

Disjunctive Duties And Supererogatory Sets Of Actions

Draft.

For citations, please refer to the published version.

ABSTRACT

I develop a 'duty-plus' approach to supererogation based on a simple intuition: if I am required to do x or y , doing x and y is a candidate for, though not necessarily, supererogation. This is an appealing view to take, located midway between two extreme positions, supererogationism and rigorism. I give a precise statement of the view through the notion of disjunctive duties, and discuss the commitments a duty-plus theorist should make, independent from the Kantian context in which this position is often discussed. I also advocate the novel claim that we should take supererogation as a property of sets of actions, rather than single actions. The latter view suffers from problems in cases of concurrent acts and accumulative supererogation.

1. Introduction

James Urmson famously claimed that all ethical theories which only operate with the three deontic categories of the required, the optional, and the forbidden were 'totally inadequate to the facts of morality'¹ because they fail to recognise a fourth category of actions, which we can call the supererogatory. Supererogatory actions should be seen as 'meritorious non-duty'², as something which is good but in no sense required.³ A number of examples can be provided to make the existence of such a realm outside duty plausible: the soldier throwing herself on a live hand grenade to save her companions, the torture victim who forgives her tormentors even though they do not regret, or the doctor who voluntarily travels to a war zone to treat the wounded.

In this essay, I will defend a position that gives supererogation a place within a three-tiered deontic theory by using the con-

¹ J. O. Urmson, 'Saints and Heroes' in A. I. Melden (ed.), *Essays in Moral Philosophy* (Seattle: University of Washington Press, 1958), 198–216, at 189–90.

² Joel Feinberg, 'Supererogation and Rules', *Ethics* 71 (1961): 276–88, at 282.

³ David Heyd, *Supererogation: Its Status in Ethical Theory* (Cambridge: Cambridge University Press, 1982), 115.

cept of disjunctive duties. The guiding intuition is that some duties can be overfulfilled: if I am required to do *x* or *y*, then doing *x* and *y* is a candidate for supererogation.⁴ On this position, it is a necessary condition for supererogation that a duty has been over-fulfilled. Put otherwise, we can go *beyond* duty, but there is no part of morality *outside* duty. Let's call this the *duty-plus* explanation of supererogation.⁵

The general approach is not new. Kantian philosophers and Kant interpreters have tried to exploit the latitude allowed by 'imperfect' duties to explain supererogation within a Kantian framework.⁶ Other authors deny that the notion of imperfect duty allows a place for supererogation in Kant's, or a Kantian, position.⁷ It is an interesting and difficult question whether and how Kant's framework leaves space for supererogation, but this will not be my question. Instead, I tackle the systematic issue of what the best version of the duty-plus theory is, and what theoretical commitments it should make. Despite the Kant-related literature, these questions are still underdeveloped.

I first describe the general appeal of the duty-plus view in the next section. I limit myself to necessary conditions for supererogation, as I explain in section 3. I develop the notion of disjunctive duty, which I will use as a simplification of the idea of 'imperfect' duty (sec. 4). On this basis, we can state a precise form of the duty-plus theory. A major claim of the paper is that we should see supererogation as a property of sets of actions rather than single actions (sec. 5). I then show how the duty-plus theory can be defended against various objections, and what some of its crucial commitments are (sec. 6).

⁴ *Ibid.*, 63.

⁵ The label is Feinberg's, though he argues against this position; *op. cit.* note 2, 282.

⁶ Thomas Hill, 'Kant on Imperfect Duty and Supererogation', *Kant-Studien* 62 (1971), 55–76; David Heyd, 'Beyond the Call of Duty in Kant's Ethics', *Kant-Studien* 71 (1980), 308–24; Richard McCarty, 'The Limits of Kantian Duty, and Beyond', *American Philosophical Quarterly* 26 (1989), 43–52.

⁷ Marcia Baron, 'Kantian Ethics and Supererogation', *Journal of Philosophy* 84 (1987), 237–62; Daniel Guevara, 'The Impossibility of Supererogation in Kant's Moral Theory', *Philosophy and Phenomenological Research* 59 (1999), 593–624; Jens Timmermann, 'Good but Not Required?—Assessing the Demands of Kantian Ethics', *Journal of Moral Philosophy* 2 (2005), 9–27.

2. Situating the Duty-Plus View

Let me start by placing the duty-plus view in the wider debate about the nature of supererogation to see why it's worth developing. We can roughly distinguish three approaches to supererogation,

Supererogationism. There are supererogatory actions, and *all* (or at least, *some*) are *outside* duty: i.e., neither required by perfect duty, nor the overfulfilment of some imperfect duty.

Duty-plus: There are supererogatory actions, and *all* are *beyond* duty: i.e., the overfulfilment of some imperfect duty.

Rigorism: There are *no* supererogatory actions, at least not in any strong sense of the word: all seemingly supererogatory actions are demanded by perfect duties or by open-ended imperfect duties.⁸

The duty-plus view is extremely appealing because it is a midway position between the two extremes of supererogationism and rigorism, and promises to combine the best features of both.

Supererogationists divide morality into two parts: Duty requires from us only a minimal contribution such that life in society does not become nasty, brutish, and short. Anything beyond these minimal duties belongs to the 'higher flights of morality'⁹. This realm is important, as it enables us to express our personality and highlights the value of autonomous choice.¹⁰ The duty-plus view incorporates some of these intuitions. Imperfect duties allow for latitude such that we can express our character, and they give space for, and embrace, autonomous choice *inside duty*.

However, against supererogationism, the duty-plus view denies that supererogation forms a wholly separate plane of higher morality. Supererogation is not seen as an area of moral achievement distinct from duty, but continuous with it. The

⁸ This corresponds roughly with the distinction made (under different labels) in David Heyd, 'Supererogation' in E. Zalta (ed.) *Stanford Encyclopedia of Philosophy* Winter 2012, 2012
<<http://plato.stanford.edu/archives/win2012/entries/supererogation/>>.

⁹ Urmson, *op. cit.* note 1, 211.

¹⁰ Heyd, *op. cit.* note 3, 172–8.

duty-plus view thus holds on to the idea that the moral realm is unified. Furthermore, supererogation in this view is a derivative idea that is secondary to duty.

To express it metaphorically, we put duty first in the moral space. In this, the duty-plus view agrees with rigorism. But it rejects the rigorist's claim that duty fills *all* of the moral space. Reaching the boundaries of duty is a logical and practical possibility, even if knowing the precise shape of these boundaries might be hard, and even if a good-willed agent might not care for these boundaries.¹¹ But unlike the rigorist, the duty-plus theorist is not revisionist about judgments of supererogation, and accepts that supererogation is an interesting and genuine moral phenomenon.

The duty-plus view is thus a theory which can account for much of the appeal of the other two theories, without accepting some of their excesses. This is why we should try to state the best possible form of it, which I will aim to do in this paper.

3. Necessary and Sufficient Conditions

Before moving on, let me note an important limitation of the paper. Imagine that Lena decides to save three children out of a burning orphanage. Once she is in the building, Lena refuses to save a fourth orphan who she leaves behind, even though doing so comes at no additional risk to her. While entering the orphanage seems supererogatory, not saving the fourth orphan is not supererogatory, and even morally forbidden. This is an example of a 'gap' or 'hole' in the scale of supererogation.¹² Even though Lena overfulfils her duty, her actions are morally prohibited.

There are other ways in which over-fulfilment is not sufficient for supererogation. One important factor in supererogation is to have the right motives or intentions.¹³ If Ramsay, in the same situation as Lena, saves the children only because he intends to abduct them, we might plausibly say that Ramsay did not act in a supererogatory manner.

¹¹ I will explain these last caveats in subsections III and IV of sec. 6.

¹² As discovered and analysed in-depth in Ulla Wessels, *Die gute Samariterin: zur Struktur der Supererogation* (Berlin: de Gruyter, 2002).

¹³ See Alfred Archer, 'Supererogation and Intentions of the Agent', *Philosophia* 41 (2013), 447–62.

Furthermore, there might be other duties relevant to a given case. Consider Ned, also in Lena's circumstances, who is a single father of five children whose lives would be greatly impaired if he died. If there are other bystanders willing to save the orphans, then Ned might be under a duty to his children *not* to risk his life, and him saving the orphans might fail to be supererogatory for that reason.

These examples should be enough to show that it's hard to state sufficient conditions for supererogation. The focus in the rest of this paper will be on stating necessary conditions for supererogation. Thus, I presume that the duty-plus view is committed to the claim that all supererogation is an instance of duty-overfulfilment, but not the inverse.

4. Disjunctive Duties

In the Kantian literature mentioned, supererogation is approached through the notion of imperfect duties. There is a great variety of ways to understand this idea: Rainbolt lists eight different meanings, Schumaker 25 which he organises under three major headings.¹⁴ This gives us an embarrassment of riches, and much of the Kant-oriented literature is an attempt to clarify what imperfect duties are. I think we can bypass much of this discussion. Our basic intuition was that doing x and y was possibly supererogatory when there is a duty to do x or y . So a natural starting point is the idea of *disjunctive duties*.

A disjunctive duty is the duty to perform (at least) one member of a non-singleton, non-empty set of actions A . For the members of this set I will assume that they are sufficiently different such that agents can distinguish between them. We should also stipulate that there is no perfect duty to perform any one of the members of A . The members of A can be, but not necessarily have to be, mutually exclusive.

The term *duty* is understood as duty proper and left as a primitive. It applies to any feasible logical connection of actions, such that there can be a duty to perform a , to perform $a \wedge b$, to perform $(a \vee b) \wedge c$, and so on. A *perfect duty* with regard to a set

¹⁴ Millard Schumaker, *Sharing without Reckoning: Imperfect Right and the Norms of Reciprocity* (Waterloo: Wilfrid Laurier University Press, 1992), 4–6; George Rainbolt, 'Perfect and Imperfect Obligations', *Philosophical Studies* 98 (2000), 233–56, at 233.

of actions A means that there is a duty to perform *all* members of A .

We can formalize the concept of disjunctive duty now as:

(D) An agent I has a disjunctive duty with regard to the set of actions A if, and only if, I has a duty to do any (i.e., one) of the members of the set A where the above constraints for A apply.

Put more formally, if we write the members of A as a_1, \dots, a_n , and 'having a duty' as $D(\cdot)$, then a disjunctive duty with regard to A is equivalent to $D(a_1 \vee \dots \vee a_n)$ ('wide-scoping').

Note that the property of 'being a disjunctive duty' only applies to *sets* of actions, not to any particular act. If there is a disjunctive duty with regard to A , let us say that each member a of A has the property of being *disjunctively required* (as a member of A ¹⁵). For clarity's sake, I will reserve the concept of disjunctive duty to sets of actions, and use 'disjunctively required' exclusively for single actions.

One sometimes encounters definitions of imperfect duties in the language of *act-types*.¹⁶ The difference between such definitions and mine is that between *classes* and *sets*. I take it that for all realistic choice situations, we can find a finite set of actions that corresponds to a class (i.e., act-type) of actions. Therefore, if one considers imperfect duties over act-types, one can equally well understand them as disjunctive duties over a set of acts that share some common characteristic. That all actions in a set required by a disjunctive duty have some shared feature will often be a convenient assumption, but we need not narrow the conceptual framework in this way.

An immediate extension of (D) is to say that a disjunctive duty with regard to a set A means that we have the duty to do any number m of the n acts in A , where $0 < m < n$.¹⁷ I will assume this generalised form throughout the paper.

As compared to the rich idea of 'imperfect' duty, disjunctive duties are a greatly simplified concept. Actual moral agents

¹⁵ One action can be disjunctively required as a member of different sets.

¹⁶ E.g., Michael Stocker, 'Acts, Perfect Duties, and Imperfect Duties', *Review of Metaphysics* 20 (1967), 507–17, at 508–9.

¹⁷ For similar definitions, see *ibid.*, 508–9, and Daniel Statman, 'Who Needs Imperfect Duties?', *American Philosophical Quarterly* 33 (1996): 211–24, at 213.

will rarely if ever encounter duties with such neat logical form in their practical deliberation. Nevertheless, disjunctive duties can be seen as the skeletal bones of more realistic duties, and are thus a helpful theoretical simplification. Take a classic example of an imperfect duty, the duty of charity. Right intentions, manner of giving, choice of beneficiary, and social conventions will play a role in an agent's deliberations about this duty. Nevertheless, reduced to its bare bones, this duty gives us a set of charitable acts, of which we are required to do a certain subset, but by no means all or the biggest possible subset.

5. Supererogation as a Property of Sets of Actions

It is the mainstream view to see supererogation as a property of actions. This is reflected in common definitions of supererogation,¹⁸ and the classic examples used to illustrate it. Against this view, I will argue that we should see supererogation as a property of *sets* of actions. (This entails the more cautious claim that we *can* see it as such.) This is a conceptual point which should be accepted even by opponents of the duty-plus view, but I will discuss it specifically in the context of that theory.

Given the definition of disjunctive duties in (D), we can formulate an *act-focussed* phrasing of the duty-plus view:¹⁹

- (A) If an act *a* is supererogatory, then
 - (i) there is a disjunctive duty with regard to some set *B*,
 - (ii) *a* is a member of *B*,

¹⁸ E.g., Heyd, *op. cit.* note 8; Thomas Hill and Adam Cureton, 'Supererogation' in H. LaFollette (ed.), *International Encyclopedia of Ethics* (Chichester: Blackwell, 2013), vol. 8, 5070–78; Gregory Trianosky, 'Supererogation' in E. Craig (ed.), *Routledge Encyclopedia of Philosophy* (London: Routledge, 1998), vol. 9, 232–5.

¹⁹ (Sets of) omissions can also be supererogatory. Assume that a homeless person starts living in Sindhu's garage. Sindhu has not invited or encouraged that person to stay, but she also does nothing to expel the homeless from her garage. This, plausibly, is a supererogatory omission. But for simplicity's sake I will only talk of 'actions' rather than 'actions and omissions' in the main text. (I thank an audience member at the Dublin conference for pointing this out.)

(iii) some other actions have been done which collectively fulfil the duty with regard to B .

I follow an intuitive usage in which we say that a disjunctive duty is *fulfilled* if the number of acts it requires (or more) have been done.

On a *set-focussed* phrasing, the duty-plus view looks as follows:

- (S) If a set of acts A is supererogatory, then
- (i) there is a disjunctive duty with regard to some set B ,
 - (ii) A is a subset of B ,
 - (iii) there is a proper subset of A which fulfils the duty with regard to B .

If we return to the idea that a disjunctive duty requires that we do m out of n actions, this means that any of the sets containing $(m+1)$ to n of those actions are candidates for supererogation. Inside the act framework, the first m actions are not supererogatory and count as fulfilling our duty, while the $(m+1)$ -th to n -th actions are candidates for supererogation.

I. Supererogatory Sets

I will argue for the superiority of (S) over (A), but let me first explain the idea of a supererogatory set of actions. We can start with two examples about other moral concepts:

It was good that Aaron failed to meet Benedict at the promised time, and instead helped Christine move house.

It was wrong for Dimitri to give Enzo a birthday present, but not Frances.

Note that it was *not* good that Aaron failed to meet Benedict, and it was *not* wrong for Dimitri to give Enzo a birthday present. In both cases, we are saying something about both acts *considered together*. It's wrong to give one good friend a birthday present *and* not to give one to another, but there is nothing wrong about either of these actions on their own.

So there is an intuitive sense in which we apply moral predicates such as 'good' and 'wrong' to sets of actions. There is no reason to think the case would be any different for supereroga-

tion, especially insofar as we define that notion in terms of these other predicates.²⁰

Another observation is that the same behaviour can be described in different ways, obscuring any hard boundaries between single actions and sets of actions. Consider the following two claims which both refer to the same conduct:

It was supererogatory that Giulia helped Henrik get his life back on track.

It was supererogatory that Giulia let Henrik live at her house, helped him get rid off his drug addiction, and encouraged him to find a job.

Sometimes we easily slide back and forth between talking about single actions and sets of actions. It would be strange if moral predicates properly only applied to the former. This should give some intuitive grip on how the notion of a supererogatory set of actions is possible.²¹ Note that the idea of a 'set' only appears in our philosophical analysis; agents need not explicitly think about sets when they deliberate.

II. Concurrent Acts

Let me now turn to objections against the act-framework. First, there is a problem with concurrent acts. Imagine that there is a disjunctive duty to do x or y and assume that an actor is doing x and y at the same time. For example, Stannis has a duty to give to one of two charities. He uses online banking, fills in the details of both charities, and by clicking 'Enter' sends money to both simultaneously.

The defender of (A) has three choices. She can say that both acts of charitable giving are supererogatory, that only one of the two is, or that none of them is. The first option is not plausible. In giving to both charities, Stannis fulfils his disjunctive duty to give to either. So one of the two actions must count as fulfilling his disjunctive duty. Even if we claim that we do not *know* which of these two actions fulfils the disjunctive duty, we are committed to the conceptual claim that there are actions which fulfil a duty *and* are supererogatory. Given that the min-

²⁰ Derek Parfit calls it a 'mistake in moral mathematics' to ignore the effects of 'sets of acts'. but he has inter-personal, not intra-personal cases, in mind (*Reasons and Persons* (Oxford: Clarendon Press, 1984), sec. 26).

²¹ I thank an anonymous reviewer for pressing me on these issues.

imal content of the notion of ‘supererogation’ is to ‘go beyond duty’, this is conceptually implausible.

The claim that none of the two actions is supererogatory is equally troubling. There is something supererogatory about giving to two charities, and so an act framework should say that at least one action being done is supererogatory (or a candidate for supererogation). Supererogation has to be *located* somewhere in the example.

The act-theorist might deny that there are two actions being done here – instead, there is one action, which is Stannis giving to two charities. But this reply makes the situation worse: if this is one action, it is both required by the duty to be charitable *and* it is supererogatory, which again leaves us in conceptually awkward territory.²²

Thus, the only feasible option for the act-theorist is to claim that one of the two actions is supererogatory, and the other one isn’t. The main problem here is that any such determination will be arbitrary. Assume that Stannis’ two charitable givings are symmetrical in all morally relevant aspects: they do the same amount of good, are equally burdensome to Stannis, and so on.

A first line of reply for the act-theorist is to claim that perfectly symmetrical cases are rare. In most realistic cases, the actions being done differ in their non-temporal features, such as how burdensome they are. We can decide on the basis of these other properties which actions to classify as supererogatory in concurrent actions – e.g., the more burdensome action is the one which is supererogatory.

But on what basis do we pick the properties which serve as tie-breakers in such cases? We could pick those properties which are associated with supererogatoriness – e.g., being burdensome. But if we do so, we move beyond theoretical resources internal to the duty-plus view. Rather, we invoke additional intuitions about the supererogatory.

The basis on which we design our tie-breaking criterion looks *ad hoc*. We are interested in explaining supererogation, so appealing to intuitions about what is supererogatory is dangerous at this point. Second, there is the comparative point: the set framework is preferable here because it does not *need* to give a

²² I will return to this case in sec. 6.

tie-breaking criterion in the first place. Thus, theoretical simplicity favours the set framework.


A second strategy is to embrace arbitrariness. In Stannis' case, we should say that one charitable giving is supererogatory, and the other one isn't. Perhaps it is unknowable which one is which, but we know that *one* of them is, and the other isn't. One motivation for this strategy is that we 'do not care' precisely which of the two actions is supererogatory. We only care that the agent acted in a supererogatory manner.²³

This reply, however, makes it hard to see why we should remain act-theorists. If we do not care which actions are supererogatory, why insist on a theory that makes arbitrary stipulations? The set framework can account for the intuition that the agent acted in a supererogatory manner, without incurring this theoretical cost. According to (S), $\{x,y\}$ is (possibly) supererogatory. Both x and y are constitutive members of a supererogatory set, but the question whether x is supererogatory, or whether y is, does not arise. Thus, we need to make no arbitrary stipulations.

III. Accumulative Supererogation

Let's turn to a second class of cases. It is plausible that having and exercising certain virtues beyond some level is supererogatory.²⁴ Most virtues express themselves in many different actions, most of which are small and insignificant. For example, kindness includes holding doors open for others. Call such cases *accumulative supererogation*: we go beyond the call of duty not in one big step, but through many small steps. Other examples of accumulative supererogation are being generous, charitable, friendly, helpful, and so on.

Let's start with a simplified case. Imagine that Renly has a duty to be kind. Consider the following two patterns,



Each of the two time lines shows a continuous pattern of kind behaviour, with the inset lines showing when a kind act has

²³ This reply has been suggested to me by David Heyd.

²⁴ For the objection that it isn't, see below.

been done. All of these are comparatively insignificant actions – being friendly with one’s barber, complimenting acquaintances, tipping the cab driver well, and so on.

Assume that pattern I describes (one instance of) a minimally duty-fulfilling pattern. If Renly does kind acts roughly this often and at roughly those intervals, he fulfils his duty of kindness. Assume that Renly acts as described by pattern II – that is, he is kind much more often than required of him.

A first objection is that each of Renly’s actions is *too morally insignificant* to deserve the label ‘supererogatory’. Without trying to give a complete account, moral significance implies praiseworthiness, and is a fitting occasion for positive reactive attitudes; it goes hand in hand with merit, and moving beyond the socially expected; and significance is correlated with conferring benefits on others, or imposing burdens on the agent.

But each single act of kindness, considered on its own, is only a small nicety which does not have any of these features. Holding open a door for someone imposes a burden which will often be not even perceived as such, as will be the benefit to the other person. It seems odd to praise an agent for holding open a door. Some reactive attitudes might be fitting in response to such acts – e.g., gratitude – but the fitting attitudes will be very ‘low-key’, and unlike the stronger reactive attitudes associated with moral significance. Lastly, small acts of kindness are common and expected.

One reply is to severely weaken the predicate ‘supererogatory’, and reduce it to merely mean ‘beyond duty’. Under this definition, when we call holding open a door supererogatory, nothing is said about merit, praiseworthiness, sacrifice, or any other sense of moral significance. We allow for ‘trifling’ supererogatory actions.²⁵

Severing the link between supererogation and moral significance would make the extent of the supererogatory both too large and uninteresting. For example, going for a long Sunday afternoon walk is ‘beyond duty’. But we exclude it from the scope of the supererogatory because it is morally insignificant.

²⁵ This might be Urmson’s response. See discussion below, and Lorenne Burchill, ‘In Defence of Saints and Heroes’, *Philosophy* 40 (1965), 152–7, at 152–3.

Even if you think that weakening the predicate 'supererogatory' is a plausible move, we should again note that an act-based theory incurs theoretical costs here which the set-theory does not have. The actions done in pattern II, *considered together*, are *highly* morally significant. Doing *all* of these actions is meritorious and praiseworthy; it does much good for others; and so on. Thus, this *set* of actions is morally significant. The set theorist does not need to weaken the predicate 'supererogatory', because she can highlight the moral significance of the set of actions.

Another reply is to deny that kind actions, *even considered collectively*, are supererogatory.²⁶ One might argue that kindness is too far away from paradigm cases of supererogation, such as Urmson's doctor volunteering to travel to the plague-ridden city. It is neither heroic nor saintly.

In reply, let me start with a point about philosophical historiography. While everyone remembers Urmson's examples of heroic supererogation, he also allows that '[i]t is possible to go just beyond one's duty by being a little more generous, forbearing, helpful, or forgiving than fair dealing demands', and he goes on to explain that he drew 'attention to the heroic and saintly deed [...] merely in order to have conspicuous cases of a whole realm of actions that lie outside the [classic deontic] trichotomy'²⁷. So the article setting the terms of the debate on supererogation allows generosity and other virtues to count as supererogation.

I think this is correct. There are various ways how we can define supererogation – e.g., as what it is praiseworthy to do but blameless to omit, or what is good to do but not wrong not to do. Whichever definition you prefer, we are interested in a peculiar 'realm' of morality. Heroes and saints are paradigm cases in the sense that they first *pointed us towards* this realm. But once we have it firmly in view, there is no *a priori* guarantee that the cases falling into this realm will look very much like our initial paradigm cases.

Kindness and other virtues can robustly fulfil various conceptions of the supererogatory. For example, exhibiting these virtues to a high degree is good but not required, and praiseworthy to do but not blameworthy to omit. Thus, we should classi-

²⁶ I thank an anonymous reviewer for raising this point.

²⁷ Op. cit. note 1, 205.

fy them as supererogatory, despite their dissimilarities with Urmson’s heroes and saints.

So much for moral significance. My second objection is again from arbitrariness. With regard to patterns I and II, the act-theorist needs to say which of Renly’s actions are supererogatory and which aren’t. But there is no non-arbitrary way to make this distinction. First, we might take temporal location as a criterion, and say that the actions done at a later time are (possibly) supererogatory, and the actions done earlier merely fulfil the relevant duty. For example, in pattern II x would fulfil a duty, while y and z would be candidates for supererogation.

But given that this is a continuous pattern of behaviour, there is no natural point at which the disjunctive duty starts to apply or ends. We might shift the time frame to the left, such that y and z are at the beginning of it; now they would cease to be candidates for supererogation. Because picking a starting point is arbitrary for many duties, determining supererogatoriness with reference to such a starting point is arbitrary as well.

A defender of (A) might again point to non-temporal features of the actions to determine their supererogatoriness. Some actions are more morally significant than others. Assume pattern III describes Margaery’s charitable giving, where duty is again described by pattern I.



The differing heights of the inset lines indicate the varying moral significance of Margaery’s acts – e.g., Margaery’s contributions to charity differ in how burdensome they are to her, or how much good they do. We might now claim that Margaery’s *big* contributions are supererogatory – e.g., v and w in pattern III.

But it’s not clear why we should pick out *those* actions and not others. Imagine the opposite view. We argue that Margaery’s big contributions fulfil her duty, and it’s the many small things she does which are supererogatory. I can see no principled, non-question-begging reason why we should not choose this way of distinguishing the supererogatory from the non-supererogatory.

The set framework avoids the arbitrariness problem. (S) classifies the *set* of actions done in patterns II and III as (possibly) supererogatory. The set theorist does not embark on a difficult search for categorising some of Renly's and Margaery's actions as supererogatory, and some as duty-fulfilling. Rather, it is *Renly's kind behaviour* and *Margaery's charitable giving*, formally analysed as the sets of actions they do, which is supererogatory.

IV. Other Advantages

There are three further advantages of a set-based over an act-based duty-plus approach to supererogation.

First, the set framework better allows us to account for standard ways of talking about supererogation. In reading the biography of a saint, we might single out no one particular action which we think was 'the saintly one'. Rather, we say that 'the saint's life', or 'the things the saint has done', were supererogatory.²⁸

Second, David Heyd has objected that any account which explains supererogation in terms of imperfect duty makes a category mistake, as the concept of imperfect duty only applies to sets of acts (or act-types), whereas supererogation is a quality of particular acts.²⁹ Assuming there is some form of category mistake, (S) avoids it (while (A) commits it), as both disjunctive duty and supererogation are defined for the same category, i.e., sets of actions.

A third advantage is more speculative. Assume that Susie's house burns down. In response, she receives an overwhelmingly supportive response from her neighbours. All of them help her out in small ways; but her neighbours, *collectively*, help her so generously that her loss is significantly reduced. I find it tempting to think that this is a case of *collective supererogation*: while no particular neighbour went beyond the call of duty, the collective support which Susie received was supererogatory. In this case, a set of actions done by different agents is supererogatory.³⁰ Whether collective supererogation exists is admittedly controversial, and the possibility has not been ex-

²⁸ See also the examples in subsection I.

²⁹ Op. cit. note 3, 63.

³⁰ This is inspired by Parfit-style cases; see op. cit. note 20.

plored to my knowledge. However, a set framework is in a good position to account for such cases.

6. Supererogation inside Disjunctive Duty

I have argued for the conceptual claim that we can and should see supererogation as a property of sets of actions. In the rest of this paper, I will sketch ways of defending the substantive claim that all supererogatory actions can be seen as over-fulfilling a disjunctive duty.

I. Individuating Duties

The defender of the duty-plus view is committed to the claim that all supererogatory sets are subsets of a set for which there is a disjunctive duty (conditions (i) and (ii) of (S)). In response, consider what Feinberg claims about Urmson's case of a doctor who volunteers to travel to a plague-ridden city:

'[the] heroic doctor is not simply doing his "duty plus more of the same." He does not travel a definite number of miles more than the total required by duty [...]. [H]e has no duty to travel one step toward the plague-stricken city or to treat one single victim in it.'³¹

We can outline the following supererogationist objection. There are certain actions that are supererogatory and which do not fall under *any* duty. These (sets of) acts are 'in *no* way obligatory'³², i.e., not even covered by any disjunctive duty. Re-phrased in the context of (S), there are some supererogatory sets of actions for which (i) is false.

There is an obvious reply: there *are* various disjunctive duties which cover the current case. First, there is the duty to help people in the plague-stricken city. If you think there is no such duty, consider the (role-based) duty the doctor has to help the sick and needy. And lastly, consider the generic duty to help others. If we just look widely enough, we will find a duty which covers the given case. Feinberg is right that there is no 'duty to travel one step toward the plague-stricken city', insofar as there is no perfect duty which requires *this particular* ac-

³¹ Op. cit. note 2, 280.

³² Heyd, op. cit. note 3, 125. Emphasis in original.

tion. But it does fall under the scope of several disjunctive duties.

This strategy generalises. For any case the supererogationist presents, it is likely that we will find some general disjunctive duty which covers it.³³ If this is the dialectical situation, it seems that the duty-plus theorist has the upper hand.

Feinberg might reply that this is a gimmicky way of describing the travelling doctor's situation. But what Feinberg needs is more than an appeal to intuition. Rather, he needs a full account which explains how duties are individuated, and how they apply to particular cases.

Consider now the following claims:

- (1) The doctor has a duty to help others.
- (2) Travelling to the plague-ridden city is an instance of helping others.
- (3) Travelling to the plague-ridden city fulfils the doctor's duty to help.

The first two propositions are natural assumptions. Denying (1) commits the supererogationist to an implausibly minimalist view of what duties we have. (2) is obvious. If we denied (3), this would have the counter-intuitive implication that, after the doctor returns from his trip, he has not fulfilled his duty to help, and we can blame him for failing to fulfil his duties, whereas we should do the opposite, and praise him. Moreover, (3) seems an implication of (1) and (2).

Here's a reply I can imagine on Feinberg's behalf. While we acknowledge (1)-(3), we might also claim that

- (4) Travelling to the plague-ridden city is not required by the doctor's duty to help others.

Consider an analogy. You have a duty to support your country. If you become a soldier, you fulfil that duty. But the duty to support your country did not require you to become a soldier. Similarly, we might say that the heroic doctor *fulfilled* his duty to help others, but that he was *not required* by that duty to travel to the city. Thus, the supererogationist can say that su-

³³ Guevara thinks that for this reason we should not define supererogation in terms of acts which do not fulfill any duties (op. cit. note 7, 595–7).

pererogation can be duty-fulfilling, but it is not required by duty.

We should note that this reply already concedes some ground to the advocate of duty-plus. In this modified supererogationist position, we acknowledge that something can be supererogatory while 'fulfilling' a duty. It's not clear that this is compatible with the supererogationist intuition that there are some areas of morality which fully fall outside the realm of duty.

More importantly, it's hard to make the current distinction in deontic logic. If some action x 'fulfils' a disjunctive duty D this is naturally read as the claim that it was one of the disjuncts of D . But if x is one of D 's disjuncts, then it is disjunctively required by D in the sense I defined. The distinction between 'fulfilling a duty' and 'being required by a duty' is thus likely to collapse; it might be merely a restatement of the supererogationist's intuitions in a question-begging way.

II. Individuating Actions

In a second set of counter-examples, an act is supererogatory, although not more than the number of acts necessary for fulfilment of a disjunctive duty have been done (condition (iii) of (S)).³⁴ Consider Ayewa, the soldier who throws herself on a live hand grenade to save others. This act is supererogatory, but it seems that (S) cannot give it supererogatory status, as *only one* act is done.

We here encounter the problem of how to individuate actions. Remember Giulia, who helped Henrik get his life back on track. We could describe her behaviour as if it were one action (help Henrik) or as if it consisted of several actions (give him a place to stay, help him overcome addiction, and so on). Which description should we prefer? Which carves the situation at its metaphysical joints? Are both admissible descriptions of the situation?

These are extremely difficult questions, and I will make no attempt at an answer. Still, there are some responses the set-theorist can give. It seems to me that the advocate of (S) is not

³⁴ This is a modification of an argument in Gregory Mellema, *Beyond the Call of Duty: Supererogation, Obligation, and Offence* (Albany: State University of New York Press, 1991), 34–6. Cf. Gregory Mellema, 'Supererogation and the fulfillment of duty', *Journal of Value Inquiry* 25 (1991), 167–75, at 171–3.

committed to the claim that on *any* possible way of individuating duties, and on *any* possible way of individuating actions, conditions (i)-(iii) of (S) are given. This would be much too demanding. Rather, she should claim that some behaviour is a candidate for supererogation if there is *some* (plausible) way to describe our duties and *some* (plausible) way to describe the behaviour such that conditions (i)-(iii) of (S) apply.

A corollary of this claim is the following. Some ways of describing our duties and actions come more naturally to us than others. The duty-plus theorist is not committed to the claim that under the best or most natural description of what is happening, the agent's acting fulfils (S). So it is no direct objection if you think that Giulia's behaviour should best be described as doing one action, as long as you agree that the alternative description is at least *plausible*. The same strategy can be applied to a wide range of cases.

Ayewa's example provides a hard case where individuating her actions as several might not look very plausible. She throws herself on the grenade; that's it. To answer this case, let's start with a description of her duties: Ayewa has a duty to tolerate a small risk to make the survival of her comrades more probable. This duty is disjunctive, as there are many slightly risky ways in which she can fulfil it. So first of all, if Ayewa has just done *anything* else covered by this duty, aside from throwing herself on the grenade, we can count her case as supererogatory according to (S). This will be the case for most realistic scenarios.

Let's now assume this is not the case – Ayewa has just entered the army, say. As a second line of reply, we might re-describe the action itself. Throwing oneself on a live hand grenade is the acceptance of a very great risk to secure one's comrades' survival. It is taking on the standard risk everyone has to accept, *and in addition* making that risk almost certain death. It is doing both of these things which is supererogatory, and so is covered by (S).

As a last resort, we might bite the bullet: no view can live without theoretical costs. If Ayewa has done *nothing* else to fulfil her duties as a soldier, and throwing herself on the grenade is the only *one* thing she has done, then perhaps we should reject that she acted in a supererogatory manner. She simply fulfilled her duty. We can still say that Ayewa acted in a very praiseworthy and heroic way – acting heroically is nei-

ther identical with supererogation, nor sufficient for it. Again, supererogation is a technical category in our moral schema: once we have discovered its existence, we should not expect that its contents are necessarily what we expect them to be, or originally intuited them to be.

III. Doing One's Supererogatory Share

Another set of objections concerns the compatibility of supererogation and thresholds. A first worry is this: According to the duty-plus view, if you do one more action than required by disjunctive duty, you act supererogatorily. But this makes supererogation too easy to achieve. It's hard to see why doing barely more than duty requires could even possibly be supererogatory.

However, nothing in the duty-plus view conceptually requires that we identify two thresholds: (1) the threshold for fulfilling a duty, and (2) the threshold for supererogation. Thus, we can introduce two thresholds, m_D and m_S , such that $0 < m_D < m_S < n$, where n is the number of disjuncts in the duty. After doing m_D acts, the disjunctive duty is fulfilled; after doing m_S acts, the conjunction of all acts becomes a candidate for supererogation.

But perhaps there is a lingering, deeper worry. If we view supererogation as fulfilling a threshold, it becomes possible to be 'minimally supererogatory'. One might object that this is self-refuting, as supererogation excludes the idea of being moral in any minimal way. Imagine Peter 'does his supererogatory share', and stops afterwards, pointing towards his supererogatory achievements. Few would say that Peter acts in a supererogatory way.

The best response for the duty-plus view is to distinguish the motivations of the agent from the (deontic) classification of the action. If it was part of Peter's motivation that he won't do more than his supererogatory share, this introduces a motive which likely undermines the supererogatory status of his actions.

Consider instead Elo's case. She has also minimally gone beyond the supererogatory threshold. But she has been motivated to do her best, and it has been an accident that she was only minimally supererogatory – perhaps she has been quite busy, and despite her best efforts, had only few opportunities and resources to give to charity. In her case, it is much less counter-intuitive to think that she acted in a supererogatory manner.

IV. Perfectible Duties

One last objection stems from the rigorist view that imperfect duties are *open-ended* or *imperfectible*: We can never fulfil them.³⁵ Statman, for example, writes that ‘charity is the sort of duty we are never really freed from [...]’³⁶. We are justified in ceasing to give to charity only if more stringent duties require us to do so.

As I have described disjunctive duties, they can be fulfilled, so the duty-plus advocate must reject rigorism. Whether we accept rigorism is likely to depend on the deeper view of morality we have, but let me offer some observations which soften the blow of rejecting rigorism.

First, the idea that duties can be fulfilled should not be mistaken for the idea that morality is undemanding. (S) can be used to describe a lenient ‘Yuppie ethics’³⁷ approach to morality, but also a very demanding theory of morality. Everything depends on where one sets the relevant thresholds. Supererogation might be common, or rare, or contingently, there might be no supererogation at all. In other words, the duty-plus view is a structural, not a substantive view of morality.

Second, I introduced the idea of disjunctive duties as a simplification. ‘Real-life’ duties will be subject to indeterminacy and uncertainty. We will often be sceptical about claims that our duties have been fulfilled. Agents who claim to have fulfilled their duties might often be mistaken, and teaching children that duties have thresholds might be harmful. The duty-plus advocate can accept this. She only claims that *as a matter of principle* duties can be fulfilled. She need not take any position on the determinacy, knowability and so on of actual imperfect duties.

Third, we should again stress the difference between motivation and deontic classification. Agents who think of their duties as annoying hurdles which are in one’s way might be morally criticisable in various ways.³⁸ The defender of the duty-plus view, however, is not committed to the claim that it’s good if agents are motivated in such ways, and she might in-

³⁵ E.g., Baron, *op. cit.* note 7; Susan Hale, ‘Against supererogation’, *American Philosophical Quarterly* 28 (1991), 273–85.

³⁶ Statman, *op. cit.* note 17, 214.

³⁷ Marcia Baron’s term.

³⁸ Statman, *op. cit.* note 17, 212–3; Baron, *op. cit.* note 7, 250–2.

clude proper motivation as a further necessary condition for supererogation.

7. Concluding Remarks

The two major points of the paper are the following. First, I advocated a novel conceptual claim that we should see supererogation as a property of sets of actions, because it better allows us to understand concurrent supererogatory actions and accumulative supererogation. This is a strong claim, and as far as I know, there is no precedent for it in the literature. But I hope to have shown that the idea cannot be easily dismissed. If we want to keep an act-based explanation of supererogation, some amendments to standard views will be necessary.

Second, I defended the duty-plus view. I provided a definition of this view through the combination of (D) and (S). I pointed towards four issues which will be crucial for a full assessment of the duty-plus view: (1) the individuation of duties, (2) the individuation of actions (perhaps the crucial issue), (3) the distinction between motivation and deontic classification, and (4) the (im-)perfectibility of duty. More would need to be said on each of these questions. But for now, I believe that the duty-plus view of supererogation emerges as a serious contender.³⁹

MATTHIAS BRINKMANN
(matthias.brinkmann@philosophy.ox.ac.uk)

³⁹ I thank Johannes Himmelreich, Jessica Laimann, and Maxime Lepoutre for comments, the audience at the Dublin ‘Supererogation’ conference, and an anonymous reviewer. I also received helpful feedback on an early version of this paper from Hasko von Kriegstein, Marco Meyer, and Andreas Schmidt. Lastly, I thank Matthew Braham without whose encouragement this paper would not exist.