F* role" is independently of an understanding of what it is for facts to *count in favor* of acts (and beliefs, etc.). This is one place where what Broome says in his book is more straightforward and

satisfactory than the manuscript we were responding to in our earlier paper (2008: 42–4). As noted above, he now takes the "for-*F* role" to simply stand in for the weight of the relevant reason involved in any case. One might still ask: is it possible to understand the idea that normative reasons have *weights*, prior to an understanding of *counting in favor*? Broome clearly thinks it is (54). We are not so sure about this.

Of course, in one way we can obviously understand "weight" independently of "counting in favor," and that is in the common, non-normative use of "weight," as when we consider physical scales. But the physical examples are meant to be metaphorical. Might we come to understand normative weight via first understanding physical weights and then making use of a metaphor? Perhaps this might work when we are considering the adoption of a *theory* of reasons, but this seems like a highly questionable hypothesis when considering how actual ordinary agents come to an appreciation of the weight of reasons early in life (say). Broome might object that it does not matter for his purposes whether agents actually come to understand that reasons have weights prior to understanding that they count in favor; what is crucial is that agents *could* do so. Perhaps they could; we are just not so sure of this.

iii. The mystery of the weight of reasons. In any case, we now think Broome's way of explaining what a "for-*F* role" is leads to a more important objection, and that is that Broome leaves it utterly mysterious what it is for a reason to have a weight (2008: 44–5). One way of putting this objection is to consider his account of reasons in relation to both the primitivist account endorsed by T. M. Scanlon and Derek Parfit (according to which no informative analysis of "reason" is possible) and our own analysis of reasons. At first blush, it may seem that Broome's account of reasons improves on the primitivist account: after all, he seems to offer an analysis where the primitivist account embraces a pessimism about the possibility of being able to do so. But if we are now to take the concept of normative weight as basic, it is not clear that Broome's analysis is really more satisfactory than the primitivist analysis.[10] Defenders of *reasons as evidence*, on the other hand, may be able to explain weights in other terms. We will say more about weights in the last section of the chapter.

[10] It might be objected that *reasons as evidence* does not really fare particularly well on this front, since it leaves *ought* unanalyzed (we have suggested that we can analyze evidence in terms of epistemic probability). Our objection is that Broome leaves both *ought* and *weights* unanalyzed. Defenders of reasons primitivism sometimes think they can at least explain oughts in terms of reasons. Regarding weights, Scanlon (2014: 111) rejects the view that weights are real, independent properties of reasons: "I conclude that there is no non-normative coin or normative coin, in terms of which the strength of reasons, considered on their own . . . can be expressed. . . . The strength of a reason is an essentially comparative notion, understood only in relation to other particular reasons" (he goes on to add that his approach to understanding the strength of reasons is "top-down," working from general principles that express normative relations concerning sufficient and conclusive reasons). Reasons primitivists (and others) often use the term "strength" rather than "weight," but we take these terms to be roughly equivalent when attributed to reasons; at least, we see no particular lessening of the element of mystery when one word is used rather than the other (and, absent the adoption of something like Scanlon's "top-down" approach, reasons primitivists are not in a better position than Broome, in this respect).

iv. Ambiguity concerning "reason," or defining "reason" disjunctively. This objection can be simply stated (2008: 45–6). According to Broome, there are two types of normative reasons: pro toto reasons, which are explanations of ought facts, and pro tanto reasons, which corresponds to *parts* of normative weighing explanations. Broome may be happy with there being two senses of (normative) "reason," but there do not seem to be any independent grounds for thinking that (normative) "reason" is ambiguous in this way. An alternative to postulating an ambiguity here is to combine both definitions of "reason" that Broome provides in order to form a disjunctive definition (2008: 35); the problem then might be said to be that such a disjunctive definition, combining reference to both weighing and non-weighing explanations, is not sufficiently unified. It certainly does not appear to cut nature at its joints (something one might hope for with definitions).

v. Too few reasons. Broome explicitly rejects the idea that facts that are (merely) evidence that you ought to *F* are, or provide, reasons to *F*, even when they are conclusive evidence that one ought to *F* (51). Here, it is fair to say, Broome is in agreement with many philosophers working on reasons. And when asked to assess, for instance, the claim that the fact that a minister accepted a bribe is a reason for him to resign alongside the claim that the fact that a reliable newspaper has reported that a minister ought to resign is a reason for him to resign (or similar pairs of claims), we have found audiences often intuitively judge that the first claim is true and the second false (in the relevant circumstances). Broome clearly shares this intuition, but he does not present much of an argument for it being correct (one may take his claim that evidence that one ought to *F* does not *explain* why one ought to *F* as a crucial step in an argument here, but this argument would rely on his account of reasons being correct, and, in any case, as we point out in (vi), he seems to say elsewhere that evidence in general *does* very often explain why one ought to *F*, so it is not clear why evidence that one ought is special in this respect). We agree with Mark Schroeder (2007: 92–7) that negative existential intuitions about reasons are often misleading, and that much can be done to dispel such intuitions by focusing on conversational pragmatics.[11]

We also think that there is a serious danger that the denial that evidence that one ought can provide reasons for action will lead to a too few reasons problem (2008:

[11] On Schroeder's Humean account of reasons, there is a reason for you to eat your car (it contains iron). The argument that Schroeder provides appeals to Gricean maxims; the defender of *reasons as evidence* might hope to also do this, but the precise form of the argument will need to take a different form (Schroeder relies on the reasons in question being very weak, but we cannot rely on that assumption, for a start). Even so, assuming Schroeder's argument is successful, it is *very* odd to think that the fact that your car contains iron *is* a reason to eat your car while the fact that a reliable book says you ought to eat spinach (one might imagine the author knows but does not say that spinach contains iron) is *not* a reason to eat spinach. Some philosophers worry that allowing for as many reasons as we do creates a counting problem (as it would unjustifiably license the continual adding together of the weights of reasons that are not independent of each other); we deny that there really is a problem of this kind (see Star 2011: 93–4, forthcoming).

48–9; 2009: 233–4). From the perspective of fundamental normative ethics, conceived of as a project that attempts to articulate and defend fundamental ethical principles, many of the mundane facts that we commonly take to be reasons are not really fundamental explanations of what we ought to do. If hedonistic utilitarianism were true (for instance), then the fact that an act would produce pleasure would be fit to be part of a fundamental explanation of why one ought to do that act, but the fact that an act would be the fulfilling of a promise would not be part of such an explanation. Assuming one thinks there is a point to the project of normative ethics (as we do), it is not at all attractive to suppose that it is only the features mentioned in fundamental explanations that get to count as reasons, since this would make too many of our ordinary reason claims false (which, as well as being very counter-intuitive, would not fit well with the deliberative role of reasons discussed above).[12]

Of course, Broome need not think that it is only facts cited in fundamental explanations of ought claims that get to count as reasons, but this is precisely why we wonder why he is so sure that evidence that one ought to *F* cannot non-fundamentally explain why one ought to *F*. If a reliable book says one ought to eat cabbage, why can that not non-fundamentally explain why one ought to eat cabbage? After all, there does seem to be something remiss about not following the advice, even if it turns out a reliable book (or a reliable doctor, say) is wrong on this occasion (and there are cases where not following testimony that provides evidence concerning what one ought to do seems morally wrong, e.g. if a reliable doctor claims a drug will save a child's life and requires the consent of a parent).

Furthermore, we take it that testimony can provide reasons for action in virtue of the fact that it can *reliably indicate* what one ought to do. Assuming Broome does not wish to claim that only (parts of) fundamental normative explanations provide reasons, general non-fundamental facts will not invariably coincide with fundamental reason facts, but will instead merely *reliably indicate* the presence of such facts. From this perspective, facts like the fact that a newspaper has reported a minister ought to resign and facts like the fact that a minister accepted a bribe seem to be on a par.[13]

[12] To be clear: this is not meant to be an objection to utilitarianism. Rather, we are simply saying that the utilitarian can and should adopt an account of reasons that allows the fact that an act would be the fulfilling of a promise to count as a reason. Star (forthcoming) argues that *reasons as evidence* best elucidates the relation between fundamental and derivative reasons in a way that respects both ordinary virtue and normative ethical theory.

[13] The difference in the extent to which each of these facts *reliably indicates* that he ought to resign is a difference in degree only. Depending on the details of one's favorite normative ethical theory, one might need to describe a fairly outlandish case to have it come out that accepting a bribe on this occasion is not accompanied by a fundamental ethical reason, but it will be possible to do so, unless the ethical theory recognizes bribes as fundamental reasons, or something always entailed by bribes, e.g., perhaps, deception, but then we could just use some other example following the same general recipe.

vi. Ought and evidence. Broome's denial that evidence provides reasons seems to be in tension with his account of *ought*. Here we must ignore many subtle features of Broome's discussion of ought, and cut to the chase: he takes the central ought (the ought that the rational requirement of Enkrasia applies to) to be the *prospective* ought, and not the *objective* ought (40–4). Broome takes it that it is easiest to illustrate the difference between these oughts by focusing on situations where consequentialist moral claims at a very general level might apply; here the objective "ought" (Broome is not sure that there really is such an ought) will be determined by actual outcomes, independent of evidence, but the prospective ought will be determined by prospects, which are dependent on the probabilities of outcomes. Broome thinks that the school of thought that the probabilities in question are to be explained in terms of available evidence gets the right results much of the time, but, for reasons that are fairly complex, it ultimately needs to give way to the main alternative school of thought, namely that the probabilities in question are purely subjective (i.e. the probabilities that Bayesian subjectivists focus on). At no point is there a parallel to his blanket refusal to allow that evidence that one ought to *F* can be a reason. He concludes, "Prospective oughts . . . are subjective only within limits. *Your subjective probabilities must be constrained by the evidence* . . . but evidence does not always determine all of the probabilities. It may leave a gap, and if it does Bayesians fill the gap with subjective probabilities" (44, emphasis added). To simplify in a way that we hope Broome will not mind (he does, after all, seem very sympathetic to the evidence-based account of probability): to a considerable although not complete extent, evidence will determine what we ought to do.

Now our criticism is fairly easy to state: given that Broome believes that, to a large extent anyway, evidence explains what one ought to do, it seems that his claim that evidence that one ought to *F* cannot be a reason is unwarranted, given that reasons are, on his own account, parts of explanations of why one ought to *F*.

vii. Fundamental ought facts. The following seems like a very plausible claim: "if one ought to *F*, then there is a reason to *F*"; otherwise, an agent might do something which they rightly take to be something they ought to do at the same time as thinking there is no reason at all to do it! It also seems reasonable to suppose that some ought claims or principles are *fundamental*, in the sense that the truth of these claims is not something that can be further explained. Whether the correct fundamental ethical theory is consequentialist or non-consequentialist, it seems likely that the principle(s) at the foundation of this theory may lack further explanation. If it is also the case that reasons are parts of explanations concerning what agents ought to do, then it will not be true of fundamental oughts that "if one ought to *F*, then there is a reason to *F*" (2008: 53).

We now think that Broome can avoid this objection by utilizing the machinery of *requirements* (provided in ch. 7) to hold on to "if one ought to *F*, then there is a reason to *F*," while holding that there are no unexplained oughts. On his view,

fundamental ethical principles are best cast in terms of normative requirement claims, rather than ought claims. For the sake of simplicity, assume moral requirements override other types of normative requirements (we assume this solution will also work with more complex theories of the way normative requirements of different types interact). Whenever one is morally required to *F*, this requirement—for example, to respect autonomy, or, say, to maximize total expected welfare—will be a pro toto reason to *F*; or, in other words, the requirement will explain why one ought to *F*.[14]

3 Weighing Explanations

As we mentioned above, it is clear that Broome thinks that the important class of reasons generally referred to as pro tanto (i.e. all reasons that may be weighed) are to be analyzed by reference to weighing explanations. It is thus crucial that, as long as we think that there are indeed pro tanto reasons (an assumption that seems very safe to hold to), there actually must be some genuine weighing explanations. We will now present a new line of argument for the conclusion that there are, in fact, no genuine weighing explanations. This may appear to be a rather strong conclusion, so it is important that we proceed by first paying careful attention to Broome's account of weighing explanations.

Mechanical weighing explanations

Broome begins his discussion of weighing explanations in general by describing mechanical weighing explanations:

> How does a mechanical weighing explanation go? Suppose a balance tips to the left. A typical weighing explanation of why it does so will go like this. There is at least one object in the left-hand pan, and there may be one or more objects in the right-hand pan. Each object has a weight. The combined weights of the objects in the left-hand pan exceed the combined weights of the objects in the right-hand pan. That is why the balance tips left. (In the theory of mechanics, the word "weight" refers to a force, which is a vector. I am not using it that way, but colloquially, to refer to a scalar magnitude.)
>
> (52)

Broome then contends that this is analogous to weighing explanations of normative facts:

> Suppose you ought to *F*. If there is a weighing explanation of why, it takes an analogous form. There is at least one reason for you to *F*, and there may be one or more reasons for you not to *F*. Each reason has a weight. The combined weights of the reasons for you to *F* exceeds the combined weights of the reasons for you not to *F*. That is why you ought to *F*. In this analogous

[14] It is no accident that Broome is careful to note in the earlier chapter on reasons that the facts that are pro toto reasons may themselves be normative facts (51).

explanation, the fact that you ought to *F* is the explanandum, and analogous to the fact that the balance tips left. A reason for you to *F* is analogous to an object in the left-hand pan, and a reason for you not to *F* is analogous to an object in the right-hand pan.

(52)

Broome's examples of normative weighing explanations

Broome provides two main examples of supposed weighing explanations of normative facts. The first runs as follows:

Suppose you ought to bring some wine to the party, and suppose the explanation of why is this. You promised to bring some wine, and you ought to keep your promises unless significant harm will result from doing so. In this case, suppose the only harm that will result is the cost to you in money and inconvenience of buying the wine, and this is not significant.

(55)

Broome provides the following gloss on why this is a good candidate for being a weighing explanation:

In favour of your bringing some wine is your promise. Against is the cost. Your promise plays the for-bringing-wine role in the explanation. It is therefore a *pro tanto* reason to bring wine. The cost plays the against-bringing-wine role. It is therefore a *pro tanto* reason not to bring wine. The promise is a reason of greater weight than the cost.

(56)

Here is Broome's second example, and his reasons for considering it a weighing explanation:

Suppose you are choosing between Montreux and Marrakesh as places to visit. Suppose that you ought to choose Montreux, on the grounds that it is a pleasant resort and not so far away as Marrakesh, though less exotic. The explanation of why you ought to visit Montreux is a weighing one. In favour of visiting Montreux are its pleasantness and proximity. Each plays the for-Montreux role in the explanation. They are *pro tanto* reasons for you to visit Montreux. Against visiting Montreux is the exoticness of the alternative, Marrakesh. That plays the against-Montreux role. It is a *pro tanto* reason for you not to visit Montreux. The reasons for visiting Montreux outweigh the one for not doing so.

(56)

Broome uses this example to illustrate his distinction between particular reasons and general reasons:

The explanation could be filled out in more detail. Montreux's pleasantness consists in its beautiful views, cosy restaurants, opportunities for lake cruises and other pleasant features. Its nearness means your journey will add less carbon dioxide to the atmosphere. On the other hand, the exoticness of Marrakesh consists in teeming bazaars, belly-dancing and opportunities for camel rides. One reason for visiting Montreux is the beautiful views. Another is the cosiness of the restaurants. A reason for not visiting Montreux is that going instead to Marrakesh will allow you to take a camel-ride. And so on.

(56)

Particular reasons figure in weighing explanations that are simply more detailed spellings-out of the weighing explanations which feature general reasons. The former explanations are not rivals of the latter, and the particular reasons to *F* do not combine weights with the general reasons to *F* to produce a stronger case for *F*-ing.

Some other explanations

Consider the following examples of explanations of non-normative facts, not owed to Broome:

- A bridge collapses. Here are some facts that may figure in the explanation of this fact. Though the bridge was well-designed and constructed, built from suitable materials and regularly (and competently) maintained, it was overloaded with cars and lorries while being buffeted by particularly strong winds.
- Here is another example. A person dies from cancer. He succumbs to the illness because it is a particularly aggressive strain of cancer and the person's immune system is compromised. He dies despite the fact that the cancer was noticed early on, and the person is treated with often-effective radiotherapy and chemotherapy.
- A third example runs as follows. Someone loses a 100 m sprint Olympic final. She loses because she faces a tough opponent (who wins) and is not very good at handling extreme stress. She had a good chance of winning, however, because she is capable of running faster than her opponent, she prepared well for the race, and she was more motivated to win than anyone.

In the above cases, explanations are offered for various non-normative facts. Interestingly, these explanations are structurally similar to Broome's own examples of supposed weighing explanations of normative facts in at least four ways.

First, in our examples, we can identify facts that, in some sense, play the "for-the-explanandum" role and those that play the "against-the-explanandum" role. Consider the first example. The fact to be explained is that a certain bridge collapses. Some of those facts mentioned in the explanation clearly contribute to the obtaining of this fact—the bridge is heavily weighed down with vehicles; the wind is extremely high. These facts can be considered the "for the bridge collapsing" facts. The other facts mentioned (that the bridge is well designed, constructed, and maintained, etc.), play the "against the bridge collapsing role". These facts might typically explain why the bridge stays up. They are also facts that create an expectation that the bridge will stay up, and which thus create the need for an explanation of why it does not.

Similar points apply to the other two examples. The cancer's aggressiveness and the man's compromised immune system play the for-dying role in the explanation of his death, while the early diagnosis and excellent medical care play the against-dying role. The athlete's speed and well-preparedness play the against-losing role, and the

nervousness plays the for-losing role. The parallels with Broome's examples are obvious—he too identifies for- and against-roles for the facts that constitute explanations of normative facts.

Second, there is also a sense is which the for- (or against-) facts may "win." The sprinter does in fact lose the race, so the for-losing facts "win out" over the against-losing facts. We may, if we like, express this by saying that the for-losing facts outweighed the against-losing facts (though we explore below how useful it really is to think of these explanations as weighing explanations). Again, this closely resembles Broome's examples.

Third, the explanations of our non-normative facts frequently concern quantities of various kinds, and these quantities often make a difference to which fact obtains. The for-dying facts in the second example include the aggressiveness of the cancer and the weak immune system of the man. The for-collapsing facts in the first example concern the strength of the wind and the weight of the load on the bridge. This parallels Broome's examples. The against-buying-the-wine facts in his first example concern the cost of the wine and the inconvenience that buying it causes. Both inconvenience and cost come in degrees (a cost may be low, high, reasonable, or exorbitant), and their degrees may affect whether one should buy the wine or not. The for-choosing-Montreux facts in Broome's second example include Montreux's relative proximity, the degree of which plays a crucial role in Broome's weighing explanation.

Fourth, Broome's distinction between general and particular reasons finds a parallel in our examples. Take the third case above. We may fill out the explanation for the sprinter's loss in more detail:

> The sprinter's training regime, as recommended by top fitness experts, consisted in 90 minute workouts that concentrated on improving her explosive power and maintaining her core strength. Her personal best in previous competitions was 0.1 seconds better than her strongest competitor. Furthermore, she very much wanted to win the race to honor the memory of her father, who was also a sprinter and had died recently. On the other hand, she has consistently found that the more prominent the competition in which she races, the more she feels the pressure. In the past this has led to false-starts. In this (highly prestigious) race, she overcompensates and has a slow start. Though she gains on the eventual winner near the end of the race, she cannot reach her and comes second.

The above explanation simply spells out the original explanation. The above explanation of the loss does not *rival* our first explanation. Nor do the more detailed facts combine weight with the less specific facts to create stronger explanations.

Are our examples weighing explanations?

We have seen that the examples we give above are similar in various ways to Broome's examples of supposed normative weighing explanations. This raises the question of whether our examples are weighing explanations. That is, are they suitably analogous to mechanical weighing explanations?

To help us discover whether they are, let's consider what Broome says about some essential elements of normative weighing explanations:

> These are some features that are essential to a weighing explanation of why you ought to *F*. The explanation must include one or more *pro tanto* reasons, either for you to *F* or for you not to *F*. Each of these reasons must be associated with something that is identified as its "weight". The reasons and their weights play a characteristic role in the explanation. The role is that the weights of the reasons on each side are combined together in some way, and whether or not you ought to F is determined by the combined weights on either side.
>
> (52)

The important points are that the facts involved in weighing explanations have weights (both for the explanandum and against it), and these weights combine with each other in such a way that the facts that favor the (true) explanandum outweigh the facts against it.[15]

The major disanalogy between weighing explanations, as Broome describes them, and our explanations is that, though the facts in our explanations often involve quantities of certain types—quantities, furthermore, that play crucial roles in these explanations (as described in the previous section)—the facts involved do not have anything readily identifiable as *weights*.

Consider our second example (the same general points apply to our other examples). The aggressiveness of the cancer, the weakness of the immune system, the earliness of detection, and the excellence of the treatment all play a role in explaining the death of the patient. However, the quantities involved in the for-dying facts do not together produce a combined weight that outweighs the quantities involved in the against-dying facts. (If they did, presumably such a special quantity would be of great interest to medical science.) Rather than having weights that combine and are then set against each other, the for- and against-dying facts in the explanation of the man's death play different roles.[16]

The job done by the for-dying facts in explaining the man's death is fairly clear. Cancer, particularly aggressive strains, greatly increase the chances of death, especially in patients with compromised immune systems. The exact physical mechanisms of how this happens are relatively well understood. The for-dying facts really do explain (bring it about, cause, etc.) the man's death.

[15] A careful reader will notice that we are using the language of facts when talking of reasons, and that Broome views this as a perhaps optional "regimentation". However, he also believes this regimentation is "harmless" (48), and he generally sticks to it, so we take it that there is nothing problematic with focusing on the regimented version of his view throughout.

[16] It might be claimed that insofar as medical science makes use of data gathered concerning the probabilities of various outcomes occurring, given various conditions, it makes use of data that can be weighed. This may be so, but understanding weights in terms of probabilities leads to *our* conception of reasons as evidence, and cannot help Broome here (assuming he wishes to reject our account of reasons, as we believe he does).

What of the (losing) against-dying facts? What is their role in the explanation of the man's death? We think there at least four main roles. First, such facts do play a background causal role in the man's death (just as do the fact the man was born, that he eats a certain diet, and that he is subject to the Earth's gravitational pull). These facts concern various physical processes. Such processes, and their interaction with the processes involved in the for-dying facts, help bring about the man's death. (Question: Is this enough to make such facts causes of the man's death? Answer: We don't know. If they *are* causes, they are to the same extent that the man's birth and attraction to the Earth are. They are certainly not *salient* causes. They also presumably do not bear the following counterfactual relation to the man's death: were these facts to be false, the man's death would not have occurred. Nor indeed, do these facts increase the probability of the man's death relative to a suitably chosen alternative scenario.)

Second, the against-dying facts (or facts of the same type) *very often do* explain why someone does *not* die of cancer. Many people do survive cancer thanks to early detection and suitably administered chemotherapy and radiotherapy. Part of what makes these facts *against-dying facts* is precisely that similar facts do prevent people dying, and, indeed, these very facts could have prevented the man from dying. They are potential explainers.

Third, the against-dying facts play an epistemic role—they lead us, when considered alone, to increase our credence that the man will live. They constitute evidence that he will. In providing such evidence, the against-dying facts create a backdrop that helps us understand the explanation of the fact that the man dies. We understand that the threat posed to him by the cancer, given his excellent medical treatment, and such like, was (for example) particularly severe, or of a certain rare type.

Fourth, and relatedly, the against-dying facts create the *need* for an explanation of the man's death. Given that such facts *can* explain the man's continuing to live, and that they constitute evidence that he will do so, this provides a context in which the for-dying facts can explain why he died. In certain contexts, some facts do not cry out for explanation. These might be facts that we would expect to obtain in the normal course of events (this is not to say, of course, that such facts *lack* explanations, but simply that we do not go out of our way to look for them). We *feel the need* to explain a fact when we are provided with information that makes that fact surprising or expected to some extent. The against-dying facts in our example help provide such a context.

To generalize then: in our examples of explanations, the against-explanandum facts play four roles in such an explanation: (1) they provide background causal facts; (2) they *can* explain (cause, bring about, etc.) the negation of the explanandum (and similar facts often *do* explain their explananda); (3) they provide evidence for the negation of the explanandum; (4) they create a need for explanation of the explanandum. Conspicuous by its absence is the idea that the against-explanandum facts have weights that combine together to compete with the combined weights of the for-explanandum facts.

If our explanations are not weighing explanations, then neither are Broome's

In the last two sections, we defended the following claims: (1) Broome's examples of (supposed) normative weighing explanations are similar in many ways to our own examples of non-normative explanations; (2) our explanations are not weighing explanations. It does not *follow* from this that Broome's explanations are not weighing explanations, but it is at least rather suggestive of it. Indeed, we think this suggestion withstands scrutiny.

First, the for-*F* and against-*F* facts in Broome's normative explanations play broadly the same four roles as the for-explanandum and against-explanandum roles in our explanations. Such roles do not call for the idea of weights that such facts have and that combine or outweigh the weights of other facts. Given this, unless there is some further reason to attribute such weights to Broome's for-*F* and against-*F* facts, we already have an adequate (if rough) account of how such facts figure in the type of normative explanations that Broome highlights. Appealing to the weights of such facts is superfluous.

Take Broome's first example (the same points apply to his second example too). That you promised to take the wine, and that your doing so is not terribly onerous, rather straightforwardly explain why you ought to do so. The role of these for-taking-the-wine facts is to explain why (make it the case that; ground the fact that) you ought to take the wine. The against-taking-wine-facts that Broome considers are the (modest) cost of the wine and the (slight) inconvenience buying it causes you. These facts play an explanatory background role—they, together with the for-taking-the-wine facts and many other facts (such as the continuing existence of off-licenses and your own birth), are part of a scenario which suffices for the fact that you ought to take the wine. The against-taking-the-wine facts too can explain, in different scenarios, why it is *not* the case that you ought to take the wine. They also serve as evidence that it is permissible not to take the wine (and even, perhaps, evidence that you ought not to take the wine), and they create a context in which an explanation of why you should indeed take the wine is called for. In short we can explain the role of these facts in Broome's explanations without invoking the idea that they have weights that combine or compete with the weights of other facts.[17]

Second, just as there is no clear way of combining the various quantities involved in the facts in our non-normative explanations into one quantity analogous to a weight, there is no clear way of translating the various quantities involved in the facts in Broome's normative explanations into one quantity analogous to a weight. Consider Broome's second example. The pleasantness of Montreux and its proximity do not plausibly combine into one quantity—a weight—that can outweigh the weight

[17] By this we mean that it is possible to give a plausible account of Broome's examples of normative explanations, according to which the facts that Broome says explain why you ought to bring the bottle, or go to Montreux, do their explanatory work without possessing weights. Of course, Broome's *own* account of how these facts explain normative facts does appeal to weights.

that is (somehow) determined by the exoticness of Marrakesh and thus make it the case that you ought to holiday in Montreux. Clearly such quantities do play an important role in explaining what you ought to do, but they do not do so via a generic one-size-fits all quantity.[18]

Might Broome here appeal to *value* as being such a quantity? It is not clear he can. First, Broome is open to the possibility that value should be understood in terms of reasons, having defended this option in print previously, and thus might not wish to appeal to value to help analyze reasons (57–8). Second, given the possibility of there being incommensurable values (something Broome is also explicitly open to), there is no singular quantity of value that can be straightforwardly measured. Third, and most importantly, value is a quantity assigned to actions and consequences of actions, not to the facts that explain why these actions ought to be performed. Given this, value cannot be the quantity called "weight" that plays the roles Broome assigns it.

Where to go from here?

We have argued that Broome's normative explanations are not weighing explanations, as they are more akin to our non-normative explanations, which are not weighing explanations. What kind of responses are open to Broome? We think there are at least four.

First, Broome may concede that his examples are not weighing explanations, but this is merely because he chose the wrong examples. Broome gave himself this wriggle-room by saying that he does "not wish to assert that any particular weighing explanations are genuine" (55). He may insist that there are indeed normative weighing explanations and that reasons can be analyzed in terms of them. The problem with this approach is that Broome's examples are paradigm cases of normative explanations (this is, no doubt, why Broome chose them). If even these types of explanations do not count as normative weighing explanations then it is doubtful that any do.

Second, Broome may say that while his cases are indeed analogous to ours, our cases are indeed (non-normative) weighing explanations. The for-explanandum facts in our cases really do possess weights that combine and are then set against the combined weight of the against-explanandum facts. But, again, aside from this idea being implausible on its face, it is superfluous given the roles we have pointed out for the for- and against-explanandum facts. What would we be adding to our understanding of such explanations by insisting upon a relatively obscure quantity possessed by these facts that is highly analogous to weight?

[18] It is true that we sometimes assign exact degrees to entities in formal contexts (such as expected value, or degrees of credence) without actually being in an epistemic position to measure such degrees, or perhaps even when we doubt there are such exact degrees (rather than something fuzzier that we are modeling). The issue here, though, is not implausible metaphysical exactness, nor our epistemic limitations. Rather, we doubt there is anything useful to measure at all (however inexact or inaccessible).

Third, Broome may simply deny that his cases and ours are analogous. While we have tried to show that they *are*, such issues are complex and relevant similarities and differences between the cases are not immediately obvious. Perhaps Broome's for- and against-*F* facts do not play the roles we attribute to our for- and against-explanandum facts. We describe our case that they do above. More plausibly, Broome might reply that appeal to weights, while superfluous to account for the roles of the facts that figure in *our* explanations, are not superfluous to account for all the roles of the facts that figure in his. While his for- and against-*F* facts may indeed play the roles we suggest, they also play roles that include their having weights.

We concede that Broome's fact may play more roles than the facts that figure in our non-normative explanations. For example, it might be that the against-going-to-Montreux facts (i.e. the for-going-to-Marrakesh facts) help explain why it is somewhat regretful if you do not go to Marrakesh, or why going to Marrakesh would be a generally fine plan, or that going to Montreux is not ideal, even if something you ought to do. What we doubt is that such roles require us to assign *weights* to the facts that figure in Broome's normative explanations. It strikes us that everything Broome wants these facts to do can be done without their having weights.

Fourth, and lastly, Broome may concede that his explanations, like ours, are not weighing explanations, but that we are still able to analyze reasons with respect to them. That is, on this proposal, a reason to *A* is a fact that plays the kinds of roles we have highlighted in this chapter. To go this route, however, would be to give up the idea that reasons have weights that can combine, compete, outweigh, and so on. This, we believe Broome, would concur, would be to give up too much.

We take none of these options. Rather, we somewhat tentatively conclude that there are no normative weighing explanations and thus that it is not possible to analyze reasons with reference to them. Like Broome, we think the idea that reasons have weights is crucial to understanding the nature of reasons. In previous work (Kearns and Star 2008, 2009, amongst others), we pointed out that evidence also has weights. We hold out hope that we can understand the fact that reasons have weights as resting on the fact that reasons are a kind of evidence (and that evidence has weights). For our latest thinking on this see Kearns (forthcoming) and Star (forthcoming). Part of our reason for thinking this is precisely our skepticism concerning the supposed weights to which Broome appeals. If there are no such weights, we all have to look elsewhere.

References

Broome, J. (2008) Reply to Southwood, Kearns and Star, and Cullity. *Ethics* 119: 96–108.

Broome, J. (2013) *Rationality through Reasoning*. Chichester, UK; Malden, MA: Wiley Blackwell.

Fletcher, G. (2013) A Millian objection to reasons as evidence. *Utilitas* 25: 417–20.

Kearns, S. (forthcoming) Bearing the weight of reasons. In Errol Lord and Barry Maguire (eds.) *Weighing Reasons*. Oxford University Press.

Kearns, S., and Star, D. (2008) Reasons: Explanations or evidence? *Ethics* 119: 31–56.
Kearns, S., and Star, D. (2009) Reasons as evidence. *Oxford Studies in Metaethics* 4: 215–42.
Kearns, S., and Star, D. (2013) Weighing reasons. *Journal of Moral Philosophy* 10: 70–86.
Scanlon, T. M. (2014) *Being Realistic about Reasons*. Oxford: Oxford University Press.
Schroeder, M. (2007) *Slaves of the Passions*. Oxford; New York: Oxford University Press.
Setiya, K. (2014) What is a reason to act? *Philosophical Studies* 167: 221–35.
Star, D. (2011) Two levels of moral thinking. *Oxford Studies in Normative Ethics* 1: 75–96.
Star, D. (forthcoming) *Knowing Better*. Oxford University Press.
Thomson, J. J. (2008) *Normativity*. Chicago, IL: Open Court.