

When does ‘Can’ imply ‘Ought’?

Stephanie Collins

This is the final draft of a paper whose definitive version will be published in

International Journal of Philosophical Studies

According to Stan Lee, the author of Spiderman, “with great power there must also come -- great responsibility.” (Lee 1962) In other words: ‘can’ implies ‘ought.’ Left on its own, this principle has the potential to be incredibly demanding. The natural move is to add a ‘at not too high a cost’ rider. Then we get something like: if you have the capacity to fulfil important interests, at not too high a cost, then you have a moral duty to do so. I will call this the ‘assistance principle’ and the duties it posits ‘assistance duties.’ The assistance principle is simple and appealing, as Singer’s famous ‘drowning toddler’ example demonstrates (Singer 1971).¹ Variants on this principle have been applied to issues as diverse as climate change (Caney 2010), humanitarian intervention (Pattison 2010), and, most prominently, global poverty (Singer 1971, 2009).

This paper will consider how the concept of sacrifice should feed into our interpretation, and implementation, of the assistance principle; that is, which kinds of sacrifice are relevant to answering the title question and limiting the demandingness of assistance duties. I will make two suggestions for how we should conceptualise sacrifice so that it appropriately limits the applicability of the assistance principle. Both suggestions challenge the view, implicit in much literature on sacrifice, that sacrifice is always a sacrifice of the *agent* or *duty-bearer*. Others’ sacrifices, and other kinds of sacrifice, are, I argue, relevant to determining the existence and strength of assistance duties.

First, I will focus on the ‘at not too high a cost’ clause. I will suggest that the costs relevant to this clause are the *net costs* to both the *capacity-bearer* and the *interest-bearer* – as opposed to the gross costs, and/or the net costs to the capacity-bearer alone. That is, we should distinguish (net) ‘agent-relative sacrifice’ from (net) ‘recipient-relative sacrifice.’ For example, if I can rescue the drowning toddler at the cost of my shoes, then this is a gross cost to me. But if I will experience a warm glow of self-satisfaction, then this is a gross benefit to me. By subtracting my gross costs from my gross benefits, we arrive at the ‘net agent-relative sacrifice’

¹ Of course, some disagree. Jan Narveson (2003), for example, claims that you have done nothing *wrong* by defaulting in Singer’s example. But even he agrees you have done something *very bad*.

of my action. Likewise, if I rescue the drowning toddler, this is a gross benefit to him. But if I give him concussion while doing so, then this is a gross cost to him. By subtracting the toddler's gross costs from the toddler's gross benefits, we arrive at the 'net recipient-relative sacrifice.' I argue that both kinds of sacrifice—net agent-relative and net recipient-relative—should feed into the 'at not too high a cost' clause.²

Second, I will suggest that we should conceptualise sacrifice (as it's relevant to the assistance principle) such that our interest-fulfilling actions can *sacrifice* the better actions of others. I call this 'ideal-relative sacrifice': in acting to fulfil an important interest, we sometimes sacrifice the more ideal actions that others would have taken to fulfil that very interest. These actions might be 'more ideal' in the sense of (i) *imposing lower costs on the interest-bearer* (i.e., the more ideal action would fulfil the interest-bearer's important interest more fully or securely, or would fulfil other of the interest-bearer's other interests, alongside the important interest that generated the duty) or in the sense of (ii) *imposing lower costs on the agent who performs them* (compared to the costs our own interest-fulfilling actions would have imposed on us). For example, suppose my rescuing action would give the toddler concussion, but your rescuing action would not. If I then go ahead and rescue the toddler (when you would have done so had I not), then I have *sacrificed* your more ideal action. In this case, your action is 'more ideal' in the sense that it *imposes lower costs on the interest-bearer*. But your action might (instead or also) be 'more ideal' in that the sense that it imposes lower costs on its performer. For example, if you are wearing cheap gumboots, while I am wearing expensive high heels, then my rescuing action involves an 'ideal-relative sacrifice' insofar as your action imposes lower costs on the agent that performs it.

If some interest-fulfilling action will sacrifice the more ideal actions of others – that is, if the action will imply 'ideal-relative sacrifice' – then the moral reason to perform that action is weaker than it otherwise would have been. Acknowledging that our interest-fulfilling actions can involve ideal-relative sacrifice allows us to see that the assistance principle demands less from us than we might initially have thought. It also gives the attractive verdict that well-meaning but uninformed agents should not rush to 'help' in contexts where their actions are not those most needed.

² When I refer to 'agent', I mean 'the entity (individual or collective) with the duty in a particular context.' This label is always context-relative: everyone who's an agent in one context has been, and could yet be, a recipient in another context; and most entities who are recipients in one context have been, or could yet be, agents in other contexts (non-human animals aside).

In sum, then, there are three kinds of sacrifice that are relevant to assistance duties: agent-relative sacrifice; recipient-relative sacrifice; and ideal-relative sacrifice. The size of the agent-relative and recipient-relative sacrifices determines whether the ‘at not too high a cost’ clause is met, and thus helps determine whether there is an assistance duty at all. Assuming there is an assistance duty, the size of the ideal-relative sacrifice helps determine the strength of the duty, that is, how easily the duty can be outweighed by other, competing duties that are held by the duty-bearing agent.

I assume that a high level of sacrifice is *sufficient* to undermine, and to weaken, *assistance* duties. Importantly, this assumption is neutral on many other issues in the philosophy of sacrifice. For example, I will not assume that a high level of sacrifice (of any, or all of, the three types) is *necessary* for any duty – including any assistance duty – to be undermined or weakened. I will also not assume that a high level of sacrifice is necessary for an act to be supererogatory. Nor will I assume that a high level of sacrifice is *sufficient* to undermine or weaken duties *other than* assistance duties. After all, some highly sacrificial acts are nonetheless required by duty, and some non-sacrificial acts are not required by duty, or are supererogatory. My question is a focused one: how should we understand the ways in which sacrifice blocks, or weakens, assistance duties? My answer is that there are three kinds of sacrifice at issue. These kinds of sacrifice circumscribe the contexts in which ‘can’ implies ‘ought.’ Factors other than sacrifice might *also* circumscribe these contexts.

To make these points, I will develop two versions of the assistance principle: the well-placed principle and the best-placed principle.³ The well-placed principle states that if you are *sufficiently* capable of fulfilling someone’s important interests, at not too high a cost, then you have a moral duty to do so. The ‘best-placed’ principle states that if you are *most* capable of fulfilling someone’s important interests, at not too high a cost, then you have a *strong* moral duty to do so. (‘Strong’ refers to how important a countervailing reason of yours would have to be to outweigh the duty in question.) Both these principles contain the ‘at not too high a cost’ clause, so I will analyse that clause in a way that applies to both principles, including explaining how both agent-relative and recipient-relative sacrifice bear on that clause.

³ Others have developed versions of the assistance principle (e.g. Singer 1971, 231; Singer 2009, 15; Unger 1996, 12; Goodin 1985, 118; Scanlon 1998, 224). These authors’ formulations tend to be brief – usually only a sentence long – and do not distinguish between the well-placed and best-placed versions. I developed these two principles in some detail in [omitted], but there I did not consider the implications for conceptualising sacrifice, did not separate the three kinds of sacrifice, and did not consider how the notion of sacrifice can explain the two principles’ differential strength.

But the two principles come apart on the question of whether, by satisfying the principle's antecedent, you sacrifice someone else's more ideal actions. If you satisfy the conditions of the 'well-placed' version, then your interest-fulfilling action might sacrifice someone else's more ideal actions (i.e., might involve ideal-relative sacrifice); if you satisfy the conditions of the 'best-placed' version, then your interest-fulfilling action does not have this sacrificial implication. I suggest that, because the best-placed version does not contain this sacrificial implication, the duties it produces are stronger than those of the well-placed principle. This suggestion accords with the common-sense idea that actions implying little sacrifice are more readily morally demanded than actions implying great sacrifice – but the sacrifice at issue here is the sacrifice *of someone else's more ideal action*, not the interests of the agent alone. Thus, the notion that we can sacrifice others' better actions through our good actions implies that the two versions of the principle are differently demanding.

The argument will take the following structure. I start by assuming that well-placed duties and best-placed duties are two species of the genus 'assistance duties.' I assume we can break down the duty-generating relation at issue in (both species of) the assistance principle into five rough components. Components 1 to 4 are the individually necessary and jointly sufficient conditions for a 'well-placed duty,' while components 1 to 5 are the individually necessary and jointly sufficient conditions for a 'best-placed duty.' The first four components are: (1) there is an *important interest* that is unfulfilled; (2) someone is *sufficiently capable* of fulfilling that interest; (3) that agent's most efficacious measure for doing so is *not too costly*; (4) the agent could do *similarly in similar circumstances*. These components generate a 'well-placed duty'. For the relation to generate 'a best-placed duty', the fifth component must be added: (5) the agent's fulfilling the interest would be *as or less costly* than any other agent's doing so. In §1-2, I briefly explain the first two components. In §3, I explain the third component, arguing that the relevant 'costs' are the net agent-relative *and* net recipient-relative sacrifices, and describing how the fourth component arises out of discussion of the third one. In §4, I explain the fifth component. Here I motivate the idea that well-placed agents can make ideal-relative sacrifices – they can sacrifice the more ideal actions of other agents – and that this sacrifice is what makes well-placed duties weaker than best-placed duties. By developing components three, four, and five, I will demonstrate that a broad and tripartite conception of sacrifice – one that includes agent-relative, recipient-relative, and ideal-relative sacrifice – should be used in interpreting and applying the assistance principle.

1. ‘An Important Interest is Unfulfilled’

First, to establish the scope of the discussion, it will help to know about the interests that are at issue in the well-placed and best-placed principles. I label these “important interests.”⁴ To picture these interests, we can imagine a continuum of interests, from trivial to important. Which interests are further towards the “important” end of the continuum will depend on one’s substantive theory of welfare, on which my formulations of these principles is neutral. For example, in an article outlining a variety of duty-generating principles, David Miller (2001, 453) focuses on “instances of deprived or suffering people—people whose basic rights to security, or subsistence, or health care are not being protected, and who as a result are in no position to live minimally decent lives.” He suggests a version of the assistance principle—a version according to which “responsibilities ought to be assigned according to the capacity of each agent to discharge them”—as one way of distributing duties to fulfil these interests. But Miller’s characterisation is just one possible characterisation of important interests. Others might refer to basic rights (Shue 1996) or basic capabilities (Nussbaum 2006).

Important interests might be indexed in various ways. For example, they might be indexed to the agent who fulfils them. I might have an important interest not just in being loved, but in being loved *by my parents* (Liao 2006; Keller 2006). This will become important for the two species of the assistance principle: if my interest is in being loved by my father, then only my father is eligible for an assistance duty to fulfil that interest. This means that his duty will be grounded in the best-placed version of the principle (not merely the well-placed version), and thus that his duty is relatively strong, since his interest-fulfilling action will not sacrifice the better interest-fulfilling actions of others – or so I shall argue in §4.⁵ Thus, the possibility of interests that are indexed to a fulfiller increases how often the best-placed principle applies.

Despite all this scope for detail, when identifying assistance duties, interests should be defined as generally as possible while capturing all that is of value in them. For example, if what matters is just “that I am fed tonight,” then we should not split hairs between my interest in “getting fed tonight by A” and my interest in “getting fed tonight by B.” This is because

⁴ The idea of X’s important interests may seem to evoke Raz’s (1986, 166) idea of “an aspect of X’s well-being (his interest) [that] is a sufficient reason for holding some other person(s) under a duty.” I will suggest that an (unfulfilled) important interest is merely necessary, not sufficient, for an assistance duty: there are other necessary conditions regarding cost to the duty-bearer, the duty-bearer’s capacities, and so on. Assistance duties are thus perhaps what Raz calls a “conditional duty” (1986, 167-168). Though what I say is consistent with Raz’s general account of interests, rights, and duties, my account isn’t meant to imply anything about non-assistance duties and the rights associated with them. It’s thus more specific than Raz’s (1986) account.

⁵ Of course, other agents might have duties to (take measures to) make my father capable of loving me. But these other people’s duties would be duties to fulfil my interest in “having parents that are capable of loving me.” This is a different interest, which might generate well-placed or best-placed duties for those people.

assistance duties include not just duties in virtue of being *well-placed* to fulfil some interest, but also duties in virtue of being *best-placed* to fulfil some interest. And if there is no morally important difference between these interests, then our account of best-placed duties should not pointlessly say A and B each have a *best-placed duty* over these different interests, based on their respectively being best-placed those agent-indexed interests. (At least, our account should not say this if we assume full compliance—a point I will address in §IV). We can say that they each have a *well-placed* duty to do so. But if they are both well-placed to meet the non-agent-indexed interest, then we do not need agent-indexed interests in order to say this. Assuming they are not equally best-placed to fulfil my interest in *being fed tonight*, only one of them should have a duty to fulfil that interest based on being *best-placed* to do so. This prevents a situation in which *all* assistance duties are best-placed duties. If all assistance duties were best-placed duties, then we would not be able to capture the intuitive differences (on which, more below) between duties in virtue of being well-placed and duties in virtue of being best-placed.

We now have the following. Subject *S* has a well-placed or best-placed duty only if:

- (1) Person *P* has an important interest, *I*, that is unfulfilled, where the interest is described in a way that includes all and only the morally relevant facts and indexes.

2. Someone is Sufficiently Capable of Fulfilling It

I will assume that assistance duties are duties to take some measure (action or omission) that has at least some likelihood of fulfilling an interest (the idea of “important interest” will re-emerge shortly). I will use “measure” assuming agents have full control over measures they take.

For an agent to have a duty to take a measure *in virtue of that measure’s propensity to fulfil an interest*, the measure must be sufficiently likely to fulfil that interest. How likely is sufficiently likely? It seems strange to believe in a duty to take some measure—in virtue of that measure’s propensity to fulfil an important interest—if the measure has only, say, a 0.0001 likelihood of succeeding. Yet in some cases—if the interest is important enough—an agent can have a duty to take measures that have only a tiny likelihood of fulfilling an interest. If the interest at stake is “all persons’ interests in being alive next week” and A could take the measure “firing a rocket at the giant asteroid that, if not hit by a rocket, will hit Earth and kill everyone this weekend,” then A might have a duty to take those measures, even if the measures are unlikely to fulfil all those interests. There are measures he can take that will *possibly* fulfil the

interests, so he can have a duty to take those measures, where the duty is grounded in the importance of the relevant interests.

This suggests that we should not a precise threshold—say, 50% or 70%—for how likely the measure is to fulfil the interest. The reason for this is that any such threshold would be somewhat arbitrary and would rule out duties in ‘high-stakes, low-likelihood’ cases like that just described. Instead, I suggest, for an agent to have an assistance duty to take a measure, the likelihood that the measure will fulfil the interest must be *proportionate* to the importance of the interest. The less important the interest, the more likely it must be that some measure will fulfil it, if that measure’s likelihood is to be proportionate to the interest’s importance. Conversely, for an extremely important interest, a measure might have a low likelihood of fulfilling it (if the measure is taken), and yet it might be proportionate. Admittedly, this brings a consequentialist element into the assistance principle. As we shall see in §3.B, though, this will not make assistance duties consequentialist through-and-through. Since my main aim in this paper is to characterise the nature of the *sacrifices* at issue in the assistance principle, I lack space to fully defend this aspect of (my characterisation of) the principle. If you would prefer to set a uniform likelihood threshold that all duty-generating interest-fulfilling measures must meet, then you can make the necessary substitutions for the rest of the paper.

For simplicity, and following earlier work (Collins 2015, 107-8), I will use “A’s likelihood” to mean “the likelihood that the interest will be fulfilled if A takes his *most efficacious measure*.” This is the measure of A’s that is most likely (of all A’s measures) to fulfil the interest, if A takes it. And I’ll use “success” to mean “the important interest is fulfilled.”

My suggestion, then, is that for A to have an assistance duty grounded in some interest, A’s likelihood of success must be in the range of likelihoods that are *proportionate* to the importance of the interest, where more important interests are proportionate to a range of likelihoods whose lower bound is lower. If A’s likelihood of success is proportionate to the interest (or if it meets the reader’s preferred likelihood threshold), then we can say A is *sufficiently capable* of fulfilling the interest to bear an assistance duty to take the relevant measures.

Thus we have the second necessary condition for assistance duties:

- (2) If: *S* takes measure *M*, where *M* is the most efficacious measure open to *S* to fulfil *I*; then *I* will be fulfilled with a likelihood that is *proportionate* to *I*’s importance, where

more important interests are proportionate to a wider range of likelihoods (with a lower bound that is lower).⁶

3. The Action is Not Too Costly

3.1 Costliness to the Agent

It is here that the notion of sacrifice makes its first appearance in (both versions of) the assistance principle. It is typical to think there is no duty to take interest-fulfilling measures if doing so would be excessively costly *for the taker of the measure*, that is, if doing so would involve excessive *agent-relative* sacrifice (Rawls 1971, 100; Fishkin 1982, 15; Dorsey 2013, 357). Below, I will argue that a high agent-relative sacrifice is not, on its own, sufficient to cancel a duty. Nonetheless, it does feed into the ‘not too costly’ condition. An agent-relative sacrifice, as I will use the term, is a sacrifice that an agent makes of her own interests. The ‘not too costly’ clause refers, in part (but only in part), to agent-relative sacrifice. In this sub-section, I will consider how we should conceptualise agent-relative sacrifice in the context of the assistance principle. (In §3.B, I will conceptualise *recipient-relative sacrifice*; in §4, I will introduce *ideal-relative sacrifice*, that is, the sacrifice an agent makes by fulfilling an interest when another agent (the ‘ideal’ agent) could more ideally fulfil that interest.)

We can think of the agent-relative sacrifice as an opportunity cost: the agent sacrifices whatever she would do instead of taking the measures. (Likewise, Overvold (1980, 108) argues that the baseline for sacrifice should not be one’s position before the sacrificial act, but should be whatever one’s position would have been had one not taken the sacrificial act.) Following Lewis (1973), we can frame this by asking what properties the agent would have in the “nearest possible world” in which she did not fulfil the interest. Whatever properties she would have in that world is the agent-relative sacrifice of her taking the measure. Notice that many things might be different between the nearest measure-taking world and the nearest non-measure-taking world: the agent might have more or less health, knowledge, friendship, love, recreation time, or access to beauty in the nearest world where the agent takes the measure, as compared with the nearest world in which she does not take the measure. All kinds of intrinsic, extrinsic, final, and instrumental values and disvalues to agent might count. Thus, the ‘opportunity cost’

⁶ If the reader prefers to have a constant likelihood threshold—rather than varying that threshold in proportion to the interest’s importance—then (2) becomes “If: *S* takes measure *M*, where *M* is the most efficacious measure open to *S* to fulfil *I*; then *I* will be fulfilled with sufficient likelihood.” I thank an anonymous reviewer for pressing me to be neutral on this issue.

framing does not reduce the range of goods that might be (agent-relatively) sacrificed; it simply provides a way of setting the baseline against which we can assess those (varied and multiple) sacrifices. Also, some opportunity ‘costs’ of taking the measure might in fact be benefits, if the nearest possible world in which the agent doesn’t take the measure is one in which she’s worse off. Again, the opportunity costs framing merely sets the baseline; it does not judge the value of that baseline. The opportunity costs might also include what Carbonell (2012) calls ‘sacrifices of self’: sacrifices that are not readily understandable as either costs or gains to agent, since they are part of the agent’s very self. It is a tricky matter whether such sacrifices of one’s very personality, values, or outlook should be understood as costs rather than benefits. Again, the point of the opportunity ‘cost’ framing does not settle that question; it simply sets the baseline for comparison.

However, when assessing whether a measure is *too* costly for the agent, we need to consider not just *which* world is our baseline for comparison, but also whether the goings-on within that world are *good or bad* for the agent. When assessing whether a given measure is or is not too costly, we should consider both the goods and the bads that the agent will undergo, in the measure-taking world as compared with the nearest non-measure-taking world. That is, when asking whether a measure has a high agent-relative sacrifice, we should compare the *net value* of these two worlds, vis-à-vis the agent.

This suggestion – to weigh up bads and goods to get an overall assessment of net value – runs counter to the view of Carbonell, who points out that some bads cannot be fully compensated for by the goods that come along with them. Suppose some interest-fulfilling measure would lead the agent to sacrifice some physical health, but increase her political freedoms. Is the measure a net good or a net bad, for her? Carbonell’s arguments imply that this question might not be answerable: “gains of wellbeing in one area of one’s life cannot be said to directly *compensate* for losses of wellbeing in another, in the most meaningful sense of ‘compensate,’ even if the gains are actually made possible by the losses.” (2012, 238, emphasis in original) Thus, she suggests, the (agent-relative) sacrifices that are relevant for the “not too costly” condition are *not* the net losses (or ‘net value,’ as I labelled the same concept). Instead, the relevant sacrifices are the *gross* losses.

This is too heavy-handed a response to the fact that some goods and bads are incommensurable. After all, some goods and bads *are* commensurable. For these goods and bads, it makes sense that we balance them off against one another, when asking whether a measure that entails goods and bads is too costly to the agent. This is not to say that a cost that comes with accompanied benefits is *not a cost*, or is in some way *irrelevant* or *insignificant* to

an overall assessment of the agent's situation (Carbonell [2012, 250] suggests that this is what a 'net' approach implies). It is just to say that when we are considering *whether a given cost can be demanded by a moral duty*, both the cost and the benefit should matter, *if* it is possible to balance them off. For example, suppose I can rescue a drowning toddler. If I did so I would receive a monetary gift from the parents. It seems reasonable to suppose that this could genuinely off-set the cost to me of muddying my shoes (assuming the shoes lack sentimental value). If so, that off-setting makes a difference to whether the rescue is too costly for me – reality of the shoe-sacrifice notwithstanding. In short, it is not *always* true that “[o]bligation is constrained by unreasonable sacrifice even when the sacrifices open doors to new sources of wellbeing.” (Carbonell 2012, 238) This is because some of those new sources of wellbeing, sometimes, can completely off-set (what would otherwise be) an unreasonable sacrifice. We should not ignore this possibility by insisting that the only sacrifice that is ever relevant is the gross sacrifice. *If* losses and gains are commensurable, then net loss (i.e., net value) is the relevant metric for agent-relative sacrifice.

Moreover, even in those cases where the goods and the bads truly are incommensurable, the non-compensability runs both ways: just as the gain doesn't compensate for the loss, likewise the loss doesn't erase the gain. The gain is real, so it's not clear why an incommensurable loss should be viewed as 'too costly,' such that it can erase an obligation to take a measure that entails the gain. For these reasons, we should use the 'net value' metric as the default metric for assessing the sacrifice an agent makes in taking interest-fulfilling measures.

That said, Carbonell is surely correct that, sometimes, an interest-fulfilling measure can realise a bad for the agent, where that bad is not commensurable with the goods the measure realises for the agent, and where the bad really does, on its own, render the measure too costly. But we can account for this by saying that measures that realise such bads have a 'trumping' power, such that they render the agent-relative value of the measure negative (or zero), irrespective of the goods that the measure also brings about. That is, Carbonell is correct that *some* bads (e.g., some that are very bad and are incommensurable with their accompanying goods) should have trumping power when it comes to assessing the agent-relative sacrifice of the measure. Perhaps the loss of the agent's basic civil rights, or the loss of their close family members, are like this. Such losses render the measures that entail them non-positively valuable for the agent, regardless of what benefits the measure would provide the agent. My point is that the presence of such cases should not deter us from assessing a measure's agent-relative sacrifice in a way that takes account of all the goods and all the bads, when it is possible to do

so. We should ask after the overall value of a measure to an agent. Sometimes, the answer to that question will be that the value is negative in virtue of one incommensurable loss alone.

The possibility of incommensurability suggests that sometimes we will be able to say only whether the agent-relative value of some measure is positive or negative, or greater or lower than the value of some other measure open to the agent—we will rarely be able to put a precise number on the agent-relative value of a measure. Sometimes, different agent-relative values (each accruing to a measure the agent could take) will only be able to be given an ordinal ranking, not a cardinal one. Sometimes, even an ordinal ranking might be impossible—there may be no fact about which measures have greatest value. For example, if *S*'s most efficacious measure for fulfilling an important interest would result in *S*'s not being able to write a great novel (that *S* otherwise would write), while *T*'s most efficacious measure for fulfilling that important interest would result in *T*'s not being able to paint a great painting (that *T* otherwise would paint), then there might be no fact about which person's measure realises highest cost to the relevant agent—because the novel and the painting are incommensurable (Anderson 1993, ch. 3; Raz 1986, ch. 13).⁷ In general, then, the idea of a measure having a 'non-negative' agent-relative value is somewhat metaphorical. However, compelling examples of assistance duties (such as easy rescue cases) suggest that this incommensurability does not always paralyse our ability to make judgments about which measures have more or less agent-relative value.

All this allows us to be ecumenical about the values that feed into the agent-relative sacrifice involved in the "at not too high a cost" constraint (including which losses count as trump cards, by rendering the measure's agent-relative value automatically non-positive). However, there are a few issues on which it is important to take sides. The first is that agent-relative value should be understood objectively (or at least, inter-subjectively), rather than subjectively. This is because this value is to be used as an input into claims about what people owe each other: what I owe you and what you owe me. For these claims to be assessed for their relative plausibility, we must have a metric of plausibility that is common between them. A crucial part of this metric is the value (negative or positive) that would be realised if the claim were upheld. An objective (or at least, inter-subjective) theory of value is needed for this metric to be properly intelligible to all parties, thus responding to their moral agency and patiency.

⁷ And as we shall see, my analysis of "best-placed" will state that the best-placed agent is that whose most efficacious measures will *not realise less* value than any other agent—it will not require that she *realise more* value than any other agent. On this analysis, all else being equal, both the novelist and the painter would have a best-placed duty.

(Here I roughly follow Carbonell (2012, 232-4), who follows Scanlon (1971).) And, as Carbonell (2015, 235) points out, there is another reason why the value at stake here must not be subjective: as Darwall (2002, 3) puts it, “if there were no difference between what a person valued and what benefited him, self-sacrifice would be impossible, except through weakness of will. [. . .] [I]t would be impossible for pursuing one’s values ever to cost one *on balance*, since realizing a value would be the same thing as benefiting from it.” Agent-relative value, then, is not to be determined subjectively.

There is a second stand we should take on how to assess agent-relative values. This stand arises from the fact that assistance duties are forward-looking imperatives, not backwards-looking assessments. The assistance principle (in both its versions) ought to be formulated to be operationalizable in practical reasoning. Yet we cannot know, ahead of taking a measure, what value (goods and bads) it will realise, vis-à-vis the agent. We should, therefore, weight the possible values by their (reasonably believed or reasonably believable) likelihoods. That is, since the assistance principle is an action-guiding and forward-looking imperative, the agent-relative sacrifice relevant to the “not too costly” constraint is agent-relative *expected* value.

To summarise: the size of the agent-relative sacrifice of a given interest-fulfilling measure is *part of* what determines whether that measure is ‘too costly’. The other determinate will be fleshed out in the next section. If the measure is too costly, then the assistance duty never arises to begin with. There is an agent-relative sacrifice just in case the cost to the agent is negative. The cost to the agent is negative just if the expected net value to the agent’s interests (i.e., the sum of her expected interest-fulfilment minus her expected interest non-fulfilment), on a non-subjectively defined understanding of interests, is lower in the nearest world where she takes the measure, than in the nearest world where she does not the measure. Some interests are so important that their non-fulfilment would render the expected net value to the agent’s interests negative, regardless of what goods-for-the-agent come along with that non-fulfilment.

3.2 Costliness to the Recipient

The agent is not the only one whose goods and bads are at stake in circumstances where the assistance principle applies. Most obviously, goods and bads for the interest-bearer (or ‘recipient’) are also at stake. As alluded to in the Introduction, a measure can be costly for the interest-bearer—either in the sense that the interest-bearer’s important interest is fulfilled, but not fully or securely; or in the sense that the interest-bearer’s other interests are frustrated by the measure. In short, at stake for the interest-bearer is not just whether the relevant important

interest is fulfilled, but also whether it is fulfilled in the right way (that is, fully, securely, or without frustrating their other interests). To illustrate with the familiar example: if the drowning toddler is saved, but is knocked around and suffers concussion, then the value of the rescue measure, vis-à-vis the recipient, are diminished relative to rescue without concussion. If the non-concussion rescue would have happened in the absence of the concussion rescue, then this is a cost (for the recipient) of the rescue scenario with the concussion.

Should the value of a measure vis-à-vis the recipient also help determine whether an interest-fulfilling measure meets the “not too costly” condition? That is, when others are fulfilling our interests, should the sacrifices they impose on us (by way of fulfilling our interest) have some bearing on whether they ought to fulfil our interests? Yes. Agent-relative expected value is not the only relevant sacrifice. In performing ‘good deeds’, well-meaning people affect those around them – and not for the best in every respect. If I am a clumsy swimmer and will cause the child to get concussion during my rescue, then, *even if my rescue is likely to be successful*, this sacrifice I impose on the child should be considered in assessing whether I should enact the rescue.

One might think that the costs to the recipient have already been taken into account in §1: it is their important interest the action is aiming to fulfil. To count their interests as also bearing upon the “not too costly” condition would, perhaps, be to count those interests twice. But this is not right. The only interest of theirs that was at issue in §1 was the specific important interest that the measure aims to fulfil. In the child drowning case, this is the important interest in being rescued. But those on the ‘recipient’ end of the assistance principle have interests that are as numerous, varied, multi-faceted and incommensurable as those on the ‘agent’ end of the principle. There is no reason why these other interests of the recipient should not be considered.

A slippery slope may appear to loom. If we’re going to consider costs *and* benefits to agent *and* recipient, then why not insist that the interest-fulfilling measure must have non-negative value regarding *the entire world*, on pain of not fulfilling the “not too costly” condition? Indeed, if one adheres to maximising, act, agent-neutral consequentialism, then it is natural to think the “not too costly” condition is met if and only if the measure produces non-negative aggregate cost regarding the entire world. But if one is a maximising, act, agent-neutral consequentialist, then one will see the assistance principle as a mere ‘rule of thumb.’ And ‘rules of thumb’ should not have the same content as the fundamental principles that justify the rules of thumb (otherwise, they would not deserve the name ‘rules of thumb’). So while it might seem natural to include ‘the entire world’ in the not-too-costly condition, this would undermine the role that the assistance principle is meant to play in one’s theory. By

limiting the relevant net value to the agent and recipient alone, we make the assistance principle operationalizable as a rule of thumb.

Meanwhile, if one is *not* a maximising, act, agent-neutral consequentialist, then one is likely to see the assistance principle as grounded in the claims (or rights) of the interest-bearer themselves. From this perspective, to consider interests beyond those of the agent and recipient would be to overlook the point that the assistance principle produces duties for a particular agent to fulfil a particular important interest (right, claim) of a particular recipient. It is the recipient—and the recipient alone—that would be wronged if his interest were not fulfilled. And they would be wronged by the agent. So, again, from a claim- or rights-based perspective, it again is natural to limit the relevant net value to agent and recipient.

We can, therefore, think of the “expected value” of a measure as being equal to: the potential weighted benefits for agent and recipient, given that the measure is taken, multiplied by the benefits’ likelihoods, minus the potential weighted costs for agent and recipient, given that the measure is taken, multiplied by those costs’ likelihoods. If this value is non-negative, then the costs are not proportionate to the benefits, so the costs do not render the measure non-obligatory, and the ‘not too costly’ condition is met. That is, a necessary condition for an assistance duty is that the measure’s expected value (regarding agent and recipient) is non-negative.

We have now arrived at the third necessary condition for assistance duties:

- (3) *S*’s taking the measure in (2) would realise non-negative expected value regarding *S* and *P*.

3.3 Aggregative and Iterative Sacrifice

Intuitively, when assessing the expected value of an interest-fulfilling measure, only that measure’s value is relevant. But this creates problems for the assistance principle. After all, small costs add up, both across time and at one time. These aggregate costs sometimes block the duty. Perhaps any one measure I might take to fulfil one person’s important interest is not prohibitively costly, but I cannot take measures that fulfil *all*, or even *many*, persons’ important interests without incurring prohibitive cost. Saving any one child is easy; saving many children requires forgoing my most important life projects and relationships.

Suppose, then, that the aggregative, but not iterative, value is non-positive—that is, the measures taken in aggregate have a non-positive expected value, but taken iteratively each has a positive expected value. For example, suppose it takes you a day to save each of 20,000 lives. While it might be proportionate (i.e. have positive value) for you to spend one day saving one life, it might be that 20,000 days (i.e., almost 55 years) is too much for morality to demand of you, whatever the cause, thus rendering the value of the 55-life-saving measure non-positive. If each duty over each life-saving measure were dependent only on the expected value of *that one measure* being non-negative, then we would be unable to cite the overall non-positive expected value as blocking any one of the interest-fulfilling duties. But this cannot be right: I should be able to cite the aggregated expected costs of all these measures as disproportionate to their aggregated expected benefits, rendering the aggregate expected value non-positive.

For example, perhaps each measure that fulfils conditions (1)–(3) is compatible with certain life-enhancing goods such as friendship, personal projects, and so on. But we cannot have these life-enhancing goods while taking measures in *all* circumstances where (1)–(3) hold. If life-enhancing goods should not be sacrificed at the altar of well-placed or best-placed duties, then we should take an aggregative, rather than iterative, approach to assessing whether the value of one act of beneficence is positive. In short, we should consider the value of our *policy* around assistance duties, not the value of discharging individual duties, if we are to give due credit to the value of, for example, life-enhancing goods (see similarly Cullity 2004, Part II; Hooker 1999; Kamm 2000, 660).

If aggregative value can be non-positive while iterative value is positive, this suggests assistance duties are not constrained by the expected (agent- and recipient-relative) value of each interest-fulfilling measure considered in isolation, but instead by the aggregated expected value of all such measures. For assistance duties, we should aggregate the value of measures that we could take to fulfil someone else’s important interest, in all cases where (1)–(3) hold. This gives us something like the following necessary condition for an assistance duty:

- (4) If *S* were to take measures in all circumstances where the other conditions of this principle hold between *S* and any individual, then *S* would realise positive *aggregate* expected value regarding *S* and all those individuals.

A problem remains. Suppose Peter is the only one capable of rescuing all of 100 children. Each of their interests is very important. It would realise positive expected (agent-relative combined with patient-relative) value for him to rescue 10 of them, but non-positive

expected value for him to rescue any more than that. (Imagine that rescuing 11 or more would cause him to develop severe hyperthermia, which serves a “trumping” function on any benefits to patients.) If (4) is necessary for an assistance duty, then Peter is off the hook *altogether*: Helping all 100 would realise negative value, so, by the lights of (4), he has no duty to help any of them.

But the right result is surely that Peter has a duty to rescue 10 children. After all, that is the threshold of positive expected value regarding himself and the recipients; the point at which the sum of Peter’s costs and the recipients’ costs does not lead to sacrifice. Peter should be let off the hook for not helping all 100, but *not* off the hook for helping 10 of them. A plausible method is for Peter to rank the 100 measures—one measure for saving each of the 100 children. Each measure is weighted according to the importance of the interest it is aimed at fulfilling, along with the other expected costs and benefits for Peter and the child in question. The weighted measures are then ranked. Peter’s duty is to work his way down the ranked measures, taking each until he reaches the limit of positive aggregate value. There might be several measures that are in a “tie” situation on the list: if Peter cannot take all the tied measures while realising positive value, then he is afforded discretion in choosing which of the tied measure(s) he takes.

Condition (4) then becomes disjunctive:

- (4) Either: if *S* were to take measures in all circumstances where the other conditions of this principle hold between *S* and any individual, then *S* would realise positive *aggregate* expected value regarding *S* and all those individuals;
- Or: when the importance in (1) and value in (3) are used to weight the measures in (2), these measures rank sufficiently high among similarly weighted measures (for which (1)–(3) also hold), such that the measures in (2), *and* all more highly-ranked measures, could be taken with positive aggregate expected value (regarding *S* and those whose important interests are thereby targeted).

4. The Action is the Least Costly of All Agents’ Similar Actions: Ideal-Relative Sacrifice

Conditions (1)–(4) get us to well-placed duties. But they do not get us to best-placed duties, understood as duties that accrue to the agent *best-placed* to fulfil some important interest. If (1)–(4) were the end of the story, then multiple agents could bear duties to fulfil exactly the same interest, if they all met (1)–(4). Such a proliferation of well-placed duties would be an

acceptable result in many instances, but sometimes we are interested in picking out one from among the many agents who are well-placed. This has a stronger duty than all the other well-placed agents. This is the agent we call on first, if it would be disastrous for all well-placed agents to act. Imagine a crowd of well-placed people, all jumping in to save the drowning child, getting in each other's way and causing further drownings. We don't want this. We want a single agent to act. Best-placed duties pick out this one agent. This agent has a stronger duty than all the other agents, even if it wouldn't be disastrous for all to act. And if it *would* be disastrous for all to act, then she has the duty in the first instance.⁸

To understand why best-placed duties are stronger than well-placed duties, it helps to develop a conception of sacrifice according to which we can *sacrifice others' actions*. The idea that an agent can sacrifice someone, or something, other than the agent of the sacrifice is familiar from the folk notion of sacrifice. The first listing for 'sacrifice' in the *Oxford English Dictionary* is "[p]rimarily, the slaughter of an animal ... as an offering to God or a deity." Taking our cue from this, it seems we can sacrifice all sorts of things – not just our interests as agents – via our interest-fulfilling actions. In this section, I suggest that one thing that you can sacrifice when performing an interest-fulfilling action is someone else's performing a similar, but better, action. This helps to explain why best-placed duties are stronger than well-placed duties: acting on (some) well-placed duties *sacrifices* a better interest-fulfilling action (an action of a better-placed person). Discharging well-placed duties thus (sometimes) involves a higher sacrifice than discharging best-placed duties. We can therefore explain the weaker strength of those well-placed duties via the familiar idea that duties involving higher amounts of sacrifice are weaker than duties involving lower amounts of sacrifice: specifically, the sacrifice of someone else's better action.

We can call this sacrifice the 'ideal-relative sacrifice': the *merely* well-placed agent sacrifices the ideal action by discharging her duty. The ideal action can be better in the sense of imposing lower costs on *the agent* who performs it (compared to the costs the merely well-placed agent's interest-fulfilling actions would have impose on the well-placed agent), or in the sense of imposing lower costs on *the interest-bearer* (in the sense that either the interest-bearer's important interest is more fully or securely fulfilled, or the interest-bearer's other interests are also fulfilled, alongside the important interest that generated the duty).

⁸ As an anonymous reviewer pointed out, there will be some situations in which it's best if a few agents help *together*, rather than one agent helping alone. I lack space to deal with such cases here, though I discuss them in Collins (2013).

To capture this, we need to add a condition to (1)-(4), such that (1)-(4) produce well-placed duties while (1)-(5) produce best-placed duties. A first pass of condition (5) is:

- (5) *S* is the agent whose measure (from (2)) is such that, if *S* takes that measure, *S* would realise no less expected value (regarding agent and recipient) than would any other agent who meets (1)–(4) if they took their measure.

That is, *S*'s interest-fulfilling measure *would not make an ideal-relative sacrifice*. This condition is different from “*S* is most capable.” The most capable agent is simply the one whose most efficacious measure is *no less likely to fulfil* the relevant interest than is any other agent's most efficacious measure. But the most capable agent might incur great agent-relative sacrifices—or impose great recipient-relative sacrifices—if she takes this measure. For this reason, we should not be concerned merely with her likelihood of success, but with other costs and benefits of her measure for her and the recipient.

But this is not quite the end of the story. The agent who is “best-placed” (by the lights of condition (5) above) might not do what her moral duty demands. If Ash has broken his leg, perhaps his housemate, Brandy, is best-placed (according to condition (5)'s first pass) to take Ash to the hospital. Yet suppose Brandy is callous, and will not doing her duty. Brandy's unwillingness to take Ash to hospital does not affect her duty to do so. After all, agents cannot get out of doing their duties by simply not wanting to. Brandy has a best-placed duty.

But what about Cara, Ash and Brandy's next-door neighbour, who knows Brandy will callously watch television while Ash suffers, and that there is nothing anyone can do to convince Brandy otherwise? What is Cara's duty? She does not meet condition (5)'s first pass: Cara's measure would realise less value than Brandy's measure; Cara's measure would involve an *ideal-relative sacrifice* according to condition (5) above. But, given Brandy's callousness, Cara surely has a strong duty, in virtue of being (second) best-placed.

We can make sense of this by noting that, in determining whether some agent has a duty, we do not assume that others will comply with morality's demands. So, when determining whether an agent's measure would involve an ideal-relative sacrifice, we should be as realistic as we can about the expected probability that others will comply with their duties, just as we are realistic about natural events that might intervene and disrupt the attempts of the agent under consideration. By contrast, when determining whether an agent has a duty, it is irrelevant whether *she* will or won't comply with morality. Agents cannot get out of assisting others by not wanting to. When assessing Brandy's duty, then, we treat her as an agent, who cannot get

out of her duty by simply not doing it. So she is best-placed, on that assumption: her measure would not involve an ideal-relative sacrifice. This gives Brandy a best-placed duty.

When determining Cara's duty, however, we treat Brandy as a feature of the environment to be worked around. Given that Brandy will not do her duty, Cara is best-placed: her measure would not involve an ideal-relative sacrifice. Because we assume that the duty-bearer will comply, but are realistic about whether others will comply, there are sometimes two (or more) agents with best-placed duties to take different measures to fulfil the same interest. This is simply because, when we are considering the duties of the different agents, we hold different things fixed where those things are relevant to the existence of an ideal-relative sacrifice. Specifically, considering an agent's duty, what matters is not that the agent's measures are not worse than anyone else's, but that they are not worse than those of any other agent *who will do their duty if they have one*. After all, it is only if those other agents did their duty that the agent's measure would involve an ideal-relative sacrifice.

We can thus slightly refine condition (5):

- (5) Of all agents that meet (2)–(4), *S*'s measure (from (2)) is such that, if she were to take that measure, she would realise no less expected value (regarding agent and recipient) than the expected value (regarding agent and recipient) that would be realised by any other agent *who would take measures to fulfil this interest if they had a duty to do so*.

Condition (5) says the measure “would realise no less expected value,” not that it “would realise more expected value.” If two agents would realise equal value, then they both have best-placed duties to take their measures, since neither's measure would sacrifice the *better* actions of another agent.

5. Conclusion

We now have two principles.

The Well-placed Principle:

If

(1)–(4)

Then

(6a) In the absence of defeaters, *S* has an all-things-considered duty to take *M*; and

- (7a) If the ranking in the second disjunct of (4) ranks multiple measures equally, and if *S* could take some but not all of those equally-weighted measures and realise positive aggregate expected value (regarding *S* and recipients), then, in the absence of defeaters, *S* has a duty to take some of the tied measures, up to the threshold of positive aggregate expected value (regarding *S* and recipients).

The Best-placed Principle

If

- (1)–(5)

Then

- (6b) In the absence of *strong* defeaters, *S* has an all-things-considered duty to take *M*;
and
(7b) If the ranking in the second disjunct of (4) ranks multiple measures equally, and if *S* could take some but not all of those equally-weighted measures and realise positive aggregate expected value (regarding *S* and recipients), then, in the absence of *strong* defeaters, *S* has a duty to take some of the tied measures, up to the threshold of positive aggregate expected value (regarding *S* and recipients).

There are three kinds of sacrifice relevant to these principles: agent-relative sacrifice, recipient-relative sacrifice, and ideal-relative sacrifice. All three of these have some bearing on whether a given agent-recipient pair meets the antecedents of the principles, in particular conditions (3) and (5). I have argued that both agent-relative and recipient-relative sacrifice should be understood as *net* sacrifice, not *gross* sacrifice. The problem of incommensurable sacrifices and benefits creates some problems for this, as Carbonell has noted. But I have suggested we can solve these problems by allowing that some sacrifices (by of either agent or recipient) can render an interest-fulfilling measure's value non-positive, without the need for that sacrifice to be traded off against the possible benefits of the measure.

The notion of ideal-relative sacrifice helps to capture why it is that well-placed duties are weaker than best-placed duties: if the agent- and recipient-relative value is held constant across two possible scenarios, but one of those scenarios contains a third party that is better-placed to fulfil the interest (and will do so if they have a duty to do so), then the agent's fulfilling the interest contains an ideal-relative sacrifice, and is thereby less valuable. This notion of

ideal-relative sacrifice also helps address a possible concern that assistance duties are, by their nature, overly demanding of agents: we have relatively weak duties to take interest-fulfilling measures that imply ideal-relative sacrifices. Thus, such duties are more easily defeated by countervailing considerations. This means that, all-things-considered, they will demand action of us less often than the assistance principle perhaps implies at first glance.

References

- Anderson, Elizabeth. 1993. *Value in Ethics and Economics*. Cambridge, Mass.: Harvard University Press.
- Ashford, Elizabeth. 2003. "The Demandingness of Scanlon's Contractualism." *Ethics* 113 (2): 273–302.
- Barry, Christian. 2005. "Applying the Contribution Principle." *Metaphilosophy* 36 (1/2): 210–227.
- Barry, Christian and Gerhard Øverland. 2013. "How Much for the Child?" *Ethical Theory and Moral Practice* 16 (1): 189–204.
- Brennan, Geoffrey and Nicholas Southwood. 2007. "Feasibility in Action and Attitude." Published online only in T. Ronnow-Rasmussen, B. Petersson, J. Jonefsson, and D. Egonsson (eds), *Hommage a Wlodek: Philosophical Papers Dedicated to Wlodek Rabinowicz*. Available at <www.fil.lu.se/hommageawlodek>.
- Chappell, Timothy (ed.). 2009. *The Problem of Moral Demandingness: New Philosophical Essays*. London: Palgrave Macmillan.
- Chisholm, Roderick. 1963. "Supererogation and Offense: A Conceptual Scheme for Ethics." *Ratio* 5 (1): 1–14.
- Collins, Stephanie. 2013. "Collectives' Duties and Collectivization Duties." *Australasian Journal of Philosophy* 91 (2): 231–248.
- Collins, Stephanie. 2015. *The Core of Care Ethics*. Basingstoke: Palgrave Macmillan.
- Cullity, Garrett. 2004. *The Moral Demands of Affluence*. New York: Oxford University Press.
- Dorsey, Dale. 2013 "The Supererogatory, and How to Accommodate It." *Utilitas* 25 (3): 355–382.
- Fishkin, James. 1982. *The Limits of Obligation*. Binghamton: Yale University Press.
- Miller, David. 2001. "Distributing Responsibilities." *Journal of Political Philosophy* 9 (4): 453–471.

- Miller, David. 2007. *National Responsibility and Global Justice*. Oxford: Oxford University Press.
- Miller, Richard. 2004. "Beneficence, Duty, and Distance." *Philosophy & Public Affairs* 32: 357–383.
- Page, Edward A. and Avia Pasternak. 2014. *Benefitting from Justice*. Special Issue of *Journal of Applied Philosophy* 31(4).
- Feltham, Brian and John Cottingham (eds). 2010. *Partiality and Impartiality: Morality, Special Relationships, and the Wider World*. Oxford: Oxford University Press.
- Goodin, Robert E. 1985. *Protecting the Vulnerable*. Chicago: University of Chicago Press.
- Hooker, Brad. 1999. "Sacrificing for the Good of Strangers—Repeatedly." *Philosophy and Phenomenological Research* 59 (1): 177–181.
- Kamm, Frances M. 2000. "Does Distance Matter Morally to the Duty to Rescue?" *Law and Philosophy* 19 (6): 655–681.
- Keller, Simon. 2006. "Four Theories of Filial Duty." *Philosophical Quarterly* 56 (223): 254–274.
- Kumar, Rahul. 1999. "Defending the Moral Moderate: Contractualism and Common Sense." *Philosophy & Public Affairs* 28: 275–309.
- Lee, Stan. 1962. *Amazing Fantasy #15*. Marvel Comics.
- Liao, Matthew S. 2006. "The Right of Children to be Loved." *Journal of Political Philosophy* 14 (4): 420–440.
- Miller, David. 2007. *National Responsibility and Global Justice*. Oxford: Oxford University Press.
- Nussbaum, Martha C. 2006. *Frontiers of Justice: Disability, Nationality, Species Membership*. Cambridge, MA: Belknap Press.
- Raz, Joseph. 1986. *The Morality of Freedom*. Oxford: Oxford University Press.
- Scanlon, T.M. 1975. "Preference and Urgency." *Journal of Philosophy* 72 (19): 655–669.
- Scanlon, T.M. 1998. *What We Owe to Each Other*. Cambridge, MA: Harvard University Press.
- Shue, Henry. 1996. *Basic Rights: Subsistence, Affluence and US Foreign Policy*. Second edition. Princeton: Princeton University Press.
- Singer, Peter. 1972. "Famine, Affluence, and Morality." *Philosophy & Public Affairs* 1 (3): 229–243.
- Singer, Peter. 2009. *The Life You Can Save*. New York: Random House.
- Sobel, David. 2007. "The Impotence of the Demandingness Objection." *Philosophers' Imprint* 7 (8): 1–17.

- Sonderholm, Jorn. Forthcoming. "World Poverty, Positive Duties and the Over-demandingness Objection." *Politics, Philosophy & Economics*.
- Wolf, Susan. 1982. "Moral Saints." *Journal of Philosophy* 79: 419–439.