Relevance and Nonbinary Choices*

Kirsten Mann

In cases where the claims of different groups of people compete, the Relevance View occupies a middle ground between aggregation and nonaggregation. It allows weaker claims to aggregate to outweigh a stronger claim just when the competing claims, compared pairwise, are sufficiently close in strength. The view has strong intuitive appeal when applied to simple binary choices, but I argue that attempts to extend it to nonbinary choices have been unsuccessful. I propose a new extension of the Relevance View to nonbinary choices based on a "binary contrastive" account of the moral reasons that obtain in the cases of interest.

I. INTRODUCTION

Suppose there are two groups of people with claims for your assistance, and you can only help one of the groups. Which group ought you to choose? One possibility is that you ought to help the group whose claims have the greatest combined strength; call this "aggregation." Another possibility is that you ought to satisfy the strongest claim; call this "nonaggregation." Both aggregation and nonaggregation are at odds with some widely held intuitions about cases. Consider, for example, Life for Headaches.

> *Case 1: Life for Headaches.* You can choose exactly one of the following options:
>
> Option A: save one person from a premature death.

* For comments and discussion, I am grateful to Christian Barry, Alan Hájek, Seth Lazar, Chad Lee-Stronach, Thomas Schmidt, Katie Steele, Alex Voorhoeve, Timothy Luke Williamson, Shang Long Yeo, and two anonymous referees and several associate editors for *Ethics*. I presented a version of this article at the Australian National University's 2019 Economics and Philosophy workshop, and I am grateful to the participants for their helpful questions and suggestions. My research is supported by an Australian Government Research Training Program Scholarship.

Ethics, volume 132, number 2, January 2022.
 https://doi.org/10.1086/716873

Option B: save *N* people from suffering mild headaches.[1]

In this case, it seems like the numbers do not matter at all. No matter how many headaches there are, you ought to save the life. These intuitions tell against aggregation. But now consider a case where the claims involved are closer in strength:

Case 2: Legs for Arms. You can choose exactly one of the following options:

Option A: save one person from losing her legs.

Option B: save *M* people each from losing an arm.

Here, it seems like which group you ought to help depends on the number of people facing the loss of their arm. For small *M*, you ought to choose Option A; if *M* is large enough, you ought to choose Option B. In Legs for Arms, then, aggregation appears the most plausible approach.

My concern in this article is to constructively explore a view which holds promise for vindicating our apparently conflicting judgements about cases like Legs for Arms and Life for Headaches. This view—the "Relevance View"—offers a criterion for when weak claims aggregate to outweigh a stronger claim and when they do not, seeking to accommodate and systematize the intuitions that favor aggregation and nonaggregation.[2] According to the Relevance View, a group of weaker claims can outweigh a stronger one, but only when the claims, compared pairwise, are sufficiently close in strength that the weak claims are "relevant" to the stronger.

My discussion of the Relevance View will focus on Alex Voorhoeve's influential version, Aggregate Relevant Claims (ARC).[3] ARC states that in a choice between satisfying the claims of one group and satisfying the claims of another, where you cannot satisfy both, you should choose the option under which you maximize the sum of the strengths of satisfied

1. This and similar cases are discussed extensively in the literature. "Life for Headaches" is introduced in Alastair Norcross, "Comparing Harms: Headaches and Human Lives," *Philosophy and Public Affairs* 26 (1997): 135–67.

2. The term "Relevance View" is introduced in Patrick Tomlin, "On Limited Aggregation," *Philosophy and Public Affairs* 45 (2017): 232–60.

3. Alex Voorhoeve, "How Should We Aggregate Competing Claims?," *Ethics* 125 (2014): 64–87. Earlier presentations of the broad view can be found, e.g., in Thomas Scanlon, *What We Owe to Each Other* (Cambridge, MA: Harvard University Press, 1998), 229–41; and Frances Kamm, *Morality, Mortality*, vol. 1, *Death and Whom to Save from It* (New York: Oxford University Press, 1993), 144–95.

relevant claims.[4] A person has a claim just in case her well-being would be higher under one of the two options—in the rescue cases Voorhoeve is interested in, a person's well-being is higher under the option in which you assist her—and her claim is satisfied if her well-being is at the higher level under the chosen option. The strength of a claim is a function of the person's well-being under the two options. A claim is relevant if it is sufficiently strong relative to the strongest claim with which it competes, and irrelevant otherwise.[5]

I will leave aside the question of how strong a claim needs to be to be "sufficiently strong" relative to a stronger claim.[6] But on plausible assumptions about where this relevance threshold is to be found, ARC delivers verdicts that accord with our intuitions in the two cases discussed above. In Legs for Arms, the two types of claims are close enough in strength that all the claims are relevant. The weaker claims have nonzero weight and so, if they are sufficiently numerous, they can outweigh the stronger claim.[7] In Life for Headaches, a headache claim is so much weaker than the claim to be saved from premature death that it is irrelevant. The only relevant claim at stake is the life claim, and so you ought to save the life, no matter how many people stand to suffer headaches.

While a range of criticisms have been leveled against the Relevance View,[8] the one that is the focus of this article concerns the view's application

4. Voorhoeve does not limit ARC to binary cases, but, as I will argue, nonbinary cases present a distinct challenge for the view, so I will deal with them separately.

5. Voorhoeve uses slightly different language here: to be relevant, a claim must be "sufficiently strong relative to *the strongest competing claim*" (Voorhoeve, "How Should We Aggregate?," 66; emphasis added). However, Patrick Tomlin has argued that this criterion is ambiguous when the sets of claims are heterogeneous. Voorhoeve limits his discussion to cases where all the claims in a set have the same relevance status (ibid., 67), but it is clear that ARC must be extendable beyond these simple cases if it is to have significant practical application. The formulation of ARC presented here is what Tomlin terms the "Anchor by Competition" interpretation (Tomlin, "On Limited Aggregation," 238–39). For a defense of this interpretation, see Kirsten Mann, "The Relevance View: Defended and Extended," *Utilitas* 33 (2021): 101–10.

6. See Voorhoeve, "How Should We Aggregate?," for discussion.

7. Alec Walen has recently argued that, rather than relevance being an on/off property, the moral weight of a claim might gradually decline as the difference in strength between it and the strongest competing claim increases, until, once the difference in strength reaches the threshold for irrelevance, the moral weight reaches zero. Thus, for example, a claim to be saved from a headache might have greater moral weight against a claim to be saved from the loss of a fingertip than against a claim to be saved from the loss of a whole finger, despite its having nonzero weight in both cases. I will assume, with Voorhoeve, that the moral weight of a claim is determined by its strength and relevance status alone, but the arguments I offer here apply equally to Walen's model. Alec Walen, "Risks and Weak Aggregation: Why Different Models of Risk Suit Different Types of Cases," *Ethics* 131 (2020): 62–86.

8. See, e.g., John Broome, *Weighing Lives* (Oxford: Oxford University Press, 2004), 56–59; and Michael Huemer, "In Defence of Repugnance," *Mind* 117 (2008): 899–933. Broome and

to nonbinary choices. In particular, critics have pointed to the intransitivity of the "is more choiceworthy than" relation and the possibility of cycles in some nonbinary choices.[9] Applying the Relevance View to successive pairwise choices from a nonbinary option set {A, B, C} can generate the following cyclic pattern of pairwise verdicts: B is more choiceworthy than A, C is more choiceworthy than B, and A is more choiceworthy than C. This can happen, for example, if the claims satisfied under Option C are relevant to the claims satisfied under B but not to the claims satisfied under A. Cycles present a problem for accounts of moral permissibility, because it is hard to see how any option can permissibly be chosen from the full option set: each seems to be rendered impermissible by the option that defeats it on a binary basis. If we accept the Relevance View, which option ought we to choose from a nonbinary option set where each option is defeated by some alternative?

The task, then, is to determine whether there is a plausible extension of the Relevance View to nonbinary choices. I argue that much rests on how we conceive of the moral reasons that are in play in the cases of interest. My argument is structured as follows. In Section II, I set out two different ways that we could think about the structure of the moral reasons: I call them "global relevance" and "binary relevance." According to global relevance, a claim is relevant or not relative to the choice context as a whole. A relevant claim's being satisfied under a given option is a global reason—a reason relative to all the alternatives in the option set—to select that option. According to binary relevance, relevance facts only obtain relative to pairs of options. A claim's being satisfied under a given option can be a "binary contrastive" reason to choose that option rather than some specific single alternative, but it is not a reason to choose the option overall in a nonbinary choice context.

In Section III, I discuss global relevance, of which I take Voorhoeve's extension of ARC to nonbinary choices to be an example. I give two arguments for why we should not accept global relevance as an account of the structure of moral reasons like "Claim c would be satisfied under Option X." Instead, I argue, if we accept the Relevance View, then we should accept a binary contrastive account of the moral reasons.

In Sections IV and V, I ask whether this binary contrastive account could be sufficient to guide choice in cyclic nonbinary choice problems. In

---

Huemer target views that reject full aggregation (of which the Relevance View is one), arguing that the intuitions that might incline someone to accept such a view are unreliable. See also Tomlin, "On Limited Aggregation." Tomlin argues that the Relevance View delivers absurd verdicts when applied to some binary choices between heterogeneous sets of claims and offers no guidance in some other, more complex cases.

9. See, e.g., Derek Parfit, "Justifiability to Each Person," *Ratio* 16 (2003): 368–90; and Alistair Norcross, "Contractualism and Aggregation," *Social Theory and Practice* 28 (2002): 303–14.

Section IV, I consider choice rules that select permissible alternatives solely on the basis of the various binary rankings of pairs of options. I argue that we should reject this version of binary relevance, because it cannot accommodate strong intuitions about permissibility in three-option cycles. In Section V, I propose a new extension of the Relevance View to nonbinary choices that is consistent with a binary contrastive account of the moral reasons. On this view, weighing up options in competing claims cases is only possible on a pairwise basis, but the morally relevant considerations in nonbinary choices are not exhausted by simple binary choiceworthiness facts like "A is more choiceworthy than B." I argue that my proposal provides a richer basis for distinguishing among options that are related by a choiceworthiness cycle than alternative interpretations of binary relevance, delivering intuitive verdicts in cyclic cases while avoiding the pitfalls of global relevance.

II. GLOBAL AND BINARY RELEVANCE

We can think of the Relevance View as an account of the moral reasons that obtain in cases of competing claims. Consider the fact "a headache claim is satisfied under Option X." In many choice situations, this fact would count in X's favor—it would be a reason to choose X. Whether it does favor choosing X seems to depend on what the agent would thereby be giving up. For example, faced with a simple binary choice between X and one other option, the headache claim arguably counts in X's favor if the alternative option is to save a finger. But it does not count in X's favor against saving a life. It is not that the life claim happens to outweigh the headache claims with which it competes, but rather that a headache claim provides no support at all for choosing X instead of saving the life. In this case, we might say that the reason favoring X is "disabled" because a sufficiently stronger claim would go unsatisfied as a result of the agent's choosing X. A reason's being disabled is fundamentally different from its being outweighed. When one reason outweighs another, both reasons remain operational. When one reason disables another in a particular context, the disabled reason ceases to have any weight in that context. In a practical sense, it ceases to be a reason for the agent at all.[10]

10. See Jonathan Dancy, *Ethics without Principles* (Oxford: Oxford University Press, 2004), chap. 3. Victor Tadros also draws a distinction between outweighing and disabling in the context of the Relevance View: Victor Tadros, "Localized Restricted Aggregation," *Oxford Studies in Political Philosophy* 5 (2019): 171–204. A similar view has been proposed by David Lefkowitz, drawing on Raz's account of exclusionary reasons, on which a life's being at stake "excludes" the headaches from the agent's consideration: David Lefkowitz, "On the concept of a morally relevant harm," *Utilitas* 20 (2008): 409–23; Joseph Raz, *Practical Reason and Norms* (Oxford: Oxford University Press, 1999), chap. 1.

The sketch just given of the moral reasons in competing claims cases needs to be made more precise in its application to nonbinary choices. If Claim c's being satisfied under Option X gives us a reason to choose X, what is this a reason for, specifically? Under what circumstances does this reason obtain, and when is it disabled? I'll discuss two ways that the moral reasons could be structured. The first approach is "global relevance."[11] On this account, a relevant claim's being satisfied under some option is a "global reason," in that it counts in favor of that option relative to the overall choice context. More formally,

> *Global Relevance.* If Claim c's being satisfied under Option X is a reason to choose X, it is a reason to choose X relative to the whole option set {X, Y, ... , Z}.
>
> Claim c's being satisfied under X is such a reason if and only if c is relevant relative to {X, Y, ... , Z}.

Voorhoeve's extension of ARC to nonbinary choices is consistent with global relevance. According to Voorhoeve's specific relevance criterion—I'll call it "Global ARC"—a claim is relevant if and only if it is sufficiently strong relative to the strongest competing claim in the option set, where "competing" is read as "mutually unsatisfiable."[12] If we assume that the sets of claims satisfied under each option are nonintersecting,[13] the criterion is as follows:

> Claim c is relevant relative to {X, Y, ... , Z} if and only if c is sufficiently strong relative to the strongest claim satisfied under any option in {X, Y, ... , Z}\{X}.

According to global relevance, a claim's being satisfied under a particular option either gives you a reason to choose that option overall, relative to the entire choice context, or gives you no reason to choose that option in the given choice context.

The rival view to global relevance is "binary relevance." On this account, a claim's being satisfied under some option can give you a reason to choose that option relative to a specific single alternative, but it is not

11. Campbell Brown uses the term "global approach" to describe Voorhoeve's extension of ARC to nonbinary choices. Campbell Brown, "Is Close Enough Good Enough?," *Economics and Philosophy* 36 (2020): 29–59, 55.

12. Voorhoeve, "How Should We Aggregate?," 78; Tomlin, "On Limited Aggregation," 239.

13. There is a question about how ARC should be applied to cases where the same claim could be satisfied under more than one option; see Tomlin's discussion of quasi-competition in "On Limited Aggregation," 251–57. I will not pursue this question here, since it does not present a problem for binary relevance.

a reason to choose that option overall. Call such a reason a "binary contrastive reason."[14] More formally,

> *Binary Relevance.* If Claim c's being satisfied under Option X is a reason to choose X, it is a reason to choose X rather than some particular alternative, call it Option Y.
>
> Claim c's being satisfied under X is such a reason if and only if c is relevant relative to Y.
>
> Claim c is relevant relative to Y if and only if c is sufficiently strong relative to the strongest claim satisfied under Y.

We can represent global reasons as increasing the value of the option they favor, and the global relevance choice rule as directing us to choose an option with the highest value. We cannot represent binary contrastive reasons as increasing the value of one option in a nonbinary choice, since such a representation would not allow for a reason to favor an option relative to one alternative and not another. Whether and how binary relevance can nonetheless guide nonbinary choice will be the subject of Sections IV and V. For now, notice that binary relevance seems to restrict us to binary comparisons of options when weighing up the various options in a competing claims case. And in cases where the various pairwise comparisons of options yield a maximal option—one which is undefeated on a binary basis by any alternative—choosing that maximal option is well supported by the moral reasons, since we are justified in choosing that option over every alternative.

III. WHY WE SHOULD REJECT GLOBAL RELEVANCE

Voorhoeve's Global ARC directs us to identify the claims that are relevant overall, relative to the choice context, and then choose the option that maximizes the sum of strengths of satisfied relevant claims. There is a single fact of the matter as to whether a given claim is relevant overall, regardless of how many options there are. In a three-way choice between nonintersecting sets of claims, then, all claims that would be satisfied under Option X and Option Y are candidates to render a claim in Option Z

14. Contrastive reasons are discussed more generally by Justin Snedegar, e.g., in "Contrastive Reasons and Promotion," *Ethics* 125 (2014): 39–63. On Snedegar's account, considerations that are reasons for an option are always reasons for that option relative to some class of alternatives, not reasons for that option simpliciter. Within Snedegar's framework, what I am referring to as global reasons and binary contrastive reasons would both qualify as "contrastive." The difference is in the contrast class: for a global reason, the contrast class is the entire option set; for a binary contrastive reason, the contrast class is restricted to a single alternative, regardless of the size of the option set.

irrelevant. To ascertain whether Claim c is relevant, we simply identify the strongest claim that would be satisfied under an alternative option and assess whether c is sufficiently strong relative to that claim.

To illustrate how Global ARC works, I'll first introduce a case that will be a focus of discussion for the rest of the article. Suppose a third option involving claims of intermediate strength is added to Life for Headaches:

> *Case 3: Expanded Life for Headaches.* You can choose exactly one of the following options:
>
> Option A: Save one person from a premature death.
>
> Option B: Save 100 people from serious injury.
>
> Option C: Save 1,000,000 people from headaches.[15]

In this and all cases to follow, we assume that

- headache claims are sufficiently strong to be relevant relative to serious injury claims, but they are not sufficiently strong relative to life claims;
- serious injury claims are sufficiently strong to be relevant relative to life claims;
- 100 serious injury claims outweigh one life claim; and
- 1,000 headache claims outweigh one serious injury claim.

In a choice between A and B, you ought to choose B: 100 serious injury claims outweigh one life claim. In a choice between B and C, you ought to choose C: 1,000,000 headache claims outweigh 100 serious injury claims. And in a choice between A and C, you ought to choose A: the headache claims that would be satisfied under C are rendered irrelevant by the life claim, and so the only relevant claim at stake is the life claim. We are thus faced with a cycle: B is more choiceworthy than A, C is more choiceworthy than B, and A is more choiceworthy than C.

Note that this example assumes that there is a claim of intermediate strength—I have referred to it as a claim against "serious injury"—such that headaches are relevant to it and it is relevant to a life claim. On some precisifications of 'sufficiently strong' in the definition of relevance, though, there will be a gap between the weakest claim relevant to a life and the strongest claim to which a headache is relevant. If this is the case, we can construct an example with the desired cyclic structure by weakening the claim in Option A or strengthening the claims in Option C to

15. Similar cases are discussed in Norcross, "Contractualism and Aggregation," 308; Parfit, "Justifiability to Each Person," 384; Voorhoeve, "How Should We Aggregate?," 76–77; Johanna Privitera, "Aggregate Relevant Claims in Rescue Cases?," *Utilitas* 30 (2018): 228–36, 232; and John Halstead, "The Numbers Always Count," *Ethics* 126 (2016): 790–802, 797–98.

yield a three-option cycle, or by adding further intermediate options between Options A and B or Options B and C to form a cycle with more than three elements. I will leave aside the possibility that there is no finite sequence of claims originating with a headache claim and ending with a life claim such that each claim is relevant to its successor.[16]

Returning now to Global ARC, the view applies to Expanded Life for Headaches as follows. The claim to be spared from premature death that would be satisfied under Option A and the claims to be spared from serious injury that would be satisfied under Option B are relevant. The headache claims that would be satisfied under Option C are irrelevant, because each is sufficiently weaker than the life claim that would be satisfied under Option A. Option B satisfies the set of relevant claims with the greatest aggregate strength, so the agent ought to choose Option B—the option to prevent 100 cases of serious injury.

A. Counterintuitive Implications of Global ARC

Global ARC seems to get the verdict right in Expanded Life for Headaches. But it yields counterintuitive verdicts in other cases. Consider the following example:

> *Case 4: Counterexample to Global ARC.* You can choose exactly one of the following options:
>
> Option A: Save one person from a premature death.
>
> Option B: Save 1,000 people from serious injury.
>
> Option C: Save 999 people from serious injury and 1,000,000 from headaches.

First, consider how ARC would order the options in a series of pairwise choices. In a choice between A and B, you ought to choose B: all claims are relevant, and 1,000 serious injury claims outweigh one life claim. In a choice between B and C, you ought to choose C: all claims are relevant, and the 1,000,000 headache claims satisfied under C are sufficient on

16. If there was no finite sequence of claims each relevant to its successor connecting a headache and a death, then either there would be some claim that was not relevant to any stronger claim, or every claim would be relevant to some stronger claim but there would be some threshold such that no claim weaker than the threshold was relevant to any claim stronger than the threshold. This would amount to a view on which claims fall into two (or more) categories and no number of claims from the inferior category outweighs any number of superior claims, including in cases where the claims are nearly indistinguishable but happen to fall on either side of the threshold. Such a view might avoid the problem of choiceworthiness cycles, but it is not clear that it could still accurately be described as a version of the Relevance View.

their own to outweigh the serious injury claims in B. In a choice between A and C, you ought to choose C: the headache claims satisfied under C are irrelevant, but the serious injury claims satisfied under C outweigh the dying person's claim. This means that there is a maximal option—Option C—which ought to be chosen over each competing option.[17]

Voorhoeve's Global ARC, as I have characterized it, does not select this maximal option. The headaches in C are globally irrelevant, and so the option that satisfies the greatest aggregate of relevant claims is Option B. Here, Option B is selected over Option C because there are claims satisfied in Option C that are rendered irrelevant by a claim satisfied in a third option—Option A.

Critics of the Relevance View have pointed to examples like this one as evidence of the view's internal inconsistency. In Case 4, Global ARC appears to violate a principle called Basic Contraction Consistency or Alpha: the most choiceworthy option overall ceases to be the most choiceworthy option when some other option is removed from the option set.[18] ARC's proponents have, however, offered a convincing response to charges of internal inconsistency. In brief, Global ARC does not violate Basic Contraction Consistency because, from the standpoint of Global ARC, changing the membership of the option set changes the morally significant properties of the remaining options. Option C chosen from {A, B, C} is a different option from C chosen from {B, C}, and so there is nothing inconsistent about B's selection from the nonbinary option set and C's selection from the binary option set.[19]

Case 4 reveals a different problem for Global ARC, though: the view is at odds with what appears to be a natural and intuitive way to treat the headache claims in our deliberation about the case. In Case 4, we compare Option A to Options B and C, and it is clear that A is not the one we should choose. The claims to be spared serious injury that would be satisfied under each of Options B and C substantially outweigh the life claim that would be satisfied under Option A. Importantly, in comparing A and C, we do not count curing headaches in C's favor—the claims in C that are relevant to a life claim are on their own sufficient to outweigh the life claim. But the headache claims' being satisfied does seem to count in C's favor against Option B. In short, given that we ought not to satisfy the A claim, we should take headaches into account in determining which set of claims we do satisfy.

17. Brown discusses a case with a similar structure ("Is Close Enough Good Enough?," 28–29).

18. Voorhoeve, "How Should We Aggregate?," 78; Amartya Sen, *Collective Choice and Social Welfare* (Amsterdam: North-Holland, 1970), 17.

19. Voorhoeve, "How Should We Aggregate?," 78–79.

This example shows, I think, that we are reluctant to set aside headache claims once and for all. We can accept that headaches give us no reason not to save a life while maintaining that they still have some role to play in the decision whether to choose Option B or Option C within the full choice context. More generally, in nonbinary cases, we might accept that weak claims do not count against alternatives that render them irrelevant while maintaining that they still count against the other alternatives in the option set.[20] If we accept global relevance, though, we must accept that any claim counts against all the alternatives in the option set, or none of them.

In response to a different case in which a claim that ultimately goes unsatisfied affects which option Global ARC selects, Voorhoeve offers a justification of the global character of his view along the following lines. Showing sufficient respect for the person with the most at stake requires in part that we change the way we evaluate the alternatives. Giving the dying person's claim its due does not necessarily require us to satisfy her claim, but it does require us not to satisfy irrelevant claims instead.[21] In cases of heterogeneous sets of claims, where failing to satisfy the dying person's claim could involve satisfying both relevant and irrelevant claims, this requirement of respect might be better expressed in terms of the weight accorded to claims, rather than the claims that are ultimately satisfied. That is, giving the dying person's claim its due might require that we not accord weight to irrelevant claims when determining which option we ought to choose. Treating the dying person's claim with respect in Case 4, then, would involve setting aside the headache claims as globally irrelevant, and so not choosing Option C. However, it is far from clear that this is what respect demands. It seems equally plausible that respect for the dying person requires that we not accord the headache claims any weight against her claim. And the pattern of deliberation about Case 4 that I described above satisfies this requirement. Headaches are set aside in the decision about whether to save the dying person's life and only have significance in weighing up Options B and C.

Proponents of global relevance might argue that Case 4 shows not that we need to abandon the global account altogether but merely that Voorhoeve's specific relevance criterion needs revision so that the headaches in this particular case are relevant overall and you ought to choose

20. Victor Tadros endorses a related position, arguing that in a binary choice between competing heterogeneous groups of claims X and Y, a claim in Group X might count against one claim in Group Y but not another. Tadros, "Localized Restricted Aggregation." Space does not permit discussion of Tadros's view here, but note that questions remain concerning nonbinary choice; see ibid., 187–91, for discussion.

21. Alex Voorhoeve, "Why One Should Count Only Claims with Which One Can Sympathize," *Public Health Ethics* 10 (2016): 148–56, 153.

Option C. That is, Case 4 could be taken to show that Voorhoeve's relevance criterion rules too many claims globally irrelevant, and we should seek a more inclusive criterion of global relevance. However, the problem with Case 4 seems to stem from the global approach itself, not from Voorhoeve's specific criterion. Intuitively, the weak claims do not count against the alternative that renders them irrelevant, but they do count against the other alternatives in the choice context. This is not possible on any view on which a claim's relevance is established relative to the overall choice context.

B. Claim Ambiguity

A further difficulty for the global account is that in some cases it is ambiguous what claims we would be satisfying by choosing a particular option. According to ARC, a person has a claim if her well-being is at stake, and the strength of a person's claim is a function of her well-being if you help her (in which case her claim is satisfied) and her well-being if you do not. The view is illustrated with examples for which this definition of 'claim' applies unproblematically, where each person can only be at two possible well-being levels. But if as a result of your decision a person could be at more than two well-being levels, then it is ambiguous what claim you would be satisfying by choosing an option in which she is better off than she could have been. For example, suppose that you have three options available to you. Under Option 1, I am in perfect health. Under Option 2, I have a headache. Under Option 3, I die. If you choose Option 1, which claim of mine have you satisfied? Have you saved me from death, or have you prevented a headache? Voorhoeve explicitly excludes such cases from consideration when introducing ARC.[22] It would be a shortcoming of the view, though, if it had nothing to say about such cases, especially because (as I will argue) we can construct cases of this type that look like they should be treated in the same way as simple cases to which ARC does apply.

Claim ambiguity is a problem for global relevance because global relevance requires claims to be defined globally, relative to the choice context as a whole. For example, to apply Global ARC to a choice between Options A, B, and C, there needs to be a single fact of the matter about which claims you satisfy by choosing A from {A, B, C}. So, if Global ARC is to apply to cases like the one described in the previous paragraph, we need to revise Voorhoeve's definition of 'claim'. In particular, we need to specify the baseline well-being level from which a claim's strength is measured. In this section, I will discuss three candidate definitions,

22. Voorhoeve, "How Should We Aggregate?," 67. For discussion, see Daniel Ramöller, "On the Possibility of Limited Weighing of Lives" (PhD diss., Stockholm University, 2020), chap. 6.

offering counterexamples to each, and argue further that no global definition will be satisfactory.

Definition 1: Lowest possible well-being as claim baseline.—The first possibility is to define a person's claim relative to the option in which she is worst off. This definition is tested in the following example, which is a variant on Life for Headaches. In all the previous cases, it has been unambiguous what claims are satisfied under each option, so I have been able to describe each option in terms of the claims that choosing that option would satisfy. Here, I first describe the options in terms of the well-being levels each person would experience. The number in parentheses after each description represents the well-being of each person in that group.

> *Case 5: Apocalyptic Life for Headaches.* There is one x person and 1,000,000 y people. You can choose exactly one of the following options:
>
> Option A: x in full health (100); y's have headaches (99).
>
> Option B: x dead (0); y's in full health (100).
>
> Option C: x dead (0); y's dead (0).[23]

Options A and B in this case have the same well-being distributions as the two options in Life for Headaches. Under Option C, you do not help anyone and both x and the y's die. Here, choosing A and B both involve helping the y people to some extent, because under Option C the y people die.

On the definition of 'claim' we are considering, the baseline well-being level for a satisfied claim is the well-being the person would have were she the worst off she could be under the choice. For both x and the y's, this is zero. We thus have the following claims satisfied, assuming for simplicity that the strength of a claim is just the difference in well-being levels:[24]

> *Case 5: Apocalyptic Life for Headaches (strengths of claims satisfied).*
>
> Option A: x: claim of strength 100 satisfied; y's: claims of strength 99 satisfied.

23. Ramöller discusses a similar example ("On the Possibility," 150–51).

24. Voorhoeve's presentation of ARC allows for prioritarian weighting of well-being, so that the strength of a claim is determined not just by the difference in well-being levels but also by absolute well-being. Following Brown's approach, we can formalize this as follows: the weight of person i's claim satisfied by choosing A rather than B might be $f(w_{i\mathrm{A}}) - f(w_{i\mathrm{B}})$ rather than simply $w_{i\mathrm{A}} - w_{i\mathrm{B}}$, where $w_{i\mathrm{A}}$ represents i's well-being under A and f is some concave strictly increasing function (Voorhoeve, "How Should We Aggregate?," 66; Brown, "Is Close Enough Good Enough?," 35).

Option B: x: no claim satisfied; y's: claims of strength 100 satisfied.

Option C: x: no claim satisfied; y's: no claims satisfied.

As already mentioned, Options A and B are identical in terms of well-being distribution to the two options in Life for Headaches. And I believe that most—even all—proponents of ARC would advocate choosing Option A and saving x's life in Apocalyptic Life for Headaches, for much the same reason that they would advocate saving the life in the original Life for Headaches. However, the claims satisfied under Options A and B look very different across the two cases. In the original Life for Headaches, we compare x's claim of strength 100 against the claims of the y's of strength 1. In Apocalyptic Life for Headaches, we compare x's claim of strength 100 and the claims of the y's of strength 99 against the claims of the y's of strength 100. Given the large number of y's, the only way that Global ARC can direct us to choose A rather than B in this case is if the claims satisfied under Option B are irrelevant. These claims are among the strongest claims at stake in the choice, though, so they must be relevant. Thus, Global ARC directs the agent to choose Option B and cure the headaches. The problem is that if we define the claims of the y's globally, relative to the option in which they are worst off, then Option A and Option B both involve satisfying very strong claims. All these claims are relevant, and so the question comes down to which option has the highest aggregate well-being. This is Option B, because of the extremely large number of y people.

If Global ARC does imply that we ought to select Option B in Case 5, this must count heavily against it. Can this implication be avoided under the present definition of 'claim'? It might be observed that Option C is Pareto dominated: alternatives to Option C exist which are at least as good for every individual and strictly better for at least one individual. Perhaps the Global ARC choice rule could be amended so that we eliminate Pareto-dominated options before claims are even identified.[25] This would deliver the desired verdict in Apocalyptic Life for Headaches; the case would reduce to Life for Headaches. But we can make a minor revision to the case so that Option C is no longer Pareto dominated, and Global ARC once again gives a wildly implausible result. We simply add one more person, z, whose life is not at stake but who stands to suffer a headache under Options A and B:

Case 5′: Apocalypse with Survivor. You can choose exactly one of the following options:

Option A: x in full health (100); y's have headaches (99); z has headache (99).

25. I owe this suggestion to an anonymous *Ethics* reviewer. Brown also raises the possibility of combining Global ARC with a Pareto rule, in the context of an example without claim ambiguity. Brown, "Is Close Enough Good Enough?," 57.

Option B: x dead (0); y's in full health (100); z has headache (99).

Option C: x dead (0); y's dead (0); z in full health (100).

Option C is no longer Pareto dominated, because z fares better in C than in A or B. Option C thus remains in contention and supplies the baseline well-being for the claims of the y's:

Case 5′: Apocalypse with Survivor (strengths of claims satisfied).

Option A: x: claim of strength 100 satisfied; y's: claims of strength 99 satisfied; z: no claim satisfied.

Option B: x: no claim satisfied; y's: claims of strength 100 satisfied; z: no claim satisfied.

Option C: x: no claim satisfied; y's: no claims satisfied; z: claim of strength 1 satisfied.

The claims of the y's in B are still relevant, and so the counterintuitive verdict that you ought to choose Option B stands. Excluding Pareto-dominated options is not enough to avoid the conclusion that in certain three-option choices you ought to cure 1,000,000 headaches rather than saving a life.

Rather than trying to avoid selecting Option B in the original Case 5, proponents of Global ARC might instead embrace the idea that the verdict should deviate from Life for Headaches. They could point to the fact that what is at stake for the y's is radically different in the two cases. In Life for Headaches, the y's will suffer headaches if you give them no assistance. In Apocalyptic Life for Headaches, the y's will die if you give them no assistance. However, embracing the global verdict would come at a large cost. First, as mentioned above, the intuition that we ought to choose Option A in this case is overwhelmingly strong, just as in Life for Headaches. Second, suppose we choose Option B from option set {A, B, C} and we are asked, "Why did you choose Option B and not Option A?" What justifies the choice of B over A, according to this version of Global ARC, is that the y's are better off under B than under A. Advocates of the Relevance View would reject such an appeal to the incremental improvement in the well-being of the y's in Life for Headaches, and it is not clear why it should be an admissible justification in this case either.

Definition 2: Next-highest well-being level as claim baseline.—To avoid the counterintuitive verdict in Apocalyptic Life for Headaches, we could define a satisfied claim relative to the person's next-highest possible well-being level, rather than her lowest possible well-being level. This would mean that in Option B of Apocalyptic Life for Headaches each y person has two claims satisfied: a claim to be in perfect health rather than suffering a headache,

and a claim to live (with a headache) rather than die. While the claims of the y's against death are relevant, their incremental claims to headache relief are not. On this definition of 'claim', then, A is the option that satisfies the greatest aggregate of relevant claims, as desired.

Making the definition more precise, for individual i, rank all her possible well-being levels as a result of your decision from lowest to highest: w_{i1}, w_{i2}, ... w_{in}. Suppose that under Option X individual i is at well-being level w_{ik} for some $k \leq n$. Then, Option X satisfies several of i's claims, as follows:

- Claim to be at well-being w_{i2} rather than w_{i1}.
- Claim to be at well-being w_{i3} rather than w_{i2}.

  . . .
- Claim to be at well-being w_{ik} rather than $w_{i(k-1)}$.

Now consider another variant on Life for Headaches. Without assistance, person x will die. You have many options for how much assistance to offer to x: you could cause her to be at any one of 100 well-being levels between dead (0) and perfect health (100). Also, as in the original Life for Headaches, you are able to cure the headaches of the y's, but this would prevent you from offering any assistance to x.

> *Case 6: Life Increments for Headaches.* There is one x person and 1,000,000 y people. You can choose exactly one of the following options:
>
> Option 1: x in full health (100); y's have headaches (99).
>
> Option 2: x has a headache (99); y's have headaches (99).
>
> Option 3: x has a more severe headache (98); y's have headaches (99).
>
> Option 4: x has a slightly more serious condition (97); y's have headaches (99).
>
> . . .
>
> Option 100: x alive (1); y's have headaches (99).
>
> Option 101: x dead (0); y's in perfect health (100).

Under Options 1–99, x has multiple claims satisfied. The claims satisfied under each option are as follows (assuming again for simplicity that the strength of a claim is the difference in well-being levels):

> *Case 6: Life Increments for Headaches (strengths of claims satisfied).*
>
> Option 1: x: 100 claims of strength 1 satisfied; y's: no claims satisfied.

Option 2: x: 99 claims of strength 1 satisfied; y's: no claims satisfied.

. . .

Option 99: x: two claims of strength 1 satisfied; y's: no claims satisfied.

Option 100: x: one claim of strength 1 satisfied; y's: no claims satisfied.

Option 101: x: no claims satisfied; y's: claims of strength 1 satisfied.

All satisfied claims are the same strength (1), so none can be rendered irrelevant by any other claim. Thus, the option you ought to choose is the one recommended by straightforward aggregation: Option 101, under which x dies and the 1,000,000 y's each avoid a headache. By breaking x's claim up into small increments, the full force of x's life's being at stake is never properly brought to bear on the choice, and the headaches of the y's remain relevant.[26]

Definition 3: Highest possible well-being as harm baseline.—Another possibility is to redescribe ARC in terms of harms avoided rather than claims satisfied. Campbell Brown frames his presentation of the Relevance View this way, with the choice rule directing you to choose the option that minimizes the sum of significances of relevant harms.[27] A person is harmed in an outcome if she is worse off than she could have been, and the significance of her harm in an outcome is a function of her well-being in that outcome and in the outcome in which she fares best—I will again assume that it is simply the difference between these two well-being levels.[28] In cases where each person can only be at two possible well-being levels, this approach is equivalent to Global ARC under the other claim definitions I have discussed. However, the different definitions can lead to different verdicts in cases where people affected by the choice can be at more than two possible well-being levels.

26. Perhaps prioritarian considerations could be brought in here so that x's claim to be at well-being 1 rather than 0, satisfied under each of Options 1–100, renders the claims of the y's satisfied in Option 101 irrelevant, since the one-unit improvement in well-being would occur off a much lower base for x than for the y's. However, it would be possible to restate the example, making x's well-being increments even smaller at low well-being levels, so that the claims of the y's were relevant.

27. Brown, "Is Close Enough Good Enough?," 41.

28. Ibid., 40, 51. Note that Brown's formulation allows for well-being levels to be weighted by a prioritarian value function. See also Marc Fleurbaey, Bertil Tungodden, and Peter Vallentyne, "On the Possibility of Nonaggregative Priority for the Worst Off," *Social Philosophy and Policy* 26 (2009): 258–85. Fleurbaey et al. use the term 'complaint' rather than 'harm' to refer to the shortfall in a person's (prioritarian-weighted) well-being relative to the option in which she fares best.

The harms approach avoids the counterintuitive verdict in Apocalyptic Life for Headaches. Presenting the three options from that case in terms of harms sustained, we have:

Case 5: Apocalyptic Life for Headaches (significance of harms sustained).

Option A: x: no harm; y's: harms of significance 1 sustained.

Option B: x: harm of significance 100 sustained; y's: no harms.

Option C: x: harm of significance 100 sustained; y's: harms of significance 100 sustained.

The harms to the y's in Option A are sufficiently smaller than the harm to x in Options B and C and are thus irrelevant. The option that minimizes the sum of the significances of relevant harms is therefore Option A, as desired, since there are no relevant harms sustained under that option.

However, the harms approach would deliver a counterintuitive result in the following example:

Case 7: Competing Small Improvements. There is one x person and 1,000,000 y people. You can choose exactly one of the following options:

Option A: x extremely badly off (9.99999); y's in full health (100).

Option B: x very slightly better off than in A (10); health of the y's somewhat impaired (90).

Option C: x in full health (100); y's dead (0).

Comparing Options A and B, Option B is very slightly better for one person, x, and Option A is better for 1,000,000 people, the y's. Under Option C, you could restore x to full health, but with the catastrophic result that you fail to prevent the deaths of all of the 1,000,000 y's.

In a binary choice between A and B, unless the weighting for harms is prioritarian to an implausible extreme, the harms suffered by the y's in Option B are relevant to the harm suffered by x in A, and so, given that the y's vastly outnumber x, you ought to choose A. In the three-option choice, we have the following distribution of harms across options, assuming as before that the significance of a harm is just the difference in well-being levels:

Case 7: Competing Small Improvements (significances of harms imposed).

Option A: x: harm of significance 90.00001 sustained; y's: no harm.

Option B: x: harm of significance 90 sustained; y's: harms of significance 10 sustained

Option C: x: no harm; y's: harms of significance 100 sustained

If the harm to x in A renders the harm to the y's in B irrelevant, as it seems it must, the option that minimizes aggregate relevant harms is B. But again, consider the question "Why should you choose Option B and not Option A?" The justification that global relevance offers for choosing B is that the only thing of relevance at stake is x's well-being, even though x has less at stake in the decision between A and B than the y's do.

As with the counterexample discussed in the first part of this section, one could read this as an example of global relevance ostensibly failing to satisfy Basic Contraction Consistency. That is, whether you ought to choose Option A or Option B depends on whether Option C is included in the option set. For the reasons given earlier, I do not believe that this would be a successful objection to global relevance. My argument here is rather that it is obvious which option we ought to choose in this case, and that is Option A. And if we did choose Option B, a natural way to criticize the choice would be to point out that we have prioritized a miniscule benefit to a single person over more substantial benefits to 1,000,000 people. In criticizing the choice of Option B in this way, we are adopting a binary standpoint, asking why Option B was chosen and not Option A.

* * *

Other definitions of 'claim' or 'harm' are possible, but I do not believe that any global definition can avoid damaging counterexamples like the ones just discussed. The idea of a claim, much like the idea of relevance, seems to make more sense under a binary comparison of options than globally. What is the moral significance for the y people of my choosing Option A in Apocalyptic Life for Headaches? It depends on which option you compare it to. If I choose Option A rather than Option B, I am failing to prevent their headaches. If I choose Option A rather than Option C, I am saving their lives.

Claim ambiguity poses a serious problem for global relevance but no problem at all for binary relevance with a binary conception of 'claim' as just outlined. There is no ambiguity about which claims are satisfied when choosing from a binary option set, because each person can only be at a maximum of two possible well-being levels. In all three cases I have used to illustrate the problem of claim ambiguity, adopting this binary definition of 'claim' delivers a single maximal option on the binary choiceworthiness relation, and this is the option that accords with our intuitions. To illustrate for Apocalyptic Life for Headaches, Options A and B

both straightforwardly defeat C (the option where everyone dies). And the pairwise comparison between A and B is just Life for Headaches, in which the Relevance View directs you to choose Option A. Thus, A is maximal. Binary relevance will not always deliver a maximal option, though; such cases are the subject of the remaining two sections.

IV. BINARY RELEVANCE AND CYCLIC CHOICES

I have argued that we should reject global relevance. If we instead accept binary relevance, accepting that a claim only counts in an option's favor relative to a single alternative, then we are forced to weigh up the options pairwise. For each pair of options A and B, we weigh our reasons to choose A rather than B against our reasons to choose B rather than A.

If the only way to weigh up options is pairwise, then it seems that the binary "is at least as choiceworthy as" relation must be what determines what you ought to choose in nonbinary choices. This is not a controversial view: it is standard in social choice theory to construct a choice function from a binary relation, on which one is required to choose a maximal option, that is, an option that is not defeated by any other option.[29] Call this approach "maximization."[30] Maximization is the obvious way to understand nonbinary choice under the binary account of reasons that I have proposed. If at least some of our moral reasons are binary contrastive, then the most natural way to understand the task of selecting the option for which you have most reason is selecting an option for which you have at least as much reason as any alternative, taken separately.

The problem for the Relevance View is that the binary choiceworthiness relation can generate cycles. Recall Expanded Life for Headaches (Case 3):

> Option A: Save one person from a premature death.
>
> Option B: Save 100 people from serious injury.
>
> Option C: Save 1,000,000 people from headaches.

In this case, B is more choiceworthy than A, C is more choiceworthy than B, and A is more choiceworthy than C. There is no maximal option. Thus,

29. Sen, *Collective Choice*, 9. Sen defines the choice set as the set of "best elements" of the option set, i.e., elements that are at least as choiceworthy as any alternative, rather than the maximal elements, i.e., elements that are not strictly less choiceworthy than any alternative (equivalently, elements that are undefeated by any alternative). The Relevance View's binary choiceworthiness relation is complete, though, and so the best elements and the maximal elements are the same.

30. Brown, "Is Close Enough Good Enough?," 50–51.

we cannot make a permissible choice if we are required to choose a maximal option.

The broad assumption underpinning maximization is that the set of permissible options is wholly determined by the pairwise rankings among options, for example, "B is more choiceworthy than A." That is, if for each pair of options you tell me how they rank with respect to each other, I can tell you which option(s) are permissibly chosen from the full option set. It is possible to retain this assumption while avoiding maximization's problem with cycles. What we are looking for is a rule whose inputs are all the binary rankings of options, whose output is the set of permissible options, and for which the set of permissible options is always nonempty, even if the options form a cycle.

Brown proposes *transitive closure maximization* as an alternative to maximization that can accommodate cycles. Under transitive closure maximization, one ought to select the (or one of the) maximal option(s) on the transitive closure of the binary choiceworthiness relation. Use $\succcurlyeq$ to represent the binary relation "is at least as choiceworthy as" and $\succ$ to represent its strict part, "is more choiceworthy than." The transitive closure of $\succcurlyeq$ on option set S, represented as $\succcurlyeq^+$, is a relation that holds for every ordered pair from S for which $\succcurlyeq$ does, and also for the smallest number of additional ordered pairs needed to make $\succcurlyeq^+$ transitive.[31]

Applying the transitive closure rule to a cycle gives the result that all options are permissible. Suppose we have a cycle, options $X_1, X_2, \ldots, X_n$ such that $X_{k+1} \succ X_k$ for all $k < n$, but $X_1 \succ X_n$. Since $\succcurlyeq^+$ holds wherever $\succcurlyeq$ does, we have $X_{k+1} \succcurlyeq^+ X_k$ for all k, and $X_1 \succcurlyeq^+ X_n$. And since $\succcurlyeq^+$ is transitive, by repeated applications of transitivity we have $X_n \succcurlyeq^+ X_1$. This means that for any two options X_k and X_j, both $X_k \succcurlyeq^+ X_j$ and $X_j \succcurlyeq^+ X_k$. Every option is a maximal element of the option set, and so you can permissibly choose any option.

Brown offers two arguments in support of this choice rule. First, in cyclic cases like Expanded Life for Headaches, the options seem to be morally "on a par": each is less choiceworthy than another alternative, each is more choiceworthy than another alternative, and there is nothing to separate them. Described in this way, ruling all options permissible seems to get the verdict right.[32] Second, the view satisfies a condition that Brown calls "Expansion Consistency." Informally, suppose our moral theory

31. Ibid., 51. Note that Brown's definition of the maximal options of S on binary relation $\succcurlyeq$ is $X \in S$ such that for any $Y \in S$, $X \succcurlyeq Y$. I earlier defined maximal as undefeated, i.e., $X \in S$ such that for any $Y \in S$, $Y \not\succ X$. As mentioned in note 29, the two definitions are equivalent in the case of the Relevance View's $\succcurlyeq$ relation because of the completeness of $\succcurlyeq$ under the Relevance View. $\succcurlyeq^+$ is also complete, and so the two definitions are also equivalent in the case of $\succcurlyeq^+$.

32. Ibid., 51.

tells us that A is permissibly chosen from our option set. And suppose that on a binary comparison B is at least as choiceworthy as A. Expansion Consistency states that B must also be permissibly chosen from our option set.[33] Expansion Consistency has the implication that if our choice rule selects any element in a cycle, it selects all of them.

A consequence of transitive closure maximization is that it is permissible to cure the headaches in Expanded Life for Headaches. Each option in the three-option cycle is maximal on the transitive closure of the binary choiceworthiness relation, so each is permissible. But it is hard to square this with the strongly held intuitions that might lead someone to accept the Relevance View in the first place. Contra Brown, it is clear that the three options are not morally on a par; it seems just as wrong to cure headaches in this case as it does in the original Life for Headaches.

Transitive closure maximization is, of course, not the only choice rule that could be adopted to guide choice when there is no maximal option. But no choice rule that takes binary choiceworthiness rankings as the only basis for choice, as transitive closure maximization does, will be able to distinguish among the options in a three-option cycle. Each option is strictly more choiceworthy than one alternative and strictly less choiceworthy than another—there is nothing to differentiate them. If we want to take seriously the intuition that curing headaches is not permissible in Expanded Life for Headaches, such choice rules will not help us.

V. DIFFERENTIATING DEFEATS

I have just argued that we should reject any choice rule which takes as its only inputs binary choiceworthiness rankings of the form "A is at least as choiceworthy as B." However, accepting that the only way to weigh up the options in competing claims cases is pairwise may not commit us to such a choice rule. There could be further moral facts—facts beyond the simple binary rankings of options—that bear on the choice. These facts could allow us to discriminate between options in cases like Expanded Life for Headaches, where the options are indistinguishable in terms of binary choiceworthiness rankings.

To see how this might be possible, consider a sporting analogy.[34] There are three boxers competing in a championship. The only way to decide the champion is to run a series of two-way competitions—the rules of boxing do not allow all three competitors to compete simultaneously. And suppose at the end of the competition, once each competitor

33. Ibid., 56.

34. See Robert Sugden, "Why Be Consistent? A Critical Analysis of Consistency Requirements in Choice Theory," *Economica* 52 (1985): 167–83, in which Sugden offers a range of sporting examples to explore different kinds of choice functions.

has fought each of the other two competitors, there is a cycle: each competitor has defeated one opponent and been defeated by another. Should all three be crowned joint winners, or no one? Maybe there are further facts about boxing that could identify a single champion. For example, what if competitor C was defeated in a knockout and the other two defeats were by decision? C has the same number of victories and defeats as the other two competitors, but her defeat seems to count against her more substantially than do either of theirs.

This example contemplates a way of choosing the winning boxer(s) when bare facts about wins and losses in pairwise matches leave us with no way to discriminate. And yet it is consistent with the idea that the only way to compare the skill levels of a group of boxers is to have them compete in pairs. This additional fact about boxing performance—that a knockout is a worse defeat than other kinds of losses—is still a fact about boxers' performance in pairwise contests. In short, the only way to identify the best boxer is to look at their victories and defeats in pairwise competition, but the nature of these victories and defeats becomes significant if the binary "defeats" relation does not deliver a maximal element (in this case, an undefeated boxer).

Let's return to Expanded Life for Headaches:

> *Case 3: Expanded Life for Headaches.* You can choose exactly one of the following options:
>
> Option A: Save one person from a premature death.
>
> Option B: Save 100 people from serious injury.
>
> Option C: Save 1,000,000 people from headaches.

As in the boxing example, there are three binary contests. B defeats A, C defeats B, and A defeats C. Each option defeats one of the alternatives and is defeated by the other. Perhaps the reason we think it would be wrong to choose C from this option set is that C's defeat by A is somehow worse than the other defeats.

Suppose that we choose Option A from {A, B, C}. This option is defeated by Option B, but there is something we can say in favor of our choice that registers against Option B: by choosing A, we have saved the life of one person. This is a reason for choosing A over B, and it is not disabled by any of the reasons that favor B over A. I will refer to the combined weight of the claims in Option A that are not rendered irrelevant by the defeating option as the "partial justification" for A. These claims are not sufficient to justify choosing A over B in a binary choice, but they provide some justification for the choice of A. Compare this to the choice of Option C from {A, B, C}. Again, we are choosing an option that is defeated by

some other option, Option A in this case. But here there is nothing we can say in favor of our choice that registers against the option that defeats it. The reasons we might have offered in favor of choosing Option C over A—the headaches cured under Option C—are disabled by the fact that a life claim would be satisfied under Option A.

C's defeat by A is worse than A's defeat by B, then, because there is less that can be said in favor of C relative to A than can be said in favor of A relative to B. There is no reason to choose C rather than A, but the fact that a life would be saved under A is a reason to choose A rather than B. The same argument can justify the choice of Option B from the cycle overall. In support of choosing B over C, the option that defeats it, we can offer the fact that 100 cases of serious injury would be prevented under B. These reasons are, together, weightier than the reasons that can be offered in support of choosing A (that one life would be saved) or C (nothing). Thus, we should choose Option B, because B has the greatest partial justification of all the options in the cycle.

The choice rule just described can be expressed as follows: when faced with a three-element cycle, we should choose the option for which most can be said relative to the option that defeats it; we should choose the option with the greatest partial justification. This is not the only way we could differentiate the options in a cycle—I will discuss another possibility below—but it does line up with a plausible account of the kind of justification that would be called for if we were to select a single option from a cycle. Choosing A from the cycle, say, invites the question "Why choose A rather than B?" where B is the option that defeats A. And so our choice of A calls for a specific kind of justification: justification relative to Option B, which would satisfy a set of claims that together outweigh the claims that would be satisfied under A. Choosing A is already justified over Option C. To put this another way, if you are choosing from a three-element cycle, how the options rank pairwise fixes a salient alternative that a given option is valued against: the option that is more choiceworthy than it on a binary basis.

The situation becomes more complicated when dealing with cycles with more than three elements, because one option in the cycle may be defeated by two or more alternatives. How should we understand partial justification in such a case, when I defined it above as the justification that can be offered for an option relative to the option (singular) that defeats it? Consider another variant on Life for Headaches:

> *Case 8: Very Expanded Life for Headaches.* You can choose exactly one of the following options:
>
> Option A: Save one person from a premature death.
>
> Option B: Save 100 people from serious injury.

Option C: Save 200 people from serious injury.

Option D: Save 1,000,000 people from headaches.

Option E: Save 2,000,000 people from headaches.

This is a five-option cycle—each option defeats the option that immediately precedes it in the list, and Option A defeats Option E. There are also further binary rankings within the larger cycle: A is defeated by C, B is defeated by both D and E, C is defeated by E, and D is defeated by A. In this case, it seems that we ought to choose Option C. This is because the case is the same as Expanded Life for Headaches, except that for each of the weaker claim types there is a further option to satisfy twice as many claims of that type. So choosing to prevent 200 cases of serious injury in this case appears to be the most choiceworthy option for the same reason as preventing 100 cases of serious injury is in Expanded Life for Headaches.

How should we understand partial justification for options that are defeated by more than one alternative? Consider Option D, which is defeated by both A and E. Two possibilities present themselves. First, we could look at reasons for choosing D that register relative to Option A. In this case there are no reasons that can be offered in partial justification for D, since the headaches cured under D are irrelevant relative to A. Alternatively, we could look at reasons for choosing D that register relative to E, in which case the partial justification for D is that 1,000,000 people would be spared headaches. Suppose we accept the latter, comparing options based on the maximum partial justification that can be offered against a defeating option. Such a rule would direct you to choose Option D in Case 8. Interestingly, it would not permit you to choose Option E and prevent twice as many headaches, since E is defeated only by A and headaches are irrelevant relative to A.

If we instead compare options based on the minimum partial justification that can be offered against a defeating option, then there are no claims that can be offered in partial justification for choosing D. This is more promising: once we know that there is nothing that can be said in support of D against A, it does not seem to strengthen the case for choosing D that we can offer some justification for choosing it relative to Option E. This suggests that in these more complex cases, we should adopt a *Maximin Partial Justification* choice rule. Applying this to Case 8, the minimum partial justification for Options D and E would be no claims, and the option which maximized the minimum partial justification for the chosen option would be Option C, as desired.

I have just set out how the proposed choice rule would apply to simple cyclic cases. Here is the rule for nonbinary choices in general:

Choice Rule: Maximin Partial Justification. Let $S = \{A, B, \ldots, C\}$ be the option set. Let $\succcurlyeq$ represent the Relevance View's binary "is at least as choiceworthy as" relation and $\succ$ its strict part.

1. If there is a maximal option(s) on $\succcurlyeq$, choose it/one of them. If there is no maximal option, proceed to Step 2.
2. Set aside all options that can be eliminated in virtue of their position in the binary choiceworthiness rankings as follows. Let $S^* = \cup_i S_i$, where the S_i are all the subsets of S satisfying the following conditions:
   i) for all $X \in S_i$ and $Y \in S \backslash S_i$, $Y \not\succ X$; and
   ii) there is no nonempty proper subset of S_i satisfying (i).[35]

   S^* may consist of S in its entirety, or it may be a proper subset of S. If S^* is a proper subset of S, set aside all nonmembers of S^*.
3. For each pair of Options X and $Y \in S^*$ such that $X \prec Y$, identify all the claims in X that are relevant relative to Y. Their combined weight (i.e., the sum of their strengths) is the *partial justification for X relative to Y*.
4. For each $X \in S^*$, select the minimum value of the partial justification for X relative to Z for any $Z \in S^*$ for which $X \prec Z$. Call this the *minimum partial justification for X*.
5. Choose the option whose minimum partial justification is greatest.

Step 1 asks if there is a maximal option (or options). If there is, then you ought to choose it (or one of them). If we accept that the key moral facts at play in our cases of interest are of the form "X is more choiceworthy than Y," then it is hard to see how any other choice could be justified. If there is a maximal option, you can justify choosing it over each of the alternatives. And if you choose a nonmaximal option when a maximal option is available, then your choice can be criticized in a way that choosing the maximal option cannot be: there is another option available that is more choiceworthy than the option you chose.

If there is no maximal option, the binary choiceworthiness rankings among options may still be informative as to which options should *not* be chosen. Step 2 narrows down our options to S^*, the "Schwartz set."

35. Thomas Schwartz, "Rationality and the Myth of the Maximum," *Noûs* 6 (1972): 97–117, 113; see also Anders Herlitz, "Non-transitive Better Than Relations and Rational Choice," *Philosophia* 48 (2020): 179–89. The simplest form S^* could take is for there to be a single S_i, and all the examples discussed in this article have this structure. Note, however, that it is possible under the Relevance View for more than one disjoint S_i to exist. For example, it is possible to have two disjoint three-option cycles such that each member of the first cycle is equally as choiceworthy as each member of the second. Each of the three-option cycles would then satisfy the definition of S_i.

We can think of S^* as a maximal subset of S. Recall that X is a maximal element of S if no other element defeats it. S^* is analogously maximal because no nonmember defeats any member. Described in this way, it seems clear that whichever option we ought to choose must be a member of S^*: no option outside S^* defeats any option within. For example, if A, B, and C formed a cycle at the top of the rankings and there were three other options, D, E, and F, which were less choiceworthy than each of A, B, and C, then the Schwartz set would be {A, B, C}. It is difficult to see how a choice from {D, E, F} could be justified. While maximization cannot tell us which option we ought to choose in this example, the binary ranking of options does narrow down the set of potentially acceptable choices.

S^* has the following properties: each element of S^* is strictly more choiceworthy than some other element of S^* and strictly less choiceworthy than some other element of S^*.[36] It might be asked at this point how the Schwartz set relates to the set of options selected by the transitive closure maximization rule discussed in the previous section. It can be shown that membership of the Schwartz set of some option set S entails maximality on the transitive closure of "is at least as choiceworthy as" on S.[37] So if an option is not maximal on $\succcurlyeq^+$, it will be excluded at Step 2. Maximality on $\succcurlyeq^+$ does not entail membership of the Schwartz set, however. For example, suppose there are four options A, B, C, and D such that A, B, and C form a cycle, D is equally as choiceworthy as A, and D is strictly less choiceworthy than both B and C. D is maximal on $\succcurlyeq^+$ but is not a member

36. To see that each element of S^* is strictly more choiceworthy than some other element of S^*, suppose the converse: there is an element, Z, of S^* such that for any $Y \in S^*$, $Z \preccurlyeq Y$. $Z \in S^*$ entails that $Z \in S_i$ for some S_i satisfying the conditions at Step 2, i.e., for any $X \in S_i$ and $Y \in S \backslash S_i$ $Y \succ X$ and there are no proper subsets of S_i satisfying the condition. But if, as we have assumed, $Z \preccurlyeq Y$ for all $Y \in S^*$, then $S_i \backslash \{Z\}$, which is a proper subset of S_i, satisfies this condition. This contradicts our assumption that S_i had no such proper subsets. To see that each element is strictly less choiceworthy than some other element, again suppose the converse holds: there is an element, Z, of S^* such that $Z \succcurlyeq Y$ for any $Y \in S^*$. Again, $Z \in S^*$ entails $Z \in S_i$ for some S_i satisfying the conditions in Step 2. But then {Z} is a proper subset of S_i that satisfies the condition that for any $X \in \{Z\}$ and $Y \in S \backslash \{Z\}$, $Y \succ X$. This contradicts our assumption that S_i had no such proper subsets.

37. As shown in Rajat Deb, "On Schwartz's Rule," *Journal of Economic Theory* 16 (1977): 103–10, the Schwartz set of option set S is equivalent to the set of maximal elements on the transitive closure of the strict part of the choiceworthiness relation restricted to S, $\succ^+$, where "maximal" on some binary relation means undefeated on that relation's strict part (note that $\succ^+$ may not be complete on S, so the two definitions of 'maximal' discussed in notes 29 and 31 above are not equivalent). This means that the Schwartz set can be defined as follows: $\{X|\ \forall Y \in S \neg[(Y \succ^+ X)\ \&\ \neg(X \succ^+ Y)]\}$. The set of maximal elements on $\succcurlyeq^+$ on S, as discussed in Sec. IV, is given by $\{X|\ \forall Y \in S X \succcurlyeq^+ Y\}$. Suppose that X is a member of the Schwartz set, and thus, according to Deb's theorem, $\neg[(Y \succ^+ X)\ \&\ \neg(X \succ^+ Y)]$ for all Y. Then, for all Y, either $X \succ^+ Y$ or $\neg(Y \succ^+ X)$. $(X \succ^+ Y)$ entails $X \succcurlyeq^+ Y$, since $X \succ Z_1 \succ \cdots \succ Z_n \succ Y$ entails $X \succcurlyeq Z_1 \succcurlyeq \cdots \succcurlyeq Z_n \succcurlyeq Y$. $\neg(Y \succ^+ X)$ entails $\neg(Y \succ X)$, which, by the completeness of $\succcurlyeq$, entails $X \succcurlyeq Y$ and hence $X \succcurlyeq^+ Y$. Thus, if X is a member of the Schwartz set, X is maximal on $\succcurlyeq^+$.

of the Schwartz set. Here, it seems appropriate to draw a distinction, as the Schwartz set does, between D and the other options: while all options are defeated in at least one binary contest, Option D does not defeat any of its competitors, whereas A, B, and C do.

Having eliminated all nonmembers of the Schwartz set, we next (in Steps 3 and 4) look at the defeats suffered by each option that remains—as just mentioned, each option in S^* will be defeated by at least one other option in S^*. For each option, X, we focus on its worst defeat: the binary contest for which the claims satisfied under X that are relevant to the defeating option have the lowest aggregate strength. In Step 5, we choose the option with the greatest partial justification relative to the option that defeats it in its worst defeat.

In summary, Maximin Partial Justification directs us to follow the binary choiceworthiness facts as far as they will take us and then look beneath the pairwise rankings of options to identify the option whose worst defeat is least severe.

The choice rule is illustrated through the following example:

> *Case 9: Another Variation on Expanded Life for Headaches.* You can choose exactly one of the following options:
>
> Option A: Save one person from a premature death.
>
> Option B: Save 100 people from serious injury.
>
> Option C: Save 200 people from serious injury.
>
> Option D: Save 1,000,000 people from headaches.
>
> Option E: Save 2,000,000 people from headaches.
>
> Option F: Save one person from a headache.

There is no maximal option—each option is strictly less choiceworthy than at least one alternative—so we move to Step 2. There is a subset of the option set for which no nonmember is more choiceworthy than any member: {A, B, C, D, E}. No proper subset of {A, B, C, D, E} satisfies this condition. We thus set Option F aside.

Next, we consider the minimum partial justification for each option that remains:

- Option A is defeated by Options B and C. The same claim is satisfied under A relative to B and C, and it is relevant relative to each, so A's minimum partial justification is the weight of one life claim.

- Option B is defeated by Options C, D, and E. The claims satisfied under B are the same relative to each of C, D, and E and are relevant relative to each, so B's minimum partial justification is the combined weight of 100 claims to be spared from serious injury.
- Option C is defeated by Options D and E. The claims satisfied under C are the same relative to D and E and are relevant relative to each, so C's minimum partial justification is the combined weight of 200 claims to be spared from serious injury.
- Option D is defeated by Options E and A. While the claims satisfied under D are all relevant relative to E, the claims satisfied under D are all irrelevant relative to A. Thus, the minimum partial justification for D is zero.
- Option E is defeated by Option A. The claims satisfied under E are irrelevant relative to A, so the minimum partial justification for E is zero.

The option whose minimum partial justification is greatest is Option C—satisfy 200 serious injury claims. Thus, you ought to choose C.

A. Objections

To conclude my discussion of Maximin Partial Justification, I will consider two possible objections to the view.

First, there is the question of whether there are alternative choice rules that are more suitable for the task at hand. Space does not permit an exhaustive survey, but one initially very plausible approach is to take defeat margins into account. If we are interested in the severity of one team's defeat by another in basketball, for example, we might look at the difference in points scored by each team. It would be strange to look only at the number of points scored by the losing team, which is analogous to how the Maximin Partial Justification choice rule works for claims.

Could defeat margins be used to distinguish more severe from less severe defeats under the Relevance View? If Option X defeats Option Y, the strength of the defeat could be equal to the difference in the combined weight of relevant claims satisfied under Option X and Option Y. However, this approach gives the counterintuitive result that curing headaches could be obligatory in Expanded Life for Headaches. Option C's defeat by A could be by a smaller margin than both A's defeat by B and B's defeat by C. To illustrate, suppose that relevant claims have the following strengths:

Claim against death: 100.

Claim against serious injury: 10.

Claim against headache: 0.1.[38]

Then, we have the following defeat margins:

B (prevent 100 cases of serious injury) defeats A (prevent one death) by a margin of 900:

$$100 \times 10 - 1 \times 100 = 900.$$

C (prevent 1,000,000 headaches) defeats B by a margin of 99,000:

$$1{,}000{,}000 \times 0.1 - 100 \times 10 = 99{,}000.$$

A defeats C by a margin of 100:

$$1 \times 100 - 1{,}000{,}000 \times 0 = 100$$

since the headaches have zero weight in this comparison.

According to this rule, C's defeat by A is the narrowest, at 100. Also, C's victory margin over B is the largest, at 99,000. Taking margins into account would thus favor curing the headaches.

A rule based on defeat margins does not capture the intuition that C's defeat by A is somehow more serious than the other defeats. The problem is that if we identify the seriousness of a defeat with the defeat margin, then the only quantity that could capture the seriousness of C's defeat by A is the strength of A's claim, and this may be substantially lower than the combined strengths of the B claims and the C claims (when relevant).

The second possible objection is that Maximin Partial Justification appeals to global reasons, not binary contrastive ones, in identifying the permissibly chosen alternatives from a cycle. That is, it could be argued that choosing the option with the greatest partial justification looks structurally the same as choosing the option with the greatest (global) value. There is an important difference between Maximin Partial Justification and global relevance, though: the global reason that the partial justification for Option X might be claimed to express is derived from more fundamental binary reasons. To return to the boxing analogy, if there is no undefeated boxer from a series of binary contests, we might award the title to the boxer who suffered the least serious defeat or the most resounding victory. Our assessment of the seriousness of each boxer's

38. These claim strengths are chosen for simplicity, but the range could be made wider with the same result if the number of people with serious injury and headache claims was adjusted.

defeat could play a similar role in selecting a winner to some global measure (if such a thing existed) of the boxer's performance in the competition overall. And yet it is still true that boxers can only compete pairwise. The method of selecting a champion just described is fundamentally different from a global version of boxing we might imagine that involves pitting all the contestants against each other in a single melee. Similarly, the partial justification for Option X is still a measure of X's performance in a binary contest with an option that defeats it, even if this quantity is used to select an option from a nonbinary option set.

Most importantly, Maximin Partial Justification avoids the problems with global relevance that I discussed in Section III. First, I argued that global relevance yields counterintuitive verdicts arising from the fact that a claim's being satisfied counts in favor of an option relative to one alternative if and only if it counts relative to all the alternatives. Maximin Partial Justification avoids these counterintuitive verdicts by allowing that a claim can have weight in some binary comparisons and not others. Second, Maximin Partial Justification avoids the issue of claim ambiguity, because it compares options pairwise and so does not require a global definition of 'claim'. Finally, in the examples I discussed in connection with both these problems for global relevance, there was a maximal option which was intuitively the only permissible choice and which the global view did not select. Maximin Partial Justification selects a maximal option where one exists and so accords with our intuitions in all these cases.

Another way of framing the objection might be that Maximin Partial Justification illegitimately treats binary reasons as if they are global reasons. Maximin Partial Justification does involve comparing the strength of binary contrastive reasons across different pairs of options—we need to be able to make sense of the idea that our reasons to choose B rather than C can be stronger than our reasons to choose C rather than A, for example. But while we would not typically weigh these reasons against one another, I do not see any reason that they should not be comparable. And as long as we adopt an interpersonally comparable cardinal well-being scale, as the Relevance View does, we can specify unambiguously the criteria for when one set of partially justifying claims outweighs another.

VI. CONCLUSION

For those seeking a middle ground between aggregation and nonaggregation, the Relevance View offers great promise. It would be disappointing, then, if it could not be applied to nonbinary choices in a satisfactory way.

I have argued that Voorhoeve's extension of ARC to nonbinary choices has counterintuitive implications. We can trace the problem back

to the global structure the view attributes to moral reasons of the form "Claim c is satisfied under Option X." Contrary to Voorhoeve's global account, I have argued that such reasons are binary contrastive: Claim c's being satisfied under Option X could count in X's favor relative to some alternatives in the option set and not others. This binary contrastive conception of the moral reasons maintains a close connection between our deliberations about binary and nonbinary cases and avoids the counterintuitive implications of the global account.

Accepting binary relevance means accepting that weighing up the options in a competing claims case can only be done pairwise. We weigh the reasons to choose X rather than Y against the reasons to choose Y rather than X. In a nonbinary choice, which option we ought to choose is determined by how the options compare pairwise. But, importantly, this is not limited to bare choiceworthiness rankings like "X is more choiceworthy than Y." Facts about each option's defeats become significant when the binary choiceworthiness relation does not deliver a maximal option. If there is a maximal option on the binary choiceworthiness relation, we ought to choose it. But when faced with a cycle, we are not forced to accept that all options are morally equivalent.