

# The Composition of Reasons

CAMPBELL BROWN

31<sup>st</sup> January, 2013

How do reasons combine? How is it that several reasons taken together can have a combined weight which exceeds the weight of any one alone? I propose an answer in mereological terms: reasons combine by composing a further, complex reason of which they are parts. Their combined weight is the weight of their combination. I develop a mereological framework, and use this to investigate some structural views about reasons. Two of these views I call “Atomism” and “Wholism”. Atomism is the view that atomic reasons are fundamental: all reasons reduce to atomic reasons. Wholism is the view that whole reasons are fundamental. I argue for Wholism, and against Atomism. I also consider whether reasons might be “context-sensitive”.

## Introduction

Imagine you have a decision to make, say, whether to drink whisky, or water. Each option is supported, let’s suppose, by exactly one reason. If you choose whisky, you’ll be happy; if water, healthy. Plausibly, what you ought to do depends on how strong each reason is, compared with the other. You ought to do whatever is supported by the *strongest* reason. Say this is drinking whisky (happiness outweighs health).

But suppose now there’s a second reason for drinking water: it’s cheap. This new reason, like the old, is less strong than the reason for drinking whisky (happiness outweighs money, too). So, as before, drinking whisky is favoured by the strongest reason (as we’ve stipulated, the two reasons on the other side are both weaker), yet this might no longer be what you ought to do. Possibly, in this new case, you ought *not* to do what is supported by the strongest reason. For although the strength of each water reason, considered by itself, is less than that of the solitary whisky reason, their combined strength might nonetheless be greater. By combining together, they might defeat an opponent which neither could defeat alone.

How does this happen? How do reasons combine together in this way? I propose we think of it in *mereological* terms. Reasons combine by composing other reasons, complex wholes of which they are parts. Roughly, the idea is as follows. Take two reasons with some bearing on a given option. Each reason, on the familiar picture, will have what might be called a “force”, consisting of both a “direction”, i.e., being either for or against this option, and a “magnitude”, i.e., being stronger or weaker, more or less weighty. Suppose, as may be the case, they also have a combined force, different from the force of either one on its own. They do this, on my proposal, by composing another reason. This third reason is the combination of the first two, and so its force is their combined force. In a slogan: the combination of the forces is the force of the combination.

Such combining may occur with greater numbers of reasons. When a reason is composed of all the reasons which have some bearing on a given alternative, we may call it an “overall reason”. What you ought to do, we may then say, is settled by overall reasons. You ought to choose whichever option, in a given situation, is favoured by the strongest overall reason (or whichever is disfavoured by the weakest overall reason, if none is favoured).

My goal here is to develop a mereological framework for reasons, and then to use this framework to investigate various structural views about reasons. The main focus will be on two views, which I call “Atomism” and “Wholism”.<sup>1</sup> With a mereological conception of reasons, we may think of reasons as occupying different levels in a hierarchy of complexity. At the bottom are *atomic* reasons, which have no other reasons as parts. At the top are *whole* reasons, which are parts of no other reasons. Atomism, as I use this term, is the view that atomic reasons are fundamental: all reasons reduce to – or better: supervene on – atomic reasons. Wholism is the view that whole reasons are fundamental, in the same way.

Though they are not primarily methodological views, it may be helpful to think of Atomism and Wholism as each suggesting a different approach to weighing reasons. Atomism suggests a “bottom-up” approach, wherein we begin at the bottom, assigning weights to the atomic reasons, and then work our way up through the hierarchy, assigning weights based on those assigned below. Wholism suggests going in the opposite, “top-down” direction, starting with the whole reasons and then working down.

Atomism and Wholism are not mutually exclusive. Some reasons are both atomic and whole. If these reasons, the atomic wholes, were fundamental, then both Atomism and Wholism would be true. However, as I argue below, combining the two is a risky policy, likely to have absurd consequences. Better, then, to accept only one. But which? I shall argue for accepting Wholism.

Section 1 sets out in more detail the rationale for the mereological approach, and defends it against an argument from Jonathan Dancy. Section 2 develops a mereological framework for thinking about reasons. Section 3 defines, using this framework, several structural views, including Atomism and Wholism and various forms of these two views. Section 4 presents an argument against Atomism. Finally, Section 5 considers whether reasons are “context-sensitive”, and rebuts two arguments for thinking that they are.

## 1 In Defence of Overall Reasons

Before discussing Atomism and Wholism, I need to respond to an objection made by Jonathan Dancy to overall reasons, since these play an important role in the mereological approach I am advocating. Dancy says that there are no such things and that to believe there are is delusional (Dancy 2004, p. 16). As I shall argue, however, there is a perfectly natural motivation for positing overall reasons, and doing so needn’t have the problematic consequences Dancy alleges.

### 1.1 Collective Plural Predication

In talk and thought about reasons, judgements of two sorts are common. We make judgements about *particular* reasons, identifying which options they favour or disfavour and estimating their strength in comparison with other particular reasons. But we also make *overall* or *final* judgements, summing up the whole situation, assessing the combined strength of all the reasons taken together, as a whole.

Judgements of the latter sort, overall judgements, typically involve what is sometimes called collective, or non-distributive, plural predication. A plural predication is a sentence of the form “The Fs are G”. Here are two examples:

- (1) The pencils are five inches long.
- (2) The pencils are scattered around the room.

---

<sup>1</sup>Why “Wholism”? Mainly to save confusion with another view which is often called “Holism”. More on Holism and its relation to Wholism below, Section 2.2.

These differ in an interesting way. In (1) a certain property, being five inches long, is ascribed to each pencil individually. Each pencil by itself is five inches long. By contrast, in (2) a certain property, being scattered around the room, is ascribed not to each pencil individually, but of all the pencils taken together, as a whole. No pencil by itself can be scattered; only a plurality or collection of pencils can have that property. Plural predication of the sort exemplified by (1) is often called *distributive*, whereas that exemplified by (2) is called *collective*.

Now consider a sentence expressing an overall judgement about reasons:

(3) Overall, the reasons in favour of buying the house are very strong.

This might be true even if each reason considered alone is quite weak. The reasons might be individually weak yet collectively strong. So (3) is a collective predication. (Indeed, the contribution of the modifier “Overall” in this context seems to be to signal a collective reading.)

Collective predications prompt an interesting question. How is it possible to ascribe a property to the Fs, e.g., the pencils or the reasons, not individually, but as a whole? A natural thought is that there must be some singular entity, which in some sense “stands in for” the Fs. Developing this thought in a mereological direction yields the view, sometimes called *mereological singularism*, according to which the singular entity is the fusion, or mereological sum, of the Fs. On this view, (2) says that the fusion of the pencils is scattered, and (3) that the fusion of the reasons is very strong.

We are thus led naturally to the view that overall judgements about reasons are judgements about fusions of reasons. And from there it is but a small step to the view that fusions of reasons may themselves be reasons. Some might resist this last step. They might accept that there are complex things composed of reasons, which favour or disfavour actions, yet be reluctant to call such things reasons. Other words might be used instead. We might, for example, call the complex thing a “case”. Overall, the reasons make a strong *case* for buying the house, we might say. I would have no objection to this alternative terminology if anyone were to insist on it.

Or you might prefer singularism without the mereology. For example, you might say that (3) attributes strength not to the fusion of the reasons, but rather to the *set* of the reasons (see, e.g., Schroeder 2007, pp. 123–45). I have no great objection to this, but it does seem to result in a less elegant framework. On the mereological approach, reasons are all of a piece: all are properties. But on the set-theoretic approach, reasons are of two kinds: some are properties, some are sets. You could perhaps say that all reasons are sets, and that when we speak of a property being a reason, what we really mean to call a reason is not the property, but the singleton set which has only this property as a member. Again, I have no great objection, but to me it seems more natural to say reasons are properties (or facts, if you prefer) rather than sets.

This is not meant to be a decisive argument for complex reasons. (For one thing, singularism is controversial; for criticisms, see McKay 2006.) What I do want to suggest, however, is that this approach yields a elegant, useful framework. Its utility, as I hope to show below, is that it enables us to clearly distinguish various views regarding the structure of reasons and to investigate the logical relations between these.

## 1.2 Dancy's Objection

The complexity of reasons is a matter of degree. Two simple reasons may combine to form a complex reason, which in turn combines with another reason to form a still more complex reason, and so on. Continuing in this way, ascending to ever higher levels of complexity, we may eventually arrive at a maximally complex reason, a reason that is composed of *all* the reasons bearing

on a given alternative.<sup>2</sup> These maximal reasons we shall call *overall* reasons; others we shall call *contributory* reasons.<sup>3</sup> As I said earlier, overall reasons enjoy a sort of priority. There may be a great many contributory reasons in favour of your doing something, but if the strongest *overall* reason is against it, then the overall reason carries the day; you shouldn't do it.

Jonathan Dancy argues that there are such things as overall reasons:

Much of our talk of reasons is about contributory reasons in this sense, reasons on one side or on the other, reasons that stack up with others to make a better or worse case for an action. But as well as talking about reasons in this way, we also speak of what there is overall reason to do. There is nothing wrong with that, of course, but it should not delude us into thinking there are such things as overall reasons in addition to the contributory ones. ... We can say we have more reason to do this than that, but most reason to do some third thing. These verdicts do not themselves specify further reasons (of an overall sort), on pain of changing the very situation on which they pass verdict. So there are no overall reasons. (Dancy 2004, p. 16)

This argument, however, does not defeat the mereological conception of overall reasons. Notice, first, that Dancy actually states two different positions in this passage. He concludes by saying that there are no overall reasons. But earlier he says something potentially weaker: that there are no overall reasons *in addition to* contributory reasons. There may be ways of interpreting this “in addition to” clause such that the weaker position is consistent with the mereological approach. It is sometimes said that, as D. M. Armstrong puts it, “mereological wholes are not ontologically additional to all their parts” (Armstrong 1997, p. 12). If so, then on my account, overall reasons do not exist in addition to contributory reasons. But that is not to say they do not exist at all.

Still, I do not want to rest my defence of the mereological approach on the “ontological innocence” of mereology. Such innocence may depend on controversial doctrines, such as “composition as identity”, which I am disinclined to accept, and, in any case, it seems wise not to give more hostages to fortune than necessary. What then of the stronger position? Does Dancy's argument support this?

It will be useful to flesh out the argument a little. As I understand it, the argument may be presented in the form of a *reductio ad absurdum*. We assume what Dancy denies, that verdicts do specify overall reasons, and show this has bad consequences. Suppose, to begin with, there are exactly two contributory reasons,  $c_1$  and  $c_2$ , bearing on some option. We weigh these up and reach a verdict regarding their overall, or combined force. This verdict, we are assuming, specifies a further, overall reason, which we'll call  $o_1$ . But  $o_1$  was, by hypothesis, not taken into account in our weighing up, and our verdict is therefore undermined. To be acceptable, a verdict must be based on *all* reasons, but ours is based solely on  $c_1$  and  $c_2$ , leaving out  $o_1$ . So we must weigh up again, this time including *three* reasons,  $c_1$ ,  $c_2$ , and  $o_1$ , and come to a new verdict. But this new verdict specifies another overall reason,  $o_2$ , on which it was not based. So we must weigh up yet again, this time including *four* reasons,  $c_1$ ,  $c_2$ ,  $o_1$  and  $o_2$ , and come to a new verdict. But ... This is getting us nowhere. We're mired in a vicious infinite regress, forever revising our verdicts, never reaching one we can accept. Every successive verdict we reach undermines itself: it specifies a reason on which it was not based, and so there is at least one reason it fails to take into account.

<sup>2</sup>This needn't be the case. It might be that all the reasons taken together have no combined force, because they cancel each other out, in which case their fusion will have no force and hence be no reason at all.

<sup>3</sup>This terminology comes from Dancy 2004, p. 16. One might alternatively say “all things considered” (for overall), and “*pro tanto*” (for contributory).

Thus our original assumption, that verdicts specify overall reasons, has the absurd consequence that no verdict can ever be acceptable.

On the mereological conception of overall reasons I propose, however, there need be no such regress. Two elements of my view are important here. First, the overall reason specified by a verdict is the mereological sum of the reasons on which the verdict is based. Let's write  $(x + y)$  to denote the mereological sum of  $x$  and  $y$ . Then, since our first verdict specified  $o_1$ , and was based on  $c_1$  and  $c_2$ , we have  $o_1 = (c_1 + c_2)$ . Likewise,  $o_2 = ((c_1 + c_2) + o_1)$ . Second, the combined force of some reasons is the force of their sum. Thus the combined force of  $c_1$  and  $c_2$ , which is what our first verdict was about, is the force of  $o_1$ . Likewise, the combined force of  $c_1$ ,  $c_2$ , and  $o_1$ , which is what our second verdict was about, is the force of  $o_2$ .

Together these two theses imply that our verdicts are about the same thing. Given the second thesis, our first verdict is about the force of  $o_1$ , and our second  $o_2$ . But given the first thesis,  $o_1$  is  $o_2$ . By definition,  $(x + x) = x$ ; the fusion of a thing with itself is the thing itself. So we have

$$o_2 = ((c_1 + c_2) + o_1) = (o_1 + o_1) = o_1.$$

There is only one overall reason,  $o_1 = o_2$ , and both verdicts are about the force of this one reason. There was therefore no need to revise our first verdict, because our second verdict was bound to be the same. The regress needn't have got started in the first place.

This might, however, suggest another objection. On my view, the combined force of  $c_1$  and  $c_2$  is the same as the combined force of  $c_1$ ,  $c_2$ , and  $o_1$ . Thus, in a sense, the overall reason  $o_1$  doesn't have any independent force: it doesn't contribute anything to the overall force of reasons which wasn't already there thanks to  $c_1$  and  $c_2$ . Some might object to this on the ground that every reason must have some independent force. This seems, however, rather like arguing that there cannot be such a physical object as, say, my computer, because every physical object must have its own independent weight, yet my computer could not have any weight over and above the weight of its parts. This is implausible, of course. I don't see how the view that all reasons must have independent force is any better.

One might worry that Dancy's overall position is somewhat inconsistent. To be a reason, he says, is to be a "favourer" (or a "disfavourer"), a thing which stands in the "favouring relation" to something (Dancy 2004, p. 29). Presumably, then, the point of his argument is to show that there can be no overall, or non-contributory, favourers. Notice, however, he is happy saying that "reasons stack up with others to make a better or worse case for an action" (Dancy 2004, p. 16). Now, what is this thing that he calls a "case"? It seems to stand in the favouring relation: it is *for* an action. And it is non-contributory: it is made of the stacked up contributory favourers, not a contributory favourer itself. It seems odd, therefore, that Dancy's hostility to overall reasons is not matched by hostility to cases.

## 2 A Mereological Framework for Reasons

I turn now to developing a more precise mereological framework for thinking about reasons. I begin by stating some assumptions about what reasons are, and about the mereological relations that hold between them.

### 2.1 Properties

I shall assume that reasons are properties. The peatiness of the whisky is a reason to drink it, the instability of the bridge a reason not to cross it, and so on. These properties I take to be "universals". If two bridges are each unstable, then there is a single property, instability, instantiated or

exemplified by both. We might then ask, if this property is a reason not to cross one bridge, must it also be a reason not to cross the other? I'll return to this question below.

I assume this mainly for convenience. An alternative view, apparently popular among philosophers, is that reasons are facts. One could, I think, develop an essentially equivalent framework using facts instead of properties, but it would be more cumbersome. Many philosophers prefer to speak as though reasons are not themselves properties, but are related to them in some other, not fully specified way: reasons are said to be variously “given by”, “provided by” or “constituted by” properties.<sup>4</sup> These ways of speaking would be fine for my purposes too, so far as I can tell, though a little less concise.

Which properties are reasons? I have in mind the following general picture. Suppose you face a choice between alternatives, objects which you may use for some purpose, e.g., houses you might buy to live in, or whiskies you might drink, or books you might read. Which of these alternatives you *ought* to choose depends on what they're like, what properties they have.<sup>5</sup> Among the properties of a given alternative, we may distinguish three sorts: those which count in favour of your choosing this alternative, those which count against, and those which count not at all, neither for nor against. In the example of a house you might buy, being spacious counts in favour, being in a dangerous neighbourhood counts against, and having an odd street number is neutral (at least for most people). The reasons are the properties in the first two categories, the favourers and disfavourers. Thus a reason is a property of an alternative which either favours or disfavors some person's choosing this alternative for some purpose.

Note, on this picture, every reason is tied to a particular person. The spaciousness of a house may be a reason for you to buy it without also being a reason for me to buy it, or for me to do what I can to get you to buy it. The issue between Atomism and Wholism is independent of the distinction between so-called “agent-neutral” and “agent-relative” reasons. Similarly, reasons are tied to particular actions. The same property of the house may be a reason for you to *buy* it, without being a reason to, say, set fire to it.

There is a natural sense in which one property may be said to be a part of another. For example, the property of *being a Polish astronomer* is naturally thought of as having the properties of *being Polish* and *being an astronomer* as parts. One part of being a Polish astronomer is being Polish, and another part is being an astronomer, and that's all there is to it. I shall assume that we are dealing with a set of properties with a certain mereological structure. This structure may be represented by a pair  $\langle P, \leq \rangle$ , where  $P$  is a finite set, and  $\leq$  is a binary relation on  $P$ .  $P$  represents some properties—i.e., those salient in some choice situation—and  $\leq$  represents the part-relation which holds between these, so  $p \leq q$  means that  $p$  is a part of  $q$ . I shall write  $p < q$  to mean that  $p$  is a proper part of  $q$ , where a proper part of  $q$  is any part of  $q$  except  $q$  itself, i.e.,  $p < q \iff (p \leq q \wedge p \neq q)$ .

It will be useful to define some functions, based on the structure  $\langle P, \leq \rangle$ .

**Definition 1.** For any property  $p$ , let  $parts(p)$  be the set of properties in  $P$  which are parts of  $p$ :

$$parts(p) = \{q \in P : q \leq p\}.$$

<sup>4</sup>For example, T. M. Scanlon states his famous “buck-passing theory of value” like this: “being good, or valuable, is not a property that itself provides a reason to respond to a thing in certain ways. Rather, to be good or valuable is to have other properties that constitute such reasons” (Scanlon 1998, p. 97).

<sup>5</sup>At least this is so when you know sufficiently well what the alternatives are like. If you're more in the dark, things may be more complicated. It may be that what you ought to do depends not on which properties the alternatives in fact have, but on which you believe them to have. We needn't resolve such complications here. So let us simply set aside such cases of ignorance.

**Definition 2.** For any set of properties  $Q$ , let  $atoms(Q)$  be the set of atomic properties in  $Q$ , i.e., the set of every property in  $Q$  of which nothing in  $Q$  is a proper part:

$$atoms(Q) = \{p \in Q : \neg \exists q \in Q (q < p)\}.$$

And let  $wholes(Q)$  be the set of whole properties in  $Q$ , i.e., the set of every property in  $Q$  which is a proper part of nothing in  $Q$ :

$$wholes(Q) = \{p \in Q : \neg \exists q \in Q (p < q)\}.$$

I assume  $\langle P, \leq \rangle$  is what might be termed a “classical” mereology. More exactly, I assume the composite function  $atoms \circ parts$  is an isomorphism between  $\langle P, \leq \rangle$  and  $\langle \wp(atoms(P)) \setminus \emptyset, \subseteq \rangle$ . This has some helpful implications. It implies that  $\leq$  is a partial order (because  $\subseteq$  is a partial order), i.e., transitive, reflexive, and antisymmetric. It also implies what is sometimes called “unrestricted composition”: for any properties, there is a mereological sum of these properties, defined as follows.

**Definition 3.** The *sum* of properties  $p$  and  $q$ , denoted here as “ $p + q$ ”, is the smallest property of which both  $p$  and  $q$  are parts, i.e., the supremum (least upper bound) of the set  $\{p, q\}$ .

Or equivalently, given that  $\langle P, \leq \rangle$  is classical,  $p + q$  is the unique property whose atomic parts include the atomic parts of both  $p$  and  $q$  and no others:

$$atoms \circ parts(p + q) = atoms \circ parts(p) \cup atoms \circ parts(q).$$

## 2.2 Contexts

A common view among philosophers is that reasons are “context-sensitive”. As Dancy puts it, “a feature that is a reason in one case may be no reason at all, or an opposite reason, in another” (Dancy 2004, p. 7). Dancy sometimes calls this view “Holism”. I should stress, however, that it is not Wholism in the sense I intend here. By “Wholism” I mean, as stated above, the view that whole reasons are fundamental: all reasons reduce to whole reasons. In fact, the negation of Dancy’s view, i.e., the view that reasons are *not* context-sensitive, which I call below “Invariabilism”, is a form of Wholism (it is equivalent to Isolationism, which I define below).

The framework I shall develop here will allow such context-sensitivity. It will allow that a property instantiated by two alternatives may be a reason for one but not for the other, or a less weighty reason for one than for the other. For example, if two bridges are both unstable, it could be that this property, instability, is a reason *against* crossing one bridge, yet a reason *for* crossing the other. Perhaps the first is, say, a pedestrian bridge over busy train tracks and has been damaged by a severe storm, whereas the second is part of an outdoor obstacle course and was designed to be unsteady so that traversing it would be more exciting.

I shall assume, however, that this can happen only when the alternatives differ in their other properties. If alternatives instantiate all the same properties, then they must be, as we might say, “normatively equivalent”, in the sense that any property that is a reason for one is an equally weighty reason for the other. This seems a safe assumption. It allows examples like the one given above. The second bridge, but not the first, was designed to be unstable; so this is one property that differentiates them. More generally, whenever it seems intuitive that a property is a reason for one alternative which has it but not for another, it will be possible, I predict, to explain why this is so, and this explanation will refer to some property had by one alternative but not the other.

Given this assumption, we may treat a context as a set of properties instantiated by an alternative. But, since the properties we are working with have mereological structure, not every set of

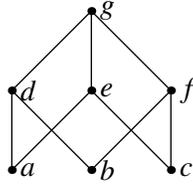


Figure 1: a property mereology

properties can be a context. A property cannot be instantiated without its parts, nor they without it. Nothing can be a Polish astronomer without being both Polish and an astronomer, and nothing can be both Polish and an astronomer without being a Polish astronomer. Thus, a context must be closed under both composition and decomposition. That is, it must contain every property all the parts of which it contains (composition), and all the parts of every property which it contains (decomposition). Thus we have the following definition:

**Definition 4.** A *context* is a set  $C \subseteq P$  such that for all  $p, q \in P$ ,

$$p + q \in C \iff \{p, q\} \subseteq C.$$

It follows that if a property is an atom in one context, then it is an atom in every context where it is present (i.e., if  $C$  and  $D$  are contexts, and  $p \in C \cap D$ , then  $p \in \text{atoms}(C)$  only if  $p \in \text{atoms}(D)$ ). However, a property may be a whole in one context without being a whole in every context where it is present. More exactly, for every property  $p$ , there is exactly one context in which  $p$  is a whole, namely,  $\text{parts}(p)$ . And every context contains exactly one whole: if  $C$  is a context then  $\text{wholes}(C)$  is a singleton.

To illustrate, suppose  $P$  contains seven properties as shown in Figure 1. Here the vertices, or “dots”, represent properties, named  $a$ ,  $b$ , and so on. Where two dots are connected by one or more lines, this indicates that the lower property is a part of the higher. Thus,  $a$  is a part of  $d$ ,  $b$  a part of  $g$ , and so on. From this structure we can construct seven contexts, as shown in Figure 2. Each diagram represents a context. The black dots represent the properties that are present in the context, and the grey dots those that aren't. Thus context  $A$  contains the single property  $a$ ,  $D$  contains three properties,  $a$ ,  $b$ , and  $d$ , and so on. Notice that  $A = \text{parts}(a)$ ,  $B = \text{parts}(b)$ , and so on. (So the function  $\text{parts}$  is a bijection between properties and contexts.) The atomic properties are those across the bottom row:  $a$ ,  $b$ , and  $c$ . These are atomic in every context where they are present. The single whole property in a context is that represented by the highest (black) dot. Thus  $d$  is the whole in  $D$ , but not in  $G$ , where  $g$  is the whole. A property may be both atomic and whole in a single context, as  $a$  is in  $A$ , for example.

### 2.3 Weighing Functions

Each property, in each context where it is present, has a certain normative *weight* (or force, strength, etc.). This weight will be represented in our framework by a “weighing function”, that is, a function  $w$  that assigns to each property-context pair  $\langle p, C \rangle$ , with  $p \in C$ , a real number, denoted  $w(p, C)$ . A positive number represents a favouring reason, a reason *for* choosing the property (or, more exactly, for choosing an alternative that has this property). A negative number represents a disfavouring reason, a reason *against* choosing the property. Zero represents a property that is no reason at all. The absolute value of the number (i.e.,  $w(p, C)$  if  $w(p, C) > 0$ , or  $-w(p, C)$  if  $w(p, C) < 0$ ) represents how weighty the reason is: the greater the value, the more weighty the reason. I assume that  $w$  represents the weight of reasons on a ratio scale. This amounts essentially

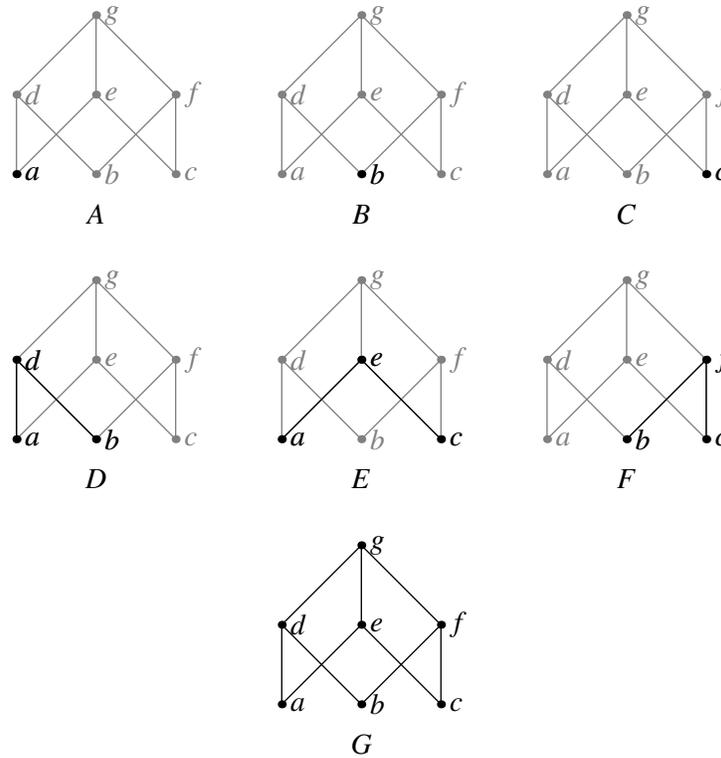


Figure 2: some contexts

to the assumption that we can compare ratios of differences between weights, e.g., we can say that the difference in weight between two reasons is, say, twice that between two other reasons.<sup>6</sup>

It is worth noting that this framework does not require that the fusion of some reasons be itself a reason. Think of a case where two perfectly matched reasons pull in opposite directions: one is a reason for some action, the other an equally weighty reason against. In this case, the *combined* force of the reasons may be zero, so that their fusion is no reason at all. So it is possible that putting together two reasons gives a non-reason. This is as it should be, I think. It is no more mysterious than the phenomenon of two poisonous chemicals combining to make something non-poisonous (as with, e.g., sodium and chloride).

### 3 Some Structural Views

With this framework in place, I may now define several structural views about reasons. I begin by defining what I mean by structural.

#### 3.1 Formal vs Substantive Conditions

Choosing is often preceded by a process of “weighing” reasons. In our framework, this process may be conceptualised as one of selecting a weighing function (or a set thereof, since weighing functions are unique only up to multiplication by a positive scalar). Or equivalently, it may be seen as a process of elimination: to select one weighing function is to eliminate all the rest. Your grounds for discarding a weighing function may be of two kinds, which I’ll call “substantive” and

<sup>6</sup>For a ratio scale, as opposed to merely a cardinal scale, we also need to fix the zero-point. We do this, as stated above, by letting zero be the weight of a non-reason, a property that neither favours nor disfavors.

“formal”. Substantive grounds concern the weights of particular properties. For example, you might think that the property of *causing pain* is always a negative reason and, therefore, reject any weighing function that assigns a non-negative weight to it. Formal grounds concern the overall pattern of weights assigned to properties. They involve requirements of consistency or coherence. For example, you might reject the sort of context-sensitivity discussed above. Consistency, you might think, requires that the same property is assigned the same weight in every context, in which case you would discard all weighing functions which assign different weights.

This distinction may be made more precise as follows. First, suppose  $w$  and  $v$  are weighing functions on the same property mereology  $\langle P, \leq \rangle$ . Then say that  $w$  and  $v$  are *isomorphic* iff there exists an automorphism  $\pi$  of  $\langle P, \leq \rangle$  such that, for all  $p \in P$ ,  $w(p) = v(\pi(p))$ .<sup>7</sup> Now suppose  $F$  is a property of weighing functions, by which I mean a *set* of weighing functions. Then say that  $F$  is a *structural* property iff  $F$  is closed under isomorphism (i.e., for all  $w$  and  $v$ , if  $w \in F$  and  $v$  is isomorphic to  $w$ , then  $v \in F$ ). The idea is that isomorphic weighing functions have a common structure, and a structural property is one that is shared by all structurally equivalent weighing functions. Finally, we may say that the reason for rejecting a weighing function is *formal* iff the weighing function is rejected in virtue of its having a structural property; otherwise the reason is *substantive*.

Here I will interested solely in formal conditions. I assume there is a set  $S$  of weighing functions that satisfy every desirable formal condition. These weighing functions are “structurally sound”, as we might say: they are internally consistent and coherent. They might not satisfy certain external, substantive requirements, but that won’t concern me here.

### 3.2 Atomism and Wholism

Atomism and Wholism, as I use those terms here, are views about  $S$ . Atomism says that  $S$  is “atomistic”, and Wholism that it is “wholistic wholistic”, where these two properties are defined, in two stages, as follows.

**Definition 5.** Let  $w$  and  $v$  be weighing functions. Then:

1.  $w$  and  $v$  are *atom-equivalent* iff, for all properties  $p$  and contexts  $C$ , if  $p \in \text{atoms}(C)$  then  $w(p, C) = v(p, C)$ ;
2.  $w$  and  $v$  are *whole-equivalent* iff, for all properties  $p$  and contexts  $C$ , if  $p \in \text{wholes}(C)$  then  $w(p, C) = v(p, C)$ .

**Definition 6.** Let  $W$  be a set of weighing functions. Then:

1.  $W$  is *atomistic* iff, for all  $w, v \in W$ , if  $w$  and  $v$  are atom-equivalent then  $w = v$ ;
2.  $W$  is *wholistic wholistic* iff, for all  $w, v \in W$ , if  $w$  and  $v$  are whole-equivalent then  $w = v$ .

Atomism is the view that all reasons supervene on atomic reasons: within the class of structurally sound weighing functions, there cannot be a difference in the weight of any property without a difference in the weight of some atomic property. Likewise, Wholism is the view that all reasons supervene on whole reasons.

It may be helpful to think of these views in the following way. Suppose two people have each considered the same practical problem. Each has finished weighing all the relevant properties and

<sup>7</sup>An automorphism of  $\langle P, \leq \rangle$  is a permutation  $\pi$  of  $P$  such that for all  $p, q \in P$ ,  $p \leq q$  iff  $\pi(p) \leq \pi(q)$ . A permutation of  $P$  is a bijective mapping from  $P$  into itself.

has settled on a weighing function. They have, however, selected different weighing functions: there is at least one property and context such that the two people assign different weights to this property in this context. Finally, suppose that each person's weighing is internally consistent. Then Atomism entails that the two people must disagree over the weight of at least one atomic property, whereas Wholism entails that they must disagree over the weight of at least one whole property.

Neither Atomism nor Wholism says anything about the consistency of particular weighing functions. Atomism does say that the set of all consistent weighing functions is atomistic, but every weighing function belongs to at least one atomistic set (e.g., for any weighing function  $w$ , the singleton set  $\{w\}$  is trivially atomistic). I turn next to defining three views that do have implications regarding the consistency of particular weighing functions. I call these "Additivism", "Isolationism", and "Marginalism". The first is a form of Atomism (i.e., it implies Atomism), and the other two are forms of Wholism.

### 3.3 Additivism

Suppose  $p$  and  $q$  are distinct atomic properties. Then Atomism says that the weight of their sum,  $p+q$ , reduces to the weights of  $p$  and  $q$ .<sup>8</sup> (Here I am, for clarity, suppressing reference to contexts.) But how? Here's a natural thought: the weight of  $p+q$  is the sum of the weights of  $p$  and  $q$ . The weight of the sum is the sum of the weights. This thought leads to the view I call "Additivism", defined as follows.

**Definition 7.** A weighing function  $w$  is *additive* iff it satisfies the following condition. For all properties  $p$  and contexts  $C$ , if  $p \in C$  then

$$w(p, C) = \sum_{q \in A} w(q, C),$$

where  $A = \text{atoms}(\text{parts}(p))$ .

Additivism is the view that all elements of  $S$  (i.e., all structurally sound weighing functions) are additive. It says that, according to any consistent weighing function, the weight of a property in a context is the sum of the weights of its atomic parts in the same context.

That Additivism implies Atomism is easily seen. Because the weight of every property (in every context) is the sum of the weights of some atomic properties, it follows that if two distinct weighing functions are atom-equivalent (i.e., they assign the same weights to all atoms), then at least one of them is not additive.

### 3.4 Isolationism and Marginalism

The two forms of Wholism I'll discuss might helpfully be introduced by analogy with weighing of a different sort, weighing physical objects. Suppose you wish to know the weight of some physical object that is part of a larger whole, e.g., the handlebar of a bicycle. One approach would be to remove the handlebar from the bicycle and then to weigh it on its own. But suppose weighing the handlebar alone is not possible for some reason: perhaps, say, it won't sit stably on the scales. Then all is not lost; there is another way. You could, first, weigh the complete bicycle, and then remove the handlebar and weigh what remains. Subtracting the latter weight from the former will yield the weight of the handlebar.

---

<sup>8</sup>Strictly, Atomism says only something weaker: the weight of  $p+q$  reduces to the weights of *some* atomic reasons, but not necessarily to  $p$  and  $q$  in particular. It seems natural, however, to think that if a property's weight reduces to the weights of any atomic reasons, then it will reduce to the weights of the atomic properties of which it is composed.

These two approaches correspond to two forms of Wholism. The first corresponds to Isolationism. On this view, the weight that a property has as part of a larger whole is the weight that it would have in isolation. In the case of an atomic property  $p$ , the weight of  $p$  in isolation is  $w(p, \{p\})$ , i.e., the weight that  $p$  has in a context where it is the only property present. In the case of a non-atomic, or complex, property, we cannot achieve full isolation. A complex property cannot be instantiated without its parts, and so there is no context in which it is entirely alone. But we can say that a complex property is isolated, in a more inclusive sense, when it is accompanied by no other properties besides those which are parts of it. In general, then, the weight of a property  $p$  in isolation is  $w(p, \text{parts}(p))$  (which is the same as  $w(p, \{p\})$  when  $p$  is atomic).

To simplify notation, I shall write  $w(p)$  to denote the weight of  $p$  in isolation; i.e.,  $w(p) = w(p, \text{parts}(p))$ . Isolationism may then be defined as follows:

**Definition 8.** A weighing function  $w$  is *isolationist* iff it satisfies the following condition. For all properties  $p$  and contexts  $C$ , if  $p \in C$  then

$$w(p, C) = w(p).$$

When a property is isolated it is a whole property, not part of any larger property. Hence the term on the right side of the above equation represents the weight of a whole property, and therefore Isolationism entails Wholism. Suppose  $w$  and  $v$  are isolationist weighing functions. If  $w$  and  $v$  are distinct — i.e., if for some property  $p$  and context  $C$ ,  $w(p, C) \neq v(p, C)$  — then for some property  $p$ ,  $w(p) \neq v(p)$ , and so they are not whole-equivalent. The set of all isolationist weighing functions is, therefore, wholistic wholistic.

The second approach corresponds to Marginalism. On this view, the weight that a property has as part of a larger whole is the difference in overall weight that would result from removing this property. Suppose  $p$  and  $q$  are non-overlapping properties (i.e., they have no parts in common). Then in the context  $\text{parts}(p + q)$ ,  $p$  is part of a larger whole,  $p + q$ . Marginalism determines the weight of  $p$  in this context as follows. First, take the overall weight of all the properties in this context, i.e., the weight of their sum,  $p + q$ . Then remove  $p$ , leaving the context  $\text{parts}(q)$ , and take the weight of all properties in this context, i.e., the weight of their sum,  $q$ . Finally, subtract the latter from the former to get the weight of  $p$  in the original context.

Thus we have the following definition.

**Definition 9.** A weighing function  $w$  is *marginalist* iff it satisfies the following condition. For all properties  $p$  and contexts  $C$ , if  $p \in C$  then

$$w(p, C) = w(p + q) - w(q),$$

where  $q$  is a property such that  $p$  does not overlap  $q$ , and  $C = \text{parts}(p + q)$ .<sup>9</sup>

Again, this is clearly a form of Wholism, because the terms on the right side of the equation are weights of whole properties.

### 3.5 Strong Additivism

These three views — Additivism, Isolationism, and Marginalism — are logically compatible. A weighing function may satisfy all three conditions: it may be additive, isolationist, and marginalist. But they are not entirely logically independent. The conjunction of any two of them entails the third. You cannot consistently accept two without accepting all three.

<sup>9</sup>Because  $\langle P, \leq \rangle$  is classical, as defined above, it follows that there exists exactly one property  $q$  satisfying these two conditions.

The conjunction of these three views (and hence the conjunction of any two of them) is equivalent to a fourth, which I'll call "Strong Additivism", defined as follows.

**Definition 10.** A weighing function  $w$  is *strongly additive* iff it satisfies the following condition. For all properties  $p$  and contexts  $C$ , if  $p \in C$ , then

$$w(p, C) = \sum_{q \in A} w(q),$$

where  $A = \text{atoms}(\text{parts}(p))$ .

Thus Strong Additivism is the view that the weight of a property, in any context, is the sum of weights of its atomic parts *in isolation*. Whereas Additivism applies only *within* a context, Strong Additivism applies *across* contexts. Additivism reduces the weight of a given property in a given context to the weights of its atomic parts *in the same context*. Strong Additivism, on the other hand, reduces this to the weights of the atomic parts in *different* contexts (i.e., the contexts in which those atomic parts are alone). Strong Additivism implies both Atomism and Wholism: it reduces all reasons to the atomic whole reasons.

The following proposition states some important logical relations between the four views defined above:

**Proposition 1.** *The following statements are equivalent:*

1.  $w$  is additive and marginalist,
2.  $w$  is additive and isolationist,
3.  $w$  is marginalist and isolationist,
4.  $w$  is strongly additive.

*Proof.* Without loss of generality, suppose  $P = \{x, y, z\}$ , with  $x$  and  $y$  atomic, and  $x + y = z$ . Then we have the following equivalences:

$$\begin{aligned} w \text{ is additive} &\iff w(z) = w(x, P) + w(y, P), \\ w \text{ is isolationist} &\iff (w(x, P) = w(x) \wedge w(y, P) = w(y)), \\ w \text{ is marginalist} &\iff w(z) = w(x, P) + w(y) = w(y, P) + w(x), \\ w \text{ is strongly additive} &\iff w(z) = w(x) + w(y). \end{aligned}$$

Now it is straightforward to confirm that the four statements are equivalent. □

So, in summary, if you accept any two of Additivism, Isolationism, and Marginalism then you must accept all three, and you must also accept Strong Additivism. As I'll argue below, however, you shouldn't accept Strong Additivism.

## 4 Against Atomism

My argument against Atomism will proceed in two steps. First, I'll argue against accepting both Atomism and Wholism. Second, I'll argue for accepting Wholism. In brief: we should choose only one of Atomism and Wholism, and we should choose Wholism.

#### 4.1 Absurd Consequences of Strong Additivism

How might you accept both Atomism and Wholism? One way to do this — perhaps the most obvious way — would be to accept Strong Additivism, which, as we saw above, entails both views. But Strong Additivism, as I'll now argue, is best avoided.

Suppose you have a choice of four drinks:

1. coffee with milk
2. whisky neat
3. black coffee
4. whisky with milk

And suppose the relevant properties are: *contains coffee* ( $c$ ), *contains whisky* ( $w$ ), *contains milk* ( $m$ ), and all mereological sums of these. If  $w$  is strongly additive then

$$w(c + m) \geq w(c) \iff w(w + m) \geq (w).$$

Strong Additivism implies, therefore, that it's okay to add milk to your coffee iff it's okay to add milk to your whisky; choosing coffee with milk is at least as rational as choosing black coffee iff choosing whisky with milk is at least as rational as choosing whisky neat. (By saying here that  $x$  is at least as rational as  $y$ , I mean simply that the overall reason for  $x$  is at least as strong as that for  $y$ ). But of course that is absurd! There is nothing irrational in both preferring coffee with milk and whisky without milk (or, for that matter, having the opposite preferences, preferring whisky with milk and coffee without).

This might not be a fatal blow to Strong Additivism. The proponent of that view might argue that I've wrongly identified the relevant properties. These are not the properties of containing coffee and so on, but some other properties. Still, examples like this do reveal a sort of trouble that Strong Additivism may lead us into if we're not careful. A less risky option is to reject Strong Additivism. It seems undesirable to constrain our choice of substantive views by the adoption of certain formal or structural conditions. But this is effectively what we would do by accepting Strong Additivism, since it is incompatible with certain substantive positions regarding which properties are reasons (insofar as it cannot be combined with such positions without leading to absurdity of the kind noted above). If we want to keep our substantive options open, we are better off without Strong Additivism. And this seems a reason, if perhaps not a decisive one, to reject it.

Suppose, then, that we do reject it. Then what? Well, then we will probably need to reject at least one of Atomism and Wholism too. Although Strong Additivism is not *implied* by the conjunction of Atomism and Wholism, it does seem the most obvious way of combining the latter two views. Combining Additivism, which seems the most obvious form of Atomism, with either Isolationism or Marginalism, which seem the most obvious forms of Wholism, yields Strong Additivism, as we've seen. Though it remains logically possible to accept both Atomism and Wholism while rejecting Strong Additivism, I cannot think of a sensible way of doing so. For example, here is one possibility. Consider the view that  $S$  contains all and only *constant* weighing functions (i.e., those functions which assign the same weight to every property in every context). On this view,  $S$  is both atomistic and wholistic.<sup>10</sup> But nearly all constant weighing functions violate strong additivity. However, this view — effectively the view that all properties have equal

<sup>10</sup>This is trivially true, because distinct constant weighing functions cannot be either atom-equivalent or whole-equivalent.

weight — seems rather implausible. If there are more plausible ways of combining Atomism and Wholism without Strong Additivism, I have not been able to discover them. That is not to say, of course, that they do not exist, but only that I doubt they do.

The question arises, then, which of Atomism and Wholism has to go? I shall argue that it is Atomism. We should keep Wholism.

## 4.2 For Wholism

Ultimately, what matter are whole reasons. What you ought to do is determined by the overall balance of reasons; or, in our mereological framework, by the balance of overall reasons.<sup>11</sup> You ought to do that which is favoured by the strongest overall reason (or that which is disfavoured by the weakest). And overall reasons are whole reasons; so what you ought to do is determined by whole reasons. If weighing functions are whole-equivalent — i.e., if they assign the same weights to all whole properties — they are, therefore, practically equivalent: they agree perfectly as to what you ought to do.<sup>12</sup>

Now recall  $S$ , the set of all structurally sound, consistent and coherent weighing functions. It seems desirable that  $S$  contain as few weighing functions as are necessary for practical purposes. (Think of this as a practical version of Occam’s Razor: do not multiply entities beyond practical need.) That is, if weighing functions  $w$  and  $v$  are both in  $S$ , and  $w$  and  $v$  are practically equivalent, then  $w = v$ . Otherwise, you could be left with pointless decisions, deciding between weighing functions that make no difference at all to what you ought to do. But this entails that  $S$  is wholistic wholistic, that Wholism is true.

Perhaps it will be objected that correct normative judgement requires weighing all the reasons, whereas Wholism implies that we need not weigh all reasons, only whole reasons. But this objection misunderstands Wholism. True, Wholism does imply that weighing whole reasons is sufficient; but that doesn’t mean that weighing non-whole reasons is unnecessary. You can’t weigh a whole reason without weighing its parts. Recall the bicycle example. The bicycle is nothing more than the sum of its parts, the handlebar, the seat, the wheels, etc. You can’t put the whole bicycle on the scales without at same time putting the handlebar and all the rest there too. Wholism in fact requires the weighing of *all* parts, including non-whole parts. It requires, however, that they be weighed, not individually, one at a time, but all together as a whole. That’s just what it is to weigh a whole reason.

## 4.3 Isolationism or Marginalism?

I have argued that we should accept Wholism. But there are many ways to do this. Two ways are by accepting Marginalism and by accepting Isolationism. Which is better? Here I am neutral. There are, I think, two things we may mean by speaking of the normative weight of a non-whole property. We could mean the weight that the property would have in isolation, which we might call its “intrinsic weight”. Or we might mean the contribution that this property makes to the overall weight of reasons, which we might call its “marginal weight”. Either way of talking is

---

<sup>11</sup>What I mean here is simply that overall reasons trump other reasons. Strictly speaking, other facts — e.g., facts about what you *can* do — may help to determine what you ought to do. More on this below (Section 5.2).

<sup>12</sup>This is not quite right. Below (Section 5.2) I suggest another practical need served by reasons: they are useful in explaining not only why we ought to do certain things, but also why we are rational to regret our inability to do certain things. The peatiness of a particular brand of whisky may help to explain why I should regret my inability to drink this whisky if, say, the pub I am in does not have any. What I need to say, then, is that rational regret is also fundamentally a matter of whole reasons. If I have overall more reason to drink a brand of whisky that *is* available, I should not regret the unavailability of the inferior brand. I’m happy to say this.

fine, so far as I can tell, provided we are clear about what we mean. When talking in the former way, it is Isolationism which is appropriate; when talking in the latter, it is Marginalism. The important thing is not to mix them up: we mustn't assume that intrinsic and marginal weight always come to the same. That way lies Strong Additivism, and the absurd consequences thereof.

Some philosophers, however, will want to reject Isolationism. As noted earlier, they think reasons are context-sensitive: the same property may have different normative weight in different contexts. But this is incompatible with Isolationism. If the weight of a property in any context (in which it is present) is its weight in isolation, then its weight is the same in every context.

In the next section, I want, therefore, to address two arguments for context-sensitivity. I shall argue that neither is decisive.

## **5 Context-sensitivity**

By "Invariabilism" I shall mean the view that reasons are not context-sensitive: a property has the same normative weight in every context where it is present. I'll address two arguments against Invariabilism.

### **5.1 Intuitions about Cases**

Arguments for context-sensitivity often involve involve cases like the following:

Suppose my seeing a movie would give me pleasure. Plausibly, the fact that seeing a movie would give me pleasure is a reason for me to see it. However, in another case that very same consideration might be no reason at all to perform the action. If, for example, I would take pleasure in seeing someone tortured in the movie (suppose it is a documentary on human rights abuses) then the fact that I would take pleasure in the film plausibly is no reason whatsoever to see it. (McKeever and Ridge 2006, p. 27)

To fit this into my framework, I'll interpret McKeever and Ridge as making a claim about a property, rather than about facts. This seems better anyway. McKeever and Ridge claim that one fact is a reason while another is not. But these facts are distinct; they are about different movies. So even if their claim is true, this would not show that the "very same consideration" is sometimes a reason, sometimes not. If "considerations" are facts, then there are two considerations, not one. This problem may be solved, I suggest, by substituting properties for facts, in the way suggested. There is a property, viz., being such as to give someone pleasure, which is present in two contexts, corresponding to the two movies. One movie is a comedy (let's say) and the other a documentary on human rights abuses. McKeever and Ridge claim this property is a reason to see the comedy but not the documentary. I assume their reason for thinking this is that, to them, this seems intuitively correct.

As McKeever and Ridge say, the Invariabilist might attempt to accommodate such intuitions by urging that what is a reason to see the first movie is a slightly different property, not the simple property of giving pleasure, but the more complex property of giving "nonsadistic pleasure" (McKeever and Ridge 2006, p. 27). The Invariabilist may count this as a reason while denying there is any reason to see the second movie, because the pleasure given by the second movie would be sadistic pleasure, being caused by seeing footage of torture. This strategy is sometimes called "expansionism", because it involves "expanding" what is counted as a reason, by building in some further normatively relevant property (See e.g. Dancy 2004, pp. 125–127). McKeever and Ridge concede "some temptation" to adopt expansionism, though apparently not an irresistible one (McKeever and Ridge 2006, p. 27).

Dancy has argued against expansionism. His main objection, as I understand it, is that expansionism is unacceptably *ad hoc*. It involves “complicating” our accounts of reasons unnecessarily. The expansionist’s substitute reasons are more complicated than the (allegedly) intuitive reasons they replace, e.g., the property of providing non-sadistic pleasure is more complicated than that of providing pleasure. But, says Dancy, “there is no reason to embark on this process of complication” (Dancy 2004, p. 127).

I want to suggest, however, a principled reason for not taking such intuitions at face value. Let me first describe a certain general phenomenon in ordinary language. Sometimes a sentence which literally ascribes some property to some object is “assertable”, a perfectly fine thing to say, even though this object does not have this property (nor is believed to have it), but is instead only a salient proper part of something else which does have it. (This phenomenon, whereby a part is used to refer to the whole, is sometimes called “synecdoche”.) Here are some examples:

- (4) My car won’t fit under the bridge; it’s towing a tall yacht.  
(Actually, the car will fit easily, though the fusion of the car and the yacht, of which the car is only a part, won’t fit.)
- (5) She got wet walking past a garden sprinkler.  
(Actually, her clothing kept her dry; but the fusion of her and her clothing did get wet.)
- (6) Your suitcase is too heavy to take on board the plane.  
(Actually, the suitcase is quite light; it’s the suitcase plus its contents that is heavy.)

Although each of these sentences is literally false, to point this out in ordinary conversation would be quite pedantic. (It’s okay, of course, in philosophy.) There may be different ways of accounting for this phenomenon, both semantic and pragmatic, but we needn’t go into details here. The essential point is just this: sometimes it’s okay to (literally) ascribe a property to something even when the thing to which the property is ascribed is merely a part of some other thing that has the property.

Now return to the movie example. Here two properties are candidates for being reasons to see the first movie: (i) giving pleasure, and (ii) giving non-sadistic pleasure. But the first is a part of the second (every non-sadistic pleasure is a pleasure). Thus, even if only the second property were a reason, it might nonetheless be okay to say the first property is a reason (or at any rate to utter a sentence which, when interpreted literally, means that the first property is a reason), this being another instance of the general phenomenon described above. This is especially likely if the first property is an especially salient part of the second, which arguably it is. Thus the Invariabilist might argue that intuitions such as those reported by McKeever and Ridge do not so clearly tell against expansionism, since even if the reason here really were (ii), we might nonetheless feel intuitively that it is (i).

## 5.2 Reasons to Do What Cannot be Done

Another argument against Invariabilism claims that it has the — allegedly implausible — consequence that you may have reasons to do things you cannot do (See e.g. Dancy 2000, p. 153; McKeever and Ridge 2006, pp. 49–50; Ridge 2007, p. 344). Here’s an example. Suppose you’re deciding where to spend your holidays. Someone recommends going to the Marlborough Sounds because of their renowned tranquility. This is, let’s grant, a reason to go there. But the moon is also famously tranquil. So, given Invariabilism, this is equally a reason to go to the moon, in spite of the impossibility of your going there. Invariabilism implies therefore that you have a reason to

do something you cannot possibly do. According to the objection, this implication is implausible. As I shall argue, however, it's not so bad.

To be sure, in a conversation about where to spend your holidays, it would be rather odd, and perhaps in some sense inappropriate, to say that the tranquility of the moon is a reason to go there. But this doesn't show that what is said is false. Its oddness might instead be explained pragmatically, as being due to its having some false implicature, e.g., that it would be fruitful to entertain the moon as a potential holiday destination (cf. Schroeder 2007, pp. 92–97). Obviously this would be a waste of your time, and so it would be pointless to talk about reasons for going there. But that doesn't mean there are no such reasons.

The objection might be thought to gain support from the familiar principle that “ought” implies “can”, that what you ought to do is a subclass of what you can do. If we allow reasons to do things you cannot do, then there's no guarantee these reasons will always be outweighed by countervailing reasons, so the strongest overall reason might turn out to be one in favour of your doing something you cannot do. Suppose the moon is the only place that instantiates the unique blend of properties — tranquility, privacy, remoteness, etc. — that you prize in a holiday destination. Then, given Invariabilism, going to the moon is what you have *most* reason to do. And so, if what you have most reason to do is what you ought to do, we have a violation of the principle that “ought” implies “can”.<sup>13</sup>

However, the Invariabilist can resist this consequence by rejecting the last step, the assumption that you ought to do that which you have most reason to do. They can say instead that you ought to do that which, *of the things you can do*, you have most reason to do. This might seem to result in needless multiplication of reasons. We need reasons, some might claim, solely for the purpose of determining what people ought to do. On the view I've proposed, however, some reasons — those which favour or disfavour impossible actions — have no bearing on what you ought to do. These reasons might seem therefore redundant, serving no purpose (cf. Streumer 2007, pp. 365–368). There are, however, other purposes reasons might serve. For example, they can explain why your inability to do certain things is *regrettable*. If the moon really is your ideal holiday destination, you might naturally *regret* your being unable to go there. But this regret would make no sense if you had no reason to go there. Why should you regret your being unable to do something you have no reason to do? Some might claim it is sufficient for your inability to be regrettable that you *would* have reason to go to the moon if you *could* go there. But this counterfactual might be false. It might be that if the moon were within your reach, it would be no longer tranquil, because overrun by “space tourists”, and so you would lack any reason to go there. So what explains your regret must be that now, although you have reason to go there, you can't go there. If there were never any reasons to do things you were unable to do, such inability would never be regrettable.

## References

- Armstrong, D. M. (1997). *A World of States of Affairs*. Cambridge University Press.
- Dancy, Jonathan (2000). ‘The Particularist’s Progress’. In: *Moral Particularism*. Ed. by Brad Hooker and Margaret Olivia Little. Oxford University Press, pp. 130–157.
- (2004). *Ethics without Principles*. Oxford University Press.

<sup>13</sup>Peter Vranas argues in the opposite direction, from the premise that one can only have reason to do what one can do, to the conclusion that one ought to do only what one can do (Vranas 2007). However, he does not argue for the premise, except to say that it has “an almost tautological ring” (Vranas 2007, p. 173). His idea seems to be that, if a reason supported something that could not be done, it could not properly be classed as a “reason for action”. But this seems wrong. Going to the moon, for example, is an action, even if not an action that you can perform.

- McKay, Thomas (2006). *Plural Predication*. Oxford University Press.
- McKeever, Sean and Michael Ridge (2006). *Principled Ethics: Generalism as a Regulative Ideal*. Oxford University Press.
- Ridge, Michael (2007). 'Anti-reductionism and Supervenience'. In: *Journal of Moral Philosophy* 4.3, pp. 330–348.
- Scanlon, T. M. (1998). *What We Owe to Each Other*. Harvard University Press.
- Schroeder, Mark (2007). *Slaves of the Passions*. Oxford University Press.
- Streumer, Bart (2007). 'Reasons and Impossibility'. In: *Philosophical Studies* 136, pp. 351–384.
- Vranas, Peter B. M. (2007). 'I Ought, Therefore I Can'. In: *Philosophical Studies* 136.2, pp. 167–216.