CrossMark

Utilitas (2021), **33**, 101–110
doi:10.1017/S095382082000028X

CAMBRIDGE UNIVERSITY PRESS

REPLY

The Relevance View: Defended and Extended

Kirsten Mann

The Australian National University
Email: kirsten.mann@anu.edu.au

Abstract

The Relevance View, exemplified by Alex Voorhoeve's Aggregate Relevant Claims, has considerable appeal. It accommodates our reluctance to aggregate weak claims in canonical cases like Life for Headaches (where one person's claim to life-saving treatment competes with millions of claims for headache relief), while permitting aggregation of claims in a range of other cases. But it has been the target of significant criticism. In an important recent paper, Patrick Tomlin argues that the view suffers from failures of internal logic, violating plausible consistency constraints and generating incoherent combinations of verdicts on cases. And in cases resembling real-world healthcare allocation problems, Tomlin argues that the view offers no guidance at all. In response, I argue that the internal logic of the Relevance View is sound, and the view's core principles, suitably clarified, support a significant extension of the view beyond the simple cases to which it is typically applied.

Introduction

You are confronted with the competing claims of two groups of people. You can only help one group, so how should you weigh their claims? According to the Relevance View, a group of weaker claims can outweigh a stronger claim, but only when the difference in claim strength, assessed pairwise, is not too great.[1] Consider Life for Headaches, in which one person's claim to life-saving treatment competes with millions of claims for headache relief.[2] The claim to be cured of a minor headache is so much weaker than the claim to be spared from death that it is not *relevant* to that claim. Hence, we should accord no moral weight to the claims for headache relief; we should save the life no matter how numerous the headaches are. In cases where the weaker claims are strong enough to be relevant to the competing claim, though, we should satisfy the set of claims recommended by straightforward aggregation.

[1]Patrick Tomlin introduces the term 'Relevance View' in On Limited Aggregation, *Philosophy & Public Affairs*, 45 (2017): 232–60. Notable presentations of the view include Thomas Scanlon, *What We Owe to Each Other* (Cambridge, MA: Harvard University Press, 1998), pp. 229–41; Frances Kamm, *Morality, Mortality, I: Death and Whom to Save From It* (New York: Oxford University Press, 1993), pp. 144–95; David Lefkowitz, On the Concept of a Morally Relevant Harm, *Utilitas*, 20 (2008): 409–23; and Alex Voorhoeve, How Should We Aggregate Competing Claims?, *Ethics*, 125 (2014): 64–87.

[2]Alastair Norcross, Comparing Harms: Headaches and Human Lives, *Philosophy & Public Affairs*, 26 (1997): 135–67.

© The Author(s), 2020. Published by Cambridge University Press

https://doi.org/10.1017/S095382082000028X Published online by Cambridge University Press

Recent discussions of the Relevance View have focussed on its influential treatment by Alex Voorhoeve, whose specific articulation of the view is called *Aggregate Relevant Claims (ARC)*.

ARC:

1. Each individual whose well-being is at stake has a claim on you to be helped. (An individual for whom nothing is at stake does not have a claim.)
2. Individuals' claims *compete* just in case they cannot be jointly satisfied.
3. An individual's claim is *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.
4. A claim is *relevant* if and only if it is sufficiently strong relative to the strongest competing claim.
5. You should choose an alternative that satisfies the greatest sum of strength-weighted, relevant claims.[3]

The Relevance View does not simply accord with our intuitions about cases like Life for Headaches; it seems to get to the heart of those intuitions: headaches just do not seem to count against a life. The view may owe its appeal, however, to the fact that it is usually illustrated with simple cases like Life for Headaches, where two homogeneous groups of claims are in direct competition. In an important recent paper, Patrick Tomlin discusses two ways in which broadening the class of cases under consideration challenges the view's plausibility.[4] Tomlin argues that:

1. applying the Relevance View to cases involving two heterogeneous groups of claims in direct competition exposes serious failures of the view's internal logic; and
2. the view fails to apply to a class of cases resembling real-word healthcare allocation problems.

Published responses to Tomlin have accepted the force of the first of these objections, either proposing substantial amendments to Voorhoeve's choice rule to avoid the objection,[5] or embracing unfettered aggregation.[6] By contrast, I will defend the Relevance View without amendment against Tomlin's first challenge. In response to Tomlin's second challenge, I will propose an extension of the Relevance View to Tomlin's problem cases.

1. The Relevance View's internal logic

How might the Relevance View apply to cases where two heterogeneous groups of claims compete? ARC tells us that 'a claim is relevant if and only if it is

[3]Voorhoeve, p. 66.

[4]Tomlin, On Limited Aggregation.

[5]See e.g. Victor Tadros, Localized Restricted Aggregation, *Oxford Studies in Political Philosophy*, 5 (2019): 171–204; Aart van Gils and Patrick Tomlin, Relevance Rides Again?, *Oxford Studies in Political Philosophy*, 6 (2020), 221–56; Korbinian Rüger, Aggregation with Constraints, *Utilitas*; and Bastian Steuwer, Aggregation, Balancing, and Respect for the Claims of Individuals, *Utilitas*.

[6]Joe Horton, Always Aggregate, *Philosophy & Public Affairs*, 46 (2018): 160–74.

https://doi.org/10.1017/S095382082000028X Published online by Cambridge University Press

sufficiently strong relative to the strongest competing claim'.[7] But consider the following case.

> CASE 1: You can save either Group A or Group B:
> Group A comprises some people facing death and others facing mild impairment;
> Group B comprises some people facing severe impairment.[8]

In this and all following cases, we are to assume that mild impairment is 'sufficiently strong', in Voorhoeve's sense, relative to severe impairment but not relative to death. This means that if a claim against death was in competition with a claim against mild impairment, the claim against mild impairment would be irrelevant. We also assume that severe impairment is sufficiently strong relative to death. Also (this will become important in later examples), it is stipulated that ten claims to be spared mild impairment outweigh one claim to be spared severe impairment, and ten claims to be spared severe impairment outweigh (although only barely) one claim to be spared from death.

The question suggested by Case 1 is whether the claims of the people in Group A facing mild impairment are relevant. If the claim of the single person facing death in Group A counts as the 'strongest competing claim', the mild impairment claims are irrelevant. If the claim of one of the people in Group B facing severe impairment is the strongest competing claim, the mild impairment claims are relevant.

Tomlin discusses the two most obvious ways to disambiguate 'strongest competing claim'. These are:

> *Anchor by Strength*: In order to be relevant, a claim must be sufficiently strong relative to the *strongest overall* claim in the competition.

> *Anchor by Competition*: In order to be relevant, a claim must be sufficiently strong relative to the strongest claim *with which it competes*.[9]

It is true that the ARC choice rule could be interpreted either way, and both interpretations would deliver verdicts consistent with our intuitions about simple cases like Life for Headaches. However, Anchor by Competition seems more consistent with our intuitions about what *explains* the non-aggregative verdict in Life for Headaches. The fact that headaches would be cured under some option, O, seems to contribute to the choiceworthiness of O differently depending on what you would be giving up by choosing O. Headaches can give you a reason to choose O rather than saving someone's finger, say, but they give you no reason not to save a life. In other words, what a claim stands to be satisfied *at the expense of* affects its moral weight. We can represent this as a property that is predicated of the claim itself: it is a property of the headache claim that it stands to be satisfied *rather than* another claim. I will refer to such properties as 'cost properties' in what follows.

Thus characterised, the core idea of the Relevance View is that the moral weight of a claim depends not just on the claim's strength, but also on whether it has a particular kind of cost property: whether it stands to be satisfied at the expense of a sufficiently

[7]Voorhoeve, p. 66.
[8]Tomlin, p. 239.
[9]Tomlin, p. 239.

https://doi.org/10.1017/S095382082000028X Published online by Cambridge University Press

stronger claim. It is not enough, then, for a sufficiently stronger claim to be merely at stake in the choice, as Anchor by Strength specifies. If a claim is to render another claim irrelevant, the two must stand to be satisfied under competing options, and not to be jointly satisfied.

Tomlin argues that whichever of the two interpretations one adopts, the Relevance View displays 'bizarre internal logic'.[10] In either case, he argues, the view would violate a plausible choice principle and generate incoherent combinations of verdicts on cases. I will set Anchor by Strength aside and defend the Anchor by Competition interpretation of the Relevance View against these criticisms, referring to this interpretation as simply 'the Relevance View' in what follows.

Equal Consideration for Equal Claims

Tomlin notes that the Relevance View violates the following principle:

> *Equal Consideration for Equal Claims*: All claims of equal strength ought to be given equal weight in determining which group to save.[11]

If two claims of equal strength are in different (competing) groups, different claims will serve as the 'strongest competing claim' for each, so one could be relevant while the other was irrelevant. The first claim would then be given non-zero weight in the decision while the (equally strong) second claim would be given zero weight.

Tomlin characterises the Relevance View's violation of Equal Consideration as a failure of the view's internal logic, which suggests that he holds Equal Consideration to have the status of a requirement of internal consistency. At the very least, he believes the principle to be so plausible that few could reject it, even the Relevance View's proponents.[12]

What status should we accord Equal Consideration? The following is very difficult to reject:

> *Weak Equal Consideration for Equal Claims*: In a given choice context, if claims are the same in *all morally relevant respects*, they have the same moral weight.

Indeed, it is hard to see how a theory that violated this principle would even be intelligible. Note, though, that the Relevance View does not violate this principle. For two claims of the same strength to be the same in all morally relevant respects, their relevance status also needs to be the same. And if it is, then under the Relevance View they have the same moral weight.

Tomlin's version of Equal Consideration says something stronger. It is equivalent to the following:

> *Strong Equal Consideration for Equal Claims*: In a given choice context, if claims have the *same strength*, they have the same moral weight.

Recall that the strength of a claim is a function of the person's well-being if she is helped and her well-being if she is not. If two claims are the same in all morally relevant

[10]Tomlin, p. 258.
[11]Tomlin, p. 241.
[12]Tomlin, p. 241.

https://doi.org/10.1017/S095382082000028X Published online by Cambridge University Press

respects, they have the same strength. But if they have the same strength, they are not necessarily the same in all morally relevant respects. According to the Relevance View, a claim's cost properties – what we would be giving up to satisfy it – also contribute to determining its moral weight.

We should view (strong) Equal Consideration as a substantive moral position. It implies that a claim's cost properties cannot make a difference to its moral weight. By contrast, the Relevance View, as I have characterised it, says that cost properties *are* morally significant.

(Strong) Equal Consideration does have intuitive appeal. However, this appeal might be due in part to its resemblance to its weaker counterpart – we might read 'same strength' as meaning the same in all morally relevant respects, for example. And the converse also has intuitive support. Above, when I discussed our intuitions about Life for Headaches, I argued that there is intuitive support for the idea that cost properties can make a difference to a claim's moral weight. In summary, we should not accord Equal Consideration the status that Tomlin does. Proponents of the Relevance View will have to give it up, but that is not a decisive reason to reject the Relevance View.

Dominance

The appeal to the independent plausibility of Equal Consideration for Equal Claims is, however, only part of Tomlin's argument. He also presents a two-stage case in which the Relevance View's verdicts appear incoherent.

> CASE 2, Stage 1: You can save either Group A or Group B:
> Group A consists of one person facing death;
> Group B consists of ten people facing severe impairment.

By stipulation, ten claims to be saved from severe impairment outweigh, though only barely, one claim to be saved from death. So, you should save Group B.

Now suppose that one claim to be spared mild impairment is added to Group A, and a billion such claims to Group B.

> CASE 2, Stage 2: You can save either Group A or Group B:
> Group A consists of one person facing death and one facing mild impairment;
> Group B consists of ten people facing severe impairment and one billion facing mild impairment.[13]

The claim of the dying person in A renders the mild impairment claims added to B irrelevant. The mild impairment claim added to A is relevant, though, since it stands to be satisfied *alongside* the dying person's claim. The addition of the single relevant mild impairment claim means (by stipulation) that the relevant claims in Group A now outweigh those in Group B. Adding one claim to Group A and a billion equally strong claims to Group B has thus tipped the scales in favour of Group A.

Tomlin describes the Relevance View's verdict on Case 2 as 'counterintuitive',[14] but his complaint seems to be that the Relevance View violates another (tacit) coherence principle. Here is his account of what goes wrong in the case:

[13]Tomlin, p. 242.
[14]Tomlin, p. 242.

https://doi.org/10.1017/S095382082000028X Published online by Cambridge University Press

> If we break this down into subgroups, the claims of those facing severe impairment trump the dying person's claim, and the claims of the one billion trump the claim of the lone person facing mild impairment. It seems to be some kind of moral alchemy, therefore, under which the claims of the dying person and the lone person facing mild impairment somehow combine to defeat the others.[15]

Tomlin is invoking a dominance principle, applied here to the moral weight of sets of claims. Here is the principle set out more formally:

> *Dominance*: For sets of claims A1, A2, B1, and B2, where A1 and A2 are non-intersecting and B1 and B2 are non-intersecting, if B1 outweighs A1 and B2 outweighs A2, then B1 ∪ B2 outweighs A1 ∪ A2.

Applying the dominance principle to the sets of claims from Case 2, we have:

A1: one life claim
A2: one mild impairment claim
B1: ten severe impairment claims
B2: one billion mild impairment claims.

Notice, though, that these are incomplete descriptions of the claim sets. They leave unspecified features of the claims that are morally significant according to the Relevance View: the claims' cost properties, and hence their relevance or irrelevance. We can redescribe Tomlin's dominance argument using superscripts R and ¬R to denote relevant and irrelevant claims respectively, as follows:

B1 outweighs A1
$B2^{R}$ outweighs A2
Therefore B1 ∪ $B2^{\neg R}$ outweighs A1 ∪ A2.[16]

The sets of claims referred to in the conclusion differ in their morally significant features from the sets of claims referred to in the premises. The Relevance View does not violate dominance, then, because the principle does not apply.

That said, something *like* dominance is violated.[17] Call the formulation of dominance just discussed 'weak' dominance, and the following principle 'strong' dominance:[18]

> *Strong dominance*: If A1, A2, B1, and B2 are sets of claims, where A1 and A2 are non-intersecting and B1 and B2 are non-intersecting; and if A1′, A2′, B1′, and B2′ are sets of claims such that A1 and A1′ are identical in all respects other than their cost properties, similarly for A2 and A2′, B1 and B1′, and B2 and B2′; then if B1 outweighs A1 and B2 outweighs A2, then B1′ ∪ B2′ outweighs A1′ ∪ A2′.

[15]Tomlin, p. 244.

[16]Voorhoeve (p. 79) employs this strategy in responding to the charge that ARC violates basic contraction consistency, drawing on John Broome's *Weighing Goods* (Oxford: Basil Blackwell, 1991), ch. 5.

[17]Tadros (p. 187) makes a similar observation with respect to transitivity.

[18]Note that this is not the standard meaning of strong and weak dominance; the terminology is used here to emphasise the structural parallels with my discussion of Equal Consideration.

https://doi.org/10.1017/S095382082000028X Published online by Cambridge University Press

How costly is having to give up strong dominance? While the weak version of the principle is difficult to reject, the strong version amounts to the claim that cost properties can be disregarded in weighing up sets of claims, and hence that they have no moral significance. The Relevance View says that there are morally significant cost properties, and, as argued above, this proposition has some intuitive support.[19]

One might say in response that we should take our intuitions about Case 2 to be judgements simply that the subsets of claims *should* be able to be weighed against each other in the manner described in the dominance principle. Our intuitions would then provide justification for strong dominance, and more generally for the idea that cost properties lack moral significance. But the reliability of these intuitions is questionable: the two-stage presentation of the case forces us to consider the added claims in isolation, directing our attention away from the cost properties the claims have within the full choice context. Since the question at hand is whether particular cost properties have moral significance, we should be reluctant to place too much weight on intuitions that may have arisen from our inattention to the same properties.

Discussion

The Relevance View violates the strong versions of Equal Consideration and dominance. But these principles are not uncontroversial requirements of consistency. Whether you accept the Relevance View or the two principles just mentioned comes down to the same question: can cost properties affect the moral weight of claims? We should take seriously the possibility of an answer in the affirmative, given how well this seems to explain the non-aggregative verdict in Life for Headaches. Appealing to the strong versions of the principles discussed here assumes at the outset that cost properties do not make a difference, and so fails to address the substantive question.

2. Quasi-competition

So far, we have been considering how the Relevance View applies to cases involving competing claims, where two claims compete just in case they are jointly unsatisfiable. Tomlin argues, however, that real-world healthcare allocation cases are unlikely to involve groups in direct competition, but rather 'quasi-competition': a group of claims such that not all claims can be satisfied within the available budget, but any two could jointly be satisfied. In quasi-competition cases, there are no pairs of claims that meet the ARC definition of 'competing claim', so there is no criterion to separate relevant from irrelevant claims. The Relevance View, then, provides no guidance.

Tomlin discusses two ways the Relevance View might be extended to apply to quasi-competition cases. On the first extension, *Relevance View (quasi-competition)*, the relevance of each quasi-competing claim is established relative to the strongest claim overall in the set. You should satisfy the subset of relevant claims with the greatest combined strength, subject to your budget constraint.[20]

On the second extension, *Relevance View (subcompetitions)*, the budget is allocated through a series of subcompetitions. At each subcompetition, the strongest claim still in

[19]A similar response could be given to Horton's argument in Always Aggregate. Horton appeals to a principle similar to strong dominance to argue that all 'partially aggregative' views must have implausible implications in cases involving heterogeneous groups of claims.

[20]Tomlin, p. 253.

https://doi.org/10.1017/S095382082000028X Published online by Cambridge University Press

contention is weighed against the claims relevant to it that could be satisfied for the same (or less) money. If the strongest claim outweighs each candidate combination of weaker relevant claims, it is satisfied. If not, it is set aside, and the candidate set of weaker relevant claims with the greatest combined strength is satisfied. The subcompetitions continue until the budget is expended.[21]

These extensions of the Relevance View are illustrated through the following example:

> CASE 3: You must choose which of the following claims to satisfy with a budget of $1 million:
> Group A: A single dying person who can be saved for $500,000;
> Group B: Twenty people facing severe impairment who can be saved for $50,000 each;
> Group C: Ten thousand people facing mild impairment who can be saved for $100 each.[22]

Relevance View (quasi-competition) sets aside all the C claims as irrelevant. Recall that ten claims to be spared from severe impairment outweigh one claim to be spared from death, so using the full budget to satisfy the twenty B claims maximises the sum of the strengths of satisfied relevant claims.

Under Relevance View (subcompetitions), the first subcompetition results in A's claim being set aside unsatisfied and ten B claims satisfied. The subsequent subcompetitions involve weighing C claims against the remaining B claims, with the result that the remaining B claims are set aside and C claims are satisfied. Ultimately, ten B claims and 5,000 C claims are satisfied.[23]

Tomlin's purpose in introducing the two extensions is not to endorse one as the correct interpretation of the Relevance View. Rather, he argues that the Relevance View is insufficiently fine-grained to justify either of the candidate extensions, or to work out their details. He challenges the view's proponents to show that it *does* have something to say about quasi-competition.[24]

A positive proposal for extending the Relevance View to quasi-competition cases

I have argued that we should understand the core idea underpinning the Relevance View as follows: a claim's moral weight depends on what satisfying it would come at the expense of. In Life for Headaches, the headache claims are irrelevant not simply because a life is at stake in the decision, but because by curing headaches you would thereby be failing to save the life. The challenge in applying this principle to quasi-competition is to articulate what it means for one claim to stand to be satisfied at the expense of another claim in cases where any two claims could be jointly satisfied.

To motivate my positive proposal, notice a feature shared by both the extensions of the Relevance View that Tomlin considers. The relevance of a claim is insensitive to whether the strongest claim – the claim which could render it irrelevant – is satisfied alongside it. Under Relevance View (quasi-competition), the relevance of each claim

[21]Tomlin, p. 254.

[22]Tomlin, p. 255. Note this is a simplified version of Tomlin's case.

[23]Tomlin discusses two variants of Relevance View (subcompetitions) (pp. 255–56), but under the simplified example discussed here, both yield the same verdict.

[24]Tomlin, pp. 256–57.

https://doi.org/10.1017/S095382082000028X Published online by Cambridge University Press

is settled at the beginning of the choice procedure. Any claim that is insufficiently strong relative to the strongest claim under consideration is ruled irrelevant, regardless of whether the strongest claim ends up being satisfied. Under Relevance View (subcompetitions), relevance for the purpose of each subcompetition is established relative only to the claims still under consideration, without reference to the results of previous subcompetitions. In Case 3, for example, a C claim has the same relevance status whether or not it would ultimately be satisfied alongside A's claim. This feature of both extensions means that they cannot give effect to the idea that relevance is sensitive to whether the claim would be satisfied *at the expense of* a sufficiently stronger claim.

I propose that we understand the relevance of a claim relative to the option in which that claim is satisfied. Quasi-competition cases can be redescribed as cases of non-binary choice among mutually unrealisable options.[25] A budget and a set of quasi-competing claims together define a set of options, with each option corresponding to a set of claims that can be jointly satisfied within the budget. While no two claims in these cases are in direct competition, the options are, at least if we restrict the option set to those options where any unspent portion of the budget is insufficient to satisfy additional claims. On this view, when we assess the relevance of claim C, we are really assessing the relevance of C *in option O*. If satisfying C *by choosing option O* comes at the expense of satisfying a sufficiently stronger claim, C's being satisfied does not count in favour of choosing option O.

Here is a more precise statement of the proposal:

> *Relevance View (options)*: Claim C in option O is relevant if and only if it is sufficiently strong relative to the strongest claim C* such that:
> – C* is satisfied in an option with which O competes; and
> – C* is not satisfied in O.
> You should choose an option that maximises the sum of the strengths of satisfied relevant claims.

What distinguishes this conception of relevance from those discussed by Tomlin is that there is no fact about whether C is relevant within the overall choice context; C's relevance may vary from one option in which it is satisfied to another.

Applying Relevance View (options) to Case 3, we first set out all the options.

1. Satisfy A and 10 Bs
2. Satisfy A, 9 Bs and 500 Cs

   …

11. Satisfy A and 5,000 Cs
12. Satisfy 20 Bs
13. Satisfy 19 Bs and 500 Cs

    …

32. Satisfy 10,000 Cs.[26]

[25]See Campbell Brown, Is Close Enough Good Enough?, *Economics and Philosophy*, 36 (2020): 29–59, for a broader discussion of possible extensions of the Relevance View to non-binary choices among options.

[26]I have not distinguished options that are the same but for the identities of the group members whose claims are satisfied.

https://doi.org/10.1017/S095382082000028X Published online by Cambridge University Press

According to Relevance View (options), C claims are irrelevant in any option in which A's claim is not also satisfied – Options 13 to 32 in this case. Recall the assumed exchange rates for different claim types: ten B claims outweigh one A claim, and ten C claims outweigh one B claim. The option that maximises the sum of the strengths of satisfied relevant claims, then, is 11: satisfy the claims of A and 5,000 Cs. And this verdict seems defensible. Given the relative costs in this case, it seems that you should satisfy C claims rather than B claims. But you should not satisfy C claims if you are thereby failing to satisfy A's claim. The verdict of Relevance View (options) satisfies both these conditions.

Tomlin's challenge to proponents of the Relevance View is to show that the view has something to say about quasi-competition cases. Relevance View (options) gives effect to the same core conception of relevance that underpins the Relevance View in its application to simple cases, and offers clear guidance in cases of quasi-competition.

Conclusion

I have argued for a version of the Relevance View which has at its core the idea that the moral weight of a claim depends on what we would be giving up to satisfy it. My proposed extension realises this same conception of relevance in quasi-competition cases. The apparent failures of the Relevance View's internal logic that Tomlin identifies are not failures of coherence, but rather cases where claims of the same strength have different moral weights because of their different cost properties. The question then becomes: do the cost properties that are central to the Relevance View's account of choiceworthiness really have moral significance? This question cannot be answered by appeal to principles that assume these properties not to be morally significant.[27]

[27]For comments and discussion, I am grateful to Christian Barry, John Cusbert, Alan Hájek, Seth Lazar, Chad Lee-Stronach, Thomas Schmidt, Katie Steele, Nicholas Southwood, Patrick Tomlin, Alex Voorhoeve, Shang Long Yeo, participants at the 2019 ANU-Humboldt-Princeton Summer Institute on Practical Normativity, and an anonymous *Utilitas* reviewer. My research is supported by an Australian Government Research Training Program Scholarship.

Cite this article: Mann K (2021). The Relevance View: Defended and Extended. *Utilitas* **33**, 101–110. https://doi.org/10.1017/S095382082000028X

https://doi.org/10.1017/S095382082000028X Published online by Cambridge University Press