




ARTICLES

Partial Aggregation for Prioritarians

Makan Nojournian 

Department of Philosophy, University of Edinburgh, Edinburgh, UK
Email: m.nojournian@ed.ac.uk

Abstract

Prioritarianism is a family of views comparing distributions of well-being. What unites prioritarians is the thought that when deciding whether a distribution is overall better than another, the worse off have priority. There are different ways of making this idea more precise. However, some of these views have extreme aggregative implications and others have extreme anti-aggregative implications. This raises the question: can prioritarians accommodate partial aggregation (aggregating in some but not all cases) and avoid both extremes? In this paper, I explore and focus on a neglected anti-aggregation condition. I identify a family of views I call ‘bounded prioritarianism’ that meet this condition by placing an upper bound on the moral significance of benefits. I argue that anyone sympathetic to partial aggregation ought to opt for a version of bounded prioritarianism.

Keywords: prioritarianism; partial aggregation; distributive ethics

1. Introduction

‘Prioritarianism’ is a family of views comparing distributions of well-being.¹ What unites prioritarians is the thought that when deciding whether a distribution is overall better than another, the worse off have priority. More precisely, prioritarian views can be characterised by the following two axioms:²

¹In this paper, I discuss prioritarianism as an axiological rather than deontic theory; it is about what is better rather than what we ought to do. For consequentialists, this implies a deontic theory as well, but I do not take a stand here on consequentialism. For an accessible introduction to prioritarianism and a survey of the literature, see Arneson (2022). A more specific way of characterising prioritarianism and its contenders is as theories of ‘social welfare’: an ‘overall better-ness’ relation over social states depending on the well-being of individuals in those states. See Adler (2019: chapter 3 and appendix) for a definition of social welfare and the landscape of social welfare theories. Even if your personal ethics is not consequentialist, there are going to be many questions about distributive justice or public policy that will hinge on questions about social welfare in particular or comparing distributions of well-being more generally. See Adler (2019: chapters 5 and 6) for a good introduction to the importance of social welfare theory in public policy. Adler and Norheim (2022) is a more detailed survey of the difference prioritarianism can make to different policy areas.

²For simplicity, I am ignoring further axioms which underpin familiar social welfare theories and are taken for granted in the literature. Since they are assumed by the contenders to prioritarianism as well, they are not distinctive characterising features of prioritarianism. These include: Welfarism, Anonymity, and the Strong and Indifference Pareto axioms. See Adler (2019: chapter 3).

© The Author(s), 2024. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

Pigou-Dalton: Transferring a fixed amount of well-being from a better off person to a less well off person that reduces the gap between their well-being levels makes things overall better.

Separability: The ranking of distributions is independent of the well-being levels of unaffected individuals.

To give examples in this paper, I represent the well-being of individuals with numbers (e.g. 20), the distribution of well-being with vectors, e.g. (20, 50, 10), and ‘overall better than’ with \succ .³ Pigou-Dalton, for example, implies that $(20, 40, 50) \succ (10, 50, 50)$. It introduces a weak constraint for giving priority to the worse off. This helps differentiate prioritarianism from a view like utilitarianism which is concerned with total welfare and does not care about the worse off more than it does about others.

Separability, for example, implies that $(20, 50, 10) \succ (30, 30, 10)$ iff $(20, 50, 100) \succ (30, 30, 100)$. Separability makes precise the idea that the relative position of an individual in a distribution does not matter. This helps differentiate it from egalitarianism which is concerned with reducing inequalities of well-being.⁴ Separability is also attractive because of the deliberational benefits it comes with: we can exclude unaffected individuals from any given comparative problem and simplify the question.

Prioritarian theories diverge on how much priority they give to the worse off. The two main families of views are ‘lexical’ (or ‘discontinuous’) and ‘continuous’ prioritarianism. ‘Leximin’ – inspired by Rawls (1971) – is the most straightforward way to give lexical priority to the worse off:

Leximin: $x \succ y$ iff the worst off individual in x is better off than the worst off individual in y ; and if they are equally well off, then the second worst off individual in x is better off than the second worst off individual in y ; and so on. If the k -th worst off individuals are equally well off for all k , then x and y are equally good.

Another idea is to think about what counts as sufficient well-being. Some people think whether people are above or below a certain threshold of sufficient well-being is morally important. You can think of that threshold as being the threshold of a reasonably good life, or you can think of it as the threshold of a life that is worth living. The sufficiency view is similar to leximin in that a group gets lexical priority, but it is different in that the group that gets lexical priority is those below the sufficiency level. More precisely:

Sufficiency: $x \succ y$ iff those below the sufficiency level in x are better off than those below the sufficiency level in y .⁵

³Strictly speaking, we need to make assumptions about the measurability and comparability of well-being for these numbers to be meaningful. See Sen (2017: chapter A3*) for a full treatment. To be able to engage with the full spectrum of prioritarian views, I assume here that well-being is cardinally measurable and that well-being levels, differences and ratios are both intra and interpersonally comparable. See Adler (2019: chapter 2) for a defence of these assumptions.

⁴See Parfit (1997) for the underlying philosophical differences between prioritarianism and egalitarianism, and arguments suggesting Separability. Arneson (2022: chapter 5) gives an introductory survey of the debate between prioritarians and egalitarians.

⁵Crisp (2003) and Brown (2005) develop different versions of this view. Whether these views will strictly speaking count as prioritarian will depend on our axiomatic definition of ‘prioritarianism’. Although both views satisfy the Separability axiom, only Brown’s sufficientism satisfies the Pigou-Dalton axiom. That said,

Any kind of lexical priority generates a problem which, following Adler (2012: chapter 5), we can call the problem of ‘Absolute Priority’:

Absolute Priority: An arbitrarily small harm or benefit to the worse off can outweigh arbitrarily large benefits to others.

Absolute Priority is an extreme anti-aggregative implication of lexical views which easily leads to counterintuitive verdicts since it is insensitive to the comparative magnitudes of harms and benefits. To give an example, consider (20, 40, 40) with 30 as the threshold for sufficient well-being. You can decide between either giving a very small benefit to the worse off (0.1 units) or a very large benefit to the others (100 units each). Both Leximin and Sufficiency say that you ought to prefer giving a 0.1 benefit to the worst off person (who is also the only person below the sufficiency level) rather than 100 units of well-being to others. What is more, even if you make 0.1 as small as you wish (say 10^{-100}) and make 100 as great as you wish (say 10^{100}), then it is still going to be the case that the small benefit to the worse off is going to outweigh the large benefit to everyone else.

Most prioritarrians adopt the Continuous Priority view which avoids this problem.⁶ Continuous priority is an aggregative view (like utilitarianism) that gives more weight to the well-being of the worse off:

Continuous Priority: $x > y$ iff the sum of well-being – transformed by the ‘prioritarian transformation function’ f – in x is higher than in y ; where f is a ‘prioritarian transformation function’ iff it is a continuous, strictly increasing and strictly concave (down) function of well-being.

The prioritarian transformation function transforms well-being by assigning different weights to different well-being levels. Since it is increasing and concave, it assigns more weight to the worse off. Since it is continuous, any distribution of well-being will receive a ‘score’ – its sum of transformed well-being – which can be used as the basis for ranking it versus others.⁷ Different versions of continuous prioritarianism use different transformation functions capturing different degrees of (continuous) priority for the worse off.⁸

A great advantage of continuous priority views is that they do not face the Absolute Priority problem because of their aggregative feature. On the other hand, they have extreme aggregative implications which lead to counterintuitive verdicts. For example:

Crisp’s sufficientism does satisfy a weaker version of the condition described in Adler (2019: chapter 3) as the Minimal Pigou-Dalton condition. If we take prioritarianism to be characterized by Separability and the Minimal Pigou-Dalton axioms, then both versions of sufficientism will count as prioritarian. Whether we should do so and which version of sufficientism is the representative version does not bear on the problems discussed in this paper.

⁶For examples of authors pressing the Absolute Priority problem, see the references in Adler (2012: 396). See Parfit (1997) for an early exposition and defence of continuous prioritarianism, and Adler (2012) for a more comprehensive book-length defence.

⁷See Adler (2019: chapters 3 and 4) for a discussion of continuity and why it imposes weighted addition.

⁸See §3 for a discussion of the two main families of continuous prioritarian functions (Atkinson and Kolm-Pollak) and illustrations. The square root function is an example of an Atkinson prioritarian function. It implies – contra leximin – that $(10, 50, 50) > (20, 40, 40)$ because $\sqrt{10} + \sqrt{50} + \sqrt{50} > \sqrt{20} + \sqrt{40} + \sqrt{40}$.

Numbers Win: An arbitrarily large harm to the worse off can be outweighed by arbitrarily small benefits to sufficiently many others.⁹

We looked at two types of prioritarianism: lexical and continuous. The lexical views have extreme anti-aggregative implications (Absolute Priority) and the continuous views have extreme aggregative implications (Numbers Win) instead. This raises the question: is there a version of prioritarianism that avoids both extremes? There is a growing literature on ‘partial aggregation’ theories which seek to aggregate in some cases (to rule out Absolute Priority) but not in others (to rule out Numbers Win).¹⁰ So another way of putting the question is: Can prioritarians accommodate partial aggregation?

Sadly the answer is no. Prioritarians face an inescapable dilemma between Absolute Priority and Numbers Win.¹¹ Adler (2012: chapter 5) argues that Absolute Priority is more counterintuitive than Numbers Win and that therefore we ought to pick a continuous priority view and accept the Numbers Win property as an unfortunate consequence of this reflective equilibrium. I agree with Adler that conceding Numbers Win is better than conceding Absolute Priority. However, I think we can do better than this in accommodating anti-aggregative intuitions. We can focus on similar but weaker anti-aggregation constraints and identify the best version of prioritarianism that accommodates them.

In what follows, I explore a neglected anti-aggregation constraint (§2) and a family of views I call ‘bounded prioritarianism’ that succeed in meeting it (§3). I then argue that anyone sympathetic to partial aggregation ought to abandon more familiar versions of prioritarianism in favour of a version of bounded prioritarianism (§4). Bounded prioritarianism is the best version of prioritarianism as far as partial aggregation is concerned.

2. A weaker anti-aggregation condition

Consider the following extreme implication of some aggregative theories:

Large Benefits Win: An arbitrarily large harm to the worse off can be outweighed by a sufficiently large benefit to another group.

An example of the problem is: one person is very badly off – say with well-being 1 – and another person is very well off – say with well-being 100. If Large Benefits Win, any large harm to the worse off person – say of magnitude 100 – can be outweighed by some large benefit to the other person who is already very well off. This seems wrong.

The problem is similar to Numbers Win but slightly different. Numbers Win meant that a large harm can be outweighed by a large benefit which is distributed as a series of small benefits across a very large number of people. Here the large harm is outweighed by a large benefit which is distributed among a fixed number of people instead, and so they are potentially going to each receive large benefits. The difference is simply between the ways you distribute the large benefit, but in both cases, a large harm has been outweighed by a benefit to a group of potentially well off individuals. (See

⁹For a classic example, see Scanlon’s (1998: 235) transmitter room case. I follow Adler (2012: chapter 5) in referring to this as the Numbers Win problem.

¹⁰See Horton (2021) for a survey of the literature on partial aggregation and its challenges. See also §2 below for some motivations.

¹¹See Adler (2012: chapter 5) for a discussion of this dilemma in the context of prioritarianism specifically. Fleurbaey and Tungodden (2010) give conditions that lead to this dilemma for social welfare theories more broadly. See Fleurbaey et al. (2009) for a more informal exposition of the same set of results.

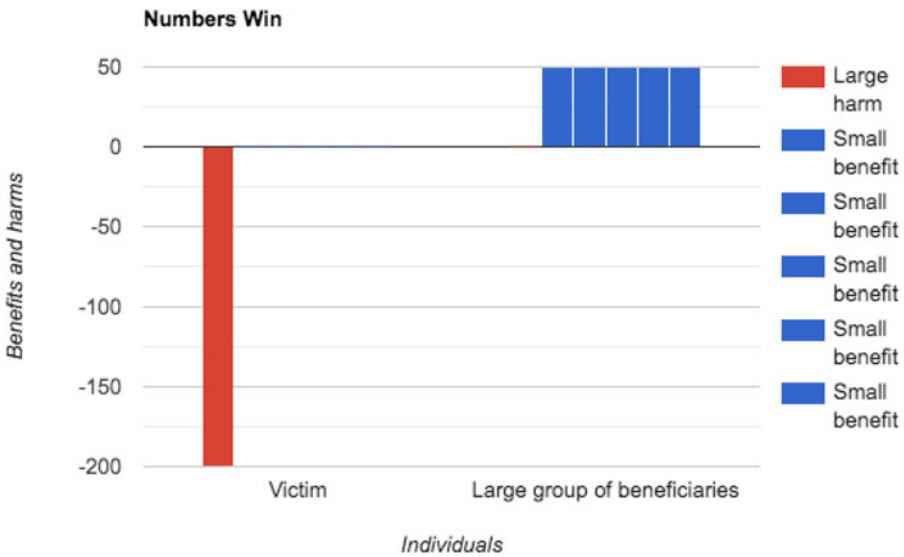


Figure 1. Numbers win.

Figures 1 and 2. These figures only illustrate the different distributions of benefits in the two problems. Whether or not the benefits outweigh the harms in these figures will of course depend on the well-being levels of the individuals and the details of the prioritarian ranking used.)

Why think that Large Benefits Win might be a problem? There are three reasons for this: independent intuitions; the implication of some prioritarian intuitions; and the implication of some key motivations for partial aggregation.

Many people might find specific examples of Large Benefits Win, like the example above, intuitively unacceptable. That said, we could be sceptical about these intuitions, given the large amounts of well-being under consideration.¹² However, interestingly in the context of prioritarianism, aversion to Large Benefits Win simply follows from much simpler low-stakes intuitions. Nebel and Stefánsson (2023) show that prioritarians who give priority to the worse off in some low-stakes cases ought to be averse to what I call Large Benefits Win cases when the stakes are high.

For example, it follows from many prioritarian views that for some level of well-being w , a 0.9% loss to someone at w cannot be outweighed by a 1% gain to another person at the same level. (The benefit is slightly greater than the loss of course, but the worse off receive greater weight in the comparison.) Nebel and Stefánsson show that it follows from this that for any well-being level w and any two groups of identical size, a 10% loss to one group cannot be outweighed by any benefit to the other group, no matter how large the benefit and how large the groups. Notice that this conclusion implies aversion to Large Benefits Win but not to Numbers Win.¹³

¹²Some have raised a similar point about Numbers Win. See §4 for discussion.

¹³See Table 3 in Nebel and Stefánsson (2023) for more examples. Strictly speaking, this result only applies to prioritarian views with ratio-scale invariance. Atkinson prioritarianism satisfies this condition. (See §3 below, and Adler (2012) for a thorough defence of the condition.) See Nebel and Stefánsson

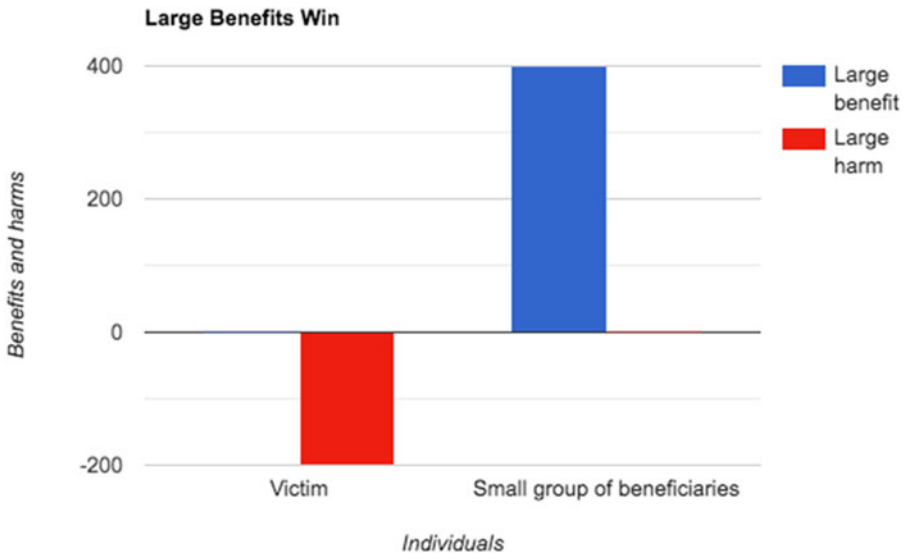


Figure 2. Large benefits win.

Putting intuitions about numbers and specific cases aside, Large Benefits Win should seem problematic to anyone sympathetic to partial aggregation. Many people who are attracted to partial aggregation think that Numbers Win is an unwelcome feature of utilitarianism and continuous prioritarianism, including continuous prioritarians like Adler (2012: chapter 5) who accept it only as a forced result of reflective equilibrium. But many of the justifications for finding Numbers Win problematic and partial aggregation desirable apply to Large Benefits Win as well, since the outweighing of large harms remains justified and the difference is simply about how to distribute the benefits.

For example, depending on what our underlying conception of well-being includes, some large harms could entail the violation of a person's integrity, personal projects or rights, and we might press for partial aggregation on the grounds that thorough aggregation leading to these results is unacceptable.¹⁴ When a large benefit to the well off outweighs an arbitrarily large harm to the worse off as in Large Benefits Win, the harm to the worse off might, for example, entail a violation of their integrity.

Another motivation for partial aggregation is a concern for fairness, either in the sense of having a fair distribution of well-being, or in the sense of ensuring everyone

(2023: Tables 1–2) for examples of their results in the absence of ratio-scale invariance. It is also worth noting that Nebel and Stefánsson use their results for very different dialectical purposes. They present them as a 'calibration dilemma' for prioritarians. I see them, however, as intuitive support against Large Benefits Win and for 'bounded prioritarianism', defined below. (Only bounded prioritarian views yield the low-stakes intuitions Nebel and Stefánsson draw on in the first place, under ratio-scale invariance.)

¹⁴See Williams (1973) on integrity and projects, and relatedly Rawls (1971) on the separateness of persons. Whether integrity or rights can be incorporated into a social welfare theory is contentious. On rights, for example, Nozick (1974) argues for a view that sees them as side-constraints, whether or not they are incorporated into our conception of value as well. Sen (1982) and others on the other hand have developed theories that incorporate rights into the value theory.

has sufficient well-being when possible, both of which can be violated by thoroughly aggregative views.¹⁵ Again, when a large benefit to the well off outweighs an arbitrarily large harm to the worse off as in Large Benefits Win, it can lead to an unfair distribution of well-being or to the worse off going below the sufficiency level.

A recently influential motivation for partial aggregation is that aggregation is only plausible when the harms and benefits are in some sense relevant or close enough, but not otherwise.¹⁶ Again, when a large benefit to the well off outweighs an arbitrarily large harm to the worse off as in Large Benefits Win, the harm might be irrelevant and not close enough to the benefit bestowed on the well off. All of these motivations for partial aggregation, depending on the details, can count against both Numbers Win and Large Benefits Win.

3. Bounded prioritarianism

Bounded prioritarianism is the version of prioritarianism that avoids Large Benefits Win:

Bounded Prioritarianism: $x \succ y$ iff the sum of well-being – transformed by the ‘bounded prioritarian transformation function’ f – in x is higher than in y ; where f is a bounded prioritarian transformation function iff it is a continuous, strictly increasing, strictly concave (down) function of well-being with an upper bound.

Bounded prioritarianism is a species of continuous prioritarianism, with the ‘upper bound’ differentiating between bounded and unbounded versions of continuous priority. Many familiar prioritarian transformation functions fall into the unbounded class, but many functions also fall into the bounded class. To give examples, consider two families of prioritarian functions familiar from the literature. These are the ‘Atkinson’ and ‘Kolm-Pollak’ functions:

Atkinson Prioritarianism: f is an Atkinson prioritarian transformation function of well-being iff $f(w) = \frac{1}{1-\gamma} w^{1-\gamma}$ for $\gamma \neq 1$, and $f(w) = \ln w$ for $\gamma = 1$.

Kolm-Pollak Prioritarianism: f is a Kolm-Pollak prioritarian transformation function of well-being iff $f(w) = -e^{\beta w}$ for $\beta > 0$.

γ and β are the priority parameters of their respective functions: the greater they are, the greater the weight accorded to the worse off. Kolm-Pollak functions are all bounded. However, only Atkinson functions with $\gamma > 1$ are bounded; the rest are unbounded.¹⁷

¹⁵Distributional fairness of course is a main concern for egalitarian social welfare theories. See also Parfit’s (1997) comments on fairness when introducing prioritarianism. See Frankfurt (1987) and Crisp (2003) on the importance of sufficient well-being.

¹⁶See Voorhoeve (2014) for an influential account. Whether or not this will count against Large Benefits Win will of course depend on how the view is fleshed out. See Brown (2020) for a survey of possible ‘close enough’ approaches.

¹⁷See Adler (2022) for an introduction to the two families of functions and their features, and Adler (2012: chapter 5) for arguments in favour of Atkinson prioritarianism. Atkinson functions are only well-defined for non-negative or positive well-being, and exhibit extreme behaviour around zero. Kolm-Pollak functions do not have these problems. On the other hand, only Atkinson functions satisfy ‘ratio-scale invariance’, which says that ratio re-scalings of well-being do not affect the ranking of distributions, implying that only well-being levels, differences and ratios are relevant to the ranking of distributions. For the purposes of this paper, I do not take a stand on ratio-scale invariance and which family of functions is preferable.




Figure 3. Unbounded prioritarian transformation functions.

Figure 3 below illustrates two unbounded Atkinson functions with well-being on the horizontal axis and transformed well-being on the vertical axis: $f(x) = 2\sqrt{x}$ and $g(x) = \ln(x)$. And Figure 4 illustrates a bounded Atkinson function and a Kolm-Pollak (and therefore bounded) function: $h(x) = -x^{-1}$ and $p(x) = -e^{-x}$.

We have seen that some prioritarian functions are bounded and some are unbounded. The bounded functions have the advantage that they do not allow Large Benefits Win. More precisely:

Large Benefits Win and Continuous Priority: If f is a continuous prioritarian transformation function: Large Benefits Win iff f is unbounded.

To see why, assume first that f is bounded with least upper bound c . Compare a distribution of well-being D with any distribution D^* in which D^* benefits a group of beneficiaries at the price of harming a particular victim. The difference between the transformed well-being levels of the beneficiaries from D to D^* cannot be greater than $n(c - w_b)$, where n is the number of beneficiaries and w_b is the well-being level in D of the worst off beneficiary. Since all three variables are fixed by D and f , there is an upper bound on the benefits of D^* relative to D . On the other hand, the difference between the transformed well-being levels of the victim can be arbitrarily large because f is strictly increasing and has no lower bound. Putting these two thoughts together with a continuous prioritarian ranking employing f , it follows that some harms to the victim




Figure 4. Bounded prioritarian transformation functions.

cannot be outweighed by any benefits to the beneficiaries. In other words, if Large Benefits Win, then f is unbounded.

Next, assume that f is unbounded. Consider D and D^* above again. It is still true that the difference between the transformed well-being levels of the victim can be arbitrarily large. But now it is also true that the difference between the transformed well-being levels of the beneficiaries can be arbitrarily large, because f is strictly increasing and has no upper bound. Putting these two thoughts together with a continuous prioritarian ranking employing f , it follows that any harm to the victim can be outweighed by sufficiently large benefits to the beneficiaries. In other words, if f is unbounded, then Large Benefits Win.

It is important to distinguish what I call bounded prioritarianism from other views that employ upper bounds in other ways. Bounded prioritarianism places an upper bound on the continuous prioritarian transformation function of well-being. To counter the extreme implications of aggregation (e.g. Numbers Win), some have suggested placing an upper bound on the total value of any number of small benefits or harms.¹⁸ This strategy can avoid Numbers Win because the total value of any number of small benefits may be less than the (dis)value of some large harms and therefore cannot outweigh it. This is a solution to Numbers Win that is not available to the bounded prioritarian. On the other hand, this strategy comes with its own problems. For

¹⁸See, for example, Carlson's (2000) exploration of the 'Moderate Trade-Off Theory', and more recently Lazar and Lee-Stronach (2019). (See especially their rejection of 'simple additivity', and further references in footnote 22). See Horton (2021) for discussion and objections.

example, it leads to the counterintuitive implication that the value of a harm or benefit to a person can change depending on how many other victims or beneficiaries are around. Since the total value of benefits is bounded, if there are millions of people in danger, saving you from this danger matters less than if there are only hundreds in danger. Bounded prioritarianism as defined above avoids this implication: the value of saving you from danger only depends on your well-being level, the magnitude of the benefit and the underlying prioritarian transformation function.

4. Arguments for bounded prioritarianism

Figure 5 below summarises the prioritarian views we have seen and their implications. In this section, I give three reasons why anyone sympathetic to partial aggregation ought to pick Bounded Continuous Priority.

Firstly, as I argued in §2, Large Benefits Win is an undesirable implication of any prioritarian theory. To remind you, Large Benefits Win is both intuitively problematic and also clashes with some prioritarian-friendly low-stakes intuitions. What is more, many arguments in favour of partial aggregation which rule out Numbers Win also rule out Large Benefits Win. If we take these considerations seriously, we will have to opt for either Lexical Priority or Bounded Continuous Priority. If we find Absolute Priority problematic too, Lexical Priority is no longer an option. Therefore, we ought to pick Bounded Continuous Priority.

Secondly, some believe that Numbers Win is a pseudo-problem and our intuitions about large numbers here cannot be reliable.¹⁹ As I argued in §2, the same cannot be said about Large Benefits Win: the large-number intuitions are about benefits to potentially well off people, and not numbers of people, and what is more, they follow from some small-number intuitions too. This means that we can concede Numbers Win because of the unreliability of our intuitions (or simply to avoid Lexical Priority), and still try to avoid Large Benefits Win. Bounded Continuous Priority does exactly that.

Finally, if we are sympathetic to partial aggregation, we would like to avoid extreme aggregative and anti-aggregative implications. Neither Lexical Priority nor Unbounded Continuous Priority can do this. Bounded Continuous Priority is the only available view that avoids both some extreme aggregative implications (Large Benefits Win) and extreme anti-aggregative implications (Absolute Priority). Therefore Bounded Continuous Priority gives us the best version of prioritarianism as far as partial aggregation is concerned.

5. Conclusion

There is a growing literature on partial aggregation and whether it is possible to avoid the extreme implications of both aggregative and non-aggregative theories comparing distributions of well-being. Partial aggregation is not easy for prioritarrians: lexical views have extreme anti-aggregative implications (e.g. Absolute Priority) and continuous views have extreme aggregative implications (e.g. Numbers Win).

I drew attention to a neglected implication of some aggregative theories which I call Large Benefits Win, and argued that it is both independently problematic and also conflicts with some of the motivations for partial aggregation. I then identified a subset of

¹⁹See, for example, the references in Horton (2021: footnote 5).

	Anti-Aggregative Implications	Aggregative Implications	
	Absolute Priority	Numbers Win	Large Benefits Win
Lexical Priority	Implies	Avoids	Avoids
Unbounded Continuous Priority	Avoids	Implies	Implies
Bounded Continuous Priority	Avoids	Implies	Avoids

Figure 5. Prioritarian views and implications.

continuous prioritarian views which I call bounded prioritarianism, and showed how they can avoid Large Benefits Win. I concluded by arguing that anyone sympathetic to partial aggregation ought to pick bounded continuous prioritarianism because it is the only aggregative view that avoids Large Benefits Win, and because it is the only view that avoids both extreme aggregative and anti-aggregative implications.

Acknowledgements. I am grateful to Matthew Chrisman, Ben Eggleston, Barry Maguire, Wolfgang Schwarz and two anonymous referees for very helpful comments on this paper, and to Campbell Brown, Nadine Dietrich and Declan O’Gara for helpful conversation on the topic. I am also grateful to the audience at the Philosophy PhD Work in Progress Seminar, University of Edinburgh, for their feedback.

References

- Adler, Matthew D. 2012. *Well-Being and Fair Distribution: Beyond Cost-Benefit Analysis* (New York: Oxford University Press).
- Adler, Matthew D. 2019. *Measuring Social Welfare: An Introduction* (New York: Oxford University Press).
- Adler, Matthew D. 2022. ‘Theory of Prioritarianism’, in *Prioritarianism in Practice*, ed. by Matthew D. Adler and Ole F. Norheim (Cambridge, UK: Cambridge University Press), pp. 37–127.
- Adler, Matthew D., and Ole F. Norheim (eds). 2022. *Prioritarianism in Practice* (Cambridge, UK: Cambridge University Press).
- Arneson, Richard J. 2022. *Prioritarianism* (Cambridge, UK: Cambridge University Press).
- Brown, Campbell. 2005. ‘Priority or Sufficiency... or Both?’, *Economics and Philosophy*, 21.2, pp. 199–220.
- Brown, Campbell. 2020. ‘Is Close Enough Good Enough?’, *Economics and Philosophy*, 36.1, pp. 29–59.
- Carlson, Erik. 2000. ‘Aggregating Harms: Should We Kill to Avoid Headaches?’, *Theoria*, 66.3, pp. 246–55.
- Crisp, Roger. 2003. ‘Equality, Priority, and Compassion’, *Ethics*, 113.4, pp. 745–63.
- Fleurbaey, Marc, and Bertil Tungodden. 2010. ‘The Tyranny of Non-Aggregation versus the Tyranny of Aggregation in Social Choices: A Real Dilemma’, *Economic Theory*, 44.3, pp. 399–414.
- Fleurbaey, Marc, Bertil Tungodden, and Peter Vallentyne. 2009. ‘On the Possibility of Nonaggregative Priority for the Worst Off’, *Social Philosophy and Policy*, 26.1, pp. 258–85.
- Frankfurt, Harry. 1987. ‘Equality as a Moral Ideal’, *Ethics*, 98.1, pp. 21–43.
- Horton, Joe. 2021. ‘Partial Aggregation in Ethics’, *Philosophy Compass*, 16.3, pp. 1–12.
- Lazar, Seth, and Chad Lee-Stronach. 2019. ‘Axiological Absolutism and Risk’, *Noûs*, 53.1, pp. 97–113.
- Nebel, Jacob M., and H. Orri Stefánsson. 2023. ‘Calibration Dilemmas in the Ethics of Distribution’, *Economics and Philosophy*, 39.1, pp. 67–98.

- Nozick, Robert.** 1974. *Anarchy, State, and Utopia* (New York: Basic Books).
- Parfit, Derek.** 1997. 'Equality and Priority', *Ratio*, **10.3**, pp. 202–21.
- Rawls, John.** 1971. *A Theory of Justice* (Cambridge, MA: Belknap Press of Harvard University Press).
- Scanlon, Thomas M.** 1998. *What We Owe to Each Other* (Cambridge, MA: Belknap Press of Harvard University Press).
- Sen, Amartya K.** 1982. 'Rights and Agency', *Philosophy and Public Affairs*, **11.1**, pp. 3–39.
- Sen, Amartya K.** 2017. *Collective Choice and Social Welfare: Expanded Edition* (Cambridge, MA: Harvard University Press).
- Voorhoeve, Alex.** 2014. 'How Should We Aggregate Competing Claims', *Ethics*, **125.1**, pp. 64–87.
- Williams, Bernard.** 1973. 'A Critique of Utilitarianism', in *Utilitarianism: For and Against*, ed. by J. J. C. Smart and Bernard Williams (Cambridge, UK: Cambridge University Press), pp. 77–151.