

# Aggregate Relevant Claims in Rescue Cases?

JOHANNA PRIVITERA

*Humboldt-Universität zu Berlin*

In *How Should We Aggregate Competing Claims*, Alex Voorhoeve suggests accommodating intuitions about duties in rescue cases by combining aggregative and non-aggregative elements into one theory. In this article, I discuss two problems Voorhoeve's theory faces as a result of requiring a cyclic pattern of choice, and argue that his attempt to solve them does not succeed.

## I. INTRODUCTION

Let rescue cases be cases in which an agent can either spare one person a certain harm or spare each of many others a less severe harm, but cannot do both, and in which the number of persons affected and the badness of the harms at stake are the only morally relevant factors. An ongoing debate in normative ethics concerns the question of how to account for intuitions regarding moral duties in these cases. In particular, it is difficult to find a rationale that accommodates two intuitions: the intuition that in those cases in which the harms at stake are similarly bad the agent ought to help the greatest possible number of persons, and the intuition that in those cases in which the harms at stake substantially differ in badness she ought to help the person facing the individually worst harm. Consider the following cases:

*Injuries* An agent has the choice between sparing one person a very painful injury or sparing a hundred others an injury that is slightly less painful, but still very bad. The pain experienced by each person would last for one week.

*Tetanus* An agent has the choice between sparing one person a very painful condition, such as the worst spasms a tetanus infection can cause, or sparing many others a hangnail. Again, the painful condition and the hangnails would each last for one week.

Many philosophers think that while in cases like *Injuries* the agent ought to help the many, in cases like *Tetanus* she ought to help the single person, no matter how many others she could spare the minor harm instead.<sup>1</sup>

<sup>1</sup> See, for example, David Brink, 'The Separateness of Persons, Distributive Norms, and Moral Theory', *Value, Welfare, and Morality*, ed. R. G. Frey and Christopher W.

Neither theories that allow for interpersonal aggregation nor ones that forbid it can typically account for both these intuitions. Aggregative theories rely on adding the harms experienced by different persons and combining them into a single measure. So, they can easily advocate helping the hundred injured by pointing out that their combined suffering outweighs the suffering of the single person. But they also seem committed to the implausible view that, if the number of those suffering a hangnail is large enough, their combined uneasiness will outweigh the suffering of the single tetanus patient. Non-aggregative theories, by contrast, determine whom we should help by considering only the severity of the harm each person faces individually. Consequently, they can easily yield the intuitively right result in cases like *Tetanus*: we ought to assist the tetanus patient because she is suffering the most severe individual harm. But by the same reasoning they also seem committed to implausible conclusions in cases like *Injuries*.

In *How Should We Aggregate Competing Claims*, Alex Voorhoeve suggests a natural solution to this problem that relies on combining aggregative and non-aggregative elements into one theory.<sup>2</sup> More specifically, he proposes a modified aggregative theory, which is constrained by non-aggregative considerations and thus avoids implausible implications in cases like *Tetanus*. In this article, I argue that while his account is supported by a compelling rationale, it also faces two problems, and that his attempts to solve them do not succeed.

The article is structured as follows. In [section II](#), I introduce Voorhoeve's theory and the rationale supporting it. In [section III](#), I discuss two problems his theory faces as a result of requiring a cyclic pattern of choice: First, the theory implausibly yields the result that we cannot but act wrongly in some specific situations. Second, it classifies as praiseworthy certain actions that seem pointless. In [section IV](#), I show that Voorhoeve's attempt to solve these problems fails.

Morris (Cambridge, 1993), pp. 252–89, at 270; Thomas M. Scanlon, *What We Owe To Each Other* (Cambridge, MA, 1998), pp. 238–41; Michael Otsuka, 'Saving Lives, Moral Theory, and the Claims of Individuals', *Philosophy and Public Affairs* 34 (2006), pp. 109–35; Frances M. Kamm, *Intricate Ethics* (New York, 2007), pp. 287–98, at 484–6; Dale Dorsey, 'Headaches, Lives and Value', *Utilitas* 21 (2009), pp. 36–58; Larry Temkin, *Rethinking the Good* (New York, 2012), chs. 2–3. Empirical research on the matter is limited, but Alex Voorhoeve reports a survey in which a substantial majority of subjects shared these intuitions in cases similar to *Injuries* and *Tetanus* in his 'Why One Should Count Only Claims with which One Can Sympathize', *Public Health Ethics* 10 (2017), pp. 148–58, at 149–50, and reviews some further, albeit imperfect, empirical evidence in his 'Balancing Small Against Large Burdens', *Behavioural Public Policy* (forthcoming).

<sup>2</sup> Alex Voorhoeve, 'How Should We Aggregate Competing Claims?', *Ethics* 125 (2014), pp. 64–87.

## II. VOORHOEVE'S THEORY

Voorhoeve calls his theory *Aggregate Relevant Claims* (ARC). His approach is phrased in terms of the claims people to be helped have on the agent instead of the harms they face – a person's claim roughly being stronger the more severe the harm she faces and the worse-off she is regardless of facing that harm.<sup>3</sup> ARC defines a claim as 'relevant' iff it is strong enough compared to the strongest claim in that rescue case, and it demands that the agent satisfies the greatest number of strength-weighted relevant claims. The approach thus allows for aggregation under the constraint that the course of action required does not depart too far from satisfying the strongest individual claim, which is what non-aggregative theories would demand. It is straightforward that this approach yields the desired results in the cases just presented. In *Injuries*, the claims of the many are almost as strong as the strongest claim and they together outweigh it, so the agent ought to help the many. In *Tetanus*, the claims at stake considerably differ in strength, so the agent ought to help the tetanus patient instead of satisfying the 'irrelevant' claims of the many.<sup>4</sup>

The rationale Voorhoeve develops in favour of ARC is this: it is permissible for each person to be substantially more concerned about her own well-being than about the well-being of strangers. Although it would, for example, not be permissible to spare oneself a very small harm at the cost of not sparing a stranger a terrible harm, it is permissible to spare oneself a certain harm at the cost of not sparing a stranger a harm that is somewhat worse. Voorhoeve calls the idealized point of view that a person would take if she were as much self-concerned as morality allows and as much other-concerned as morality requires the 'permissible personal perspective' of that person. A claim now is 'relevant' iff, from the permissible personal perspective of the

<sup>3</sup> Voorhoeve's examples and his statement of what defines the strength of a claim slightly differ from mine. The persons in his examples are already burdened by certain diseases that lower their well-being and the agent can increase their well-being by helping them. Accordingly, Voorhoeve defines a person's claim as being stronger 'a) the more her well-being would be increased by being aided; and b) the lower the level of well-being from which this increase would take place' (Voorhoeve, 'How Should We Aggregate Competing Claims?', p. 66). I adjusted the definition to fit my speaking of harms instead of levels of well-being and to fit the situation in my examples in which the agent can *prevent* harms instead of reacting to their having already occurred.

<sup>4</sup> Similar strategies of combining aggregative and non-aggregative elements have been discussed before. Derek Parfit considers one in his 'Justifiability to Each Person', *Ratio* 16 (2003), pp. 368–90, at 378–85, but ultimately rejects it because of a problem I will discuss in [section III](#). Kamm endorses one in her *Intricate Ethics*, pp. 297–8, at 484–6. Voorhoeve, however, gives what I think is the most compelling rationale and the best-worked-out overall account of such an approach. For a recent critique of Voorhoeve's theory, see John Halstead, 'The Numbers Always Count', *Ethics* 126 (2016), pp. 789–802.

person holding it, it is at least as important as the strongest competing claim – competing claims being claims that cannot be jointly satisfied.<sup>5</sup>

In some rescue cases, unanimity as to which claim should be satisfied emerges among the permissible personal perspectives of everyone affected. More specifically, such unanimity arises in cases like *Tetanus*, in which the strongest claim competes with much weaker – that is, irrelevant – claims. Since in these cases the strongest claim is objectively so much stronger than the weaker claims, it will be seen as having priority not only from the permissible personal perspective of the person holding the strongest claim, but also from the permissible personal perspectives of those holding the weaker claims. Cases in which the strongest claim competes with *relevant* claims – that is, cases like *Injuries* – are different. Here, since the objective differences between the claims are small, from the permissible personal perspectives of those holding the weaker claims, their own claims take priority. So, no unanimity emerges. According to ARC, if such unanimity arises, it ought to be respected and the agent ought to satisfy the strongest claim. If no such unanimity arises, aggregative considerations should be decisive in determining what to do.<sup>6</sup>

The most compelling feature of this rationale, I think, is that it roughly follows an idea that seems very natural in dealing with rescue cases:

[In cases like *Tetanus* it] would be wrong to allow the claim of the one to be saved from death [or, in our case, the tetanus spasm] to be outcompeted by any number of claims to be rid of the very minor harm, because given the difference in what is at stake for each person, someone facing the very minor harm should withdraw her claim if she is aware of the situation.<sup>7</sup>

This thought is very much to the point. It just seems plausible to think that if the difference between the claims is so significant that not even someone holding the weaker claim herself would be justified in satisfying her own claim, a neutral agent would not be either.

### III. TWO PROBLEMS FOR VOORHOEVE'S THEORY

ARC can accommodate intuitions about rescue cases, but it also yields two problematic implications which result from its requiring a cyclic pattern of choice between options like the following:

<sup>5</sup> Voorhoeve, 'How Should We Aggregate Competing Claims?', pp. 70–2.

<sup>6</sup> This unanimity-based consideration is first introduced as one among many reasons that speak in favour of ARC (see Voorhoeve, 'How Should We Aggregate Competing Claims?', pp. 73–5). In later sections, however, it becomes clear that this is the consideration ultimately supporting it (see Voorhoeve, 'How Should We Aggregate Competing Claims?', p. 77).

<sup>7</sup> Voorhoeve, 'How Should We Aggregate Competing Claims?', p. 75.

- (a) Satisfy the claim one person, *A*, has.
- (b) Satisfy the claims of a number of persons, *B*, that are less strong than *A*'s claim but still relevant to it, and that would together outweigh *A*'s claim.
- (c) Satisfy the claims of a larger number of persons, *C*, that are less strong than the claims of the *B*-persons, so that they are not relevant to *A*'s claim, but that are still relevant to the claims of the *B*-persons and would together outweigh them.

Between such options ARC requires the agent to make the following choices: She should choose (b) from {(a), (b)} because the claims of the *B*-persons are relevant in this situation and together they outweigh *A*'s claim. She should choose (c) from {(b), (c)} because the claims of the *C*-persons are relevant in this situation and together they outweigh the claims of the *B*-persons. And she should choose (a) from {(a), (c)} because the *C*-persons' claims are irrelevant in this situation. So, she ought to prefer (c) over (b), (b) over (a), and (a) over (c).

Drawing on John Broome,<sup>8</sup> Voorhoeve nevertheless argues that ARC does not violate the transitivity axiom of rational choice. The axioms of rational choice only require that the choices a theory mandates be rational given the values the theory sets. So, to find out whether ARC violates such an axiom, the options in question must be individuated by characteristics that are, according to ARC itself, the relevant grounds for choice. That means including relevance and irrelevance in the description. So, ARC mandates one to choose (relevant-b) from {(relevant-a), (relevant-b)}, (relevant-c) from {(relevant-b), (relevant-c)}, and (relevant-a) from {(relevant-a), (irrelevant-c)}. And this does not violate transitivity.<sup>9</sup> I grant this point to Voorhoeve. However, since (relevant-c) and (irrelevant-c) both consist in helping the *C*-persons, ARC still requires a cyclic pattern of choice. It still mandates that helping the *C*-persons must take priority over helping the *B*-persons, which must take priority over helping *A*, which must take priority over helping the *C*-persons. In what follows I will show that this pattern of choice, although not being irrational in light of ARC's values, yields two very implausible normative implications.

Voorhoeve is aware of the first of these implications. Call a case in which an agent is presented with three options like (a), (b) and (c) at the same time a 'three-option-case'. As Derek Parfit notes, according to theories like ARC, when presented with a three-option-case, the agent

<sup>8</sup> John Broome, *Weighing Goods* (Oxford, 1991), ch. 5.

<sup>9</sup> Alex Voorhoeve, 'Vaulting Intuition: Temkin's Critique of Transitivity', *Economics and Philosophy* 29 (2013), pp. 409–23, at 414–15.

cannot but act wrongly.<sup>10</sup> For if (c) is to be chosen over (b), (b) is to be chosen over (a), and (a) is to be chosen over (c), no matter what the agent chooses, there will always be something she should have chosen instead. This seems very implausible.

The second implausible implication flows from a modified version of the classic money pump argument. The *classic money pump argument* is employed within theories of rational decision-making to show that agents should not have cyclic preferences because this makes them exploitable.<sup>11</sup> The idea is this: assume an agent has cyclic preferences between holiday destinations such that she would be willing to pay some money (1) for exchanging a seaside holiday for a city trip, (2) for exchanging the city trip for a hiking tour, and (3) for exchanging the hiking tour for a seaside holiday again. The problem is that, were she to accept such a series of bargains – as she should, according to her preferences – she would lose money and end up in the same situation she was in at the beginning. Thus, having cyclic preferences, the argument goes, makes her an easy victim for potential exploiters.

The classic money pump argument is controversial. Critics have pointed out that it is not the cyclic preferences alone that make an agent exploitable, but rather the combination of having cyclic preferences and lacking foresight. Foresighted agents would, after all, realize that they are being exploited and reject the bargains.<sup>12</sup> This criticism seems justified. However, I will argue that, when applied to an ethical context, a *modified money pump argument* yields an independent problem for ethical theories requiring cyclic patterns of choice. And this independent problem is not threatened by the foresight objection.

So, let us assume ARC were true and a moral agent, who had one of the options (a), (b) or (c) discussed above available to her at the outset, was repeatedly offered to swap, for a small fee, the available option for the one that ARC ranks to be preferable. For example, imagine an agent who had the possibility to do (a) at the outset was offered to exchange the possibility to do (a) for the possibility to do (b), then (b) for (c), and then (c) for (a), each exchange coming at a small cost. As

<sup>10</sup> Parfit, 'Justifiability to Each Person', p. 384.

<sup>11</sup> The argument was first sketched in Frank P. Ramsey, 'Truth and Probability', *The Foundations of Mathematics and other Logical Essays*, ed. R. B. Braithwaite (London, 1931), pp. 156–98, at 182. Donald Davidson, J. C. C. McKinsey and Patrick Suppes developed it further in their 'Outlines of a Formal Theory of Value, I', *Philosophy of Science* 22 (1955), pp. 140–60, at 145–6.

<sup>12</sup> The foresight-based objection was first raised in Frederic Schick, 'Dutch Bookies and Money Pumps', *The Journal of Philosophy* 83 (1986), pp. 112–19. The view that an agent having cyclic preferences can be pumped even if she is foresighted is defended in Wlodek Rabinowicz, 'Money Pump with Foresight', *Imperceptible Harms and Benefits*, ed. Michael J. Almeida (Dordrecht, 2000), pp. 123–54.

in the non-moral scenario, the agent would have reason to accept each independent bargain (ARC ranks the option offered to be preferable), but by accepting the bargains she would lose money and end up with the same option she had in the beginning. The problem, now, is that advocates of ARC are committed to the view that by accepting these bargains the agent does something *morally good*. If ARC were true, the agent would be more praiseworthy the more bargains she accepts – although more bargains just lead her to lose money without achieving different results. This seems very implausible.<sup>13</sup> And unlike the case of the classic money pump, this problem cannot be solved by appealing to foresight. Whether or not the agent foresees that accepting the bargains will burden her unnecessarily does not matter – according to ARC she would still be praiseworthy for accepting them.

So, as a result of requiring a cyclic pattern of choice, ARC faces two objections: it implausibly yields that one cannot but act wrongly in three-option cases and it classifies actions as praiseworthy that just seem pointless.

#### IV. WHY VOORHOEVE'S SOLUTION DOES NOT WORK

Voorhoeve is aware of at least the first of these two objections – that is, the one about the impossibility of not acting wrongly in three-option-cases – and argues, in reply, that according to ARC the agent ought to choose (b) when presented with (a), (b) and (c) simultaneously. If this were true, it would not only counter the first objection, but also the second one. For, it could then be argued that in the modified money pump scenario, although only being presented with two options at a time, the agent actually has the choice between all three of them. From this and the duty to choose (b) in three-option-cases, it would result that it is only praiseworthy for the agent to accept the bargains offered until she secures the possibility to do (b). So, it would not be praiseworthy to incur more and more vain losses. In what follows, however, I will argue that the argument Voorhoeve gives for ARC's yielding a duty to choose (b) in three-option-cases is not sound and his theory is therefore still confronted with both problems.

Obviously, the principle of satisfying the greatest number of those strength-weighted claims that are not too much weaker than the strongest claim does yield a duty to choose (b) in three-option cases, and

<sup>13</sup> Voorhoeve does not say anything about how demanding his theory is. But if we assume that morality requires agents to do the right thing even if this comes at some (non-extreme) cost to themselves – an assumption that seems easy to justify – an even stronger point can be made. ARC would then in certain situations *require* agents to bear unnecessary costs.

the Broome/Voorhoeve-strategy of redescribing options does so, too.<sup>14</sup> The crucial question, however, is whether the rationale supporting ARC also yields this result. Voorhoeve thinks it does. His reasoning is this. When comparing A's claim, which is the strongest in this rescue case, to each of the competing claims, the following picture emerges: in comparing A's claim to the claims of the C-persons, unanimity arises that A's claim should take priority over each of the claims of the C-persons. So, the claims of the C-persons can be set aside. In comparing A's claim to each of the claims of the B-persons, no unanimity emerges as to which claim should take priority. So, in order to decide between A's claim and the claims of the B-persons, we should resort to aggregative considerations. Since the claims of the C-persons have been set aside and the claims of the B-persons together outweigh A's claims, ARC requires the agent to do (b).<sup>15</sup>

The problem with this argument is that the claims of the C-persons are set aside although there is only unanimity that they have lower priority than A's claim, and not that they also have lower priority than the claims of the B-persons. It is true that in both *Tetanus* and *Injuries* it seemed plausible just to compare the strongest claim to each of the competing claims and set aside the claims that were unanimously judged to have lower priority. But in those cases, the strongest claim was also the *only* claim that each of the weaker claims competed with. So, it was obvious that the weaker claims could be set aside.

In three-option-cases, however, setting aside the claims of the C-persons just because they are unanimously judged to have lower priority than the strongest claim is not justified. Recall the compelling rough rationale that Voorhoeve gives for ARC and that I cited in [section II](#) of this article. The rough rationale says that if competing claims considerably differ in strength we can expect from the person holding the weaker claim that she withdraws her claim, if she is aware of the situation. The idea, here, seems to be that the person holding the weaker claim should understand that instead of helping *her*, the agent should rather help the person holding the strongest claim. She should withdraw her claim, *in order that* the person holding the strongest claim be helped. But this reasoning does not apply to the C-persons in three-option-cases. For, they have no reason to think that, if they withdraw their claims, the agent will help person A (whose claim takes priority over each of theirs). Rather, if the C-persons withdraw their claims, the agent will help the B-persons, whose claims do not take priority over

<sup>14</sup> The Broome/Voorhoeve-argument even yields that there is at least one best option for *any* given feasible set of alternatives – the condition ‘feasible’ ruling out sets that, for example, contain both (relevant-x) and (irrelevant-x).

<sup>15</sup> Voorhoeve, ‘How Should We Aggregate Competing Claims?’, pp. 77–8.



the *C*-persons' claims. So, in three-option-cases it cannot be expected from the *C*-persons that they withdraw their claims and consequently it cannot be inferred that (b) is the required course of action. ARC therefore remains vulnerable both to the objection that an agent cannot but act wrongly in three-option-cases and to the praiseworthiness of the vain losses objection.<sup>16</sup>

[j.privitera@gmx.com](mailto:j.privitera@gmx.com)

<sup>16</sup> I thank Kirsten Meyer and Gabriel Wollner for their very helpful comments on earlier versions of this article. I also benefited a lot from the comments of two anonymous reviewers and from extensive discussions with Jonas Harney, Roland Hesse, Lukas Tank, Stefanie Thiele, and with the participants of the Normative and Applied Ethics Colloquium at Humboldt-Universität zu Berlin.