Foundations for Non-Aggregative Ethics

Word Count: 10,975

Intuitively, the numbers can matter. If we can build a playground at one or another park, the fact that more children would go to one park surely bears on where to build it. The number of beneficiaries can affect what we ought to do.

But there are also cases where, intuitively, the numbers do not matter. If a stagehand collapses just before the curtain rises, it seems wrong to continue with the opera simply to avoid a short delay for the viewers. We should not allow someone to risk serious complications simply to spare many others from mild inconvenience, no matter how many there are. We must help the stagehand immediately. Here, the number of beneficiaries does not affect what we ought to do.[1]

Some people will doubt this last judgment. They may direct our attention to the truly jaw-dropping numbers of people we may imagine to be threatened with mild inconvenience. Would we really insist on inconveniencing a trillion people? A trillion trillion? A googol, which is 10^{100} and outnumbers the atoms of the observable universe? Such numbers defy comprehension; we lack any visceral sense of their enormity. Our minds picture much the same thing when asked to imagine a trillion grains of sand as when asked to picture a googol grains of sand even though the difference between them is about a googol grains of sand. But if we had any intuitive sense of these very large numbers, perhaps we would see that it is morally disastrous when a googol people are mildly inconvenienced, just as we see that it is morally disastrous when even a thousand children starve.

---

[1] Scanlon 1998: 235.

I see the appeal of this view, and it may well be that these intuitions reflect our failure to fully grasp how very many little things can add up to something big.[2] If this is right, then many of our intuitive moral reactions are mistaken and need to be retooled in the light of guidance from theory. But I also see the appeal of its polar opposite: the moral significance of what happens to separate persons does not sum. What happens to me matters morally and so does what happens to you, but these bits of moral significance do not sum, combine, or aggregate. For when I consider the googol audience members waiting a few extra minutes for the show to start, I find myself curiously unmoved by the delay. It just does not seem very important. Each of an unfathomably large group of people was mildly inconvenienced. Nothing more. These delays do not collectively amount to a moral catastrophe, and if the same thing were to happen tomorrow, I would neither tear out my hair nor loose bitter tears.

I believe that we must choose between these extreme possibilities, and this paper will go on to develop a version of the latter. But you may doubt that we face such a stark choice. Faced with two extremes, why not seek compromise? Perhaps the moral significance of what happens to separate individuals sometimes sums and sometimes does not. When does it sum? There are many possibilities but one promising answer goes: when each individual in the group has interests that are sufficiently strong relative to the strongest competing interests. Individuals whose interests are not sufficiently strong are simply set aside rather than summed. Call this the Compromise View (or, for the cognoscenti, the Relevance View).[3] It can explain why the numbers matter in some but not all choices. Let me sketch how.

---

[2] For discussion, see Broome 2004: 55 – 59, Mogensen 2022, and Pummer 2013.

[3] Voorhoeve 2014 calls this Aggregate Relevant Claims, and it has become fairly influential. For some examples of its further development see Brown 2020, van Gils & Tomlin 2020, Hart 2022/2024, Mann 2021, Rüger 2020, Steuwer 2021, and Tadros 2019.

When deciding where to build the playground, we may suppose that the children all have similarly strong interests. So the moral significance of adding a given playground sums across different children. We can compare this sum for one playground to the sum for the other playground. Since these sums grow as the number of children who would use the respective playground increases, the numbers affect where we should build the playground.

By contrast, the unresponsive stagehand has far more at stake than any audience member. So we set aside, rather than summing, the interests of the audience members. No matter how many audience members there are, all of their interests are simply set aside. So their numbers do not affect whether we should save the stagehand.

The Compromise View thus offers an explanation of these cases that aligns with our instinctive judgments about what to do. Moreover, it differs from the extreme possibilities above where the moral significance of what happens to separate individuals always sums or never sums. It seems to offer an attractive alternative.

But I believe that the Compromise View, or any alternative to the two extreme possibilities, faces a serious challenge. Even before addressing the question of what to do in these cases, there is a prior question worth considering: how morally important is a few minutes delay for each of a googol people?

It may seem silly, but please bear with me. Imagine an unfathomably large group of people each waiting a few extra minutes for the show to start. Think of families scattered across the cosmos: trillions of trillions ... of trillions of families scattered across trillions of trillions ... of trillions of planets. Each family, settling in to watch the broadcast, must wait a few extra minutes for it to start. How much does this matter?

I envision three possible answers. First: there is no fact of the matter absent specification of the alternatives. Was it perhaps the case that the sole alternative to any delay involved sacrificing a young stagehand? Then the delay matters little. But if the delay resulted merely from some producer's laziness or carelessness, then it is unfathomably bad.

I find this answer baffling. I seem able to consider the delay in itself and to form judgments about its importance without knowing the alternatives. This seems true more generally. The wobble of my kitchen table matters less than the leak under my sink. This remains true even if the table is easy to fix but the leak requires an expensive and hazardous excavation into the aquifers below. Even if the leak is not worth fixing, it continues to matter more than the wobble of my table. The alternatives do not change how much these facts matter; why should they do so for the delay above?

Second: the delay matters enormously. I doubt that this is anyone's intuitive reaction, but it certainly seems possible that our intuitions systematically neglect how little things can add up to something big. If true, though, this poses a problem for the Compromise View. Given that our intuitions about the importance of the delay are deeply mistaken, why trust the intuition that we should save the stagehand despite the delay? How could our moral intuitions be reliable judges of what to do in such cases when they are such poor detectors of how much things matter? This possibility thus undermines our evidence for the Compromise View. Insofar as one finds this possibility attractive, one should doubt our intuitions about cases where many people are mildly affected. This should lead one to reject the Compromise View in favor of the extreme view that the moral significance of what happens to separate people always sums.

Third: the delay matters little. This, I believe, is how things strike us intuitively. If this is true, then it is presumably because the moral significance of what happens to each person never sums. For if moral significance sometimes sums across persons, why not in this case? When we consider the delay in itself, we specify nothing (e.g. weighty countervailing interests) that would prevent the interests of separate persons from summing. We can go further and explicitly rule out such obstacles. Suppose that the delay happens not in order to rescue a stagehand, but simply because the stage manager, chatting with their colleagues, forgot to call the maestro to the pit. This is sheer carelessness, not a weighty consideration. Yet I continue to judge the importance of the delay much the same as before. The stage manager acts unprofessionally and should be reprimanded or perhaps fired, but their conduct does not strike me as morally abhorrent, nor does the delay now come to seem extremely morally important. The delay just does not seem to matter very much. This supports the view that people's interests never sum, pace the Compromise View.

When faced with diametrically opposed possibilities, often the wisest course involves combining the grain of truth in each. That is what the Compromise View attempts. But I believe that, unusually, this strategy does not succeed here. The problem, in brief, is that regardless of what we say about what to do, we still need to adopt some attitude towards the bare event of a mild delay for each of very many people, and whatever we say there seems to force our hand later on. The delay may matter enormously, or it may barely matter; either possibility favors one of the extreme views over the Compromise View.

No doubt this is all much too quick and neglects many possible responses, complications, and alternatives.[4] But I hope it does enough to make pressing the choice between these extreme possibilities.

Ultimately, I want to know which possibility is correct. Does the moral significance of what happens to separate individuals always sum or never sum? But while generations of utilitarians have given me a good grasp of the former, my sense of the possibilities for the latter remains weak. Most non-utilitarians incorporate both aggregative and non-aggregative moral considerations; they seek to circumscribe utilitarianism, not to banish it. Some prominent philosophers do propose theories that eschew all aggregative moral considerations; perhaps the best-known are Nozick 1974, Scanlon 1982/1998, and Taurek 1977, while related ideas appear in Kamm 1985, Nagel 1978, and Rawls 1971. Several contemporaries also defend views that claim to eschew all aggregation.[5] The view I will propose differs from each of these in ways that, I believe, enable it to better meet certain challenges and illuminate the space of possibilities. But it will be easier to explain the differences and why they are helpful after presenting the view.

Here is the plan. Building on insights from Kamm, Kumar, and Scanlon, this paper sets out a novel theory on which the moral significance of what happens to separate individuals never sums (§1). Predictably, this theory implies that we must save the stagehand no matter how many audience members there are. More laudably, it can also explain why, all else equal, we should build the playground that more children would use (§2). The view may thus allow

---

[4] One may instead, for example, accept a bounded or asymptotic view on which the moral significance of what happens to separate people sums only up to some limit (Lazar & Lee–Stronach 2019). Alternatively, one may favor the views of Henning 2023, Scanlon 2021, Zhang 2024, or others. I apply the challenge to these views elsewhere.

[5] See Dougherty 2013, Gustafsson 2014, Henning 2024, Hirose 2001, Kumar 2001, and Munoz-Dardé 2005.

I exclude Scanlon 2016 and Zhang 2024 from this list because although their theories limit the role of aggregative considerations in important ways, they clearly also include such considerations.

us to maintain these judgments alongside our inarticulate but, I believe, deep-seated sense that morality is somehow about individuals considered as individuals.

§1: An Alternative to Aggregation

We seek a theory that does not sum the moral significance of what happens to distinct individuals. Call these non-aggregative theories. Before examining particular candidates, though, we need to understand this condition more clearly. The issue is finicky but critical; stringent constructions can wrongly deprive non-aggregative theories of pertinent moral distinctions; lax constructions can wind up admitting aggregation through the back door.

So when does a theory avoid aggregating the moral significance of what happens to distinct individuals? To answer this question, it will help to develop the picture behind our non-aggregative moral judgments a bit further. In what follows, I present an argument for this view that helps to clarify its content (§1.1), extract some guidance for determining when theories aggregate across individuals (§1.2), and finally set out the theory of this paper (§1.3 – 1.6).

§1.1: Why Not Aggregate?

We have already seen one reason to doubt that the moral significance of what happens to distinct individuals has a sum: a mild delay for each of very many people does not seem disastrous, even though the sum of the moral significance of this delay for each person would be enormous if it existed. But this is not the only argument against aggregation worth our attention. For one, this intuition competes with another: a thousand deaths may seem disastrous even when a single death would not, suggesting that the moral significance of what

happens to separate people does sum. Our intuitions thus seem to equivocate on whether there is a sum of the moral significance of what happens to separate individuals. For another, this intuition on its own gives us little positive sense of the non-aggregative worldview.

A more illuminating argument against aggregation builds on insights from a posthumous publication by John Taurek.[6] Developing this argument will help us see how the normative world might look without aggregation.

Some properties, such as weight, clearly sum across persons. You and I together have some combined weight which is the sum of your weight and mine. In general, the sum of the weight of two (disjoint) objects is the weight of their mereological fusion. One could easily multiply examples: height, volume, age, net worth, citation count, and so on.

Other properties, Taurek observes, seem not to sum. Consider sprinting ability. Suppose for simplicity that this supervenes on one's best time in the hundred meter dash, and that my best is twenty seconds, yours is nineteen, and Usain Bolt once ran it in nine. Intuitively, you are a bit better at sprinting than me while Bolt is much better than either of us. But how does Bolt compare to the sum of your sprinting ability and my sprinting ability? Would we together deserve gold? Or if Bolt has more sprinting ability, how many ordinary runners would it take to exceed him? Such questions seem to have no good answer because the sprinting ability of distinct people does not sum. There seems to be no way to add your sprinting ability to mine; a fortiori there is no way to compare this sum to Bolt's sprinting ability. Individuals to varying degrees can sprint, but these degrees do not sum across individuals.

This raises a question: to which category does the moral significance of what happens to distinct individuals belong? Does it behave like weight, which sums across persons, or like

---

[6] Taurek 2021 was read at the APA but not published for decades thereafter (McMahan 2021: 4). For other arguments against aggregation, see Korsgaard 2013, 2018: 157 – 160 and Nebel 2023.

sprinting ability, which does not? This is hard to answer directly, but when we struggle to categorize some phenomenon, we can often gain epistemic leverage by considering how to categorize similar phenomena. I propose that we apply this strategy here and ask: what other properties are akin to the moral significance of what happens to an individual? And do these properties sum across persons?

I take it that the moral significance of what happens to an individual crucially involves their point of view, where a point of view minimally includes some conscious states and evaluative attitudes towards the world. Entities without conscious states, like pencils, lack this kind of moral significance, and the same seems true about entities with some conscious states (e.g. the qualia associated with seeing blue) but no evaluative attitudes.[7]

Moreover, the moral significance of what happens to an individual depends on their perspective in complex ways. This seems true even if we consider only simple ways of making people better or worse off. Whether giving my sister a graduation gift makes her better off partly depends on what the gift means from her point of view. Although the details do not matter for many practical or even theoretical purposes, the benefit to her in fact depends on complex features of her perspective, such as her attitudes towards this gift and our relationship, along with her understanding of the relevant social norms. Analogous remarks apply more generally. We may often be entitled to ignore these complications for the purposes of efficient policy-making, but the benefit one derives from some act of charity implicates complex attitudes on the recipient's part towards the donor and gift.[8] And in calling these features complex, I just mean to distinguish them from far simpler properties of a perspective,

---

[7] For similar views, see Korsgaard 2018: 17 – 21 and Scanlon 1982: 113 – 114. For related discussion, see Lee 2022 and Mogensen 2025.

[8] For an influential discussion, see Mauss 2001.

such as the duration of some conscious experience or whether it includes the qualia associated with the color blue.[9]

Let us suppose, then, that the moral significance of what happens to an individual partly depends on complex features of the perspective of that individual. What else is like this? Here are some candidates: curiosity, cleverness, funniness, surprise, reasonableness, rationality. Like moral significance, these depend on complex features of an individual's perspective and admit of comparisons between individuals. We can discuss whether someone is more curious than someone else, and if so whether they are much more curious than the latter. And while we cannot make these comparisons to a high degree of precision, comparisons of the moral significance of what happens to different people also exhibit considerable imprecision. Moreover, these properties also correspond to standards that are not purely up to us. Believing that one is incredibly curious does not make it so, just as believing that something is incredibly morally important fails to make it so.

Turn the second stage of our inquiry and ask whether these analogues sum across persons. Start with curiosity. As a child, Kurt Gödel was apparently termed der Herr Warum (Mr. Why) for his endless questioning.[10] Suppose that the child Gödel was far more curious than my sister, who was in turn a bit more curious than myself. How does the sum of my

---

[9] What about simple animals? Some, including the UK government, maintain that shrimp feel pain (UK Animal Welfare (Sentience) Act of 2022). If so, then the morally significant interests of shrimp may be exhausted by relatively simple pain experiences disanalogous to the above. Will this argument thus leave untouched the possibility that the interests of shrimp sum?

I think not. The question is whether the moral significance of some event for some entity sums across entities. If the moral significance of what happens to separate people does not sum, then moral significance as such does not sum across entities. To show that some quantity fails to sum, it suffices to show that it fails to sum in some case; one need not separately show that it fails in every case. Conversely, to show that some quantity sums, one must show that it sums in every case, not just in some of them. This is why, for example, fixed-population comparisons cannot alone show that there is a sum of wellbeing; one must also show that this sum applies to variable-population comparisons (Nebel 2023: 1077 – 1078).

[10] Dawson 1997: 1.

sister's curiosity and my curiosity compare to that of the child Gödel? This question seems hard to answer because it seems hard to see how one can sum the curiosity of two people. Individuals can be curious to varying degrees, but these degrees do not seem to sum. Similar remarks, I believe, apply to the other examples. One person may be far more surprised than another by some jumpscare. How great is the sum of their surprise? How does this compare to the surprise of some third person? Again, there seem to be no good answers.

Ultimately, whether these properties admit of a sum requires further investigation; these intuitive judgments are not decisive, and it is certainly worth sharpening up these questions and carefully considering candidate definitions of the sums of these quantities. What seems right, though, is that we should not consider the moral case in isolation. If similar properties do not sum, that is reason to doubt that the moral significance of what happens to an individual sums.

This argument is also worth distinguishing from another, to my mind worse, alternative. The alternative argument uses the premise: there is a sum of the moral significance of what happens to separate people only if there is someone that undergoes the events that happen to each individual. Since there is no such entity, there is no such sum. But this argument offers little help epistemically since anyone tempted to reject its conclusion would presumably also reject its major premise. By contrast, the argument above uses no such premise. Rather, it tries to gain epistemic leverage by considering whether similar, including non-moral, examples sum. Proponents of aggregation can accept this strategy, which could have turned out to favor their view if there had been, for example, a sum of curiosity.

Finally, this argument offers some sense of what the non-aggregative view looks like. The moral significance of what befalls separate people could be like their curiosity, cleverness,

funniness, surprise, reasonableness, or rationality. They could be like features of their outlook, of the world from their point of view, which often fail to sum. We will draw on these analogies in what follows.

§1.2: What Counts as Aggregation?

It might seem easy to distinguish the views that sum the moral significance of what happens to separate people from the views that do not. Just ask: does the view mention such a sum? If it does, then it is an aggregative view; if it does not, then it is a non-aggregative view.

But because addition and subtraction are so closely related, aggregative views can be reformulated to avoid explicitly mentioning any sum. Sums can be replaced by differences.[11] Schematically, the problem is that an aggregative theory might consider whether:

$a > b + c + d$

a exceeds the sum of *b, c,* and *d*

or it might consider whether:

$(a - b) - c > d$

The difference of (the difference between *a* and *b*) and *c* exceeds *d*.

Prohibiting only explicit summation will prohibit the former while allowing the latter, but it is not clear how much they differ. For example, I might want to know whether my atlas weights more than my backpack, cat, and dog. Instead of directly weighing the former and the latter and then comparing, I could first determine by how much my atlas outweighs my backpack,

[11] The possibility that aggregative theories can be given purportedly non-aggregative reformulations is not an idle concern; to my mind, it is exemplified by Gustafsson 2021's intriguing argument for utilitarianism from purportedly non-aggregative premises.

then subtract the weight of my cat, and finally compare the result to the weight of my dog. Should the latter procedure count as non-aggregative?

In view of this problem, one might categorize subtraction as aggregative. Perhaps any view that mentions sums or differences with respect to some dimension is aggregative, while any view that does neither is non-aggregative.

But this expands the category of aggregative theories and shrinks the space for non-aggregative theories too far. As §1.1 discusses, some properties that do not sum across persons nonetheless seem to "subtract" across persons. For example, although the curiosity of my sister and my curiosity do not sum, we can nonetheless consider the degree by which her curiosity exceeds mine, or the "difference" between her curiosity and mine. She might be a bit more curious than me or far more curious than me. Moreover, we can compare these differences across comparisons. It is plausible that the difference between kid Gödel's curiosity and my sister's curiosity exceeds the difference between my sister's curiosity and mine.

A theory of curiosity that accepted these observations would not seem ipso facto to be an aggregative theory. Moreover, classifying it as such would dramatically constrain the resources available to non-aggregative theories. Non-aggregative theories of curiosity would then be confined to using at most ordinal comparisons, such as that my sister is more curious than me. But even a complete ranking of everyone by their degree of curiosity (e.g. Gödel > my sister > me) would not capture all the curiosity facts, since it omits facts about the size of the gaps in curiosity between people in the ordering.

So non-aggregative theories must be allowed some kinds of subtraction (else they are wrongly deprived of theoretical resources) but not all kinds of subtraction (else any aggregative view can be reformulated to count as a non-aggregative view). Which kinds should be allowed?

Taking the example of curiosity as a model, I suggest that a theory is non-aggregative with respect to some degreed property (e.g. curiosity) if it uses only:

(i) comparisons between things with respect to this property,

EG: my sister is more curious than me

(ii) comparisons between the degree by which x exceeds y in this property and the degree by which z exceeds w, and

EG: the degree by which my sister is more curious than me is less than the degree by which Gödel is more curious than my sister

(iii) logical combinations of the above.

EG: my sister is more curious than me and less curious than Gödel

This is only a sufficient condition, but that is all we need to later show that the theory of this paper does not aggregate the moral significance of what happens to separate individuals.

We also need to rule out theories that seem aggregative even though, officially, they employ no summation. At a minimum, non-aggregative theories cannot allow iterated subtraction of the kind found in our first example. Although they can take the difference of two things, they cannot generally consider the difference of this difference. They can consider $x - y$ but not $(x - y) - z$. This may seem ad hoc, but our example of curiosity can again help to make sense of this:

Let $x - y$ be the degree by which x exceeds y in curiosity. Then,

Gödel – My Sister = the degree by which Gödel exceeds my sister in curiosity

(Gödel – My Sister) – Me = the degree by which (the degree by which Gödel exceeds my sister in curiosity) exceeds me in curiosity

While there is plausibly some degree by which Gödel exceeds my sister in curiosity, it is not so clear that there is some degree by which (the degree by which Gödel exceeds my sister in curiosity) exceeds me in curiosity. One might suspect the latter of a category error. People can exceed other people in their degree of curiosity, but degrees of curiosity cannot do this. Compare the following example:

> Let $x \mathbin{\hat{-}} y$ be the number by which x exceeds y in citations. Then:
>
> Gödel $\hat{-}$ Tarski = the number by which Gödel exceeds Tarski in citations
>
> (Gödel $\hat{-}$ Tarski) $\hat{-}$ Lindenbaum = the number by which (the number by which Gödel exceeds Tarski in citations) exceeds Lindenbaum in citations

It seems doubtful that numbers have citation counts, so this last term looks like a category error. Even if we assume that all numbers have zero citations, this last term differs starkly from the proceeding. It may be sensible for a theory to admit all terms of the form $x \mathbin{\hat{-}} y$ where x and y are people, without also admitting all terms of the form $(x \mathbin{\hat{-}} y) \mathbin{\hat{-}} z$, even when x, y, and z are all people.

To summarize, non-aggregative theories can compare the degrees by which individuals exceed each other in some property. Curiosity, which does not aggregate across persons, admits of such comparisons. But non-aggregative theories cannot admit iterations of subtraction, for doing so threatens to collapse the distinction between aggregative and non-aggregative theories. And, of course, non-aggregative theories cannot admit of summation. More work is needed to explore and organize the logical space here, but these guidelines will suffice to categorize the theories we encounter in the remainder of the paper.

§1.3: The Separateness of Persons

We build up to the theory of this paper in stages. The theory follows from three principles, one of which we have already discussed:

The Separateness of Persons (Separateness)

The moral significance of what happens to separate people does not sum.[12]

I adopt this name because I think the principle offers a promising way to understand Rawls's famous remark that "Utilitarianism does not take seriously the distinction between persons."[13] Those who disagree may prefer another name.[14]

By definition, the non-aggregative theories are those that satisfy Separateness, while the aggregative theories are those that violate it. This means that non-aggregative theories cannot aggregate a particular property: the moral significance of what happens to an individual. It is worth keeping in mind that this does not prohibit non-aggregative theories from aggregating other quantities. Consider the silly view that one should always help the heaviest group of people. When faced with the option of helping one group or another, one should help whichever group weighs more. This view is non-aggregative in our sense since it does not sum the moral significance of what happens to separate people, even though it does

---

[12] I intend "moral significance" to include a person's wellbeing, but also such considerations as being treated fairly even when this does not otherwise redound to one's benefit. See Scanlon 1998: 202 – 206, 216 – 218 and Wallace 2019: 160 – 162 for relevant discussion.

[13] Rawls 1971: 27.

[14] Some take the separateness of persons to be the claim that persons have rights, that we have stronger reason to refrain from harming others than we do to benefit them, or more generally that interpersonal and intrapersonal trade-offs differ in morally significant ways (Brink 1993). Such claims characterize a logically independent dimension of the separateness of persons from the one that interests me here.

To see this, consider threshold deontology, which accepts both that people have rights and that people's interests sum. Or consider "maximin welfarism," which claims we ought to maximize the minimum level of welfare. On this view, there are no rights and yet people's interests do not sum. Separateness is a claim about the kinds of morally significant things that exist. It does not itself imply that people have rights or, more generally, that certain interpersonal tradeoffs are unacceptable even though similar intrapersonal tradeoffs would be acceptable.

sum their weights. Or consider the view that one should maximize the sum of the number of people that one helps and the number of people that lead happy lives. This is also a non-aggregative view in our sense. It sums the *number* of people helped by some action with the *number* of people who will lead happy lives, but Separateness does not claim (implausibly) that *numbers* do not sum. More generally, theories can use numbers without violating Separateness. Such theories may, of course, face other objections. For example, they may be completely implausible in the absence of additional commitments that presuppose an aggregative theory.

Non-aggregative theories cannot appeal to certain considerations, such as the sum of people's interests. So what can they appeal to? How could they work? We can start to address this question by considering some examples.

Suppose that you finally decide to throw away an old pair of jeans. At some point you meant to fix them up, but after several years you have finally accepted that this will never happen. As you lift the lid of the trash bin, someone passing by notices and asks if they could use your old jeans for a quilting project. Clearly, you should give them the jeans; it would be churlish to refuse. Why is this? One explanation is that while giving them the jeans needs no justification, throwing away the jeans does call for some justification. The perspective of this stranger could call you to account, demanding some justification for your decision whereupon there would be nothing for you to say.[15] Throwing away the jeans cannot be justified to them, whereas giving them the jeans can be. Thus, you should do the latter.

It will also help to consider a variation on this example. Suppose that, as you head out to throw away the jeans, you bump into two of your neighbors, Alice and Bob. As chance

[15] This need not imply that this stranger would be within their rights to actually insist that you provide them with some justification. Such insistence may be rude and may overstep their prerogatives. Nonetheless, it may be true both that we should take others into account in these ways and that others cannot actually insist that we do so.

would have it, each of them is working on a project that could use an extra pair of jeans. Clearly, you may give the jeans to either neighbor, and you should not just throw them away. Why is this? As before, giving the jeans to one of your neighbors, say Alice, seems to call for some justification from the perspective of the other neighbor, say Bob. But here, there is a justification ready at hand: giving the jeans to Alice will allow her to use the jeans for her project. You can ask Bob to consider Alice's perspective and to recognize that they have similar interests at stake. As a result, Bob cannot insist on receiving the jeans. Giving the jeans to either neighbor is permissible because doing so can be justified to the other.

None of this reasoning, which I hope sounds sensible enough, requires aggregation. It involves individuals recognizing what other individuals have at stake, no sums needed. Non-aggregative moral theories can accept such reasoning, and the main role for our theory is to extend such reasoning to more complex cases. The heart of our theory, as in these explantions, is the idea of justifying an action to another. In this, of course, we take inspiration from the vision of Scanlon 1982.

What kinds of considerations can justify an action to another? Given Separateness, we cannot justify an action to someone by appealing to the sum of the interests of separate people. It is natural, then, to think that only the moral significance of what happens to someone can justify an action to someone else. We assume that this is true, a point of agreement with Scanlon 1998: 229.

§1.4: A Problem for Other Non-Aggregative Theories

Our commitments thus far leave open a range of possible views, including the influential views of Kamm 1993/2006, Kumar 2001/2016, and Scanlon 1998.[16] Before continuing to develop our theory, I want to present a problem for these views that, as §2.3 discusses, our theory avoids. (Readers uninterested in competing views may skip this section and §2.3).

One major challenge for any non-aggregative theory is to explain why we should save two lives rather than one. (Aggregative theories, of course, have an easy time of this). They struggle to explain what is wrong with reasoning like the following:

> Suppose that you can either save Alice or save Bob and Carol. If you save Alice, then both Bob and Carol can demand some justification. You can justify saving Alice to Bob by asking him to consider Alice's perspective. You can justify saving Alice to Carol by asking her to consider Alice's perspective. Thus, you can justify saving Alice to everyone. Thus, it is permissible to save Alice.

There are two influential non-aggregative accounts of why we must save the larger number: the Tiebreaking Argument (Scanlon 1998: 230 – 234) and the Balancing Argument (Kamm 1985: 182/1993: 114 – 119/2006: 57 – 61, Kumar 2001/2011: 138 – 143). Let me explain how these accounts work and why they reject the reasoning above.[17]

According to the Tiebreaking Argument, this reasoning fails at the first step: we cannot justify saving Alice to Bob by asking him to consider Alice's perspective. This is because Bob has two kinds of reasons to reject the option of saving Alice: (i) doing so allows him to die, and

---

[16] For Kamm, I have in mind the views she proposes which seem to avoid aggregation, such as her Sobjectivity 1 and Sobjectivity 2 (Kamm 1993: 162 – 163, 167 – 168) along with the Context-Aware View (Kamm 2006: 58 – 59). Kamm also discusses views that accept aggregation, such as her Sobjectivity 3 and Sobjectivity 4 (Kamm 1993: 175 – 176, 180 – 181) along with the "Method of Virtual Divisibility" (Kamm 2006: 66 – 69).

[17] I have taken some liberties with the presentation of Scanlon's Tiebreaking Argument, since in Scanlon's framework, individuals reject principles, rather than token actions.

(ii) doing so fails to appropriately value his life. By contrast, Alice has only the former kind of reason to reject the option of saving Bob and Carol: (i) doing so allows Alice to die. Thus, Bob's reasons are stronger than Alice's, so we cannot justify saving Alice to Bob. On the other hand, we can justify saving Bob and Carol to Alice by asking Alice to consider, say, Bob's perspective. Thus, we must save Bob and Carol.

(ii) is the key step because it introduces an asymmetry between Alice and Bob. Where does this reason come from? The idea is that even without considering Bob, one is permitted to save Carol rather than Alice. Taking Bob into account should then suffice to turn this permission into a requirement. Any policy that permits us to save Alice rather than save Bob and Carol must therefore fail to take due account of the value of Bob's life. That is the additional objection that (ii) records.[18]

According to the Balancing Argument, this reasoning fails at the second step: after asking Bob to consider Alice, you cannot also ask Carol to consider Alice. The idea is that after Alice and Bob are considered, they "balance" (Kamm 1993: 101, 2006: 58) or "neutralize" (Kumar 2001: 167, 2011: 141) one another. We must then set both of them aside when deciding what to do and cannot compare them with anyone else. Only Carol then remains to be considered, and her interests favor the option of saving Bob and Carol. Thus, we must save Bob and Carol.

Here is the problem. Both views, for structural reasons, cannot explain the following combination of judgments:

1. One should give Alice $1,000 rather than give Carol $900.
2. One should give $900 to each of Bob and Carol rather than give Alice $1,000.

[18] One may wonder why Alice cannot similarly complain, if one is required to save Bob and Carol, that her presence made no difference. After all, one would be required to save Bob and Carol even if Alice were absent (Raz 2003: 359 – 363). I reply to this objection in an unpublished manuscript ("A Defense of Tiebreaking"). The idea is that, among the sensible rules, saving the greater number maximizes the number of cases in which one's presence makes a difference. Given the aim of making individuals matter, majority rule beats unanimity.

To my mind, $100 is a large enough difference to make (1) plausible. Even so, just as one should save two lives rather than one, one should also benefit two people rather than provide one person with a slightly larger benefit. Readers who disagree with this example may nonetheless agree with other examples of the same structure. For example: one should give one person both health and dental insurance rather than give another person only health insurance, but one should give two people health insurance rather than give another person both health and dental insurance.

There are versions of both views can explain either judgment on its own, but no version of either view explains their conjunction. Let me explain why.

Start with the Tiebreaking Argument. Suppose that we can either give Alice $1,000 or give Bob and Carol each $900. We might argue, in parallel to the above, that we cannot justify helping Alice to Bob by asking him to consider Alice's perspective. As before, this is because Bob has two kinds of reasons to reject the option of helping Alice: (i) doing so forsakes his interests, and (ii) doing so fails to appropriately value him. By contrast, Alice has only the former reason to reject the option of helping Bob and Carol: (i) doing so forsakes her interests. The rest of the argument could proceed as before.

(ii) remains the key step that introduces an asymmetry between Alice and Bob. Can we defend it? Before, we appealed to the fact that, even without considering Bob, one is permitted to save Carol's life rather than save Bob's life. Thus, taking Bob into account should suffice to turn this permission into a requirement. One must treat Bob as a tiebreaker. But this argument does not apply here because without Bob, one is not permitted to help Carol rather than help Alice. Without Bob, there is no tie. Thus, choosing to help Alice rather than Bob and Carol need not neglect the value of Bob's life. There is a lacuna in the argument for (ii). Moreover,

one might argue that a permission to help Alice or Bob and Carol does take due account of the value of Bob's life. After all, Bob's addition turns the requirement to help Alice (rather than Carol) into a permission to help Bob and Carol. Why does this fail to appropriately value Bob?

Turn to the Balancing Argument. At first, this argument may seem inapplicable because none of the opposed interests perfectly balance each other. But proponents of the Balancing Argument need not and do not insist that balancing requires perfectly equal interests; rather "roughly equal" (Kumar 2011: 140) considerations can be balanced against one another. Even granting this, though, there is a dilemma for the Balancing Argument: either one's interests in receiving $1,000 are roughly equal to one's interests in receiving $900 or they are not.

If they are, then consider the choice between giving Alice $1,000 and giving Carol $900. Their interests balance or neutralize one another so we set them aside when deciding what to do. There are no other interests to consider. Thus, we are permitted to choose either option (or, perhaps, required to flip a coin). This conflicts with the obligation to help Alice.

If they are not, then consider the choice between giving Alice $1,000 and giving Bob and Carol each $900. By hypothesis, we cannot balance or neutralize anyone's interests. The Balancing Argument then does not tell us what to do, and it is not clear how a non-aggregative moral theory should extend its reasoning to cover this case. Thus, in either case, the Balancing Argument does not explain the conjunction of the above judgments.

The Tiebreaking and Balancing Arguments fail to explain the above judgments. This need not refute them, but it does suggest that they are incomplete. Moreover, in view of the similarity between the above judgments and the original case of saving one or two lives, one

may reasonably expect a unified explanation of these cases. Our theory, to which I now return, will be able to provide one.

§1.5: No Multiplicity

Recall that one challenge for non-aggregative theories is to explain why the following reasoning errs:

> Suppose that you can either save Alice or save Bob and Carol. If you save Alice, then both Bob and Carol can demand some justification. You can justify saving Alice to Bob by asking him to consider Alice's perspective. You can justify saving Alice to Carol by asking her to consider Alice's perspective. Thus, you can justify saving Alice to everyone. Thus, it is permissible to save Alice rather than Bob and Carol.

The next step in developing our theory is to answer this challenge.

Our theory locates the fault at the second stage: the attempted justification to Carol fails. This is because Alice's perspective can justify helping her to one person, e.g. Bob, but her perspective cannot justify helping her to multiple others, e.g. Bob and Carol. It is unfair and smacks of double-counting to allow Alice's perspective to justify this to both Bob and Carol. The second principle of our theory generalizes this thought:

> No Multiplicity
>
> One cannot justify an action to multiple people by repeatedly appealing to the moral significance of what happens to the same person.

This principle prevents Alice's perspective from being an inexhaustible font of justification, available for address to as many people as one likes. This captures what intuitively goes awry in the above reasoning. Each counts for one. That is why one cannot count for many.

Moreover, there is something odd about a view on which an agent can appeal to someone's perspective to supply justifications to multiple others, but nobody can demand justifications from multiple others. Bob cannot, for example, insist on hearing from multiple people whose lives are at stake; his objection is answered upon hearing from Alice. Given Separateness, an agent cannot appeal to the interests of multiple people to justify an action to one person. One might think, in parallel fashion, that an agent cannot appeal to the interests of one person in order to justify an action to multiple people. That is what No Multiplicity claims.

One might object to No Multiplicity on the grounds that it is perfectly fine to repeatedly offer the same justification to as many people as one likes. After all, if some reason makes an action permissible, then one's justification to anyone for performing that action could be that reason. If I decide to go on a hike because it will be relaxing, then in response to anyone who asks me why I went on a hike, my answer could be to relax. There seems to be nothing illicit about this, pace No Multiplicity.

But this conflict is merely apparent and arises from an ambiguity in the phrase "to justify an action to someone." There is an important difference between the justifications to others that (i) make some action permissible and those that (ii) explain to someone why an action is permissible. Often, we have in mind the latter sense: a student might justify their decision to study philosophy to skeptical friends by appealing to the thrill of thinking through first principles; I might justify going for a hike to anyone who asks by explaining that I find them relaxing. In this sense of justification to others, it is quite plausible that the same justification can be provided to multiple people.

These justifications, though, are quite different from those that *make* the relevant actions permissible. The speeches above do not *themselves* count in favor of majoring in philosophy or going on a hike; rather, they cite other considerations that favor these choices. No Multiplicity applies to the justifications to others that make actions permissible; it does not constrain the explanations we provide to others for why some action is permissible. As a result, it does not rule out appealing to the same consideration when explaining, to multiple people, why one went for a hike. This limited scope for No Multiplicity corresponds to a more general and intuitively significant distinction between (i) repeatedly using a reason to make an action permissible and (ii) repeatedly appealing to a reason to explain oneself to different people.[19] These two senses of repeatedly raising the same consideration must be distinguished.

Two final observations. First, No Multiplicity implies that the justifications that one can provide to different parties are connected. Whether you can justify some choice to one person by appealing to Alice's perspective depends on whether you also justify that choice to someone else by appealing to Alice's perspective. In light of this, we cannot consider the justifications we provide to each person separately; we need to consider them together to ensure they cohere. What is morally basic, on this view, is not the individual justifications we provide to particular individuals, but the complete package of all justifications that we provide to every individual. Call this a justificatory narrative. In general, there are many possible justificatory narratives, and we will soon explain how they bear on what we ought to do.

Second, No Multiplicity is similar to Kamm and Kumar's proposal from §1.4 that roughly equal interests "balance" or "neutralize" one another. It differs in two ways. First, it applies even when people have starkly different interests at stake. In this respect, No

[19] Utilitarianism, for example, restricts (i) by insisting that no one's interests can be counted more than once, but it places no restrictions on (ii).

Multiplicity can be seen as a generalization of balancing or neutralizing. Second, and relatedly, it does not require us to set aside the interests involved. It simply bars us from repeatedly appealing to those interests. We will see, in §2.3, how this allows our theory to close the explanatory gap raised in §1.4.

§1.6: Minimax Complaint

Thus far, we have only considered cases where an agent justifies an action to someone by appealing to the equal moral significance of what happens to someone else. What happens when an agent justifies an action to someone by appealing to someone with less at stake?

Suppose that you are the only person in our friend group who has a car. You can spend the morning driving your beat-up minivan to help me move or to drop off another friend at the airport. Without a car, I would have to schlep my books to Brooklyn over a series of subway trips. Without a car, the other friend could take the subway and bus to the airport. As it happens, you end up driving the other friend to the airport. When I complain, you try to justify our choice by asking me to consider the perspective of the other friend.

How much this should succeed in diminishing my complaint seems to depend on the importance of this choice for the other friend. If this other friend is so bad at navigating public transit, they cannot reliably get to the airport, then this may be a justification strong enough to make any complaint inappropriate. But if this other friend is perfectly competent and just dislikes waiting for the subway, this justification should do very little to diminish the force of my complaint. We can also consider intermediate possibilities: perhaps the friend has several suitcases and it would be challenging (but feasible) to transport them all on the subway in one trip. This is a stronger justification than before and should do more, intuitively, to diminish my

complaint, but I may still have some complaint even after considering their perspective. After all, I may have had to make three such trips to transport the books.

This suggests that even when one person has more at stake than another, the latter can still provide partial justification for some action to the former. Consideration of the latter's perspective should still diminish the complaint that the former has against the action. The strength of the former's complaint should then depend on the moral significance of the choice for both parties.

Our theory accepts these claims. On this picture, when you justify (to me) the decision to leave me stranded by appealing to the perspective of your other friend, I am left with some complaint whose strength depends on the significance of this choice for me and for this other friend. This raises the question of how, in general, the strength of such complaints is determined, but for this paper, we will use intuitive judgments about the strengths of such complaints.

We now turn to the third and final principle of the theory. The need for this principle arises from Separateness and No Multiplicity, which jointly imply that there are cases where no matter what we do, some people are left with some complaint.

For example, recall the choice between giving Alice $1,000 (A+) or giving Bob and Carol each $900 (B∧C). On any possible justificatory narrative (i.e. complete specification of the justifications we provide), we cannot justify A+ to everyone. This is because both Bob and Carol can demand some justification. By No Multiplicity, we cannot ask both of them to consider Alice's perspective. One of them, say Carol, then receives no justification. Carol then has a complaint against A+. Moreover, Carol's complaint is fairly weighty, as it corresponds to her interest in receiving $900.

We also cannot justify B∧C to everyone. This is because Alice can demand some justification. We can ask Alice to consider Bob's perspective or we can ask her to consider Carol's perspective, but by Separateness, we cannot ask her to consider their combined interests. If, say, we ask Alice to consider Bob's perspective, Bob provides some partial justification but Alice is left with some residual complaint against B∧C. This complaint is relatively weak, though, as Bob has nearly as much at stake as Alice does.

Ideally, we would choose an action that, on some justificatory narrative, leaves no one with any complaint. Sometimes, as we have just seen, we cannot do so. Nonetheless, even in such cases, we can still approximate this ideal. We can pick the action that, on some justificatory narrative, minimizes the maximum complaint of any individual. In the above case, for example, we can choose B∧C. This will leave Alice with some complaint, but her complaint will be weaker than the complaint that choosing A+ would leave Carol with. Generalizing this thought yields the final principle of our theory:

> Minimax Complaint
>
> One may choose an action iff relative to some justificatory narrative, the strongest individual complaint against this action is no stronger than the minimum possible maximum individual complaint relative to any action and justificatory narrative.

To parse this principle, think of it as operating in steps. First, by varying the action and justificatory narrative, how weak can we make the strongest individual complaint against what we do? Call this the theoretical minimum for the maximum individual complaint. Second, which actions are such that, given some justificatory narrative, the strongest individual complaint they raise is no stronger than the theoretical minimum? Exactly those actions are

permissible, according to Minimax Complaint. On some justificatory narrative, they come as close as possible to the ideal of an action that raises no complaints.

One might be perturbed by the thought that sometimes, no matter what we do, somebody has a legitimate complaint. It might seem especially worrying to think that this can happen in cases as mundane as that above. Why not, instead, think that individuals with conflicting complaints should be again directed to consider one another's perspectives? In the above example, perhaps Alice, with her weak complaint against B∧C, should be invited to consider the perspective of Carol, who has a strong complaint against A+. In light of this, Alice should drop her weak complaint against choosing B∧C, making this option complaint-free.

I see the appeal of this thought, but it leads us back to aggregation. Complaints in our theory are like "differences" between the moral significance of some choice for one person and the moral significance of this choice for another person. Asking somebody with a complaint to consider some other person's complaint as justification is like taking the difference of these differences. But views that consider such iterated differences, as §1.2 discusses, seem no different from aggregative moral theories, and classifying such views as non-aggregative would threaten the significance of the distinction, allowing aggregative theories to be reformulated with iterated differences. Despite appearances, therefore, I do not believe the iterated complaints view to be a non-aggregative moral theory.

Our theory is the conjunction of Separateness, No Multiplicity, and Minimax Complaint. Could this also be, despite appearances, an aggregative moral theory? To assuage this concern, let us think about the kinds of information this theory employs. It considers individual complaints, which are like "differences" in moral significance, and it compares these

complaints in strength. As §1.2 discusses, this is the kind of structure that curiosity, a non-aggregative property, exhibits. So our theory seems to be non-aggregative after all.

§2: Applications

As mentioned in §1.1, our intuitions about whether there is a sum of the moral significance of what happens to separate individuals seem to equivocate, with different cases raising different judgments. I therefore do not believe that these issues should be settled primarily by considering which theory best matches our intuitive reactions to particular cases. Nonetheless, such considerations can serve an auxiliary role, and this section shows how our theory, the conjunction of Separateness, No Multiplicity, and Minimax Complaint, yields the intuitively correct verdicts in the cases considered above.

§2.1: Playground

Suppose that the town council can build a playground in one of two neighborhoods. For ease of illustration, suppose that one child, Una, lives in Einhaus while two children, Dora and Tressa, live in Twyford, and that each child will benefit to the same degree.

Consider the option of building the playground in Einhaus. No matter how we try to justify this action, some child will be left with a complaint. This is because, by No Multiplicity, we can ask one of Dora and Tressa to consider Una's perspective but not both. So either Dora or Tressa will be left with a complaint against choosing Einhaus.

Alternatively, consider the option of building in Twyford. This can be justified to everyone. The only one who can demand a justification is Una, and we can ask Una to consider the perspective of, say, Dora. Since Una and Dora have similar interests at stake, Una

is left with no complaint. By Minimax Complaint, we must choose to build in Twyford. Similar reasoning applies to any case where one can bestow sufficiently similar benefits on a smaller or larger group.

§2.2: Theatre

Now consider the theatre case. Suppose Una is the stagehand who has just fainted and there are a trillion audience members waiting impatiently. We can help Una, thereby delaying the show by a few minutes, or we can start immediately, thereby putting Una's health at risk.

Helping Una cannot be justified to everyone on any justificatory narrative because, as before, No Multiplicity bars us from appealing to Una's perspective to justify this act to multiple audience members. So nearly all the audience members will have an unanswered complaint. But each of their complaints will be fairly weak, since it corresponds to their interests in avoiding a short delay. Thus, the maximum individual complaint will be fairly weak.

By contrast, starting the show immediately leaves Una with a weighty complaint on any justificatory narrative. We can ask Una to consider the perspective of some audience member as (extremely slight) justification, but the stakes are so much higher for Una that she retains an weighty complaint. Separateness bars us from asking Una to consider the combined significance for multiple audience members.

Thus, on any justificatory narrative, Una's complaint against starting the show immediately is far stronger than the strongest complaint from any audience member against helping Una. By Minimax Complaint, we must help Una.

§2.3: Alice, Bob, and Carol

Our theory can also answer the challenge in §1.4 to explain both of the following judgments:

1. One should give Alice $1,000 rather than give Carol $900.

Giving Alice $1,000 can be fully justified to Carol by appealing to Alice's interests. By contrast, giving Carol $900 cannot be fully justified to Alice by asking her to consider Carol's perspective given how much less Carol has at stake. So Alice retains some complaint against giving Carol $900. Thus, by Minimax Complaint, one must give Alice $1,000.

2. One should give $900 to each of Bob and Carol rather than give Alice $1,000.

Consider the option of helping Bob and Carol. We can partially justify this action to Alice by asking her to consider, say, Bob's perspective. This leaves Alice with a weak complaint, since she has only a bit more at stake than Bob does.

Alternatively, consider the option of helping Alice. As before, one cannot justify this action to both Bob and Carol on any justificatory narrative since No Multiplicity bars this. So either Bob or Carol will receive no justification. That person will have a fairly strong complaint against helping Alice, since they will have about $900 at stake.

Thus, helping Bob and Carol leaves Alice with a weak complaint whereas helping Alice leaves someone with a fairly strong complaint. By Minimax Complaint, one must help Bob and Carol.

Conclusion

It is at best unclear whether the moral significance of what happens to separate people has a sum.[20] Just as there is no sum of the curiosity of separate people, there may be no sum of the

[20] Nebel 2023.

interests of separate people. And some of our moral reactions, such as our sense that a short delay for very many people matters little, suggest the same.

This paper develops a moral theory that forgoes such sums. The theory consists of three principles that, I believe, represent simple and natural generalizations of the ideas underlying many non-aggregative views. The theory also explains why we ought to save two lives rather than one and overcomes an explanatory gap that arises on other non-aggregative views. It remains to be seen, though, whether this theory, or any non-aggregative moral theory, can be extended to address the far more complex cases that we actually face.

Bibliography

Brink, David (1993). The Separateness of Persons, Distributive Norms, and Moral Theory. In Raymond Gillespie Frey & Christopher W. Morris, *Value, Welfare, and Morality*. pp. 252-289.

Brown, Campbell (2022). Aggregation and Self-Sacrifice. *Ethics* 132 (3):730-735.

Dawson, John W. (1997). *Logical Dilemmas*. Routledge.

Dougherty, Tom (2013). Rational Numbers: A Non-Consequentialist Explanation Of Why You Should Save The Many And Not The Few. *Philosophical Quarterly* 63 (252):413-427.

Geach, P. T. (1956). Good and Evil. *Analysis*, Vol. 17, No. 2: 33-42.

van Gils, Aart & Tomlin, Patrick (2020). Relevance Rides Again?. *Oxford Studies in Political Philosophy Volume 6.* Edited by David Sobel, Peter Vallentyne & Steven Wall. Oxford University Press.

Gustafsson, Johan E. (2021). Utilitarianism without Moral Aggregation. *Canadian Journal of Philosophy* 51 (4):256-269.

Hart, James (2022). Tie-breaks and Two Types of Relevance. *Ethical Theory and Moral Practice* 25 (2):1-20.

Hart, James (2024). Limited Aggregation's Non-Fatal Non-Dilemma. *Australasian Journal of Philosophy*.

Henning, Tim (2023). Numbers without aggregation. *Noûs* (3):755-777.

Hirose, Iwao (2001). Saving the greater number without combining claims. *Analysis* 61 (4):341-342.

Horton, Joe (2017). Aggregation, Complaints, and Risk. *Philosophy and Public Affairs* 45 (1):54-81.

Horton, Joe (2021). Partial aggregation in ethics. *Philosophy Compass* 16 (3):1-12.

James, Aaron (2012). Contractualism's (Not So) Slippery Slope. *Legal Theory* 18 (3):263-292.

Kamm, Frances Myrna (1984). Equal treatment and equal chances. *Philosophy and Public Affairs* 14 (2):177-194.

Kamm, Frances Myrna (1993). *Morality, Mortality Volume I*. Oxford University Press.

Kamm, Frances Myrna (2006). *Intricate Ethics*. Oxford University Press.

Korsgaard, Christine M. (1993). The Reasons We Can Share. *Social Philosophy and Policy* 10 (1):24-51.

Korsgaard, Christine M. (2013). The Relational Nature of the Good. In *Oxford Studies in Metaethics*, Volume 8, ed. Russ Shafer-Landau, 1-26. Oxford: Oxford University Press.

Korsgaard, Christine M. (2018) *Fellow Creatures: Our Obligations to the Other Animals*. Oxford University Press.

Kumar, Rahul (2011). Contractualism on the Shoal of Aggregation. *Reasons and Recognition: Essays on the Philosophy of T.M. Scanlon.* Edited by R. Jay Wallace, Rahul Kumar, and Samuel Freeman. Oxford University Press.

Lazar, Seth and Chad Lee-Stronach. (2019). Axiological Absolutism and Risk, *Noûs* 53: 97–113.

Lee, Andrew Y. (2022). Speciesism and sentientism. *Journal of Consciousness Studies*, 29(3-4), 205–228. https://doi.org/10.53765/20512201.29.3.205

Mann, Kirsten (2021). Relevance and Nonbinary Choices. *Ethics* 132 (2):382-413.

Mauss, M. (2001). *The Gift: The Form and Reason for Exchange in Archaic Societies* (2nd ed.). Routledge. https://doi.org/10.4324/9781003572350

McMahan, Jeff. (2021). "Introduction," in *Principles and Persons: The Legacy of Derek Parfit*. Oxford University Press.

Mogensen, Andreas L. (2025). Once more, without feeling. *Philosophy and Phenomenological*

Research, 111, 343–365. https://doi.org/10.1111/phpr.70018

Munoz-Darde, Veronique (2005). The Distribution of Numbers and the Comprehensiveness of Reasons 1. *Proceedings of the Aristotelian Society* 105 (2):207-233.

Nagel, Thomas (1979). Equality. In *Mortal Questions*. Cambridge University Press.

Nebel, Jacob M. (2023). The Sum of Well-Being. *Mind* 132 (528):1074–1104.

Nozick, Robert (1974). *Anarchy, State, and Utopia*. Basic Books.

Rawls, John (1971). *A Theory of Justice*. Harvard University Press.

Raz, Joseph (2003). Numbers, With and Without Contractualism. *Ratio* 16 (4):346–367.

Reibetanz, Sophia (1998). Contractualism and aggregation. *Ethics* 108 (2):296-311.

Rüger, Korbinian (2020). Aggregation with Constraints. *Utilitas* 32 (4):454-471.

Scanlon, Thomas. (2021). Contractualism and Justification. *Reason, Justification, and Contractualism: Themes from Scanlon*. Edited by Markus Stepanians & Michael Frauchiger. De Gruyter.

Scanlon, Thomas. (1998). *What We Owe to Each Other*. Harvard University Press.

Scanlon, T. M. (2021). Contractualism and Justification. In Markus Stepanians & Michael Frauchiger, *Reason, Justification, and Contractualism: Themes from Scanlon*. Berlin, Boston: De Gruyter. pp. 17-44.

Smart, J. J. C. & Williams, Bernard (1973). *Utilitarianism: For and Against*. Cambridge University Press. Edited by Bernard Williams.

Steuwer, Bastian (2021). Aggregation, Balancing, and Respect for the Claims of Individuals. *Utilitas* 33 (1):17-34.

Tadros, Victor (2019). Localized Restricted Aggregation. *Oxford Studies in Political Philosophy*

Volume 5. Edited by David Sobel, Peter Vallentyne and Steven Wall. Oxford University Press.

Taurek, John (1977). Should the Numbers Count? *Philosophy and Public Affairs* 6 (4):293-316.

Taurek, John (2021). Reply to Parfit's 'Innumerate Ethics.' *Principles and Persons: The Legacy of Derek Parfit*. Edited by Jeff McMahan and others. Oxford University Press.

Voorhoeve, Alex (2014). How Should We Aggregate Competing Claims. *Ethics* 125 (1):64-87.

Voorhoeve, Alex (2017). Why One Should Count only Claims with which One Can Sympathize. *Public Health Ethics* 10:2:148-56.

Wallace, R. Jay (2019). *The Moral Nexus*. Princeton University Press.

Zhang, Erik (2024). Individualist Theories and Interpersonal Aggregation. *Ethics* 134 (4):479-511.