

Bearing the Weight of Reasons

Stephen Kearns, Florida State University

0. Reasons, the magical animal?

Lisa: I'm going to become a vegetarian.

Homer: Does that mean you're not going to eat any pork?

Lisa: Yes.

Homer: Bacon? Lisa: Yes Dad. Homer: Ham?

Lisa: Dad, all those meats come from the same animal.

Homer: Right Lisa, some wonderful, magical animal!

(The Simpsons, 1995, "Lisa the Vegetarian")

My attitude to reasons is similar to that of Homer's to this magical animal. Just as he finds it hard to believe that all these different types of meat come from the same creature, I find it hard to believe that reasons can do all that has been claimed of them. Reasons *count in favour of* actions, *justify* actions, *motivate* actions. Reasons make actions *right* (Schroeter and Schroeter 2009), or *explain* why we ought to do them (Broome 2013). Some claim that we may analyze *value* in terms of reasons (Scanlon 1998). Others claim that what we *ought* to do can be understood as what we have most (or perhaps conclusive) reason to do (Evers 2010). Perhaps *evidence* should be understood as reason to believe.

I suspect that no one type of thing plays all of these roles. More carefully: one type of thing might play all these roles (facts, or true propositions, for example), but do not do so in

virtue of *being reasons*. A fact counts as a reason because it bears the reasons-relation to an action (or belief, desire, etc.).¹ When it is said, then, that reasons play the above roles, the suggestion is that these roles either reduce to the reasons-relation (or to something that includes the reasons-relation), or that reasons play these roles *in virtue* of bearing the reasons-relation to an action (or to something else). For example, a simple analysis of evidence is as follows: A fact F's being evidence that P is F's being a reason to believe that P. In this case, evidence is *analyzed* in terms of the reasons-relation. Another example—one might hold that a fact, F, justifies S in Aing *because* F is a sufficient reason for S to A. In this case, the justifies-relation is said to obtain *in virtue of* the (sufficient) reasons-relation (rather than reduce to it).

My suspicion is that, whether we want to reduce such relations *to* or say they obtain in virtue *of* the reasons-relation, we will fail to find one such reasons-relation (or unified class of reasons-relations) to do the job. At best what we will find are various disparate relations, some of which play some of the roles assigned to reasons, and others of which play other roles. No one relation plays all these roles. If we are feeling strict, we might claim that this means there are no such things as reasons. If we are less strict, we have at least two options. First, we might become pluralist about reasons—they are multiple relations all of which deserve the term “reasons-relation”. Second, we might think that one relation (or one class of relations) deserves the term more than others.

The latter is my tack in this paper. In particular, I claim that one class of relations—certain evidential relations—realize most of the most important roles for reasons. In fact, my aim is a bit narrower than this. As is befitting the theme of this volume, I shall claim that

¹ This is not to suggest the reasons-relation is merely two-place. It likely involves also an agent, a time, and perhaps other entities.

understanding reasons-relations as evidential best captures the roles we give to the weights (or strengths) of reasons.

1. Roles of the weights of reasons

Reasons-talk brings with it an ideology shot through with the idea that reasons have weights. We talk of one reason's being stronger than other. We talk of sufficient reasons, conclusive reasons, *pro tanto* reasons, having most reason, weighing reasons against each other, summing the weight of reasons, etc. Why have we developed such talk? What are the main claims that we are committed to in buying into this ideology? There is no single answer to these questions. In what follows, I will set out some of the roles we assign the weight of reasons (and related notions, such as conclusive reason).

We may divide these roles into two broad categories: internal and external. The internal roles for the weights reasons are those that relate reasons (and their weights) to other reasons (and their weights). The external roles for the weights of reasons are those that relate reasons (and their weights) to other phenomena, such as reasoning, value, oughts, and justification.

1.1 External roles of the weights of reasons

1.1.1 Oughts

How are reasons and oughts related? In particular, how do the weights of reasons (and related notions) interact with what we ought to do? Here are some natural suggestions:

(O1) Weighing reasons helps agents discover what they ought to do.

(O2) Necessarily, if an agent, S, has conclusive (or decisive) reason to A, S ought to A.

(O3) It is not the case that, necessarily, if an agent, S, has pro tanto reasons to A, S ought to A.

(O4) Necessarily, if an agent, S, has most reason to A, S ought to A.²

Let us take each of these in turn. When in the process of deliberating about what she ought to do, an agent weighs reasons. It is difficult to describe this process in a theory-neutral way (see 2.2.1 for my understanding), but the basic idea is that an agent considers various facts that are relevant to her situation, weighs them in some way (both by summing the weights of reasons for the same actions and comparing the weights of reasons for different actions) and concludes (often in such a way that produces knowledge) what she ought to do on the basis of such weighing. Any account of reasons should either show how weighing reasons produces such knowledge, or explain away the idea that reasons really play this role.

O2 introduces the idea that reasons can be *conclusive*. I sometimes get the impression that “conclusive reason” (and other phrases, such as “decisive reason” or even sometimes “sufficient reason” and “most reason”) is being used as a synonym for “ought”. When it is not being so used, O2 is making a substantive but intuitive claim. It seems we have a pretheoretical idea that there is a property reasons can have (conclusiveness) such that their having it is sufficient for ought. In 2.2.1 I argue that my understanding of reasons can cast light on the nature of conclusiveness.

O3 highlights a property of reasons that is related to conclusiveness. A pro tanto reason to A is one that, though it speaks in favour of Aing, need not do so decisively. In effect, O3 points

² I take “most reason” to mean: more reason to do than any other alternative.

out that reasons do not always suffice for oughts. Any account of reasons must leave room for this fact. Consider the simple idea that a fact, F, is a reason for S to A iff F explains why S ought to A. This idea, though perhaps initially plausible, fails to account for the existence of pro tanto reasons. According to this simple account, it is necessary that, if F is a reason for S to A, S ought to A. (This is because, necessarily, if F explains why P, then P.) O3 rules out such accounts of reasons.

O4 makes an explicit connection between the weight of reasons and ought. Unlike O2, which posits the existence of a property, conclusiveness, that relates reasons to ought, yet does not explicitly express how this property is itself related to the strength of reasons, O4 posits a straightforward link—what one ought to do is simply what one has most reason to do. One might think what one ought to do reduces to what one has most reason to do, or that what one ought to do is otherwise set by what one has most reason to do. Either way, many have found O4 a plausible claim.

1.1.2 Permissibility

One connection between the weight of reasons and permissibility mirrors O1:

(P1) Weighing reasons helps agents discover what they are permitted to do.

The rationale for it is also the same. Agents who weigh reasons discover not only what they ought to do, but what they are permitted to do.

As mentioned above, the idea of a conclusive reason is intimately connected with the idea of ought. If O2 is true, so is P2:

(P2) Necessarily, if an agent, S, has conclusive (or decisive) reason to A, S is permitted to A.

Having conclusive reason entails not only ought, but also permissibility.

There is another type of tie between the weight of reasons and permissibility—one that is particular to permissibility. This says:

(P3) Necessarily, if an agent, S, has sufficient reasons to A, S is permitted to A. It is not the case that necessarily, if an agent, S, has sufficient reasons to A, S ought to A.

I take it that the term “sufficient” as used in P3 has no special significance. The main idea that P3 expresses is that reasons have a connection with permissibility. There is some property reasons might have such that their having this property is sufficient for the relevant action’s being permissible (just as there is a property reasons might have (being conclusive) that is similarly related to ought).

Again, what exactly this relationship is up for debate. Just as with the phrase “conclusive reason” and “ought” I sometimes wonder if “sufficient reason” is simply being used as synonymous with “permitted”. When it is not being so used, P3 makes a substantive claim about the relationship between reasons and permissibility. An account of reasons is better off if it can either capture this relationship or explain it away.

1.1.3 Value

As well as their being obligatory or permissible, actions (beliefs, intentions, etc.) can also be valuable (both positively and negatively). Actions can be good and bad to various degrees.

The value of an action can come apart from its deontic status—a good action might not be one an

agent ought to perform (it might be not good enough, it might be supererogatory, etc.), and a bad action might be obligatory (there are necessary evils). One obvious type of external role the weight of reasons might have, then, relates to the *value* of actions. Here are two ways in which reasons arguably do this:

(V1) Weighing reasons helps agents discover what it is valuable to do, and the degree to which various actions are valuable.

(V2) The degree of value an action has is set by the weight of reasons both for and against performing that action.

Claim V1 parallels O1 (and P1). Just as weighing reasons for and against an action can help us discover whether we ought to perform it, it seems we also weigh reasons to help us determine the degree of value actions have.

V2 makes a bold claim—that the degree of value an action has is set by the weights of reasons both for and against it. I use the somewhat slippery word “set” here because there are some rather different ways of understanding just what V2 comes to. The degree of value an action has may simply *reduce to* the combined weights of the reasons for and against it. A less reductive option is to claim that the weights of reasons for and against an action *explain why* the action has a certain degree of value (the weight of the reasons *for* the action contributing the degree of positive value of the action, and the weight of reasons against contributing to the negative value).

To be fully explicit, the kind of value I have in mind is not expected value, or subjective value, or moral value, whatever exactly these notions come to. V2 makes a claim about what we

might term *objective* value (though not necessarily intrinsic value—an action might be objectively valuable in virtue of its valuable consequences). V2 posits a link between the weights of reasons to perform an action and how genuinely valuable that action is. To give one example: that John is in severe pain is a reason to give him painkillers. This fact, in virtue of being a reason, helps set the value of giving John painkillers. One might think, for instance, that this fact is a very strong reason to give him painkillers. Let's say there are no other (independent) reasons for or against this action. According to V2, it follows that there is high positive value in giving John painkillers.

1.1.4 Justification

Actions can also be justified or unjustified. Being justified comes apart from being obligatory, permissible, or valuable.³ Consider the following example. Sally's house is on fire, though she is completely unaware of this (indeed, we may suppose she has evidence to the contrary). There is conclusive reason for her to leave the house, and thus she ought to leave. She is certainly permitted to leave. However, Sally has promised Holly that she will stay in the house. In such a case, Sally is not justified in leaving the house, even though she is permitted to, and she ought to, and there is no value in staying in the house. Consider also the case of Jill. Jill has been informed by extremely reliable authorities that charity X is a very worthy cause (and she possesses no evidence to the contrary). In fact, however, X is a scam that has done an excellent job of hoodwinking even the experts. The money they receive funds terrorism. In such a case, Jill is justified in giving money to X, even though there is no value in her so doing and she ought not to do so. I take these two cases to show that (a) an action can be valuable,

³ In saying this, I am contrasting justification with *objective* ought/permissibility. In 2.3.1, I allow that justification goes together with subjective ought/permissibility.

something that ought to be done, permissible, and yet not justified, and (b) an action can have no value, is not something that ought to be done, and is impermissible, and yet justified.

Justification comes in degrees—actions are more or less justified. Furthermore, the degree of justification an action possesses is related to the reasons for (and against) performing it. Roughly, it seems that the more reason (to which Jill has access) Jill has to give money to X, the more justified she is in doing so. The more she hears about the types of good X allegedly does, from multiple independent sources, the more warranted she is in donating to them. If there are reasons she knows against donating to the charity, the less warranted she is in doing so.

According to this idea, reasons are *justifiers*. That is:

(J) The degree to which an agent is justified in A_{ing} is set by the weight of reasons both for and against A_{ing} (to which the agent has access).

Again, I use the word “set” to cover a whole range of possible reductive or non-reductive relations. The idea behind J is simple. An agent is justified in A_{ing} by the reasons she possesses for and against A_{ing} .

J, or something like it, is highly plausible. It also makes sense of various practices to which justification is related. We may be called on to publically justify our actions. We do so by citing the reasons we have to perform them. We may hold various reactive attitudes (of resentment, for example) if we believe an agent’s behaviour is unjustified. If we learn they had good reasons for such behaviour, we cease having such attitudes. In general, the stronger the reasons someone has to do something, the more we treat them as justified in their doing it (as reflected in the strength of our reactive attitudes, the severity of our punishment (or generosity of our reward) and the degree of our admonishment (or praise)). We can also be called on to justify a possible course of action α

yet to be undertaken. In doing this, we cite the reasons which favour it and argue that they outweigh those against. Justification is tied up, then, with being criticizable by others (and ourselves) and being subject of certain demands to defend past, present or future actions. We meet these criticisms and demands precisely by appealing to our reasons for action, and their weights.

Justification and reasons are also arguably linked through virtue. A virtuous agent is one whose actions are both justified and responsive to the reasons for these actions. A virtuous agent will act on those reasons that are sufficient or conclusive, or (in cases of ignorance), those that present the best available case for a certain action. The better the case for performing the action, the weightier the reasons and greater the justification (and, in the case of a virtuous agent, the stronger the motivation).

1.2 Internal roles for the weights of reasons

1.2.1 Aggregation

The basic idea that the strengths of reasons aggregate is simple. If F and G are both (distinct) reasons to A, then in some manner, the strengths of F and G can combine to produce stronger reason to A. Just as two distinct physical objects can combine their masses (a mass weighing 1g and another (distinct) mass weighing 2g have a combined mass of 3g), so can two distinct reasons combine their weight. For example, let's say I have two reasons for going to the grocers: I am out of potatoes and my friend who works there wants to talk to me. Each reason on its own might mount a modest case for going to the grocers. Together, however, they might mount a considerably stronger case for doing so. The strengths of both reasons combine.

Care must be taken with the idea. Sometimes reasons do not combine their strengths.

John is in pain. This is a reason for me to give him painkillers. John is also in severe pain. This is also a reason for me to give him painkillers. Still, these reasons do not combine their strengths so that they together make a stronger case for my giving John painkillers than each does independently. To fully account for the way in which reasons aggregate, then, we must have some understanding of why and when to rule in or out the summing up of reasons' weights. I briefly address this problem in 2.2.5.

1.2.2 Competition

Reasons compete in the sense that reasons for different actions weight against each other. The reasons for Aing might outweigh the reasons for Bing (or simply the reasons for not Aing), or they might themselves be outweighed. For instance, in the classic trolley case, it is plausible that one has a reason to refrain from switching tracks (because switching the tracks will lead to a person's death), but one has stronger reasons to switch tracks (because doing so will save the lives of five people). In this case, the latter reason *outweighs* the former reason.

Indeed, it does more than simply outweigh. That reasons compete also suggest that some reasons win, and some lose. These notions come apart from the idea of outweighing. Reasons that outweigh others might still not be winning reasons, perhaps because they are themselves outweighed, or because, though they provide a stronger case for a certain course of action, they fall short of being sufficient or conclusive. An account of reasons does better if it can spell out the notions of outweighing, winning and losing.

2. The weight of reasons and the strength of evidence

In this section I posit that a particular class of evidential relations between facts and actions

can fill many of these roles mentioned above.

2.1 Reasons as evidence

In a series of papers co-authored with Daniel Star, I make the following claim about reasons:

RA: Necessarily, F is a reason for an agent S to A iff F is evidence that S ought to A.⁴

One defense we made of this claim stemmed from the fact that our account of reasons helps us understand how reasons can have different strengths, or weights. As we put the argument (in Kearns and Star 2009):

- (1) Reasons can have different strengths.
- (2) RA is the best explanation of how this is possible.
- (3) Therefore, RA is true. (*inference to the best explanation*)

We took (1) to be relatively uncontroversial. Our main defense of (2) was that, first, it does well explaining how reasons aggregate and compete (see 2.2.5 below) and second, it makes the fact that reasons have weights much less mysterious than it otherwise would be. Our idea was that we have an independent grip on the notion that evidence has strength, and that, furthermore, this notion can be spelled out in relatively straightforward ways (one way we suggested was evidential probability—the weight of evidence equals the degree to which it increases the probability of a proposition).

⁴ See Kearns and Star 2008, 2009, 2012, 2013a, 2013b

While I still stand by the main elements of that argument, one important element that was missing was a defense of the claim that the strength of evidence plays the kinds of roles that the weight of reasons plays. Why think that the weight of a reason and the weight of evidence are so closely related? This paper attempts to answer this question.

Given my skepticism that there is one unique reasons-relation that fills all the roles touted for reasons, I now think that RA might turn out to be too restricting. Once we have given up on their being a single all-encompassing reasons-relation, we are somewhat free to suggest various relations that play the reasons-roles. Being evidence that one ought is one of a united class of relations between a fact and an action. The members of this class are related roughly by being evidential relations between a fact and the deontic or evaluative status of an action (or belief, intention, etc). Consider the following such relations:

E1: F is evidence that S ought to A.

E2: F is evidence that it is valuable (or good) that S As.

E3: F is evidence that it is permissible that S As (or that it is not the case that S ought not to A).

E4: F is evidence that that S ought not to A (or that it is impermissible that S As). E5: F is evidence that it is bad that S As.

We might understand reasons *for* an action in terms of E1-3—reasons for Aing are evidence that one ought to A, that one is permitted to A, or that it is good (or best) to A.⁵ Reasons against Aing can be understood in terms of E4-5—i.e. evidence that one ought not to A, or that it is bad to A.

The above relations concern only evidence, not the *weight* or *type* of evidence. Given that evidence comes in different strengths, can be conclusive, sufficient, outweighed, etc., there are whole host of other relations evidence might bear to deontic claims, including:

E6: F is conclusive evidence that S ought to A.

E7: F is conclusive evidence that it is permissible that S As.

E8: F is stronger evidence that S ought to A (or that S is permitted to A) than G is that S ought not to A.

In what follows, I suggest that such relations can play (most of) the roles we have given to reasons and their weights.

2.2 Roles these evidential relations play well

2.2.1 Ought and the weight of evidence

To start with one of the most promising, consider again O1:

(O1) Weighing reasons helps agents discover what they ought to do.

Understanding the reasons-relation (or reasons-relations) as evidential casts considerable light on how O1 can turn out true. What better way is there of discovering what one ought to do than by

⁵ As Gert (this volume) suggests, evidence for oughts and evidence for permissibility may correspond nicely with his distinction between reason's having requiring and justifying strength.

considering the *evidence* concerning what one ought to do? In fact, this strikes me as the *only* possible way one can do this.

An evidential conception of reasons provides us with the following sketch of weighing reasons. An agent considers facts, some of which constitute evidence that she ought to A, others evidence that it is permissible that she As (or, perhaps equivalently, that it is not the case that she ought not to A), and others still evidence that she ought not to A. She has similar evidence with regards to the other options open to her—Bing, Cing, etc. (indeed, some facts which are evidence that she ought (not) to A (or that she is permitted to A) might also be evidence that she ought (not) to B, C etc. (or that she is permitted to B, C, etc.)). Each of these facts is evidence of a certain strength concerning what she ought to do (and what it is permissible for her to do). On many occasions the combined strength of such evidence suffices to produce knowledge in the agent of what she ought to do. She may, for instance, be in possession of various facts that together conclusively show that she ought to A.

The process of weighing reasons, according to this proposal, is the process of attempting to determine what one ought to do (or what it is permissible, or what is best, or what one ought not to do, etc.) by considering the evidence concerning what one ought to do, and coming to a relevant conclusion based on the strength of evidence for and against such deontic and evaluative claims. The plausibility of the evidential conception of reasons increases (or decreases) relative to how accurate this picture of weighing reasons is. I think it captures it well.

O2 says:

(O2) Necessarily, if an agent, S, has conclusive (or decisive) reason to A, S ought to A.

There is a very natural understanding of O2 on an evidential conception of reasons. We may first understand a conclusive reason as follows:

(CR) F is a conclusive reason for S to A iff F is conclusive evidence that S ought to A.

We can further understand conclusive evidence that P as evidence that is sufficient to produce knowledge that P. One can come to know P on the basis of this evidence. What's more, knowledge is factive (necessarily, if one knows that P, then P). On this (highly natural) understanding of conclusive evidence it follows that, necessarily, if F is conclusive evidence that P, then P. This result, combined with CR entails O2.

In other words, my evidential conception of reasons identifies the property “conclusiveness” that reasons to A may have with an independently significant property of evidence—that of being sufficient to produce knowledge that one ought. This identification does at least three nice things. First, it helps explain why O2 is true. Being conclusive evidence that P is factive, and thus there being conclusive reason that one ought to A does indeed necessitate that one ought to A. We need not appeal to a somewhat mysterious kind of property of reasons that suffices for ought. Second, it also explains why we might be *interested* in the notion of a conclusive reason—the idea of facts that are sufficient to produce knowledge of what we ought to do is obviously an important one. In seeking out such facts, we may come to know our obligations.

Third, it provides at least the start of an account of how the weights of reasons are related to their conclusiveness. Evidence, just like reasons, comes in various strengths. Facts are (together) conclusive when they are strong enough to produce knowledge. On one popular way of understanding this idea, propositions are related by evidential probabilities. Some evidence, F, for P is conclusive just in case the evidential probability of P given F is 1. The weight of evidence

(and thus reasons) for P can be understood as the evidential probability of P given the evidence, and the conclusiveness of evidence (and thus reasons) for P can be understood as the evidential probability of P given the evidence's equaling 1.⁶

Recall O3:

(O3) It is not the case that, necessarily, if an agent, S, has pro tanto reasons to A, S ought to A.

The main thrust of O3 is that something may be a reason to A without it's necessitating that one ought to A. The evidential conception of reasons captures this idea very elegantly. Just as not all reasons are conclusive, not all evidence is conclusive. A fact can be evidence for the proposition that one ought to A, or that it is permissible to A, without being sufficient to produce knowledge that one ought, or that it is permissible. Non-conclusive evidence that P *makes a merely partial case* that P. Evidence concerning what one ought to do is thus well placed to capture the nature of pro tanto reasons.

2.2.2 Permissibility and the weight of evidence

P1 says:

(P1) Weighing reasons helps agents discover what they are permitted to do.

P1 can be accounted for by an evidential conception of reasons in the same way as O1 can (see

⁶ One worry one might have with this proposal is that if conclusive evidence of a proposition requires that this proposition has an evidential probability of 1 (given this evidence), and evidence is knowledge-producing only if conclusive, then there can be no knowledge-producing *inductive* evidence. This worry assumes that a proposition cannot have an evidential probability of 1 given purely inductive evidence. I am happy to reject this assumption—how weighty some evidence is may in part depend on whether what it is evidence *for* is true or not. If one is not happy to accept this, we can instead claim that evidence may be conclusive (and thus knowledge-producing) if it raises a proposition's evidential probability above a certain threshold.

above).

P2 is explained by my evidential conception of reasons in basically the same way as O2 is. Recall, P2 says:

(P2) Necessarily, if an agent, S, has conclusive (or decisive) reason to A, S is permitted to A.

A conclusive reason to A is conclusive evidence that one ought to A (on my account). That F is conclusive evidence that one ought to A entails that one ought to A. That one ought to A entails that one is permitted to A. Thus that F is a conclusive reason to A entails that one is permitted to A.

How to understand P3 proves slightly less obvious. It says:

(P3) Necessarily, if an agent, S, has sufficient reasons to A, S is permitted to A. It is not the case that necessarily, if an agent, S, has sufficient reasons to A, S ought to A.

I suggest that we understand the notion of a sufficient reason to A as follows:

(SR) F is a sufficient reason for S to A iff F is conclusive evidence that S is permitted to A.

Whereas a conclusive reason to A is conclusive evidence that one ought to A, a sufficient reason is conclusive evidence that one is permitted to A. This is the property, on my evidential conception of reasons, that links sufficient reasons to permissibility. SR is attractive for the same reasons that CR is attractive—it explains the truth of P3, it explains our interest in sufficient reasons, and it elucidates the link between a reason's being sufficient and its weight (for how exactly it does these things, see the discussion of CR above). Furthermore, SR captures the nature of sufficient reasons

without having to posit a new property of reasons that is completely distinct from conclusiveness—sufficient reasons are conclusive of permissibility, while conclusive reasons are conclusive of obligation.

2.2.3 *Value and evidence*

V1 states:

(V1) Weighing reasons helps agents discover what it is valuable to do, and the degree to which various actions are valuable.

The understanding of V1 recommended by my evidential conception of reasons is identical, *mutatis mutandis*, to those of O1 and P1 (see above).

2.2.4 *Justification*

(J) says:

(J) The degree to which an agent is justified in A_{ing} is set by the weight of reasons both for and against A_{ing} (to which the agent has access).

On the evidential conception of reasons, this translates to:

(JE) The degree to which an agent is justified in A_{ing} is set by the weight of (available) evidence that she ought or is permitted to A , and the weight of (available) evidence that she ought not to A .

In effect, how justified someone is in doing something depends on how strong their evidence is that they are permitted to do it (this overall strength itself being calculated from the strength of

evidence for its permissibility and the strength of the evidence for its impermissibility). A corollary of JE suggest itself:

(JCE) Necessarily, if an agent has conclusive (available) evidence that she ought or is permitted to A, she is fully justified in Aing.

If true, JCE helps explain the truth of the following:

(JCSR) Necessarily, if an agent has conclusive or sufficient (available) reason to A, she is fully justified in Aing.

But are J and JCR true?

Why yes! Evidence justifies (Williamson 2002). True, it is most renowned for justifying *beliefs*, but it also justifies action. Thomson 2008 contends that a fact justifies Aing to the extent that it justifies believing that one ought to A (see also Kearns and Star, 2008, 2009). I'm (now) tempted to say that a fact justifies Aing to the extent that it justifies believing either that one ought, *or that one is permitted*, to A. In essence, what one is justified in doing is what one is justified in believing one is permitted/ought to do. As evidence justifies the belief, it also justifies the action.

Examples seem to corroborate this. Those facts that justify someone's doing something are those facts that serve as evidence for its permissibility. For instance, that John is in severe pain (a fact I know) is both evidence that I should give him my painkillers and serves to justify this action. Furthermore, it justifies this action *in virtue* of being evidence that I ought to perform it. Had it been no evidence at all that I ought to do so (or that I am permitted to do so), then I would not be criticizable for failing to perform the action.

The examples in 1.1.4 were used to show that justification can come apart from obligation, permissibility and value. Sally is unaware that her flat is on fire and so not justified in leaving. Jill has excellent but misleading evidence that X is a worthwhile charity and is thus justified in giving to it, despite the facts that she shouldn't and that there is no value in so doing. These examples also nicely bring out the fact that justification is linked to evidence. Why does Sally lack justification for leaving the flat, given that she ought? Because (1) she *lacks evidence* that she ought to, or is permitted to, and (2) an agent lacks justification for doing something that they ought to do, or are permitted to do, precisely when they lack evidence that it is something they ought to do, or are permitted to do. Why is Jill justified in giving to X, given that she shouldn't? Because (1) she *has* excellent evidence that she ought to and (2) an agent has justification for doing something that shouldn't do precisely when they have good (on balance) evidence that they should do it. Jill can easily defend her decision to give to X by citing her available evidence; her actions are entirely consistent with her being virtuous; indeed, her actions might be considered morally worthy.⁸

All of this is just to say that the role of reasons as justifiers highlighted in 1.1.4 is played by evidence concerning what one ought and is permitted to do. The practices concerning justification, such as criticizing actions, demanding justifications of actions, having certain reactive attitudes, rewarding or punishing, etc. are most sensibly based on the evidence available to agents when they act.

Hume says “a wise man...proportions his beliefs to the evidence” (1999, 10.4). I suggest

⁸ Does this mean *wrong* actions may be morally worthy? Well, this rather depends on what we mean by “wrong”. Actions that one ought *objectively* not to perform might still be the *morally right* thing to do, on the assumption that what is morally right is what one *subjectively* ought to do. (The labeling here is not ideal: subjective oughts are not merely what agents think they (objectively) ought to do, but are (roughly) what their evidence suggests they ought to do.) On this assumption, actions that one objectively ought not to do are not necessarily morally wrong, and thus may even be contenders for being morally worthy.

that a virtuous person proportions her *actions*, and motivations for those actions, to the evidence. One is justified in Aing only if one's evidence that one ought or is permitted to A is sufficiently strong. If one's evidence is *conclusive*, then one is fully justified in Aing. Below (in 2.3.1), I suggest that if one has most evidence that one ought to or is permitted to A then one is (outright) justified in Aing. A virtuous person, given their epistemic position, will do what they can *know* is right, and, failing that, what they have most evidence is right.

2.2.5 Aggregation and Competition

Just as reasons aggregate and compete, so does evidence. Two pieces of evidence for a proposition can combine strengths to provide a stronger case for that proposition than either does independently. That I am out of potatoes and that my friend who works there wants to talk to me are both pieces of evidence that I ought to go to the grocers. Each independently might mount a modest case for the claim that I ought to go to the grocers. Together, however, they might mount a considerably stronger case that I ought.

Conflating the strength of reasons and of evidence also explains those cases in which the strength of two reasons does not combine. That John is in pain and that John is in severe pain are both evidence that I ought to give him painkillers. However, they do not combine their strengths to produce a collectively stronger case that I ought so to do. Why is this? Because these pieces of evidence are not *independent* (in this case, the conditional probability that John is in pain, given that John is in severe pain, is 1—the latter proposition increases the probability of the former). It is just when evidence is independent that it aggregates (as it does in the grocer case).

Evidence, just as with reasons, competes. F may be evidence that P, while G is evidence that \sim P. F might outweigh G in the sense that F is stronger evidence that P than G is that \sim P. (We

may again think in terms of evidential probability—the evidential probability of P given F is higher than the probability of \sim P given G.) The evidence that P outweighs the evidence that \sim P. Evidence also wins or loses. The evidence that P might be conclusive (sufficient to produce knowledge). Conclusive evidence that P wins out against evidence to the contrary. Evidence outweighed by conclusive evidence can be said to lose. Given this, I suggest we think of winning reasons to A as conclusive evidence either that one ought to A, or that one is permitted to A, or that it is good to A. Losing reasons to A are evidence that one ought to A, or that it is permitted to A, or that it is good to A that are outweighed by conclusive evidence either that one ought not to A, or that one ought to B (where B is incompatible with A).

2.3 Two roles these evidential relations play not so well

The evidential conception of reasons I advocate doesn't play the role of reasons perfectly. In this section, I shall look at two claims about reasons that such a conception fails to capture.

2.3.1 Most reason and ought

O4 states:

(O4) Necessarily, if an agent, S, has most reason to A, S ought to A.

If by “most reason” we simply mean *conclusive reason*, then O4 causes my account no difficulties.

If, on the other hand, we interpret “most reason to A” as “more reason to A than any alternative”, and by *this* we also do not also simply mean “conclusive reason”, but rather intend to make an explicit claim about the relative weights of reasons, then O4 does not seem to fit well with the evidential conception of reasons. This is because the most straightforward understanding of O4 given how I think of reasons is as follows:

(ME1) Necessarily, if an agent, S, has most evidence that she ought to A (or is permitted to A), S ought to A.

The problem with MO is that it is, most probably, false. This is because it seems perfectly possible to have most evidence that one ought to A, while this evidence is misleading. A detective might have the most evidence that Carl is the culprit, but really, Carl is innocent. Similarly, then, she might have the most evidence that she ought to arrest him, but it is not the case that she ought. Plausibly, then, I must reject MR.

I can, however, accept something like it. I have been using “ought” in its objective sense. This is the sense in which what one ought to do can come apart from what one is justified in doing. There is also a subjective sense according to which one ought to do only what one is justified in doing (the same distinction can be made regarding permissibility). This suggests the following variation on O4:

(O4j) Necessarily, if an agent, S, has most (available) reason to A, S is justified in Aing (S subjectively ought/is permitted to A).

Based on my evidential conception of reasons, this further suggests a variation of ME1:

(ME2) Necessarily, if an agent, S, has most (available) evidence that she (objectively) ought/is permitted to A, S is justified in Aing (S subjectively ought/is permitted to A).

Two points: first, the plausibility of O4 may rest in its similarity to (or its interpretation *as*) O4j. That one is justified in doing something is naturally put as one’s being permitted or obligated to do it. O4j claims that having most reason suffices for ought/permissibility *in this sense*. Once we see this, it is no longer clear that we should be committed to O4 (when interpreted as concerning

objective ought).

Second, given that my account spells out O4j as ME2, we can understand why O4j is (quite possibly) true. As I argued in 2.2.4, evidence is what justifies. Either one's evidence is conclusive, in which cases one's action is fully justified, or it's not, in which case the best we can be expected to do (barring sheer luck) is to do what we have the most evidence that we are permitted/ought to do. This is the option available to us that garners the least warranted criticism and that we can most reasonably defend (though there may be exceptions to this—consider, for example, Parfit's mine-shaft case (2011)). I propose, then, that, understood one way, O4 specifies a role for reasons that evidence does *not* play, but understood in another, O4j, it does.

2.3.2 Reasons, evidence and value

The other role an evidential conception of the weight of reasons does not play well is V2. V2 says:

(V2) The degree of value an action is set by the weight of reasons both for and against performing that action.

It is, I suspect, a commitment to (something like) V2 that gives rise to the most common objection to the Kearns and Star account of reasons. Indeed, this objection would apply to any evidential account of reasons. We (Kearns and Star) once illustrated the objection using the following example:

A newspaper says that there are people starving in Africa. This is evidence that one ought to give money to Oxfam. However, the fact that the newspaper says that people are starving in Africa is not a reason to send money to Oxfam...Rather, it is the fact (if it is a fact) that there are people starving in Africa that is a reason to send money to Oxfam.

It follows, if the above claim is right, that something can be evidence that one ought to do something without being a reason for them to do it. We replied by claiming that the fact that the newspaper says there are people starving in Africa (henceforth, “the newspaper fact”) is indeed a reason to send money to Oxfam.

Based on what I have said in this paper, my case for this claim should be relatively clear. The newspaper fact plays many of the significant roles that are attributed to reasons. It helps us discover what we ought to do (and what is valuable); it *justifies* action (someone in possession of the newspaper fact would be justified in sending money to Oxfam); it’s being conclusive (if it is) suffices for what one ought to do; it’s being sufficient (if it is) suffices for what it is permissible for one to do; it can combine with other facts to produce a stronger case for giving to charity, and compete with facts that favour other actions.

Still, one role the newspaper fact certainly does *not* play is that of making the action (of giving money to Oxfam) *valuable*, or of determining the degree of value it has. This action has value because, presumably, it would alleviate the suffering, and perhaps save the lives, of a certain number of people. Exactly how valuable the action is will depend on exactly how much suffering would be alleviated, how many people would be positively affected, and what could be done with the money instead. The newspaper fact is irrelevant to such facts and thus plays no role in the degree of value the action has.

My diagnosis of why people find the above type of (alleged) counterexample compelling, then, is that they are committed to V2, or something very much like it. So—I must reject V2 and

my objectors accept it.⁹ What now? Here's how I'm tempted to proceed. I have argued that an evidential conception of reasons plays many of the most significant roles of reasons. I have also expressed skepticism that there is one united class of reasons-relations that can fill all the roles we want of reasons. I propose, then, that V2 is a genuine desideratum of an account of reasons, but one we must give up. Evidential relations between facts and normative (and evaluative) properties of relations provide so much of what we want from reasons that they deserve to be thought of as the genuine reasons-relations. One role they don't play is that suggested by V2. Unless, however, we find some other types of relation that do play this role PLUS the other roles, then we'll either have to accept that the evidential relations I highlight count as reason-relations (because they are the best, but not perfect, fit for the roles we want reasons to play), or become eliminativists about reasons. This latter strikes me as an overreaction.

Of course, *some* things explain the (degree of) value of actions. These things, however, are *not* reasons (or rather, not reasons in virtue of playing this role), nor the weights of reasons. Indeed, those things that set the value of actions do not necessarily do so by having weights. Rather, we have a splendid pre-made type of relation between facts and the value of actions that can explain the value of actions. It is simply the *explanatory* relation. Facts can explain why actions have value, and explain why they have certain degrees of value. Such facts explain why actions have value not by themselves having a distinctive kind of weight or strength (a value-giving strength, as it were), just as the facts that explain why someone is a certain height don't have a distinctive height-giving strength. There is simply nothing whose *weight* determines the value of an action (for a defense of a related claim (that reasons are not parts of weighing explanations of

⁹ In rejecting V2, I take it that I also must reject buck-passing accounts of value. Very well then—I hereby reject such accounts! Further, other attempted counterexamples may require different treatment.

oughts), see Kearns and Star (forthcoming)).

It is not terribly surprising that reasons do not play all the roles suggested above. The idea that what *justifies* action is also what provides *objective value* to action strikes me as confused. The newspaper fact does indeed provide justification (whether sufficient or not) for one's donating to Oxfam. It does not make the action valuable. There are thus facts that are *justifiers* and not *value-imbuers*.¹⁰ We are faced with a choice: are reasons the latter or the former? Two thoughts: first, the answer may not matter too greatly—whichever reasons are, the other still exists. Second—conceiving of reasons as justifiers better fits our other desiderata of an account of reasons, and particularly their weights. This paper is a step in the direction of defending this claim.

3. Conclusion

Perhaps there is a special, unanalyzable relation that can play all the roles I have mentioned in this paper. After all, Homer was wrong—pork, ham, and bacon all come from the same magical animal. Still, until such a time as we come across our pig (perhaps it is flying just out of view), we'll have to make do with what we've got. And though it's not perfect, in my estimation, the strength of evidence can bear the weight of reasons.¹¹

References

Broome, J. (2013) *Rationality Through Reasoning* (Chichester: Wiley-Blackwell)

¹⁰ Similarly, there are value-imbuers that are not justifiers. Consider the unknown fact that Sally's flat is on fire. This imbues leaving the flat with value, but does not justify her doing so, as she is unaware of it. There are even plausibly value-imbuers that *cannot* be justifiers. Some incredibly complex fact may imbue an action with value but be too intricate for us to grasp, and thus unavailable to justify our actions.

¹¹ Thanks to Errol Lord, Barry Maguire and a anonymous referee for very helpful comments on this paper. Thanks also to Daniel Star, Julia Markovits, Josh Gert and the audience at the Weighing Reasons workshop at Princeton.

Evers, D. (2010) "The End-Relational Theory of 'Ought' and the Weight of Reasons" *Dialectica* 64:405-417

Gert, J. (2014) "The Distinction between Justifying and Requiring: Nothing to Fear" *Weighing Reasons*, ed. Errol Lord and Barry Maguire, Oxford University Press.

Hume, D. (1999) *An Enquiry concerning Human Understanding*, ed. Tom L. Beauchamp (Oxford: Oxford University Press)

Kearns, S. and Star, D. (2008) "Reasons: Explanations or Evidence?" *Ethics* 119: 31-56

(2009) "Reason as Evidence." *Oxford Studies in Metaethics* 4: 215-42

(2011) "On Good Advice: A Reply to McNaughton and Rawling." *Analysis* 71: 506-8

(2013a) "Reasons, Facts-About-Evidence, and Indirect Evidence." *Analytic Philosophy* 54: 237-43

(2013b) "Weighing Reasons." *Journal of Moral Philosophy* 10:70-86

(forthcoming) "Weighing Explanations" *Weighing and Reasoning: A Festschrift for John Broome*, Iwao Hirose & Andrew Reisner (eds.), Oxford University Press

Parfit, D. (2011) *On What Matters Volume 1* Oxford: Oxford University Press

Scanlon, T. M. (1998) *What We Owe to Each Other* Cambridge, MA: Harvard University Press

Schroeter, L. and Schroeter, F. (2010) "Reasons as Right-Makers" *Philosophical Explorations* 279-296

Thomson, J.J. (2008) *Normativity* Peru, Illinois: Open Court Publishing

Williamson, T. (2002) *Knowledge and its Limits* (Oxford: Oxford University Press).