

'ALL THINGS CONSIDERED'*

Ruth Chang
Rutgers University

One of the most common judgments of normative life takes the following form: With respect to *some* things that matter, one item is better than the other; with respect to *other* things that matter, the other item is better; but *all things considered*—that is, taking into account *all* the things that matter—one item is better than the other. For example, in deciding which of two career paths to pursue, I might judge that the deep-sea diving career is more exciting, while the career selling auto insurance affords greater financial security, but that, *all things considered*—that is, taking into account excitement, financial security, and whatever else is relevant to the judgment at hand—the diving career is better. Or a society choosing between two schemes for distributing benefits and burdens might judge that one scheme is better with respect to utility, and the other is better with respect to maximizing the welfare of the worst off, but that, overall—taking into account both ideals of utility and maximin—the one scheme is better. Such judgments range over questions not only about what we should *do* but what we should *want* or *believe*. Sitting in a restaurant, scanning its menu, I might find myself trying to decide what I should want to eat. With respect to cost, the grilled cheese sandwich is best; with respect to healthfulness, the broccoli medley is to be preferred; and with respect to taste, it's the steak *frites* hands down. But I might judge that, *all things considered*—taking into account cost *and* healthfulness *and* taste—the broccoli medley is the best option. Or in deciding which of two theories to believe, I might ascertain that the one theory is simpler, the other explanatorily more powerful, but conclude that *all things considered*—taking into account simplicity and explanatory power—the one theory is better than the other.

In this paper, I want to explore *how* all-things-considered judgments are possible, assuming that they are. In particular, I want to examine the question of how the different considerations relevant to an all-things-considered judgment come together in a way that gives each relevant consideration its proper due. How, for example, do considerations as seemingly disparate as cost, healthfulness, and taste, or utility and maximin, or explanatory power and theoretical

simplicity, come together in a judgment that, with respect to *all* of the relevant considerations, the one item is better than—or otherwise normatively related to—the other?

An all-things-considered judgment requires that the things considered stand in normative relations that are capable of sustaining the judgment. These normative relations give the relative importance of the things considered to the all-things-considered judgment and include not just aggregative relations such as outweighing but also nonaggregative relations such as trumping, silencing, canceling, excluding, being more stringent than, and so on. For simplicity, we can say that one item is all-things-considered 'better' than another if the items stand in a normative relation that in some sense 'favors' the one over the other.

Suppose, for example, that I judge that policy #1 is all-things-considered better than policy #2. If utility and maximin are the considerations relevant to the judgment, there must be some normative relation between them that supports that judgment. Suppose that policy #1 is marginally better with respect to utility but significantly worse with respect to maximizing the welfare of the worst off. If utility is significantly less important to the all-things-considered judgment than maximin, then policy #1 cannot be all-things-considered better than policy #2. Or consider the judgment that, all things considered, the broccoli medley is best. If cost, taste, and healthfulness are the considerations relevant to that judgment, there must be normative relations among them that support that judgment. Suppose that the broccoli medley is marginally better than the other options with respect to healthfulness but significantly worse with respect to cost and taste. If healthfulness is significantly less important than cost and taste, it cannot be true that, with respect to all three considerations, the broccoli medley is best. All-things-considered judgments presuppose that the things considered stand in normative relations that support the judgment. But how are these normative relations determined? In virtue of what does utility count more or less than maximin? What determines how important healthfulness is as against cost or taste? In general, how is it that the things considered in an all-things-considered judgment stand in normative relations that make such judgments possible?

I want to propose an answer to this question that, if correct, provides a unified account of all-things-considered judgments and highlights a deep connection between values and reasons. My suggestion is that 'all things considered' is, in effect, a placeholder for a more comprehensive *value* that includes the things considered as parts, and that this more comprehensive value determines how the things considered normatively relate to one another.¹ So, for example, when I judge that policy #1 is all-things-considered best, I am judging that, with respect to some more comprehensive value that includes the values of utility and maximin as parts, that policy is best. It is in virtue of this more comprehensive value that utility and maximin have their relative importance. And when I judge that the broccoli medley is all-things-considered better than the grilled cheese,

I am judging that, with respect to some more comprehensive value that includes the values of cost, taste, and healthfulness as parts, the broccoli medley is better. This more comprehensive value determines the importance of healthfulness vis-à-vis cost and taste.

This suggestion will strike some as curious. To make it seem less so, it is worth noting right away that the 'comprehensive value' view is certainly correct for *some* all-things-considered judgments. Sometimes, 'all things considered' quite clearly stands for a more comprehensive value that includes the relevant considerations as parts. For example, in the all-too-familiar task of evaluating philosophers for a job, one might judge that philosopher #1 is more original and insightful than philosopher #2, that philosopher #2 is clearer and more historically sensitive than philosopher #1, but that, *all things considered*, philosopher #1 is better. Here, 'all things considered' is a placeholder for a value—philosophical talent—that has the things considered—(philosophical) originality, insightfulness, clarity, historical sensitivity—as parts, and it is in virtue of this more comprehensive value that originality counts for so much as against historical sensitivity, and so on. My suggestion is that even when the things considered appear to be very different—cost, taste, and healthfulness; utility and maximin; simplicity and explanatory power—an all-things-considered judgment that gives each of these considerations its proper due does so in virtue of a more comprehensive value that has the things considered as parts. If my proposal is correct, then, as a general matter, 'all things considered' is a placeholder for a more comprehensive value that determines the relative importance of the things considered.

Now one might naturally wonder: if, as I claim, there is a more comprehensive value that includes, say, cost, taste, and healthfulness, as parts, *what is it?* Here my proposal may seem to get curiously and curiously. I believe that, in many cases in which the considerations relevant to the all-things-considered judgment are very different, the more comprehensive value that accounts for their normative relations *has no name*. Lest their missing moniker suggest that they are a philosopher's fantasy, consider (dis) values, such as 'sexually harassing' or 'tubular', 'rad', and 'phat' which were not long ago nameless and yet were referred to in everyday conversation by expressions such as 'behaving like a first-class jerk' and 'wicked, man', before we gave them their names. The namelessness of a value is an accidental product of our naming practices.

My case for the comprehensive value approach has three parts. In the first, I examine the usual answers to the question of how all-things-considered judgments are possible. Although this question is rarely considered explicitly, many philosophers say enough to imply one of two broad views that I will call, respectively, the 'simple' and the 'sophisticated' orthodoxies. I will argue that the simple orthodoxy must be rejected and that the sophisticated orthodoxy is problematic in ways that give us reason to cast about for an alternative view. If the usual ways of understanding all-things-considered judgments are

problematic, one very large obstacle to believing that the comprehensive value approach might be correct is thereby removed. But we must also explain how an appeal to more comprehensive *values* helps matters. The second part of the paper explores what it is that makes certain values 'hang together' as parts of some more comprehensive value, while other values do not hang together to form a more comprehensive value at all. As I will suggest, it is because more comprehensive values have a 'unity' that they can account for the normative relations among their component values in different situations; more comprehensive values can explain how all-things-considered judgments are possible because they have a unity in virtue of which their components have the normative relations they do. I end by considering some of what I take to be the most serious objections to this proposal. These have to do with doubts about whether there are such (nameless) values, for if there are, then it is plausible to think they do the normative work I suggest they do. I try to dispel the main sources of doubts about them. My aim in this paper, however, is not to argue that the comprehensive value view *must* be correct; I only want to say enough in its favor to suggest that, all things considered, it is a view that should be taken seriously.

I.

A few clarifications may help to head off some worries. The first worry has to do with the scope of my claim. I have said that "in general" 'all things considered' holds the place of a more comprehensive value, and that the comprehensive value approach offers a "unified" account of all-things-considered judgments. If these claims are understood to imply that 'all things considered' *always* stands for some more comprehensive value, then the claim is clearly false. For we might *stipulate* a weighting of considerations that we deem relevant to an all-things-considered judgment. I might say to myself that the best evening for me is one filled with one part higher learning to two parts puerile humor, and accordingly judge that an evening that begins with a chapter of *The Critique of Pure Reason* and quickly segues into reruns of *Fawlty Towers* is all-things-considered best. When we stipulate what matters in an all-things-considered judgment, 'all things considered' holds the place of the stipulation, not of a more comprehensive value, and it is the stipulation, not a more comprehensive value, that determines the normative relations among the things considered.

My claim also excludes all-things-considered judgments that follow from a 'rule of comparison' such as the Pareto Rule, which says that, if one item is at least as good as the other on each of the relevant criteria (and the converse does not hold), then that item is all-things-considered better. In this case, the all-things-considered judgment proceeds without there needing to be any normative relations among the things considered; 'all things considered' stands for the rule of comparability, and thus, so far as that judgment goes, there are no normative relations among the things considered that need explaining. In both stipulations

and rules of comparison, the all-things-considered judgment does not involve any *substantive* judgment about how the things considered normatively relate. My claim about all-things-considered judgments, then, is meant to cover only those judgments that 'put together' the things considered in this substantive way.²

The question about all-things-considered judgments as I have posed it assumes that even when the things at stake are very different, a correct all-things-considered judgment can nevertheless be made. The second worry involves doubt about this assumption. One might think that when the relevant considerations are very different, *no* all-things-considered judgment can be made. Sidgwick believed that this was true of egoistic and utilitarian values: such considerations cannot be normatively related, and hence the duality of practical reason. But Sidgwick offered no *argument* for this conclusion; it just struck him that there was no way to put together prudential and moral considerations.³ He never contemplated the possibility that, just as there is a more comprehensive value that puts together originality and historical sensitivity, there might be a more comprehensive value that puts together the values of, say, helping someone in need and protecting one's own interest. Indeed, despite Sidgwick's claim to the contrary, there can plausibly be normative relations among moral and prudential values. For example, taking into account both moral and prudential considerations, it is plausibly better that I save the drowning baby than avoid getting my expensive Italian shoes wet. My claim is not that every consideration stands in some normative relation to every other consideration, so that there is always an all-things-considered judgment no matter which considerations are at stake. How, for example, could there be an all-things-considered judgment when the things considered are the beauty of the number nine and the efficiency of a corkscrew? As we will see later, an appeal to more comprehensive values helps to explain why some considerations can be put together in an all-things-considered judgment while others cannot. For now, however, let me assume that in many cases in which the things at stake are very different, there can be an all-things-considered judgment about which is better.

This brings us to a third worry. Sure, a skeptic about our question might say, all-things-considered judgments are possible, but there is no *mystery* as to how they are possible. Such judgments merely report the overall strength of our preferences, and we can always, as it were, *look and see* which thing we most prefer. When I judge that the grilled cheese is better than the broccoli medley with respect to cost, I am reporting a preference for the grilled cheese that has a certain strength; when I judge that the broccoli medley is better than the grilled cheese with respect to healthfulness, I am reporting a preference for the broccoli medley of a certain strength; and when I judge that all things considered, the broccoli medley is best, I am reporting that my preference for the broccoli medley is the strongest. However, as Derek Parfit and others have persuasively argued, a normative claim of betterness of reasons cannot be given by a non-normative claim about one's preferences.⁴ What I most strongly prefer is one

thing, what is all-things-considered best or what I have most reason to do or have is quite another. For our purposes, we might put the point like this: Given that I have certain preferences with different strengths, the normative question remains: how are those preferences put together to yield an all-things-considered judgment about what I *should* most prefer? The question of how disparate considerations can be normatively related so that all-things-considered judgments are possible cannot be dismissed by attempting to reduce all-things-considered judgments to facts about the strength of one's preferences.

A final point of clarification. For simplicity, I assume that the things considered are *values*, understood in the broadest sense of that term to include rights, duties, obligations, and standards of excellence, as well as their 'negative' counterparts such as ugliness and badness, although I will for the most part stick to the 'positive' values. So long as reasons are always associated with values in this sense, this assumption is harmless. Although I will refer to these values as if they were abstract values like utility, healthfulness and simplicity, it is not, strictly speaking, abstract values whose normative relations need to be explained but their particular instantiations as borne by the items about which the all-things-considered judgment is made. For example, in judging that, with respect to liberty and equality, policy #1 is better than policy #2, we need to account for the normative relations between the particular liberty and the particular equality instantiated by each policy, not that between the abstract values of liberty and equality. Moreover, the normative relations of interest are those that hold not in the abstract but relative to a set of circumstances. Our task is not to explain how a particular liberty, for example, normatively relates to a particular utility *regardless* of the circumstances but to explain how those particular values are related in a given set of circumstances. The puzzle about all-things-considered judgments is thus more precisely stated as follows: What determines how the particular values relevant to an all-things-considered judgment normatively relate in a given set of circumstances?

II.

Many philosophers assume that, even when the values at stake in a situation are very different, an all-things-considered judgment is possible. Indeed, much of normative theorizing is an attempt to figure out just which such all-things-considered judgments are correct. Although few philosophers have explicitly considered the question of what determines the normative relations among values that makes such judgments possible, they say enough to suggest one of two answers. The 'simple' answer holds that the values at stake themselves determine their own normative relations with one another. The 'sophisticated' answer holds that the values themselves are not sufficient; some additional normative item—usually a principle or purpose—is needed to account for the normative relations among them. Both views deny that there need be any more

comprehensive value to account for the normative relations of the values. I'll consider each view in turn.

So far as I know, the simple view has never been explicitly defended, but it is probably the most widely held.⁵ Its apparent plausibility may trade on its being analogous to a view about forces in the physical sciences.⁶ Physics has an easy answer to the question, In virtue of what does one physical force physically relate to another? It's *just a basic fact* about physical forces that they interact in the way that they do. If you put a certain electromagnetic force in the same room as a certain gravitational force, you will get a certain interaction between them. The way the forces interact is determined not by some third, more comprehensive force that includes both electromagnetic energy and gravity as parts but by the nature of the individual forces themselves.

The simple view similarly holds that it's *just a basic fact* about values that they normatively relate in the way that they do. If you put a certain liberty in the same circumstance as a certain equality, those values will normatively relate in a certain way. The way they relate is not determined by some third, more comprehensive value that includes that particular liberty and equality as parts, but by the nature of the individual values themselves.

I want to raise three doubts about the simple view. The first is a challenge that, strictly speaking, applies to any alternative to the comprehensive value view. Why should our understanding of all-things-considered judgments be fragmented in the way that these views suppose? We have already seen that *some* all-things-considered judgments invoke a more comprehensive value that has the values at stake as parts. Suppose you are on the admissions committee of your philosophy department and your task is to select an incoming class of graduate students. You are down to your last slot, and the choice is between Alyssa and Bertrand. Alyssa shows flair and originality but is somewhat unclear in expressing her thoughts. Bertrand shows not much flair and originality, but he is slightly more lucid than Alyssa in his writing. In all other relevant respects they are equally good or they cancel each other out. It seems clear that, with respect to flair, originality, and clarity of expression, Alyssa is better than Bertrand. That is, Alyssa is better than Bertrand *all things considered*. What determines the normative relations among Alyssa's merits and Bertrand's? The natural, intuitive answer is that these normative relations are given by a more comprehensive value, 'philosophical promise', which has flair, originality, and clarity, among other values, as parts. Alyssa is better than Bertrand all things considered because the particular values of flair, originality, and clarity that Alyssa bears make her more philosophically promising than the particular values of flair, originality, and clarity that Bertrand bears.

Compare this case with another. It's April 15th, and you are trying to figure out how to spend your afternoon. You could either do your taxes or go for a stroll in the park. Doing your taxes is better with respect to avoiding a hefty fine and doing your duty to your country; taking a walk is better with respect to having an enjoyable afternoon. Suppose with respect to avoiding financial

penalty, doing one's patriotic duty, *and* having an enjoyable afternoon, that doing one's taxes is better; that is to say, it is better *all things considered*. But if, as the simple orthodoxy assumes, no more comprehensive value has these disparate considerations as parts, we need an explanation of how this case differs from the one above. Why is there a more comprehensive value in the admissions case but not in the tax preparation case? If in some cases the values at stake normatively relate in virtue of being the values that they are, but in other cases, their normative relations are determined by a more comprehensive value, then there should be some explanation of what makes the former cases different from the latter. The challenge is to provide such an explanation.

One obvious candidate explanation is that the values at stake are 'very different' from one another in one kind of case and not so different from one another in the other. When values are not so different, they can be parts of some more comprehensive value, and when they are 'very different', they can't. But how is this notion of being 'very different' to be explained? Sometimes one thing is 'very different' from another when an attempt to relate them normatively would lead to conceptual confusion. A certain girth may be more important than a certain volume with respect to the more comprehensive criterion of bulkiness. But there is no more comprehensive criterion according to which it makes sense to say that mass is more important than color.⁷ The same, arguably, might hold for some values. Perhaps practical values cannot be normatively related to epistemic values. the value of the truth cannot coherently be said to be more or less important than the value of, say, achieving the good life.⁸ However, since values that are 'very different' in this sense can stand in no normative relations, this sense of 'very different' cannot explain why some values supposedly put themselves together while others are put together by a more comprehensive value.

Another sense of 'very different' might be thought to do the trick. Perhaps values that belong to intuitively different evaluative 'genres' or 'categories' cannot be parts of any more comprehensive value but can nonetheless put themselves together. But this suggestion, too, is problematic. Literary and musical values are very different in this sense: they belong to different 'genres' of value. But such values *can* be normatively related by a more comprehensive value, such as 'artistic excellence'. A Nancy Drew mystery is all-things-considered worse than a Beethoven sonata, where 'all things considered' is a placeholder for the more comprehensive value of artistic excellence. Perhaps there are other genres for which a more comprehensive value seems to be lacking. In this case, we need an account of 'genres' that explains why values from some 'genres', like the literary and the musical, can be put together by a more comprehensive value while values from other 'genres' supposedly cannot. And an appeal to an intuitive idea of 'genres' does not help us to make progress on this question.

The challenge to the simple view—and to any alternative to the comprehensive value view—then, is to provide a plausible way to differentiate cases in which

values supposedly put themselves together and cases in which they are put together by a more comprehensive value. In the absence of such an account, the view looks to be an unmotivated hodge-podge of intuitions about all-things-considered judgments. Sometimes it *seems* that all-things-considered judgments involve appeal to a more comprehensive value, and in these cases there is such a value; while at other times it *seems* that they make no such appeal, and in these cases there is no such value. We need an explanation for this lack of unity.

My second doubt can be stated briefly. According to the simple view, the values at stake determine their own normative relations in any given set of circumstances. 'Circumstances' are nonnormative facts like that I have a dollar in my pocket, that the average person is five feet tall, or that it is raining outside. Now consider two particular values: the low cost of the grilled cheese sandwich and the healthfulness of the broccoli medley. How those values normatively relate will depend on the circumstances. If I have only twenty dollars in my pocket, the low cost of the grilled cheese sandwich will matter more relative to the healthfulness of the broccoli medley than it would if my billfold were bursting. Is there some general explanation of the relative importance of these values in different circumstances? The simple view says no. It is just a basic normative fact that if I have twenty dollars in my pocket, the low cost of the grilled cheese sandwich counts so much as against the healthfulness of the broccoli; and that if I have thirty dollars in my pocket, it counts a little less; and if I have fifty dollars, it counts a little less still. We might naturally wish to explain these normative relations in terms of a general principle, such as 'the more money you have, the less cost matters', but the simple view denies that there is any normative consideration beyond the values at stake that accounts for their normative relations. (Indeed, the appeal to principles leads us to the sophisticated view, which we will examine in due course). In this way, the simple view amounts to an objectionable particularism, objectionable because it is commits us to an embarrassing surfeit of particular basic normative facts that seem more naturally explained in terms of more general normative considerations.⁹

The third doubt takes a bit more explaining. On what is perhaps the most natural elaboration of the simple view, all-things-considered judgments are 'one-tier': there is no intermediary that determines that only some values and circumstances are relevant to the situation. The things considered are *all the values there are*, and the circumstances relevant to the judgment are *all the extant circumstances of the universe*.

As I write this sentence, I might wonder what I should do next, all things considered. Should I write the next sentence? Should I join the Peace Corps? Take a stroll in the park? According to the one-tier conception, the considerations relevant to answering this question include *every* value there is: the value of writing the next sentence, the value of the work I could do as a relief worker, the value of petting my dog, the value of seeing a particular painting at the local museum, and so on. And the circumstances that are relevant to my decision

include *every* extant circumstance in the universe: the fact that I have a deadline, the fact that millions of children are starving around the world, the fact that the London stock exchange is up, and so on. What I should do, all things considered, is determined by the interaction among all the values in all the extant circumstances. This, the simple view says, is how it is for all-things-considered judgments.

This way of understanding all-things-considered judgments is of a piece with explanations in the physical sciences. In explaining the interaction of physical forces, every extant circumstance in the universe is part of the relevant circumstances in which those forces interact. And while a particular explanation might focus on the interaction of two particular physical forces, every force in the universe is 'at stake' in the explanation. If, for example, the earth's gravitational pull changed, the way the two particular forces interact would change. In describing how the world is, 'all things considered', *everything* is in principle relevant to the explanation.

The 'one-tier' conception of all-things-considered judgments is, however, problematic. Most obviously, it cannot account for the distinction between situations in which a value is irrelevant and situations in which it is relevant but does not affect the normative relations of the other values relevant to the judgment. Suppose you are a clothing designer and are choosing between two models to showcase your new collection. Physical beauty is relevant to the choice, but it turns out that the models are equally beautiful (perhaps they are twins). In this case, although beauty is relevant to the choice, the beauty of each model cancels out the beauty of the other, and so their beauty makes no difference to how the other values at stake—such as their poise, charisma, and photogenicity—normatively relate. Now compare the case in which you are on an appointments committee and must choose between two philosophers to fill a vacant chair in your department. Here, physical beauty is irrelevant to which you should choose. The simple orthodoxy cannot distinguish these cases. It maintains that, in both cases, beauty is relevant but fails to affect the normative relations of the other values at stake.

The same problem arises for circumstances. There is an intuitive distinction between circumstances that are 'internal' to and constitutive of a situation and those that are external and irrelevant to that situation. If I am trying to decide what to have for breakfast, the fact that I feel like having cereal is plausibly a circumstance internal to and constitutive of the situation, while the fact that the price of silk in Hong Kong has doubled overnight is external and irrelevant to the situation. The simple view cannot account for this distinction because it claims that all extant facts are relevant in each concurrent situation.

It is worth noting as an aside that the one-tier conception may help to explain why so many philosophers have, by my lights, a distorted sense of the significance of morality. If moral values are always relevant to any situation and if, as it seems, moral values have special or overriding force, then it is easy to see

how every situation can seem like a moral one. If the one-tier conception is mistaken, then whether moral values are relevant in a situation becomes an open question.

The fundamental problem with the one-tier conception is that it leaves no room for the normative question, What values *should* be at stake in an all-things-considered judgment? As I write this paragraph, I may wonder whether I should take myself to be in a situation in which I am writing a philosophy paper, where what matters is trying to make some small progress on a philosophical question, or whether I should take myself to be in a situation in which I am working to combat world hunger, where what matters is alleviating the suffering of others. Insofar as this normative question makes sense, the one-tier conception must be rejected. For according to that conception, there is no question of which values *should* matter; all values do matter in every situation, and nothing further needs to be said.

If all-things-considered judgments presuppose two normative questions, then they are 'two-tier'. First is the question, Which things *should* be considered in the all-things-considered judgment? Second is the question, Given an answer to the first question, how do the values at stake normatively relate in the circumstances? That there are two distinct normative questions here can be seen by considering that different mistakes can be made in answering each. We might criticize a colleague for being 'insensitive' when he fails to see that what should matter in discussions with students is not only his own intellectual enjoyment but serving the needs of the student. This is not the criticism that, having correctly seen that the needs of his student is at stake, he mistakenly concludes that those needs are outweighed by the enjoyment he can get from the discussion. The first error involves a failure to appreciate the appropriateness of taking certain things to matter. The second involves a mistake in discerning the normative relations among the things that do matter.

We can now see that the relation between values and physical forces is not one of analogy but the reverse. In the case of normative relations among values, there is always an intermediary that determines which values matter in the first place. The explanation of the interaction of physical forces, by contrast, involves no intermediary; there is nothing that determines that only some forces 'matter' and not others. In the explanation of physical reality, 'everything' in principle matters. Of course, sometimes physical explanation is 'idealized' by being made relative to a 'closed system'. In explaining the interaction between the earth and its moon, for example, scientists might assume that there are no other forces, such as those from the sun or distant stars, at work. But it is understood that an idealized explanation of this sort gives only an approximation to physical reality. What physicists take to be an idealization, namely relativization to a 'closed system', is the reality in the normative case. Values are different from physical forces because the normative relations among values are always relativized to a 'closed system', that is, to 'what matters' in the situation in which they are related.

A different elaboration of the simple view might take on board the idea that all-things-considered judgments are 'two-tier'. According to this alternative view, *something*—we leave open what this is—determines which values are at stake and the relevant 'internal' circumstances in which they figure. Then, once the relevant values and circumstances have been fixed, the values deemed relevant put themselves together in the given circumstances.

This 'two-tier' version of the simple view, however, must also be rejected. It assumes that once the relevant internal circumstances and values are fixed, the normative relations among those values are fixed. But the *same* particular values can be at stake in the *same* particular internal circumstances, and yet their normative relations can differ, depending on the *external* circumstances.

A toy example will illustrate the point. Suppose that, right now, there is a child in Bangkok whom only you could save from an unpleasant death. You are faced with a choice. Carry on with the routine of your life or leave the room and hop on a plane to Bangkok. Fix whatever nonnormative facts seem plausible as circumstances of the choice situation. Suppose that the values at stake are the prudential value of carrying on as usual and the moral duty to aid innocents whose life is threatened. According to the modified simple view, the duty to aid and the prudential value of carrying on as usual put themselves together in the circumstances as you have fixed them. But how those values normatively relate cannot be determined simply by appeal to the values themselves and the given circumstances. This is because the *external* circumstances can affect how the values normatively relate even if the internal circumstances are taken as fixed and given.

We might naturally think that, other things equal, the greater the prudential cost, the less normative significance the duty to aid has in relation to it. Whether this is so, however, depends on 'what matters' in the situation. Suppose that the world so far has never seen a supererogatory act, and that after you either choose to stay or succeed in saving the child, the earth will be destroyed by a wandering meteor. In these external circumstances, it might be that 'what matters' in the situation is doing something supererogatory rather than simply acting in accordance with one's moral duty. If doing something supererogatory is what matters in the situation, then the greater the prudential cost, the *greater* the relative significance of the moral duty to aid. How the moral and prudential values at stake normatively relate in the given circumstances depends on 'what matters' in that situation. In this way, the normative relations among the values at stake cannot be determined by the values themselves; something with further content is needed.

Now it might be suggested that this further content is given by the external circumstances. In this case, we need not appeal to a normative intermediary in order to account for the normative relations among the values at stake. But this suggestion is implausible. Suppose, right now, I am told all the extant circumstances of the universe. I cannot thereby determine which values are at stake in which internal circumstances, let alone how those values are normatively

related. This is because, compatible with the extant 'external' circumstances, there are many different normatively possible situations. In one possible situation so-and-so's philosophical originality is at stake; in another, the market value of the Mona Lisa; in still another, the suffering of strangers. Given the extant circumstances of the universe, we have, in some sense, a 'choice' as to which situation we 'find ourselves in'. Exactly how this 'choice' proceeds is something I set aside here. The point I want to emphasize is that the external circumstances underdetermine which values might be relevant and thus cannot themselves determine the normative relations among those values. If this is right, there must be a *normative* intermediary between the external circumstances and the normative relations among the values at stake. The simple view, which denies that anything normative beyond the values at stake themselves is required to put those values together, must therefore be rejected.

III.

So far we've seen that the values at stake cannot determine their own normative relations in a given set of circumstances. Something with further content—*normative* content—is needed to do this normative work. What could this further content be?

It might be suggested that a general normative concept like 'value' or 'what one has most reason to do' can determine the normative relations among the values at stake.¹⁰ Such general normative concepts, however, are 'category concepts', that is, formal concepts that denote groupings under which substantive normative considerations can be collected, and do not have the normative content required to *structure* the normative relations among the values that they group together. Take, for instance, the concept of *value*. This notion cannot itself provide the normative relations among particular values; it only collects certain considerations together as belonging to the category of *value* as opposed to, for example, the category of *reason* or of *color*. Similarly, *what one has most reason to do* demarcates a category of the normative in contrast to, for example, the normative category of *what one has most reason to believe*.¹¹ And just as the nonnormative category concept, *color*, does not itself determine any ranking among colors, the normative category concepts, *value* and *what one has most reason to do*, do not determine any ranking among values or reasons for action. Thus the further consideration that puts together the values at stake cannot be a category concept but must itself have substantive normative content.

If this is so, the question then becomes, Is this consideration given, as the sophisticated orthodoxy would have it, by something such as a principle or purpose that is *not* a more comprehensive value, *or* is it given by a more comprehensive value? I want to suggest that no normative consideration could plausibly account for the normative relations of the values at stake without presupposing a more comprehensive value.

We might begin by asking how a principle, purpose, or other non-value *could* determine the normative relations among the values relevant to an all-things-considered judgment. In particular, what content must it have in order to do this normative work?

One suggestion is that this consideration has as its content the particular normative relations among the values at stake that it is supposed to determine. Our purpose in choosing between two philosophers, for instance, might be to get a particular weighting of originality, insightfulness, clarity, historical sensitivity, and so on. Or a principle governing duties to aid might simply be a collection of particular normative relations among the duty to aid and competing values in different circumstances. The normative consideration is itself nothing but a collection of the particular normative relations it determines.

This suggestion, however, will not do. It precludes the possibility of genuine disagreement about how the values at stake normatively relate given agreement on what principle, purpose, or other normative consideration applies to the situation. On this view, if you and I disagree over which philosopher to appoint, our disagreement cannot be a genuine disagreement over whether one candidate's originality should count so much as against another candidate's historical sensitivity; it can only be a disagreement about which normative consideration—e.g., principle or purpose—applies in the situation. But you and I should be able to agree that a certain principle or purpose applies to the situation and still *disagree* about which of two items we should choose according to that principle or purpose. This would not be possible if the contents of these considerations were simply given by the very set of normative relations among the values at stake. If a normative consideration is to determine the normative relations among the values at stake, it cannot simply have as its content the relations it is supposed to determine.

Indeed, a normative consideration such as a principle or purpose could plausibly determine the normative relations among the values at stake only by presupposing a more comprehensive value. Consider purposes. On the face of it, a purpose seems capable of determining which values are relevant as well as what the normative relations are among them. Suppose I put you in a room with Aye and Bea and tell you that your purpose is to appoint a philosopher. It may seem that you now know that their originality and historical sensitivity are relevant to the choice and that an impressive originality beats a banal historical sensitivity. Strictly speaking, however, you don't know any of this. That is because the purpose, 'to appoint a philosopher', has no normative content. As we have already seen, for a normative consideration to do the required determining work, it must itself have normative content. Once purposes are understood as normative goals, such as 'to get the *best* philosopher', it seems that we can determine not only which values are relevant but how they are normatively related.

But it is unclear how a normative purpose, such as 'to get the best philosopher', can determine the relative importance among the competing values of

originality, historical sensitivity, and so on, *without* presupposing a more comprehensive value, such as philosophical talent or goodness as a philosopher, that unifies them. A normative purpose can determine the values relevant to achieving the purpose, but if it is to account for the normative relations among those values, it must presuppose something that can put those values together. Purposes operate as 'pointers' to more comprehensive values that do the work of putting together the component values at stake. And it is in virtue of these more comprehensive values that one alternative better achieves the purpose than another.

The same goes for principles. Take the principle, 'Other things equal, one ought to keep one's promises.' How can such a general slogan determine the normative relations of very particular values at stake in very particular circumstances? Those who appeal to principles allow that the operation or content of a principle depends on complex background claims about when different circumstantial features affect the relative importance of the values at stake. So, for instance, Tim Scanlon's contractualism holds that it is a background 'structure of understanding' that determines when the cost of keeping a promise gets high enough to count against keeping it, and Frances Kamm's exploration of moral principles governing permissible harming relies on fine-grained background claims about when particular circumstances affect the moral strength of a duty not to harm.¹² If a principle is to be capable of accounting for the normative relations of particular values across many different circumstances, it must rely on these 'background' claims. But how are these background claims determined? Why, for example, should the background claim about promises be that when the cost reaches a certain degree, one has less reason to keep the promise rather than more reason? Are these background claims just part of the content of the principle itself? In this case, the principle would be given by a collection of particular normative relations among various values in particular circumstances, and disagreement over what these particular normative relations are, given agreement on the content of the principle, would be impossible. The problem is that a normative consideration that is capable of determining the normative relations among values in a variety of circumstances must have content apart from these relations that determines them. It is hard to see how a principle could have this content without presupposing a more comprehensive value. More comprehensive values have a 'unity' in virtue of which the normative relations among their component values can be determined in different circumstances. This unity is that in virtue of which, say, a particular cost counts more than a particular moral duty in one set of circumstances, while the very same cost counts less—or less than it did—in another set of circumstances. A similar argument can be made for any kind of normative consideration that is not itself a more comprehensive value.¹³

Now it might be suggested that there is no need to presuppose a more comprehensive value if it is just a *general* basic normative fact that certain values stand in certain normative relations across different circumstances. That is, the

sophisticated orthodoxy might be understood as maintaining that a principle, purpose, and so on, amounts to a general basic normative fact that plays exactly the same determining role as a putative more comprehensive value. In answer to the question, how do principles, purposes, and so on, account for the normative relations among values at stake?, it might be insisted that it is a general basic normative fact that certain values stand in certain relations in a variety of circumstances, and that is all there is to be said. This view does not amount to an objectionable particularism because the basic normative facts are at the same level of generality as more comprehensive values and, indeed, are supposedly that in virtue of which the particular facts about how particular values relate in particular circumstances hold. But the view does nevertheless collapse into a version of the simple orthodoxy, for if it is just a basic normative fact—at whatever level of generality—that values stand in certain normative relations with one another, then it is ultimately the values at stake that put themselves together in different circumstances. And we have already argued that the simple view is mistaken: either it is committed to the problematic one-tier conception of all-things-considered judgments or it mistakenly assumes that two particular values will always have the same normative weights in the same internal circumstances, regardless of ‘what matters’ in the choice. In any case, claiming that the normative relations among values at stake is a matter of a basic normative fact throws in the towel too early. It has us hitting normative bedrock when there are further philosophical resources—more comprehensive values—available to do the required explanatory work.

The sophisticated orthodoxy, therefore, is problematic on two counts. It needs to explain why some all-things-considered judgments proceed by appeal to a more comprehensive value while others do not. Without such an explanation, it leaves us with an unjustifiedly fragmented account of such judgments. More importantly, it fails to recognize that an explanation of the relations among values requires more than merely gesturing at principles, purposes, or other nonvalues; there is the further question as to *how* such considerations put values together, that is, what contents they must have in order to do this normative work. If we ask what familiar normative consideration could do this work, it seems that only *values* can fit the bill.

Values can play the role of determining which values are at stake and the normative relations among them because they have a *unity* in virtue of which some values and not others are their components and in virtue of which those component values hang together in the way that they do. Take philosophical talent again. It is in virtue of the unity of philosophical talent that physical attractiveness is irrelevant to philosophical talent, that a particular originality makes one more philosophically talented than does a particular historical sensitivity, and that you and I might have a genuine disagreement about whether technical prowess makes someone more or less philosophically talented than someone with an understanding of the historical sweep of philosophical ideas. You and I can agree that philosophical talent gives what matters in a situation,

but disagree about the relations among its component values in various circumstances. It is this *unity* of values that explains how the values at stake are normatively related as they are and thus how all-things-considered judgments are possible.

IV.

I suspect that most of the resistance to the idea that more comprehensive values unify the values at stake has its source in the difficulty of explaining what it is about a value in virtue of which its component values are structured as they are—*what is this 'unity' in virtue of which its components hang together in the way that they do?* Put another way, what makes a *value* different from a *mere collection* of normative relations among values? I think this question raises a profound mystery, one that is probably connected to the mystery of what makes certain concepts hang together while others do not. I am not going to attempt to say in what the unity of a value consists. Instead, I'll do what philosophers always do when they get stuck: offer a metaphor and hope that it's illuminating.

Values, I believe, are like pieces of a jigsaw puzzle, but not any old jigsaw puzzle. Some jigsaw puzzles are put together by a unifying 'picture'; the puzzle, when completed, depicts a pastoral scene, or a frigate on the high seas, or a pair of playful puppies. One piece of the puzzle fits next to another in virtue of the picture which determines how the pieces are put together. When values come together in virtue of a 'picture' that relates them, they form a more comprehensive value, and it is in virtue of this 'picture' that they are normatively related as they are. Other jigsaw puzzles are put together differently; there is no 'picture' determining how the pieces fit together but only a depiction of something homogenous like the color red or a bumble bee's honeycomb. The pieces fit together simply in virtue of their shapes interlocking in the right way—they put themselves together. If my arguments in this paper are correct, values cannot be like the pieces of a jigsaw puzzle of this second kind. Values cannot put themselves together simply by their own 'shape'; something with further content—a 'picture'—is needed.

The 'picture' that puts values together is the unity of a more comprehensive value. Although it is hard to explain just what this 'picture' is, it is important to emphasize that the mystery of what makes values hang together is not peculiar to the more comprehensive value approach. Any normative theory that recognizes values is saddled with the problem of explaining their unity; the 'problem of the unity of values' is a problem even for ordinary values such as 'beauty' and 'philosophical talent'. We have no account of what it is about such values in virtue of which their component values hang together in the way that they do. But this does not block the thought that they are values nonetheless, and so it should not ground skepticism about the more comprehensive values of interest.

Nor should the fact that many of these more comprehensive values are *nameless* be seen as providing evidence against their existence. As we noted at the outset, many ordinary values that have names today were nameless not long ago. A value can do normative work even without the benefit of a name.

Indeed, the idea that some values are nameless can be traced back to Aristotle. In identifying the virtue and vice concerned with the pursuit of small honors, for example, Aristotle held that the mean between the extremes of ambitiousness and unambitiousness has no name. Similarly, he thought that the regulation of feelings of anger involved a nameless virtue and vice; at one extreme is an excess of anger, which is nameless, and at the other extreme is a deficiency in anger, which is nameless, and the mean between these two extremes is also nameless.¹⁴ I believe that Aristotle was right in pointing out that many perfectly ordinary virtues and vices have no names; my proposal extends Aristotle's insight to values generally.

Another worry about nameless value might be that, because we cannot explicitly articulate their contents, they must be an illusion. But the fact that we cannot explicate the content of a concept does not mean that there is no such concept or that we do not possess it. Newton worked with the concept of the limit of series, which remained nameless for some two centuries, without being able explicitly to give its content.¹⁵ So just as Newton possessed the concept of a limit, ordinary thinkers today may possess and employ concepts associated with nameless values

Nameless values typically suffer from a high degree of epistemic and semantic indeterminacy, but this should not be thought to show that they do not properly exist as values. Many named values suffer from both kinds of indeterminacy, e.g., 'justice', and yet we have no doubt that those values exist. And although nameless values will typically be more indeterminate than named ones, it is not clear how this difference in degree can be parlayed into an argument that nameless values are not on all fours with their named siblings.

Perhaps the most potent doubt about nameless values lies in the suspicion that they are in some way fake—they are 'Frankenstein' values, artificially stitched together to satisfy the mad cravings of those seeking unity where none is to be found.¹⁶ Or, to put the worry differently, once we allow nameless values, the floodgates are open for values to be cobbled together any-old-how. But this worry overlooks the fact that values cannot be merely stipulated combinations of any values whatever. As we have seen, nameless values have content beyond a mere collection of the relations among their component values, and it is in virtue of this content that its component values are normatively related as they are. I have suggested that this content is given by the unity of the value: the 'picture' in virtue of which its components hang together in the way that they do. Sometimes values come together to form a unity that is a more comprehensive value, and sometimes they do not. Why this is so remains a deep axiological mystery, but the fact that they sometimes do is not subject to doubt.

There is a test that helps to distinguish more comprehensive values from these Frankenstein imposters. Suppose I put you in a room with a corkscrew, a quadratic equation, and plate of sea urchin sushi. Your task, I go on to explain, is to judge which item is all-things-considered best, where the values that matter to the judgment are the utility of the corkscrew, the mathematical beauty of the equation, and the taste of the sushi. You would not even know where to begin your deliberations. This is because there is no more comprehensive consideration that includes the values at stake. If you then say to me, 'The equation is best because its abstract beauty is ten times more important than the particular taste of sea urchin sushi or the particular efficiency of the corkscrew', you will have created a Frankenstein imposter. Some considerations look like Frankenstein imposters but are perhaps best regarded as attempts to codify a genuine nameless value. Suppose you are in a room with Anastasia and Beatrice, and your task is to judge which of the two is best, all things considered. The things considered turn out to be talent, community spirit, articulateness, beauty, and.... poise in a swimsuit. At first blush, it seems that you wouldn't know how to begin your deliberations. But, of course, every year the judges of the Miss America beauty pageant face this very task. They can engage in deliberations in part because there is an apparent rule of comparison that provides a formula for judging what it is to be best with respect to these criteria. But what looks like a stipulated formula may instead be an attempt to codify a certain (dated) notion of femininity that unifies beauty, poise, talent, and so on.¹⁷ We know that there is an underlying value at work if, for instance, someone who believes that poise in a swimsuit should count for 90 percent of the overall judgment can be said to be making a normative mistake rather than simply stipulating a different weighting. The difference between a Frankenstein imposter and a genuine value, then, is that only the latter has a unity that determines how its component criteria *should* be put together.

Our situation today with respect to more comprehensive nameless values is, I believe, akin to our situation in the 1950s with respect to sexual harassment. Imagine that the year is 1950, and two secretaries, Ann and Betty, are whispering around the office water cooler about the behavior of their bosses. They make various all-things-considered judgments about the badness of their bosses' behavior. They might agree, for example, that, *all things considered*, demanding sexual favors in exchange for keeping one's job is worse than creating a work environment that treats sexually predatory behavior as a joke, and disagree about whether putting pinups in the common room is all-things-considered worse than patting women on the head for a job well done. The things considered in their all-things-considered judgments include sexual exploitation, lack of respect, sexual domination, condescension, chauvinism, and so on. They would be skeptical of the claim that there is a more comprehensive value that unifies these considerations, for sexual harassment did not then have a name; indeed, they would insist that it *just seems* to them that, for example, *this*

instance of sexual exploitation is worse than *that* instance of chauvinism, and there is nothing more to be said.

Fast-forward to today where Arthur and Bob are whispering around the department coffee machine about candidates for a philosophy post. They might agree that, *all things considered*, Cassandra is better than Dirk, but disagree about whether Dirk is better than Eliza. The things they consider in their all-things-considered judgments include philosophical talent, teaching ability, curricular 'fit' with the department, community spiritedness, research productivity, collegiality, and so on. Like Ann and Betty, they will be skeptical of the claim that there is a more comprehensive value that unifies philosophical talent, teaching ability, community spiritedness, and so on. But just as Ann and Betty were mistaken in failing to see that there is a more comprehensive value of sexual harassment in virtue of which their judgments are correct or incorrect, so too Arthur and Bob may be failing to see that there is a more comprehensive value with a unity that structures philosophical talent, teaching ability, community spiritedness, and so on, in virtue of which their judgments are correct or incorrect. I believe that once we begin to focus our attention on such more comprehensive values, we will be in the same position with regard to these nameless values as we are today with regard to what were once nameless values.

Let me close with a general remark about the upshots of the comprehensive value view for our understanding of the normative domain. If the comprehensive value approach to understanding all-things-considered judgments is correct, then we have an important conclusion about the connection between values and reasons. All-things-considered judgments about what is better than what form the basis of what we should do, want, and believe. And if, for those judgments to be possible, there must be a more comprehensive value that determines how the things considered normatively relate, then it is *values* that tell us what reasons we have, all things considered.

Notes

* Many thanks to Kit Fine for very helpful comments on an earlier draft, to Derek Parfit for lengthy and incisive discussion of several key points, and to audiences at Harvard University and the University of California at Davis for useful discussions. Members of those audiences whose comments I remember as being especially helpful include Arthur Applbaum, Frances Kamm, Leonard Katz, Erin Kelly, Niko Kolodny, Mike Martin, Alison McIntyre, Karen Neander, Derek Parfit, Tim Scanlon, and Dennis Thompson. Some of the themes of this paper are discussed in a different setting in the companion paper, 'Putting Together Morality and Well-Being', in Peter Baumann and Monika Betzler, eds., *Practical Conflicts: New Philosophical Essays* (Cambridge: Cambridge University Press, 2004).

1. Talk of 'parts' should not be taken to signal any heavy metaphysical weather. One value is 'part' of another if bearing the one contributes constitutively to bearing the other.

2. There is an argument that rules of comparison *do* involve a substantive judgment of the sort of interest here. See my 'What is Incomparability?', draft ms. Insofar as they do, they fall under the scope of the present argument. My concern here is not with instrumental all-things-considered judgments. In such cases, 'all things considered' does not strictly refer to the things considered—the values at stake—but to their instrumental role in achieving something else. I do, however, mean to include cases of 'symmetrical conflict' in which the same exact value is at stake but in conflict with itself because, for example, both instantiations cannot be achieved given the way the world is. In this case, we have a degenerate case of the 'more comprehensive' value as being identical with the values at stake.
3. For an argument that Sidgwick was mistaken, see my 'Putting Together Morality and Well-Being'.
4. See e.g., Joseph Raz, 'Incommensurability and Agency', in R. Chang ed., *Incommensurability, Incomparability, and Practical Reason* (Cambridge: Harvard University Press, 1997), pp. 110–128; Tim Scanlon, *What We Owe to Each Other* (Cambridge: Harvard University Press, 1998), ch. 1; and for the most detailed and thorough defense, Derek Parfit, *Rediscovering Reasons*, draft ms, chs 1–2. Cf. my 'Can Desires be Reasons for Action?', in *Reason and Value: Themes from the Moral Philosophy of Joseph Raz* (Oxford: Oxford University Press, 2004).
5. It has roots in both classical utilitarianism and deontology.
6. Many thanks to Shelly Kagan who first pointed out this analogy to me in the course of defending what I am calling the simple view.
7. Recall that, as I am understanding all-things-considered judgments, 'importance' is a constitutive, not an instrumental notion.
8. I owe this suggestion to Derek Parfit.
9. Cf. Jonathan Dancy for a defense of this sort of particularism, most recently in his 'The Particularist's Progress', in Hooker and Little eds., *Moral Particularism* (Oxford: Oxford University Press, 2000), pp. 130–156. The more comprehensive value approach provides an explanation of the particularistic facts that particularism holds are bedrock. For the related point that the way in which a particular circumstantial factor can affect the normative relations among values in different ways in different circumstances, see Frances Kamm, *Morality, Mortality*, vol. II, pp. 51–68 and "Killing and Letting Die: Methodology and Substance," 64 *Pacific Philosophical Quarterly* 1983: 297–312; and Shelly Kagan, "The Additive Fallacy," 99 *Ethics* 1988: 5–31. The more comprehensive value approach offers an explanation of *why* circumstantial factors can affect the normative relations among values depending on what other circumstances are present, *viz.*, because the comprehensive value that gives what matters may vary from one circumstance to another.
10. James Griffin moots this suggestion in thinking about the normative relations between particular rights and the social good. See his forthcoming book on human rights.
11. This is not to say that there are no relations among these categories but only that the domain of the normative can be divided into different categories in various ways. It is worth pointing out that this distinction between category concepts and more comprehensive value concepts shows that *values* are doing important

work in the proposed view. If my argument is right, it is not enough to explain how all-things-considered judgments are possible to say that there is some general concept, such as *what one has most reason to do*, that 'covers' the things considered. There must be a concept with the *unity* of a value at work. Thanks to Derek Parfit and Tim Scanlon for raising an objection that led me to see the need for this distinction.

12. See Scanlon, *What We Owe to Each Other*, 197–201, 50–54; and Frances Kamm, *Morality, Mortality* (Oxford: Oxford University Press, 1996), 51–68.
13. I focus on principles and purposes because I believe that other sorts of normative consideration—policies, a theory of value, and so on—ultimately rely on principles or purposes as they are intuitively understood.
14. Aristotle, *Nicomachean Ethics*, tr. Terence Irwin (Indianapolis: Hackett, 1985), Book IV, 1125b. Indeed, Aristotle repeatedly underscores the point that *many* virtues and vices have no names (1107b, 1108a, 1115b).
15. I take this example from Christopher Peacocke, "Implicit Conceptions, Understanding and Rationality", *Philosophical Issues* 9:43–88, who argues that an inability to explicitly articulate the content of a concept does not show that one lacks the concept.
16. Thanks to Stephen Robert Grimm who pressed me to say more about this worry in the companion paper, 'Putting Together Morality and Well-Being' and who offered, in the course of expressing his skepticism, the term 'Frankenstein values'.
17. Judgments about the best decathlon athlete might be another case in which there is an underlying more comprehensive value, such as 'best all-around athlete', at work; such judgments seem to proceed according to stipulated formula of comparison, but this formula may be guided by a notion of all-around-athleticism. Thanks to Frances Kamm for suggesting that the beauty pageant example is one in which, like the decathlon case, there is an underlying value at work.