

- Lazar, S., 'Necessity in Self Defense and War', *Philosophy & Public Affairs* 40/1 (2012): 3-44.
- McMahan, J., 'The Just Distribution of Harm Between Combatants and Non-combatants', *Philosophy & Public Affairs* 38 (2016): 342-79.
- Parfit, D., *On What Matters*, vol. iii (Oxford: Oxford University Press, 2017).
- Scheffler, S., *The Rejection of Consequentialism* (Oxford: Oxford University Press, 1982).
- Thomson, J. 'Killing, Letting Die, and the Trolley Problem', *The Monist* 59/2 (1976): 204-17.
- Thomson, J. 'Turning the Trolley', *Philosophy & Public Affairs* 36/4 (2008): 359-74.

More Supererogatory

Thomas Hurka and Evangeline Tsagarakis

In one of its uses the concept of supererogation is binary, or doesn't admit of degrees. If the moral duties or reasons an act would fulfil don't outweigh the costs it would involve for you, the act is simply supererogatory; if they do outweigh those costs, it's simply not supererogatory but instead is required. In another use, however, this concept does admit of degrees, and one of two supererogatory acts can be more supererogatory than the other. Supererogatory acts are commonly said to be 'beyond the call of duty'; surely one act can be further beyond duty than another. These acts are associated with saintliness or heroism, and one can be more saintly or heroic than another; they also call for admiration, and can be more or less admirable. An act's degree of saintliness, heroism, or admirability may not depend only on how supererogatory it is; other factors, such as ones relating to its motivation, may be, and we'll later argue are, relevant. But the extent to which it has these qualities depends in part on its degree of supererogation, or on how far beyond duty it is.

This paper will first give examples where one of two acts is more supererogatory than the other, and propose an initial account of what makes that so. It will then connect this account to a more general theory of supererogation, as well as to an account of degrees in other deontic concepts. And it will contrast its theory of supererogation with the one Derek Parfit defends in *On What Matters*; one weakness in Parfit's theory, it will argue, is that it doesn't allow degrees of supererogation.

1. Degrees of Supererogation: Examples

An act is supererogatory if the moral duties it will fulfil don't outweigh the costs or sacrifices it will involve for you. But of two supererogatory acts that fulfil the same moral duty, one can impose greater costs than the other. Consider two acts

The idea of a gap treatment of degrees of supererogation in Sections 1 and 2 is due to Tsagarakis; the general defence of the permission-based theory in Section 2 and the critique of Parfit in Section 3 are due to Hurka; and Section 4 is a joint product. For discussion we're grateful to Brendan de Kenessey, and for helpful comments to James Goodrich.

that will save another person's life. One requires you to sacrifice one of your arms, while the other requires you to sacrifice both your arms. If the duty to save another person doesn't outweigh the cost of one arm, both acts are supererogatory. But surely the second act, which requires you to sacrifice two arms rather than one, is more supererogatory than the first. It's more heroic to make the larger sacrifice than to make the smaller one, and also more admirable; you go further beyond duty when you sacrifice more. And we can give an initial, relatively theory-neutral account of what makes this so.¹

The moral duty to save the other person turns on the benefit to him of being saved, or of being spared the harm of death. And the reason both acts are supererogatory is that this benefit, though greater than the costs the acts impose on you, isn't vastly greater. To put it another way, the gap between the size of the benefit to him and the size of the cost to you isn't large enough to make saving him a duty. But this gap isn't the same for the two acts: it's smaller for the second act, given its larger cost for you. And we can take this difference to explain why the second act is more supererogatory. In cases of benefiting beyond duty, we can say, the degree of supererogation depends on the size of the gap between the benefit to the other person and the size of the cost to you. More specifically, the smaller this gap is, the more supererogatory the act, and the larger the gap, the less supererogatory. A smaller gap means the act is further from being your duty; it would require a greater increase in the benefit the act gives the other to make it your duty. But when the gap is larger, or when the benefit exceeds the cost by more, the act is closer to being your duty and less far beyond duty.

In this first example the two acts fulfil the same duty but differ in their costs for you. But we can also consider a case where two acts fulfil different duties but impose the same cost. Imagine two acts, both of which require you to sacrifice one of your arms but one of which will save another person's life while the other will prevent him from losing both his arms. Here the second act seems more supererogatory, and we can again explain why by looking to the gap between the benefit it provides and the cost it imposes on you. In the first act this gap is large, because the harm of death is much greater than the harm of losing an arm. This means this act is not that far from being your duty; not that much would have to

¹ We assume that among the duties that bear on these acts is a duty to promote the good of all, including yourself, impartially. This implies that if an act's cost to you is vastly greater than its benefit to another—for example, if it involves sacrificing an arm to spare him a mild headache—it doesn't fulfil your moral duties; it violates them and can't be supererogatory. We also assume that in these acts the cost—either one arm or two—is the smallest that will produce the relevant benefit. If you sacrifice two arms to save another's life when you could do so equally well by sacrificing one, your act doesn't fulfil your impartial duty as well as the less costly alternative and may therefore be wrong. Whether the act is supererogatory then depends on whether supererogatory acts must be all-things-considered permitted; we leave this issue open. And even if the act is supererogatory, whether it's more so than its alternative depends on a more complex measure we introduce in Section 2; on the simplest reading of this measure the more costly act isn't more supererogatory.

be added to the benefit to make the act required. In the second act, however, the gap is smaller. And if a smaller gap between benefit and cost makes for a higher degree of supererogation, by making the act further from your duty, the second act is more supererogatory. This verdict may not seem quite as compelling as the one in the first example, because another factor, to be discussed later, may differentiate this pair of acts that didn't distinguish the earlier pair. But if the size of the gap between benefit and cost makes one of two acts more supererogatory in the first example, it should do the same here.

In both these examples, one end of the relevant gap is the same for both acts. In the first, both acts provide the same benefit to the other; in the second, both involve the same cost for you. This makes comparative judgements of supererogation relatively easy to make, since they require only a comparison of costs in the one example and of benefits in the other. These judgements are harder to make when neither end of the gap is the same for two acts, as when we compare saving another person's life at the cost of two arms for yourself with preventing him from losing two arms at the cost of one arm for yourself. To determine whether one of these acts is more supererogatory than the other we need to compare the size of two gaps that differ at both their ends, that is, to compare the gap between a life and two arms with the gap between two arms and one arm. Because these comparisons are harder to make, our intuitions about which if either act is more supererogatory may be less firm. But if the gap between a life and two arms is greater than the gap between two arms and one, sacrificing one arm to save another person two arms is more supererogatory than sacrificing two arms to save his life. Since the second act is closer to being your duty, the first is further beyond duty. (Again another factor, to be discussed later, may complicate our judgement here.)

2. Supererogation and Prima Facie Permissions

The above account of degrees of supererogation, though persuasive as far as it goes, is in two respects incomplete. In taking one end of the gap to be the benefit an act gives another person, it assumes that the duty a supererogatory act fulfils is always the duty to promote the good, which we're assuming is the duty to promote the good of all, including yourself, impartially. The account therefore doesn't extend to supererogatory acts, if there are such, that fulfil other duties, say by keeping a promise, paying compensation, or refraining from directly causing harm when, given the costs of doing so, that isn't required. And though the account is right to make the other end of the gap depend on an act's costs for you, it doesn't explain exactly how or why these costs have the normative effect they do. Different theories of supererogation give different explanations of this, by

linking the costs to different normative factors that weigh against, and can at times outweigh, the requiring force of the duties a supererogatory act would fulfil. A complete account of degrees of supererogation should be embedded within some such theory. We'll now propose one, which we'll argue is both attractive in itself and well suited to accommodate degrees of supererogation.

This theory takes the duties relevant to supererogation to be what W. D. Ross called *prima facie* duties,² and can also be called *pro tanto* duties or duties other things equal. They include a *prima facie* duty promote the good, and perhaps also similar duties to keep promises, pay compensation, and so on. But it supplements these duties with an independent *prima facie* permission, or permission other things equal, to promote your own good and, as part of that, to avoid harms to yourself. This *prima facie* permission functions in the same way as a *prima facie* duty. If there's a *prima facie* duty to do acts of type *F*, then an act's being *F* tends to make it, or counts in favour its being, what Ross called your duty proper; this can also be called your duty all things considered or *simpliciter*. Analogously, if you have a *prima facie* permission to do acts of type *G*, an act's being *G* tends to make it, or counts in favour of its being, all-things-considered permitted, or permitted *simpliciter*. In each case the *prima facie* or conditional factor contributes to an all-things-considered or unconditional one of the same type: *prima facie* duties to duties *simpliciter*, and *prima facie* permissions to permissions *simpliciter*.

Just as one *prima facie* duty can conflict with another, so can a *prima facie* permission conflict with a duty. This will be so if there's a *prima facie* duty to do acts of type *F* and a *prima facie* permission to omit acts of type *G*, and some act is both *F* and *G*. And just as *prima facie* duties can have differing weights, so can *prima facie* permissions: in each case it's a matter of the strength of the tendency to make for duty or permission *simpliciter*. The *prima facie* duty to do *F* is stronger than the *prima facie* duty to do *G* if there's some other *prima facie* duty, say to do *H*, that the duty to do *F* outweighs but the duty to do *G* does not. (A simpler test says the duty to do *F* is stronger than the duty to do *G* if, in a direct conflict between them, what you ought all things considered to do is *F*.) Analogously, the *prima facie* permission to do *F* is stronger than the *prima facie* permission to do *G* if there is some *prima facie* duty, say to do *H*, that the first permission outweighs but the second does not. (Here the simpler test isn't applicable because *prima facie* permissions don't conflict: it can be both all-things-considered permitted to do an act and all-things-considered permitted not to do it.) Thus the *prima facie* permission not to sacrifice two arms is stronger than the *prima facie* permission not to sacrifice one arm if there is some duty that the first would outweigh but the second would not, for example some benefit to

another that you would be required to provide at the cost of one arm but not at the cost of two arms. There presumably is some such benefit.

A theory of supererogation that uses these two *prima facie* factors can say the following: an act is supererogatory if the *prima facie* duties that favour doing it are outweighed by a *prima facie* permission not to do it, where that permission derives in some way from the more general permission to promote your own good. That the *prima facie* duties are outweighed by the *prima facie* permission means not doing the act is all-things-considered permitted, but doing the act is also all-things-considered permitted. The *prima facie* duties to do the act entail a *prima facie* permission to do it—'ought' in general entails 'may'—and that permission isn't undermined by the *prima facie* permission to avoid its costs. As noted above, *prima facie* permissions don't conflict; it can be all-things-considered permitted to do an act and also all-things-considered permitted not to do it. So supererogatory acts are ones it's both simply permitted to do and simply permitted not to do. And what grounds the simple permission not to do them is a *prima facie* permission to avoid costs to yourself that, in these cases, outweighs whatever duties the acts would fulfil.³

This theory can be extended to allow degrees of supererogation. These will again depend on a gap, now that between the strength of the *prima facie* permission to avoid an act's costs and the strength of the duties to do it that the permission outweighs. But here the size of the gap has the reverse effect to the one described above: now a larger, not a smaller, gap makes for a more supererogatory act. Return to the first example, where you can save another's life either by sacrificing one of your arms or by sacrificing two. Since the *prima facie* permission to avoid losing two arms is stronger than the permission to avoid losing one, the *prima facie* permission in the first act is stronger than the one in the second. The first permission therefore outweighs the duty to save the other's life by more, which makes the gap in the first act larger. And that is the act that is more supererogatory. It would take a greater increase in the strength of the *prima facie* duty an act with this cost fulfils to make it simply your duty; maybe it would need to save ten lives, whereas only five would be needed to make sacrificing one arm your duty. The first act is therefore further from being simply your duty than the second, or further beyond duty. And in this context, where we are comparing *prima facie* permissions with *prima facie* duties rather than benefits with costs, a larger gap makes that the case.

As so extended, the theory corrects the incompleteness of the initial account of degrees of supererogation. Because the factor on one side of the gap is now the strength of a *prima facie* duty rather than the size of benefit, it can allow supererogatory acts, if there are such, that fulfil duties other than beneficence,

² W. D. Ross, *The Right and the Good* (Oxford: Clarendon Press, 1930), ch. 2.

³ For this theory, see T. Hurka and E. Shubert, 'Permissions to Do Less than the Best: A Moving Band', *Oxford Studies in Normative Ethics* 2 (2012): 1–27.

such as promise-keeping and compensation. And because the factor on the other side of the gap is a prima facie permission to pursue your good that's present in contexts other than that of supererogation, it gives a more general explanation of why an act's costs to you can make it not required. Moreover, its account of degrees of supererogation connects with accounts of degrees in other deontic concepts such as wrongness.⁴

In one use the concept of wrongness doesn't admit of degrees. If the prima facie duties an act violates outweigh the prima facie duties and/or permissions it fulfils, the act is simply wrong; if they don't, it's simply not wrong. But in another use one act can be more wrong than another or, if we want to reserve the word 'wrong' for the binary concept, more seriously wrong. And the seriousness of a wrong again depends at least in part on a gap, now that between the strength of the prima facie duties an act violates and the strength of any duties and/or permissions it fulfils. Imagine two wrong acts, one that gives a small benefit to one person rather than a somewhat larger benefit to another, and one that gives the same small benefit to one person rather than a vastly larger benefit to another. Here the second act is more seriously wrong because the gap between the strength of the prima facie duty it violates and the strength of the one it fulfils is larger. (When an act that violates a prima facie duty fulfils no duty or permission, its seriousness as wrong depends just on the strength of the violated duty; but that is just the gap between the strength of that duty and nothing.) Derek Parfit implicitly took this view at one point in *On What Matters*. He argued that if killing one person to save several others was permitted only if the number saved was at least a hundred, then killing one to save ninety-nine would be only 'slightly wrong' or something there's only a 'weak moral reason' not to do.⁵ And the explanation is presumably that, when the number saved is so close to the threshold for permissibility, the gap between the strength of the duty the act violates and that of the one it fulfils is small. His claim about this kind of case is actually contestable. It can be argued that when the duty a wrong act violates is deontological, like the duty not to kill, rather than just about giving benefits, the seriousness of the wrong depends not only on the size of the gap but also on the strength of the violated duty on its own. But even on that view the gap is relevant, so killing one to save 99 is less seriously wrong than killing one to save 50 or 25 or two. An account of degrees of supererogation based on the gap in strength between a prima facie permission and the prima facie duties it outweighs parallels an attractive account of degrees of seriousness of wrong.

The above theory of supererogation, which weighs prima facie duties against a prima facie permission, isn't the only or even most common such theory. But it has several attractive features.

First, it gives a straightforward explanation of why the costs to you that make an act supererogatory aren't just those of some precise size but all those within a range. If you're permitted not to save another's life at the cost of one of your arms, you're also permitted not to save it at the cost of two arms, two legs, or your life. You have the first permission because the prima facie permission to avoid losing an arm outweighs the prima facie duty to save a life, but if that is so, the prima facie permissions to avoid losing two arms, two legs, and so on also outweigh that duty, since they're stronger. The result of weighing your general prima facie permission against the duty to benefit others isn't an absolute permission to weigh your own good more heavily than others' in just some precise ratio, say 5:1 and no more or no less. If you're permitted to prefer one unit of your own good to five units of good for others, you're also permitted to prefer it to four, three, or two units for others.

Second, the theory can be extended to accommodate a related but contrary phenomenon. Standard cases of supererogation involve what's been called an 'agent-favouring permission', since they permit you to prefer a lesser good for yourself to a somewhat greater good for someone else. But common-sense morality also grants 'agent-sacrificing permissions', since it sometimes permits you to forgo a greater good for yourself to give a somewhat smaller one to someone else.⁶ Just as it allows you to care somewhat more about yourself, so it allows you to care somewhat less. (It may do this most when the other person is someone close to you, like your child or spouse. But it also allows it to some degree with strangers.) The theory we've proposed can grant agent-sacrificing permissions if it supplements its prima facie permission to pursue your own good with a parallel prima facie permission not to pursue it, or to omit pursuing your good, where that can involve accepting a harm to yourself. This second prima facie permission can likewise conflict with prima facie duties, such as an impartial duty to promote everyone's good, including your own; and it can sometimes outweigh them. When it does, giving a greater good to yourself and giving a lesser one to another are both all-things-considered permitted. The same type of normative factor, a prima facie permission, that allows acts with less than the best overall consequences in cases of supererogation can also allow them in cases of agent-sacrifice.

Finally, the theory we've proposed can explain how doing a supererogatory act, rather than permissibly omitting it, is in that deontic respect morally superior. On the one hand, the absolute status of the two is the same: both doing the

⁴ See T. Hurka, 'More Seriously Wrong, More Importantly Right', *Journal of the American Philosophical Association* 5 (2019): 41–58.

⁵ Parfit, *On What Matters* (Oxford: Oxford University Press, 2017), iii, 346, 341.

⁶ M. Slote, *Common-Sense Morality and Consequentialism* (London: Routledge & Kegan Paul, 1985), ch. 1.

supererogatory act and not doing it are all-things-considered permitted but not required. On the other hand, there seems to be something especially laudable about going beyond duty. This is often in part a matter of the motive from which a supererogatory act is done. Someone who sacrifices his arm to save another's life usually has a more benevolent, and therefore a more virtuous, motive than if, in the same circumstances, he preferred to preserve his arm. But if you sacrifice your arm to save another person only because he's a billionaire who you hope will give you a large reward, your act, even though less admirably motivated, is still beyond duty and therefore supererogatory. The concept of supererogation is deontic rather than evaluative, one in the same category as right and wrong rather than like morally good or worthy, which depends entirely on motivation. A supererogatory act should therefore be superior deontically as well as, when that is the case, motivationally, and it can be so on the permission-based theory. Though both a supererogatory act and its alternative exercise a prima facie permission, in one case to do the act and in the other not to, the first act also fulfils one or more prima facie duties. And a crucial feature of prima facie duties is that they don't disappear when they're outweighed. They leave what have been called 'traces,' such as the appropriateness of feeling what Ross called 'compunction' about not fulfilling an outweighed duty, as when you break a promise because you're helping accident victims, and sometimes a further duty to compensate the person to whom that duty was owed.⁷ An additional trace, we can now say, is to make a supererogatory act in this respect deontically superior to its equally permissible alternative. Though no better in its absolute deontic status, and independently of any value in its motive, it's better in relation to prima facie moral factors because it fulfils prima facie duties as well as exercising a prima facie permission.⁸

3. Parfit on Supererogation

Whereas this theory weighs normative factors of two different types, prima facie duties and a prima facie permission, the more common theories of supererogation see your permissions simpliciter to do and not do a supererogatory act as arising from a conflict between two factors of the same type, which they often call 'reasons' and take to count, like prima facie duties, in favour of an act's being one you have most reason, or ought simply, to do. The competing reasons, however,

⁷ Ross, *The Right and the Good*, 28.

⁸ Imagine that one person sacrifices his arm when this is his duty, while another does so when his act is supererogatory. The first person's motive may be just as good as the second's; his desire to promote the good impartially may be just as strong, so he too would have done the supererogatory act had that been his situation. Even so, we think the second person's act is all things considered more admirable, because it is beyond duty whereas the first person's isn't.

reflect different 'points of view'⁹ or have different contents and sources. Parfit gives one version of this theory. He recognizes two fundamental categories of reasons: impartial reasons to care about everyone's good equally, and partial reasons to care about your own good or that of people close to you such as your children or spouse. The two kinds of reason can conflict, and in some cases one of them outweighs the other. If you can save thousands of people's lives at minimal cost to yourself, your impartial reason outweighs your partial one and you're required to do the saving. But in other cases neither kind of reason outweighs the other, nor are their weights exactly equal. They then in effect cancel each other out, and you're permitted all things considered both to do the act favoured by impartial reasons and to do the one favoured by partial reasons, so the former is in that way beyond duty. Though the two kinds of reason are to some degree comparable, they're only imprecisely rather than precisely comparable, since sometimes no determinate judgement of their relative strengths can be made. When that's the case, the impartially better choice isn't required but is supererogatory.¹⁰

This is an initially attractive result, but theories of this type have difficulty matching the attractive features of the permission-based one above. Consider first the fact that supererogatory acts occupy a range. This wouldn't follow from Parfit's theory if impartial and partial reasons could be compared precisely. If partial reasons were, say, exactly five times as strong as impartial ones, you would be permitted both to give five units of good to another and to give one unit to yourself. But you would be required to give six units to her rather than one to yourself, since then the impartial reason would outweigh the partial one, and you would be forbidden to give four, three, or two units to her rather than one to yourself, since then the partial reason would weigh more. Hence Parfit's insistence that the two categories are only imprecisely comparable, which implies that a certain kind of transitivity fails. If B is exactly equal to C, then if A is greater than B it's also greater than C. But if B and C aren't determinately ranked, so neither is greater than the other nor are they exactly equal, A can be greater than B but also unranked with respect to C. In our context this means that if the partial reason to retain an arm and the impartial reason to save another's life aren't determinately ranked, the reason to retain two arms can be stronger than the reason to retain one arm but likewise unranked with respect to the reason to save a life. You can be equally permitted to choose between saving a life and keeping one arm and between saving a life and keeping two arms; in both cases the saving is supererogatory.

⁹ S. Scheffler, *The Rejection of Consequentialism* (Oxford: Clarendon Press, 1982), ch. 2.

¹⁰ Parfit, *On What Matters* (Oxford: Oxford University Press, 2011), i. 137–41. Parfit doesn't use the word 'supererogation', but what he describes is equivalent to it.

Parfit's claim of imprecise comparability is overwhelmingly plausible if read epistemically, as a matter of what we're able to judge. If we're asked to specify an exact level of cost to you at which saving another switches from being permitted to being required, we can't do so. But epistemic imprecision is compatible with real or metaphysical precision. It's possible that the normative truth is completely determinate, with exact weights for all normative factors, but we just can't know them. (Ross seems to have taken this view about his *prima facie* duties: he thought we can never know which of two duties is stronger even though there's always an objective truth of the matter.) So if Parfit's theory is to yield a range of objective permissions, where a series of acts are actually supererogatory rather than just ones whose determinate status we can't know, it needs metaphysical imprecision, where the moral truth itself is only partly determinate. That is a much more contentious claim. Some philosophers may deny that any truths can lack full determinacy; reality is always completely precise. But even apart from that it seems problematic to rest the intuitively compelling claim that we're not always required to sacrifice for the greater good, and have a range of permissions not to do so, on an abstruse and contestable thesis in moral metaphysics. Would we really have to abandon our belief in supererogation if we learned that the normative truth is completely determinate? The metaphysical thesis seems of the wrong kind to ground the moral one.

At one point Parfit suggests that the imprecise comparability of his two categories of reason follows from the fact that they differ in kind, the one involving a reference to yourself and the other not.¹¹ But while some philosophers have argued that goods distinct in kind can't be compared at all, that's not his view; he thinks the impartial reason to save thousands of lives outweighs the partial reason to avoid some minimal cost to you. And it's hard to see why a difference in kind that doesn't entirely block comparison shouldn't allow it entirely, as it would given metaphysical precision. Parfit's theory can allow an objective range of supererogatory acts only by assuming a thesis of metaphysical imprecision or indeterminacy that's both some distance from our belief in supererogation and less securely grounded than that belief. Our permission-based theory, in contrast, requires no such thesis. Even if there's a precise point at which the *prima facie* permission to promote your own good is outweighed by the *prima facie* duty to promote another's, say the point where her good is five times as great as yours, you're permitted to prefer a unit of your own to four, three, or two units of hers.

Second, theories like Parfit's that weigh partial against impartial reasons have difficulty accommodating agent-sacrificing permissions. Consider a case where you can give either one unit of good to another person or two units to yourself. Here your impartial reason tells against giving the one unit to the other, since this

will result in less total good than if you had the two. Your partial reason also tells against giving the other the one, since this will result in less good for you. So both kinds of reason tell against giving the other the one, and doing so is all-things-considered or simply wrong: there's no permission to prefer another's even slightly lesser good. Moreover, on a natural reading these theories also deny the permission to prefer another's equal or slightly greater good, as in some cases of supererogation. If the two types of reason aren't determinately ranked, then just as sometimes an impartial reason outweighs a partial one, so sometimes a partial reason should outweigh an impartial one. This can't occur when an act's costs to you are greater than its benefits to others, since then your impartial reason, too, condemns the act. Nor can it occur when the costs are much or only moderately smaller, since then the impartial reason outweighs the partial one or the two aren't determinately ranked. But it can occur when the costs to you are just a little smaller than the benefits, for example if you can either give either five units of good to another or four units to yourself. If partial reasons can ever outweigh impartial ones, they do so in cases like this, and the upshot is that some acts of preferring another's greater to your lesser good are not only not required but positively forbidden. Some acts that would normally be considered supererogatory are wrong.

Parfit doesn't himself draw these conclusions. He acknowledges that some theories of the same general type as his say we're always required to give our own good more weight than a stranger's, but he rejects this claim as 'too egoistic.' He says that if you're in pain and have the only dose of morphine, you're permitted to use the morphine on yourself but are also permitted to give it to a stranger in the next bed if he's suffering the same pain as you, and even if his pain is somewhat less bad, so the benefit to him is somewhat smaller. More generally, he expresses sympathy for a view that allows you to give somewhat more weight to a stranger's good than to your own, though not as much more weight as you may give your own.¹² He therefore allows some agent-sacrifice and presumably also allows, as supererogatory, any act whose cost to you is a little smaller than the benefit it gives another.

These are attractive claims, but it's hard to see how Parfit can accommodate them in the theory he initially proposes, where all-things-considered permissions result from conflicts between normative factors of the same kind, namely reasons that count in favour of an act by tending to make it one you have most reason, or ought simply, to do. This is certainly the way he understands his impartial reasons. When your reason to save thousands of lives outweighs your partial reason to

¹² Parfit, *On What Matters*, I, 139–40. This last claim follows in the permission-based theory if the *prima facie* permission not to promote your own good is weaker than the *prima facie* permission to promote it. Then the first permission is more likely to be outweighed by the duty to promote the good impartially, and permissions all things considered to agent-sacrifice more limited than ones to agent-favour (see Hurka and Shubert, 'Permissions to Do Less than the Best', 13–14).

¹¹ Parfit, *On What Matters*, I, 138–9.

avoid some minor cost, it makes the saving positively required rather than something that's merely allowed. And how can a partial reason of this kind, one that positively favours promoting your good, prefer an act that leaves you with less good than some alternative would? In allowing agent-sacrifice Parfit seems tacitly to understand the relevant partial reason not as one that tends to make for a requirement, as his impartial reason does, but as having merely permissive force, like what we've called a *prima facie* permission. It's not that, in his case of permitted agent-sacrifice, you have a positive reason to give the stranger with the less bad pain your morphine. It's that you lack a positive reason to keep it for yourself, or have a permission other things equal not to keep it. The all-things-considered agent-sacrificing permission isn't grounded in a normative factor of the same kind as an impartial reason, but in one from a different, merely permissive category. And that category is needed because, given only factors with the same requiring force, acts of agent-sacrifice are doubly condemned.

Finally, Parfit's theory has difficulty explaining why a supererogatory act is deontically superior. If you save another's life at the cost of your arm, your act better fulfils impartial reasons than if you preserved your arm, and so is impartially or, as Parfit would say, morally better. But if you don't sacrifice your arm and let the person die, your act better fulfils partial reasons and is therefore partially, or prudentially, better. Each is therefore better by some standard, and there's no ground to value one kind of betterness more than the other, or to say one of the two acts is in some neutral way superior. From a neutral standpoint the two acts are on a par. The permission-based theory, in contrast, weighs factors with different kinds of force but doesn't distinguish them as moral and prudential. It can say a supererogatory act is, though no different in its final deontic status than the alternative, nonetheless preferable because it both exercises a *prima facie* permission and fulfils a *prima facie* duty rather than just exercising a permission.

Relatedly, and returning to our initial topic, a theory like Parfit's has difficulty explaining how one supererogatory act can be more supererogatory than another. In the permission-based theory this difference depends on a difference in the relations between the factors that make for supererogation: an act is more supererogatory when a *prima facie* permission outweighs *prima facie* duties by more. But in Parfit's theory there is and can be no such difference: the relation between the factors it considers is the same in all supererogatory acts. This follows from the way, in his theory, imprecise comparability blocks transitivity. In our first example, your impartial reason to save the other's life isn't determinately ranked against your partial reason to keep one arm, but it also isn't ranked against your partial reason to keep both your arms. That the second partial reason is stronger than the first has no effect on the relation between the two types of reason in these acts, which is the same. And it's in fact the same in all cases of supererogation: in all such cases the relation between the competing reasons is just that of not being determinately ranked, so no differences in these relations

are ever possible. The very feature Parfit needs to generate a range of supererogatory acts, namely imprecise comparability, prevents him from recognizing degrees of supererogation. If it's intuitive that some acts are more supererogatory than others, the fact that theories like his can't allow that is a further strike against them.

4. Non-Supererogatory Admirability

When one of two acts is more supererogatory, they don't differ in their final deontic status: each is one you're all-things-considered permitted to do and also to omit. They differ in certain accompanying properties, so the first is further beyond duty and therefore more heroic or saintly, where these properties in turn make different responses to them appropriate. Thus a more supererogatory act calls, other things equal, for more admiration or praise and may warrant more pride on the part of the agent. This last difference can even be a test of the phenomenon: it can be evidence that one of two acts is more supererogatory that more admiration or praise for it is fitting, and we may tacitly rely on this evidence when we compare the acts. The test must, however, be used with caution, because an act's degree of supererogation isn't the only factor that affects how much admiration it calls for; others can too.

Imagine that you sacrifice your arm to save not one but five people's lives, where the number of lives at stake makes your doing this your duty all things considered, or that a soldier sacrifices his life to save enough of his comrades that, given his special responsibilities to them, the sacrifice is morally required. Though these acts aren't supererogatory, they're surely still to a considerable degree heroic and admirable; the soldier, for example, would surely still be given a medal for bravery. This suggests that another factor relevant to an act's admirability is just the size of the cost you accept in order to act on a moral consideration, regardless of whether that consideration is outweighed, as in supererogation, or outweighs competing considerations, as when an act is obligatory. (This factor may increase the act's admirability either by making its motivation more virtuous or in some purely deontic way; we leave this issue open.) But this can complicate our judgements of degrees of supererogation, and may do so in particular in the three examples in Section 1.

The first example involves one act of saving another's life at the cost of one arm and another act of saving him at the cost of two arms. The second act is more supererogatory because there's a larger gap between the strengths of the *prima facie* permission to avoid its costs and the *prima facie* duty it fulfils; on that basis it's more admirable. But this act is also more admirable because it accepts a larger cost, independently of that cost's role in determining the size of a gap. We suggested in Section 1 that this example may seem especially compelling, and we can now see why. We may tacitly assess the two acts' degrees of supererogation by seeing how far each calls for admiration, and here the second act's greater

admirability has two grounds, with that due to its degree of supererogation reinforced by that of its involving a greater cost.

In the second example, however, there's no such reinforcement. Here one act saves another's life at the cost of one arm for you while the other saves him from losing two arms at the cost of one for you. The second act is more supererogatory because it involves a larger gap between permission and duty, but the greater admirability this makes for isn't reinforced by any greater accepting of costs; the two acts have the same cost. As a result the second act isn't more admirable by as much as in the first example, so an intuitive test for degrees of supererogation yields a less compelling verdict here. The second act still is somewhat more admirable, because it's further beyond duty. But a natural indicator of greater supererogation isn't as clearly present, and the same may be true in the third example. It compares saving another's life at the cost of two arms for you with saving him from losing two arms at the cost of one arm for you. Here gap considerations make the second act more supererogatory—since it's further from being required—and in that respect more admirable, but considerations of accepted costs make the first more admirable, and it may be unclear which act calls on balance for more admiration. Here the second factor not only fails to reinforce the first but positively weighs against it.

There may also be a third relevant factor. It's arguably more important to care about weightier moral considerations, or stronger prima facie duties, and some may say it's therefore more admirable to fulfil these duties. (A contrary view would say it's less admirable to fulfil them, because doing so is more expected or unquestionably demanded.) Accepting this view wouldn't affect our judgements in the first example, where both acts fulfil the same prima facie duty. But it would be relevant to the second, where the less supererogatory act fulfils a stronger duty; there it would generate a ground of admirability that competes with that based on degree of supererogation. It would do the same in the third example, where saving a life at the cost of two arms fulfils a significantly stronger duty than saving two arms at the cost of one. There it might combine with the factor of costs to make the first act on balance more admirable. Even then, however, the second would call in one respect for more admiration, since it goes further beyond duty. It's just that, in this example, the other two factors would make the effect of this particular one especially hard to see.

5. Conclusion

If supererogatory acts are beyond the call of duty, it seems that some can be further beyond duty than others and therefore more supererogatory. We've given an account of this phenomenon within a general theory that sees supererogation occurring when a prima facie permission to promote your own good outweighs a

prima facie duty, for example to promote the good of all impartially. This account uses the same tools as accounts of how other deontic concepts such as wrongness admit of degrees, and therefore is continuous with those. It also strengthens the claims of the permission-based theory of supererogation, in particular against ones like Parfit's that seem unable to allow degrees of supererogation.

Bibliography

- Hurka, T., 'More Seriously Wrong, More Importantly Right', *Journal of the American Philosophical Association* 5 (2019): 41–58.
- Hurka, T., and Shubert, E., 'Permissions to Do Less than the Best: A Moving Band', *Oxford Studies in Normative Ethics* 2 (2012): 1–27.
- Parfit, D., *On What Matters*, vol. i (Oxford: Oxford University Press, 2011).
- Parfit, D., *On What Matters*, vol. iii (Oxford: Oxford University Press, 2017).
- Ross, W. D., *The Right and the Good* (Oxford: Clarendon Press, 1930).
- Scheffler, S., *The Rejection of Consequentialism* (Oxford: Clarendon Press, 1982).
- Slote, M., *Common-Sense Morality and Consequentialism* (London: Routledge & Kegan Paul, 1985).