University of Reading

The Aggregation Problem for Scanlonian Contractualism

An Exploration of the Relevance View, Mixed Solutions, and Why Scanlonian Contractualists could be, and perhaps should be, Restricted Prioritarians

Abstract In this thesis, I discuss the aggregation problem for T. M. Scanlon's (1998) "contractualism". I argue that Scanlonian contractualists have the following two options when it comes to the aggregation problem. First, they can choose to limit aggregation directly via a specific version of the Relevance View, "Sequential Claims-Matching". Second, Scanlonian contractualists can adopt a so-called "mixed solution" of which I propose a specific version. My mixed solution does not limit aggregation. Rather, it either avoids some of the (for Scanlon at least) counterintuitive results in certain cases, or at least makes these same results look a lot less counterintuitive by putting them into perspective thanks to a plausible plurality of precisely specified values. The stakes between these two options are as follows. The first option justifies Scanlon's intuition in one vs. many cases of which his "World Cup Case" is one example, but it comes at the cost of a seemingly quite limited range of application: any version of the Relevance View only works in cases in which we have no more than two groups of competing claimants. The second option has the exact opposite implications: my mixed solution cannot justify Scanlon's intuition in one vs. many cases, but it could be applied to any case with any number of groups of competing claimants. In this choice between pre-theoretical intuition and (practical) feasibility, I would urge Scanlonian contractualists to choose feasibility.

Keywords Aggregation • Contractualism • Scanlon

Aart van Gils
PhD (Philosophy)
September 2019

Table of Contents

Declaration

I confirm that this is my own work and the use of all material from other sources has been properly and fully acknowledged.

§5-7 of my third chapter, "Sequential Claims-Matching: A (More) Precise Version of the (Local) Relevance View", are part of a larger 50-50 co-authored article with Patrick Tomlin (forthcoming) "Relevance Rides Again? Aggregation and Local Relevance" in Volume 6 of *Oxford Studies in Political Philosophy*.

Aart van Gils

Acknowledgements

Many people have most kindly helped me in realising my goal of writing a PhD thesis. I am grateful to all. Several people deserve a special mention.

First, my father, Paul, my mother, Ellen, and my brother, Jan. Without their unconditional love and support, I would never have been able to study at university, let alone write a PhD thesis. They have instilled in me the love for philosophy before I had been acquainted with it as a subject. Furthermore, they have inspired me to go all in for what has become my passion: studying philosophy. And studying philosophy has given me the highest level of fulfilment I have found in life so far. For these and countless other invaluable life lessons I am eternally grateful to them.

Second, I am indebted to my supervisors, Brad Hooker and Patrick Tomlin, for their infinite patience, support, and countless constructive suggestions for helping me to continuously improve what has become the final draft of this thesis. Without their help, I could not have done this. Furthermore, I am grateful for all the valuable ways in which they have taught and shown me in leading by example how to be better at philosophy, to be at least as invested in the success of another's projects as one's own, and for having cemented my belief that academic philosophy should be about finding out the truth together (where we can) rather than having the final say. Truly, I could not have asked for more wonderful supervisors.

Third, a special thanks to my great philosophy friends Jamie Draper, Alex McLaughlin, and Josh Wells for all our discussions and their kind support during our time together of writing our theses at Reading.

Fourth, my thanks to all members of the University of Reading philosophy department. The impact of one's working environment on one's work is beyond any reasonable doubt. The philosophy department at Reading truly shines in terms of having a wonderfully collegial atmosphere. The weekly Graduate Research Seminars in which each of us philosophy PhD

students would take turns in presenting work in progress, be the 'official' or appointed respondent for that week's presentation, or contribute to the general discussion have been essential in helping me to become a better student of philosophy in so many ways. Furthermore, I have learned so much on such a variety of topics of which I would have otherwise known next to nothing.

Fifth, I firmly believe that in any aspect of one's (professional) life, to have the right role models makes a watershed difference to one's (career) outlook and fortune. Obviously, my supervisors have both been my two most important role models during my time as a PhD student and they will remain so for the rest of my time that I am fortunate enough to do academic philosophy. However, another one of their supervisees, Julia Mosquera, who started her third year as a PhD student when I started my PhD, has also played an invaluable role in this regard. Julia's professionalism, ambition, and overall approach to her (philosophical) career served as the perfect model for me to aspire to. I am most grateful to her for being an inspiration to me and for having given me many invaluable pieces of advice regarding writing, conferencing, networking, and all that comes with making the most of one's time as a PhD student.

Relatedly, I would have never been ready as a philosophy student to take on the challenge of writing a PhD thesis without the guidance and support from the following four people, each of whom I hold in the highest regard. Alan Thomas, who introduced me to both analytical (political) philosophy and American Football, supervised my BA thesis at Tilburg University, and encouraged me to always work on my (philosophical) weaknesses and never use italics for emphasis. Bart Streumer, who supervised my MA thesis at the University of Groningen, kindly pushed me to be evermore precise in my arguments, and who encouraged me to apply to Reading to work with Brad. Kent Hurtig, who during my time as a visiting student on the MLitt programme at St Andrews/Stirling would use deontic logic to illustrate whether I was right or (more commonly) wrong in my argument about either philosophy or football. And, last but certainly not least, Jonathan Dancy. In my first year as a PhD student, I was fortunate enough to catch Jonathan's 'farewell tour' as a member of staff at Reading, as part of which he organised a reading groups/seminar series on (the then still draft version of) his (2018) *Practical Shape* and an aptly labelled 'masterclass' on three articles that had influenced him most and he thought were must-reads for anyone studying philosophy.[1] Although I love his writing, it is specifically Jonathan's ability to get to the heart of the matter as a participant to any philosophy discussion – to

[1] For those curious, these three articles are Christine M. Korsgaard's "Realism and Constructivism in Twentieth-Century Moral Philosophy", John McDowell's "Values and Secondary Qualities", and David Wiggins's "Truth, Invention, and the Meaning of Life"

philosophize ‘on the go’, if you will – which has made a lasting impression on me and which has inspired me ever since.

Finally, without the generous funding from the University of Reading and the Aristotelian Society, I would not have been able to start and finish this thesis. My thanks to all members of each board for having faith in me and awarding me with a research stipend.

University of Reading

Scanlonian Contractualism: What it is and Why it's all about Aggregation

Abstract T. M. Scanlon's contractualism as defended in his (1998) *What We Owe to Each Other* has been one of the most important contributions to moral philosophy. That is not to say that it is without its shortcomings or has not given rise to a whole host of objections to each of its several distinctive features, such as its "reasons primitivism", its "buck-passing" account of goodness, or its "trustee" account regarding non-human animal interests. But another central problem for Scanlon's contractualism is the focus of this thesis: the problem of interpersonal aggregation. In this first chapter, I introduce Scanlon's contractualism, focusing on the core elements that, together, give rise to the aggregation problem. I will start discussing aggregation and the problem it presents for Scanlon's contractualism in chapter 2.

Key words Aggregation • Contractualism • Scanlon

Introduction

Mark Timmons (2004) describes the contrast between "formal contractualism" and "substantive contractualism" as follows:

> Although contractualism in ethics is associated with non-consequentialist moral theories, a consequentialist might propose a contractualist justification of her consequentialist moral principles and a contractualist account of what reasons there are to comply with a moral principle. The contractualist ideal of agreement does not favour one normative moral theory over another. Let us call this kind of contractualism 'formal contractualism'. However, Scanlon and others often take contractualism to be opposed to certain normative moral theories, especially versions of consequentialism. So what we may call 'substantive contractualism' represents one way of developing formal contractualism according to which a moral principle is justified only if *each individual* (who is part of the relevant justifying agreement) either does or would agree to the principle in question. This kind of 'distributive justification' is opposed to the kind of aggregative or collective justification characteristic of versions of consequentialism. (Timmons, 2004, p. 93).[1]

[1] This is why I understand both Stephen Darwall's (2006) 'second-personal' contractualism and Nicholas Southwood's (2010) "deliberative" contractualism as different versions of formal contractualism. Darwall does not offer an explicit definition of 'second-personality'. However, we can get an idea of what he understands by 'second-personality' by looking at requirements and consequences he identifies as following from this notion. Darwall states that "second-personality (…) requires empathy or the capacity to put oneself in another's shoes." (2006, p. 44). According to Darwall, there are four "interdefinable, irreducibly second-personal notions" under the heading of 'second-personality' (p. 246). These four notions are: (1) addressable demands; (2) the (equal) authority to address these demands; (3) second-personal reasons that are implicit in their address; and (4) accountability to others (2006, p. 246). Darwall comments on these four second-personal notions that "any one of these four notions implicitly involves the other three, and there is no way to reduce the distinctive normativity of any one of them to norms and values that can be appreciated fully without the second-person standpoint." (2006, p. 246). Given Scanlon's aim of having his substantive contractualism as an alternative to utilitarianism, and because both Darwall's and Southwood's version of formal contractualism by definition cannot rule out utilitarianism or any other first-order moral theory, I will not further discuss these two accounts. Regarding John Rawls's (1999) 'political' contractualism, John Harsanyi (1975), in my view rightfully, objects to Rawls that the parties in the Original Position should at the very least be permitted to choose average utilitarianism rather than maximin. The fact that Rawls does not allow the parties to pick average utilitarianism from behind the veil of ignorance shows that Rawls was wrong to at least give the impression that his account was a formal contractualism. Jonathan Quong (2017) states that "Contractualism is a method or tool, not a first-order moral theory." (2017, p. 65). In light of the distinction between "formal" and "substantive" contractualism, I read Quong as claiming that the only 'proper' kind of contractualism is "formal" contractualism. If that is a correct interpretation of Quong, then I think this is too strong a statement by Quong. Contractualists are permitted to take a (substantive) stance on a (substantive) moral question without being required to abandon their position as "contractualists".

Although (the discussion of) formal contractualism certainly has plenty to offer, in this thesis, I focus on the aggregation problem for T. M. Scanlon's (1998) "substantive" contractualism.[2] From here on out until the end of my thesis, when I talk of "aggregation" I refer to "interpersonal aggregation" only unless specified otherwise. Interpersonal aggregation I understand as follows:

> Interpersonal Aggregation: The benefits (or burdens) that come to different individuals may be taken together (i.e. they may be 'aggregated') and so considered as one total benefit (or burden).[3]

In this first chapter, I will only briefly indicate what the aggregation problem is. Chapter 2 of this thesis marks the start of the discussion of the aggregation problem.

The main aim of this chapter is to provide the wider context of Scanlon's contractualism – what it is and why it faces the aggregation problem. For now, we simply need take note of the stakes involved in this discussion. As Joseph Raz (2004) puts it: "The success of [Scanlonian] Contractualism depends on the success of its treatment of aggregation." (2004, p. 59). The importance of Scanlon's treatment of aggregation for his contractualism is underscored by the fact that "Scanlon's interest in aggregation was the motive for his initial interest in contractualism." (Raz, 2004, p. 60).[4]

The first step is to acknowledge that Scanlonian contractualism faces many problems regarding several of its distinctive features. To name a few problems, there are redundancy, indeterminacy,

---

[2] For a neat bullet-point overview by the man himself of the main claims of his book, see Scanlon (2003). For three excellent overviews of Scanlon's contractualism by other authors, see Matt Matravers (2003a), Thomas Nagel (1999), and especially Philip Stratton-Lake (2004).

[3] Interpersonal aggregation should be contrasted with "intrapersonal" aggregation, which is the aggregation of different benefits or burdens that come to one and the same individual. Intrapersonal aggregation is used to determine the quality of an individual's life 'as a whole' and then compare this to the quality of another individual's life 'as a whole'. Scanlon accepts intrapersonal aggregation (1998, pp. 237–238). For a detailed discussion of intrapersonal aggregation, see Larry Temkin (2009; 2012, Ch. 4, pp. 96-128). John Broome (2004) provides a convincing argument against 'temporal discounting' of well-being – the idea that well-being that comes earlier in time would be of greater value than well-being that comes later (Broome, 2004, pp. 70–71). Temporal discounting is a common shortcoming of people's rationality: it seems of far greater value to buy a new car now than to save for your retirement.

[4] Scanlon, in earlier work, comments in a footnote on the importance of finding "the solution to the problem of how aggregative considerations and considerations of relative urgency can be combined in a systematic way, i.e. whether and when less urgent interests of many people can outweigh more urgent interests of a few. This is an important problem which must be faced in any event but which I have not attempted to deal with here." (Scanlon, 1975, p. 667, fn. 8). For further context, formulated as such, Scanlon's earlier discussion closely resembles F. M. Kamm's (1998) discussion in her chapter 13 (1998, pp. 268–283). But by the time of *What We Owe to Each Other*, Scanlon appears to have changed his view as being more distant from Kamm. In *WWOtEO*, "urgency" or prioritarian concerns play a most minimal role, as I will discuss in detail in chapter 3 of this thesis.

ex ante vs. ex post, meta-ethical, and non-identity/future generations worries or objections too.[5] Although each of these must be addressed by Scanlonian contractualists, strictly speaking, they are irrelevant to the question my thesis addresses.[6] My thesis question is whether Scanlonian contractualists could come up with a plausible position regarding aggregation. My thesis shows that there is indeed a plausible contractualist position regarding aggregation.

Nevertheless, contractualism might be torpedoed by (at least) one of the other objections. Equally, even if all the other objections could be convincingly answered, contractualism would not be acceptable if it cannot plausibly deal with aggregation. Hence, I focus on the aggregation problem for Scanlon's contractualism in this thesis. Let me now start to set the stage.

§1. Scanlon's Contractualism: its Kind and its Distinctive Features

The canonical statement of Scanlon's contractualism is as follows:

> Contractualism (…) holds that an act is wrong if its performance under the circumstances would be disallowed by any set of principles for the general regulation of behavior that no one could reasonably reject as a basis for informed, unforced general agreement. (Scanlon, 1998, p. 153).

The best way of understanding this central statement of Scanlon's contractualism is by looking at each of its constitutive features mentioned, such as "principles" and "reasonable rejection", in turn.[7] But before I will discuss these constitutive features of contractualism, let me briefly reflect

[5] For an excellent overview of what (Scanlonian) contractualism is and the most important objections raised to it, see Elizabeth Ashford and Tim Mulgan (2018). For exemplary discussions of redundancy and indeterminacy as relating to Scanlon's contractualism, see Brad Hooker (2003) and Mark Timmons (2004). For exemplary discussions of the ex ante vs. ex post distinction – mostly as applied in the context of risk – as relating to Scanlon's contractualism, see Johan Frick (2015), Barbara Fried (2012), Sune Holm (2018), Joe Horton (2017), Aaron James (2012), James Lenman (2008), Sophia Moreau (1998), and Korbinian Rüger (2018). For Scanlon's focused discussion of his meta-ethical view, see Scanlon (2014). For a critical discussion of Scanlon's meta-ethics, see Sarah Marshall (2003). For the non-identity problem, see Derek Parfit (1992, pp. 351–379). For Parfit's more recent discussion why Scanlonian contractualists, being advocates of a "person-affecting" view, have a hard time answering the non-identity problem and so to plausibly account for our obligations towards future generations, see (Parfit, 2011b, pp. 213–243).

[6] Apart from Scanlon himself, Rahul Kumar (1999; 2001; 2003a; 2003b; 2009; 2015; 2017) has been one of the most prominent defenders of Scanlonian contractualism, having argued for a contractualist response to several of the objections raised by the authors mentioned in the previous footnote 5. For a recently defended contractualist account of risk, see John Oberdiek (2017).

[7] Here I would like to especially thank each of my Part 3 students whom I had the pleasure of teaching a module on "Contemporary Consequentialism and Contractualism" during the Autumn term of 2018 at the University of Reading. Their excellent questions on and objections to Scanlon's contractualism have helped me tremendously in further improving my understanding of Scanlonian contractualism.

on the bigger picture by reflecting on what kind of contractualism Scanlon proposes by using a distinction between three kinds of contractualism drawn by Jonathan Quong (2017).

Here are the first two kinds, "consent contractualism" and "rationality contractualism":

> Consent contractualism: [Applies to cases in which] it is presumptively wrong to do something, ϕ, to a person or group of persons unless those people freely consent to it. But in circumstances C, there seem to be good reasons to ϕ, and yet obtaining the actual consent of the relevant people is not possible. When faced with a problem of this kind, asking whether suitably situated contractors would consent might be a helpful way of deciding whether it is permissible to ϕ. (Quong, 2017, p. 67).[8]

> Rationality contractualism: [A]ll forms of contractualism which are designed to consider what people would accept if they were rational, either in a narrow self-interested sense, or in some broader sense[.] (Quong, 2017, p. 69).[9]

Having presented these first two kinds of contractualism, the point is to leave these two kinds of contractualism to the side for the remainder of the thesis. My focus is on the third kind of contractualism as so identified by Quong, "fairness contractualism", of which he does not provide a strict definition but rather explains it using the two most famous versions of fairness contractualism:

---

[8] Quong's "it" in the first sentence refers to an act. But Scanlon's contractualism, as shown in the canonical statement above, applies the 'reasonably rejectable' test to sets of principles, not to acts. This quotation from Quong construes Scanlon as an act-contractualist, whereas Scanlon's own canonical statement sounds much more rule-contractualist. I will discuss the act-contractualism vs. rule-contractualism distinction in §2.

[9] Most versions of rationality contractualism are more often referred to as 'contractarianism'. The original contractarian account has been provided by Thomas Hobbes (2009). The most important contemporary contractarian account is defended by David Gauthier (1986). For an excellent overview of the distinction between 'contractualism' and 'contractarianism', their origins, and comparative advantages, see Stephen Darwall (2003). However, there are 'non-contractarian' versions of rationality contractualism. The primary example is Southwood's (2010) "deliberative contractualism". Southwood's version of rationality contractualism faces Jussi Suikkanen's (2014) "conditional fallacy". Southwood (2018) responds by differentiating, following Michael Smith (1995), between an "example model" and an "advice model" of his deliberative contractualism (Southwood, 2018, pp. 9–10). What seems to me both a natural fit and a, as far as I am aware, fully neglected area of research is to combine (either version of) Southwood's deliberative contractualism with Jonathan Wolff's (2003) discussion of contractualism and the virtues. Specifically, looking at Southwood's third kind of "deliberative norms", "reflexive norms", his advice model which would show us how we come up short in the situation as compared to our idealised counterparts, and Wolff's discussion of the contractualist-compatible virtues as to how we might better ourselves (i.e. how to narrow the gap between ourselves and our idealised counterparts) seems a promising area for fruitful philosophical investigation. Furthermore, bringing in the idea of virtues and self-improvement offers an alternative response to "the concessional fallacy" (Southwood, 2018, pp. 10–11). Similarly, for Scanlon's contractualism I think that as part of "what we owe to each other" we owe it to others to continuously better ourselves and so talk of virtues is most appropriate.

> Of course, everyone in philosophy is familiar with the most famous version of this model: Rawls's use of the original position and the veil of ignorance as a method for determining the best principles of justice to regulate the distribution of the burdens and benefits of social cooperation amongst citizens (Rawls, 1999, pp. 102–168). We can call this *fairness contractualism*. We can also envisage a more general version of contractualism, where the contractual device is used to determine what would be morally right or permissible in a given context, and not merely what would be fair (Scanlon, 1998). Although such versions of contractualism are concerned with more than fairness, they can without too much distortion be included as variants of fairness contractualism since they share the same basic structure: the contractual device is designed to reflect some of our existing moral judgements in order to help us answer further moral questions about what is fair or right. (Quong, 2017, pp. 67–68).[10]

Although Quong does not provide a strict definition of fairness contractualism, he comes close to providing a formal definition by pointing to the aim of fairness contractualism:

> In fairness contractualism, the aim is to situate contractors in such a way that the choice of a rule or a set of rules, R, to govern a situation gives us good reason that R is fair, just, or morally right (I will focus on fairness, for simplicity). (Quong, 2017, p. 74).

Here it is important to pause for a moment because Scanlon does not leave it open whether we understand "R" as either "a rule" or "a set of rules". He argues for the evaluation of one principle at a time whilst holding others fixed – where Scanlon understands principles as "general

[10] Adam Hosein (2013) argues that Scanlon cannot use the Rawlsian central justification of cooperative fairness to social justice to defend a contractualist theory of rightness. However, I think that Hosein's objection to Scanlon fails because of straw-manning: for the objection to stick Hosein has made the not-so-innocuous dual revision to Scanlon's contractualism that (1) reasonable rejection could equally be explained in terms of reasonable acceptance and that (2) the contractors are motivated by self-interest (Hosein, 2013, p. 501). The impact or importance of the first revision is debatable, although Scanlon himself argues there is a difference: "What is basic to contractualism as I understand it is the idea of justifiability to each person (on grounds that he or she could not reasonably reject). Unanimous acceptance is a consequence of this condition's being fulfilled, but it is not itself the basic idea." (Scanlon, 1998, p. 390, fn. 8). However, the second revision is significant: we are now no longer talking about Scanlonian contractualism. The psychological profile of Scanlonian contractors as being motivated to find and act in accordance with principles that no one could reasonably reject is necessary for Scanlonian contractualism to be or remain "Scanlonian". Fortunately, given my focus on aggregation, Hosein's objection to Scanlon's contractualism and my criticism of Hosein are irrelevant because of the following caveat Hosein makes: "Sometimes individuals must distribute benefits among others. And perhaps the idea of justifiability to each (along with associated ideas of reasonable complaints and so on) can help us figure out what those individual[s] should do to treat the recipients fairly." (Hosein, 2013, p. 507).

conclusions about the status of various kinds of reasons for action. So understood, principles may rule out some actions by ruling out the reasons on which they would be based, but they also leave wide room for interpretation and judgment." (Scanlon, 1998, p. 199):[11]

> [I]t is misleading to suggest that when we are assessing the "reasonable rejectability" of a principle we must, or even can, set aside assumptions about other rights and entitlements altogether. (…) [A] sensible contractualism, like most other plausible views, will involve a holism about moral justification: *in assessing one principle we must hold many others fixed*. This does not mean that these other principles are beyond question, but just that they are not being questioned at the moment. (Scanlon, 1998, p. 214). [*my italics*].

I will come back to the question of whether the contractualist test should be applied to whole sets of rules or each rule while holding others fixed in a moment. But, first, going back to Quong's description of the aim of fairness contractualism, Scanlon does not leave it open whether "R" stands for "a rule" or "a set of rules". Given the holistic approach Scanlon advocates for his contractualism, Scanlon appears to say that "R" primarily stands for "a rule" and only secondarily for "a set of rules". However, Brad Hooker (2003) presents a persuasive argument as to why Scanlonian contractualists should instead commit themselves to always evaluating the set of rules as a whole:

> Moral principle A might be best given that moral principles B, C, D, E and F are fixed, B best given that A, C, D, E and F are fixed, C best given that A, B, D, E and F are fixed, and so on. In other words, each principle in the set can look superior to every possible alternative to that principle, when the principle and the alternatives are judged in a way that takes the set's other principles to be fixed. So we might think that no one principle in this set should be changed unless others are changed as well. But this hardly establishes the optimality or justifiability of the whole set. For all we have said, the package of moral principles A, B, C, D, E and F might be far inferior to a package of alternative moral principles, such as G, H, I, J, K and L. (Hooker, 2003, p. 63).

[11] I am using "rules" and "principles" interchangeably here, but more needs to be said on their apparent difference in Scanlon's contractualism. I will leave this for my next and final §2 of this chapter, in which I will discuss Hanoch Sheinman's (2011) distinction between act-contractualism and principle-contractualism (i.e. rule-contractualism).

However, evaluating the set of rules might not be possible on Scanlonian contractualism because Scanlon explicitly raises and answers the following question: "How many valid principles are there, then? An indefinite number, I would say." (Scanlon, 1998, p. 201). As I read "indefinite" here, it suggests that there is an "infinite" number of principles.[12] But, as Hooker notes as well, if there is an infinite number of principles, then we cannot evaluate the set of principles as a whole (Hooker, 2003, p. 65). Hooker presents the choice for Scanlonian contractualists as follows:

> [T]he contractualist test [of reasonable rejection] cannot be run on an infinite set of principles. If contractualism must hold that it endorses an infinite set of principles, then contractualism must evaluate principles individually, rather than as sets. And I suspect that another possible diagnosis of why Scanlon favours evaluating principles individually rather than as sets is that he thinks contractualism must endorse an infinite set of principles. But obviously, rather than accepting Scanlon's conclusion that principles are to be evaluated individually rather than as sets, we might reject the view that contractualism must endorse an infinite set of principles. (Hooker, 2003, p. 66).

I argue that Scanlon should revise his contractualism by rejecting the endorsement, if Scanlon intends this endorsement, that there is an infinite set of principles. The set of principles must always be evaluated (as well). Because they defend a version of fairness contractualism as defined by Quong, Scanlonian contractualists therefore cannot be ambivalent between the evaluation of one rule at a time or the entire set of rules altogether: R must always (also) represent the set of rules, and never (just) one rule at a time.

Having discussed the point on how the evaluation of principles ought to proceed in Scanlon's contractualism, I return to Quong's discussion on the contractualist methodology. For all three kinds of contractualism, Quong identifies five key questions that will help design the form of contractualism to be used. I will use these five questions to further illustrate Scanlon's contractualism. Quong's five questions for designing any version of contractualism are as follows:

i. Constituency: Whom do the contractors represent?
ii. Primary Question: What question do the contractors confront?

[12] Hooker considers the possible interpretation of Scanlon's contractualism as Scanlon denying that within contractualism there is a finite number of principles (Hooker, 2003, pp. 65–66). However, Hooker considers this on the basis of Scanlon's remark that "There is no *fixed list* of 'morally relevant considerations' or of reasons that are 'morally excluded'." (Scanlon, 1998, p. 157) [*my italics*]. In my view, this is circumstantial evidence at best for the interpretation that Scanlon might advocate an infinite number of principles.

iii. Motives and Interests: What are the motives and interests of the contractors?
iv. Information: How much information should be provided to the contractors?
v. Non-agreement: What will happen to each contractor in the event of non-agreement? (Quong, 2017, p. 70).

Regarding the first question of constituency, Quong notes the following:

> There is nothing in the logic of contractualism itself that dictates the constituency of persons who ought to be represented within a model of fairness contractualism. The appropriate constituency for any fairness contractualism must be determined by independent ideas regarding who has the standing to be wronged or treated unfairly by some proposed rules regulating conduct or the distribution of advantages. (Quong, 2017, p. 74).

For Scanlon's contractualism, the answer to this first question comes in two steps. In the first step, we should note that Scanlon makes it clear that his contractualism is only supposed to apply to a subdomain of morality. That subdomain is the domain of interpersonal morality, or as Scanlon calls it: "what we owe to each other". Scanlon directly addresses the scope of contractualist morality by raising the following question: "if the account that contractualism offers is correct, what does this imply about the range of creatures toward whom one can behave wrongly in this central sense?" (Scanlon, 1998, p. 178). Scanlon answers the question by considering toward which beings we might act wrongly:

(1) Those beings that have a good; that is, those for which things can go better or worse
(2) Those beings in group (1) who are conscious, and capable of feeling pain
(3) Those beings in group (2) who are capable of judging things as better or worse and, more generally, capable of holding judgment-sensitive attitudes.[13]
(4) Those beings in group (3) who are capable of making the particular kind of judgments involved in moral reasoning

[13] "Judgment-sensitive attitudes" is a technical term invented by Scanlon to denote "attitudes that an ideally rational person would come to have whenever that person judged there to be sufficient reasons for them and that would, in an ideally rational person, 'extinguish' when that person judged them not to be supported by reasons of the appropriate kind." (Scanlon, 1998, p. 20).

(5) Those beings in group (4) with whom it is advantageous for us to enter into a system of mutual restraint and cooperation. (Scanlon, 1998, p. 179).

To work out where the dividing line is between the kind of beings toward whom we can act wrongly in the contractualist sense and those beings who cannot be wronged in this sense, we start at the top of the hierarchy – group (5) – and work our way down. I will not rehearse Scanlon's arguments for why a group is included or not.[14] My aim is merely to show where Scanlon draws the dividing line, so that I have in turn shown how to understand the scope of contractualism on this point.

Scanlon says that only and all members of group (3) are included and he explains why he does so as follows:

> Although it is morally objectionable, in the broad sense, to fail to take account of the pain and distress of nonrational creatures, we do not have the reason we have in the case of rational creatures to accept the general requirement that your conduct be justifiable to them. (Scanlon, 1998, p. 184).

The plausibility of where Scanlon draws the line between the kind of creatures whom we can wrong from those we cannot may certainly, not just as such but even within the parameters of his own contractualism, be questioned. However, I will not do so here given the purposes of my thesis. All that matters here is that I have presented how and where Scanlon draws the boundaries of the subdomain of morality to which his contractualism and its 'non-reasonably-rejectable' principles apply.

But this was only the first step in answering Quong's first methodological question of constituency. The second step is to ask whom the contractors represent in Scanlonian contractualism within the domain of "what we owe to each other". The answer? Everyone. Most importantly, Scanlon argues that for his contractualist principles we should not just be concerned with those directly affected by the principles in question, but both those who might be affected on a future occasion and even those of whom we would somehow know that they would never be directly affected by this particular principle (Scanlon, 1998, pp. 202–203). Why? Scanlon argues as follows:

[14] For Scanlon's further arguments as to how and why he includes and excludes the different classes of beings as he does, see Scanlon (1998, pp. 179–187).

> First and most obviously, widespread performance of acts of a given kind can have very different effects from isolated individual instances. Slightly less obviously, perhaps, the general authorization or prohibition of a class of actions can have significance that goes beyond the consequences of the actions that are performed or not performed as a result. This can be seen both from the point of view of agents and from that of the people who may be affected by these actions. As agents, if we know that we must stand ready to perform actions of a certain kind should they be required, or that we cannot count on being able to perform acts of another kind should we want to, because they are forbidden, these things have important effects on our planning and on the organization of our lives whether or not any occasions of the relevant sort ever actually present themselves. If, for example, I lived in a desert area and were obligated to provide food for strangers in need who came by my house, then I would have to take account of this possibility in my shopping and consumption, whether or not anyone ever asked me for this kind of help[.] (…) The same is true from the point of view of those affected by actions. Our need for privacy, for example, is not met simply because, as a matter of fact, other people do not listen in on phone calls and go through our personal files. In order to have the benefits of privacy we need to have assurance that this will not happen, and this is something that general acceptance of a principle can provide. (Scanlon, 1998, p. 203).[15]

In an admittedly vague summarising formulation, Scanlon's justification for his answer that everyone, and not just those directly affected (on this occasion) by the principle in question, should be part of the constituency seems to be that having "everyone" included within the constituency helps determine the shape of our lives.[16] In a less vague formulation, everyone should be included in the constituency because we all have plans and should be able to make plans. A smoother

[15] Although on the basis of Rawlsian contractualism, rather than Scanlonian contractualism, Tommie Shelby (2007; 2016) — in the spirit of thinking about on whom the ghetto poor can reasonably count and what (natural) justice (still) requires of them — discusses how the content of individual duties of justice are different for those who, like the ghetto poor, find themselves living in unjust conditions due to systematic neglect by the state and its institutions. I cannot go into discussing Shelby's account here, but I have mentioned it because of the attention I believe it deserves and to have an example of a non-idealised version of fairness contractualism as applied to a contemporary relevant issue for a particular society, namely that of the USA. Scanlon's (2018) emphasis on the right background conditions or institutions in the justification of (some degree of) inequality makes his account of inequality much closer to Rawls's, although differences remain.

[16] For an insightful defence of the idea that a different narrative structure to the same amount of total well-being in a life can make a life better or worse, see David Velleman's (2010) Chapter 3, "Well-Being and Time" (pp. 56–84). For arguably the most extensive and sophisticated discussion of time and narrativity, see Paul Ricoeur (2009a; 2009b; 2008).

organisation, integration, and/or alignment of all individuals' plans and so the organisation of community and society would also be in the interest of everyone.

In Scanlon's contractualism, the second and third key questions – the "primary question" of what question the contractors confront and the question of what the motives and interests are of the contractors – are tightly linked. Exhibit A is the canonical statement of Scanlon's contractualism, which, as a reminder, is as follows:

> Contractualism (…) holds that an act is wrong if its performance under the circumstances would be disallowed by any set of principles for the general regulation of behavior that no one could reasonably reject as a basis for informed, unforced general agreement. (Scanlon, 1998, p. 153).

So, the contractors are concerned with the primary question of finding a set of principles for the general regulation of behaviour that no one could reasonably reject. But why are Scanlonian contractors motivated to find these non-reasonably-rejectable principles?

> The contractualist ideal of acting in accord with principles that others (similarly motivated) could not reasonably reject is meant to characterize the relation with others the value and appeal of which underlies our reasons to do what morality requires. This relation, much less personal than friendship, might be called a relation of *mutual recognition*. Standing in this relation to others is appealing in itself – worth seeking for its own sake. A moral person will refrain from lying to others, cheating, harming, or exploiting them, "because these things are wrong." But for such a person these requirements are not just formal imperatives; they are aspects of the positive value of a way of living with others. (Scanlon, 1998, p. 162).[17] [*my italics*].

So, there are two points to note regarding the contractors' motives and interests. First, to repeat, the contractors are motivated to find non-rejectable principles. In other words, each is motivated to ensure that her conduct is justifiable to each who is or could be affected by her choices or actions. Second, there is the further question as to why each person would be motivated in this

[17] This could be taken as Scanlon's reply to Glaucon's "Ring of Gyges" story presented in Book II of *The Republic* which poses arguably the ultimate question for moral philosophers: why be moral if you could perfectly and always get away with being immoral? (Plato & Jowett, 1892, pp. 108–110). In later work, Scanlon (2012) labels this as "the question of *practical significance*: why we should regard moral demands as ones we have strong reason to accept as guides to conduct." (Scanlon, 2012, p. 226).

way. In other words, why would or should any single one of the contractors care about 'mutual justifiability'? They should and do care because of (the promise of) the positively valuable relationship of "mutual recognition" that acting in accordance with principles that no one could reasonably reject brings.

To make Scanlonian contractors' motivation clearer, let me compare their motivation with the motivation of the contractors on Rawls's contractualism.[18] Rawlsian contractors, instead of being motivated by "reasonableness", are "mutually disinterested" (in each other's fates) and rational "in the narrow sense, standard in economic theory, of taking the most effective means to given ends." (Rawls, 1999, p. 12).[19] The way that Rawls ensures that the contractors in his "Original Position" come up with reasonable results despite being motivated 'merely' by mutual disinterest and rational means-ends maximisation is via his answer to the fourth question in Quong's model.

This fourth question, how much information the contractors should have, marks the starkest contrast between Rawlsian contractualism and Scanlonian contractualism.[20] Whereas Rawls strips his contractors of any specific information regarding their plan of life, sex, race, or social or economic position in society by placing the contractors behind a (thick) "veil of ignorance", Scanlon not only permits but encourages contractors to make use of the general information from the different standpoints – "the [individual] points of view that the question of reasonable rejection requires us to take into account" – that is available (Scanlon, 1998, p. 203).[21] There is an emphasis on "general" in the phrase "general information" in the previous sentence though. Like Rawls, Scanlon requires contractors to ignore the aims or life-plans that are particular to the individuals (directly) affected. However, unlike Rawls, Scanlon does not artificially impose a veil of ignorance but rather makes use of the 'natural' veil of ignorance:

> Since we cannot know, when we are making this assessment, which particular individuals will be affected by it in which ways (who will be affected as an agent

[18] I am especially grateful to Josh Wells for his many insightful explanations of Rawls's contractualism to me.
[19] For an excellent critical discussion of the motivation of Rawlsian contractors and the Original Position, see Samuel Freeman (2007, pp. 142–166, 253-262).
[20] Scanlon thinks that not having the veil of ignorance, or "only a different and more limited kind", also has the consequence that the contrast between contractualism and utilitarianism "stand[s] out more clearly" on Scanlon's contractualism. (Scanlon, 2003a, p. 125).
[21] For Scanlon's full discussion of standpoints, see Scanlon (1998, pp. 202–206). Peter Timmerman (2015) argues that Scanlon's contractualism, better than both Kantianism and utilitarianism, accounts for the relevance of perspective-taking to moral judgment. I am indebted to Peter for having introduced me to Scanlon's contractualism for his MA tutorial on Scanlon's (1998) *What We Owe to Each Other* in the Fall semester of 2015 at the University of Groningen.

> required to act a certain way, who as a potential victim, who as a bystander, and so on), our assessment cannot be based on the particular aims, preferences, and other characteristics of specific individuals.[22] We must rely instead on commonly available information about what people have reason to want. I will refer to this as information about generic reasons. (…) Generic reasons are reasons that we can see people have in virtue of their situation, characterized in general terms, and such things as their aims and capabilities and conditions in which they are placed. (Scanlon, 1998, p. 204).

Rawls's contractualism and Scanlon's contractualism will often converge on what justice or "what we owe to each other" requires. Take, for example, a racist or sexist principle. On Rawls's contractualism, such a principle will not be accepted by the contractors because once the veil of ignorance is lifted, they might find themselves a member of the discriminated and/or oppressed group.[23] On Scanlon's contractualism, such a sexist or racist principle could be reasonably rejected because there will most likely be an alternative principle which, if accepted, would impose smaller burdens on everyone.

The last paragraph shows that in both forms of contractualism what we morally ought to do is always determined based on a comparative question. Scanlon makes this "comparative nature of the question of reasonable rejection" especially clear for his contractualism:

> According to contractualism, in order to decide whether it would be wrong to do X in circumstances C, we should consider possible principles governing how one may act in such situations, and ask whether any principle that permitted one to do X in those circumstances could, for that reason, reasonably be rejected. In order to decide whether this is so, we need first to form an idea of the burdens that would be imposed on some people in such a situation if others were permitted to X. Call these *the objections to permission*. We then need, in order to decide whether these objections provide grounds for reasonably rejecting the proposed principle, to consider the ways in which others would be burdened by a principle forbidding one to do X in these

[22] If we cannot know such facts, then we aren't assessing particular actual acts. For with particular actual acts, we do know who probably will benefit or lose. So, regarding the question of whether Scanlon's contractualism is a version of act-contractualism or rule-contractualism, this suggests his view is rule-contractualist.

[23] A, for me at least, puzzling feature of Rawlsian contractualism to keep in mind is that it would be, strictly speaking, incorrect to say that Rawlsian contractors could not or would not 'run the risk' on the sexist or racist principle. Part of the information that is stripped from contractors behind the veil of ignorance are "the special features of [their] psychology such as [their] aversion of risk or liability to optimism or pessimism." (Rawls, 1999, p. 118).

> circumstances. Suppose that, compared to the objections to permission, *the objections to prohibition* are not significant, and that is therefore reasonable to reject any principle that would permit one to do X in the circumstances in question. This means that the action is wrong, according to the contractualist formula. Alternatively, if there were some principle for regulating behavior in such situations that would permit one to do X and that it would not be reasonable to reject, then doing X would not be wrong: it could be justified to others on grounds that they could not reasonably refuse to accept. (Scanlon, 1998, p. 195). [*my italics*].

Scanlon's proposal of making the question of "reasonable rejection" of any principle comparative via the weighing of the "objections to permission" against the "objections to prohibition" is a first crucial constitutive feature of his contractualism because, in combination with a second constitutive feature of contractualism, it is what lands Scanlon in hot water regarding aggregation.

This second constitutive feature is one that Scanlon considers to be one of contractualism's central commitments, which Derek Parfit (2011b) has labelled as "the Individualist Restriction":

> The Individualist Restriction: In rejecting some moral principle, we must appeal to this principle's implications only for ourselves and for other *single* people (Parfit, 2011b, p. 193).[24]

So, in Scanlon's contractualism, for any principle we need to apply the test of 'reasonable rejectability'. The 'reasonable rejectability' test is comparative in nature: we should weigh the "objections to permission" of the principle under consideration against the "objections to prohibition" of that same principle. How should we understand either kind of these objections? Here I should highlight a sentence from the longer passage from Scanlon above:

> In order to decide whether [a principle which permitted X could reasonably be rejected], we need first to form an idea of the *burdens* that would be imposed on some people in such a situation if others were permitted to X. (Scanlon, 1998, p. 195). [*my italics*].

[24] In Scanlon's own words: "the justifiability of a moral principle depends only on various *individuals'* reasons for objecting to that principle and alternatives to it." (1998, p. 229). This restriction stems from Thomas Nagel's (2012) idea of "pairwise comparison", which is the idea that in assessing the rightness of a particular way of acting, we should compare everyone's reasons one-on-one (i.e. "pairwise") (Nagel, 2012, p. 125).

First, to be clear, we need to consider the very same question for the objections to prohibition. There is nothing special or different about considering the permissibility of a principle as opposed to its prohibition. For either kind of objection, we need to consider the burdens that would be imposed on people and determine which give rise to the stronger objection – those that permit the principle or those that prohibit it. But "people" in the previous sentence should be read in one way only, given the Individualist Restriction: only burdens that apply to each person as an individual should be considered for the 'reasonable rejectability' test of any principle. The 'individualist' understanding of sizing up objections to be made against the permission or prohibition of a principle should then be understood according to what Parfit calls the "*Greater Burden Claim*" (Parfit, 2011b, p. 192):

> Greater Burden Claim: 'it would be unreasonable…to reject a principle because it imposed a burden on you when every alternative principle would impose much greater burdens on others.' (Parfit, 2011b, p. 192).[25]

Parfit importantly adds, I think correctly, that "Scanlon uses the phrase 'impose a burden' in a wider sense, which covers not only harming someone but also failing to give someone some possible benefit." (Parfit, 2011b, p. 192). Given the "comparative nature of the question of reasonable rejection" as defended by Scanlon, the following inverse principle of the Greater Burden Claim seems entailed by Scanlon's contractualism:

> Veto: it would be reasonable to reject a principle because it imposed a burden on you when at least one alternative principle would impose smaller burdens on everyone else.

The aggregation problem for Scanlon's contractualism arises because Veto is such a strong claim, at least in the kind of cases where we can benefit one of either two groups of people, but we cannot benefit everyone. The combination of the fact that we cannot benefit everyone and Veto generates

[25] Parfit's single quotation marks suggest that he is quoting Scanlon (1998). However, he cannot be because this phrase cannot be found in Scanlon (1998). But the claim is nevertheless well-supported and certainly entailed by Scanlon's account. For the clearest textual support for this Greater Burden Claim, see Scanlon's comments on (1) the 'water rights' example (1998, pp. 192–194) and both (2) the "Rescue Principle" and the "Principle of Helpfulness" (1998, p. 224). Elizabeth Ashford (2003) uses Scanlon's "Rescue Principle" to persuasively argue that the structure of contractualism commits Scanlonian contractualism to a much more demanding picture of morality than Scanlon himself finds intuitively plausible. I will discuss the topic of (moral) demandingness in relation to Scanlonian contractualism in §2.1. of chapter 3.

some very implausible results for Scanlon's contractualism. There are four important issues which I will mention here but will leave until the next chapter to discuss. First, the details of these kinds of cases. Second, why more precisely they put questions of aggregation at the front and centre of the discussion of contractualism. Third, why Scanlon defends his Individualist Restriction in response to what aggregation appears to entail in some of the cases of the kind described. And, fourth, how his Individualist Restriction (and Veto as entailed by it) turns out to be a cure seemingly worse than the 'disease' of always permitting aggregation.

Having presented the key elements of Scanlon's contractualism which give rise to the aggregation problem, this brings me to the fifth and final key question: non-agreement. For Scanlonian contractualism in general the consequences of non-agreement seem that the contractors would deprive themselves of standing in the valuable relation of mutual recognition with each other. Apart from being intrinsically valuable, standing in this fundamental moral relationship with other people is likely to bring further benefits, and so is also instrumentally valuable. For example, if I recognise your motivation to have your choices and actions be such that I could not reasonably reject them, then I will likely come to trust you (more), be (more) willing to help you out (with increasingly arduous projects), and so on. In short, the relation of mutual recognition works, amongst other things, as a fertilizer to prudential considerations regardless of the particulars of your or my life-plans. Whether my life-plan is to be a philosophy professor and yours is to be an investment banker, we both need a supporting network that goes beyond the purely self-interested cooperative schemes of *quid pro quo*.

Many more interesting things could be said regarding non-agreement and Scanlon's contractualism, but that is beyond the scope of this thesis. In Scanlonian terms, I think we could quite plausibly say that refusing to find any agreement on interpersonal aggregation can be reasonably rejected in its own right.

Having presented the key features of Scanlon's contractualism which give rise to the aggregation problem for contractualism and having identified Scanlon's contractualism as a version of "fairness contractualism" using Jonathan Quong's methodological questions, I turn to a final important question in the next and final §2 to further specify the kind of contractualism that Scanlon's contractualism is. The main question of the next section is whether Scanlon's contractualism is a version of act-contractualism or rule-contractualism.

§2. Act-Contractualism vs. Rule-Contractualism

The distinction between act-consequentialism (or, utilitarianism) and rule-consequentialism (or, utilitarianism) is well-known.[26] Surprisingly, the parallel distinction has barely been discussed in relation to contractualism.[27] I say "barely" because thanks to Hanoch Sheinman (2011) there is a distinction between act-contractualism and rule-contractualism that we can work with.

Unfortunately, Sheinman labels the distinction as being between act-contractualism and "principle contractualism". But, as I have discussed in §1, principles play a special role in Scanlonian contractualism with a formal definition as "general conclusions". Therefore, I will use the label of "rule-contractualism" instead. Nothing is lost by this change. If anything, I think it is now even clearer that what I consider here for Scanlon's contractualism, following Sheinman, is the analogous distinction as drawn within the rival views of consequentialism and utilitarianism.

Having clarified this point, I now turn to the different kinds of contractualism that Sheinman identifies. First, there is "rule-contractualism":

> Rule-Contractualism: A has an obligation to do X (and A's failure to do X is wrong) just when principles no one could reasonably reject as public standards of behavior

[26] I will follow Elizabeth Ashford and Tim Mulgan (2018) in their definition of utilitarianism as the moral theory that "takes persons' moral status to be grounded in their capacity for well-being and suffering, and takes well-being to be the sole moral value. It takes the appropriate response to this value to be to promote it. Utilitarianism is thus a *consequentialist* moral theory – morality is concerned with bringing about valuable outcomes" (Ashford & Mulgan, 2018, Section 3). Commenting on his own motivations for developing his account, Scanlon says that he "look[s] to views such as contractualism specifically as ways of avoiding utilitarianism." (1998, p. 215). However, Scanlon extensively criticises "teleological" accounts of morality – those accounts which take "states of affairs" or "ways the world might go" as the bearers of intrinsic value and look "to realize those states of affairs that are best—that is, have the greatest value." (1998, pp. 79, 80). For his full criticism of teleology, see Scanlon (1998, pp. 79–87). So, Scanlon objects to both the broader target of consequentialism as well as the more specific target of utilitarianism. For Scanlon's most illuminating discussion of the relation between utilitarianism and consequentialism, see his imaginatively and aptly called section "The Shadow of Hedonism" (1998, pp. 100–103).

[27] Although Scanlon certainly hints at the possible relevance and importance of this distinction: "Taking familiar controversies about act and rule utilitarianism as a background, it would be natural to ask why justification of our actions should proceed by way of principles at all. Why not consider individual acts instead?" (1998, p. 197). But Scanlon dismisses (the importance and/or relevance of) this question as follows: "Put in this way, the question is misconceived. To justify an action to others is to offer reasons supporting it and to claim that they are sufficient to defeat any objections that others may have. To do this, however, is also to defend a principle, namely one claiming that such reasons are sufficient grounds for so acting under the prevailing conditions." (1998, p. 197). I find this reply unpersuasive as it would reduce the scope of any principle to apply only to one most specific act and acts very similar – if not identical – to it. In other words, I would understand Scanlon as arguing that he defends a version of act-contractualism, although he confusingly claims principles to be most important. Furthermore, I think that Scanlon here conflates his explanation of what a principle is with his own "universality of reason judgments" – the idea that any set of factors, G, which give me reason to X in circumstances C entails the universal claim that G gives anyone else a reason to X in circumstances C (1998, p. 73). Defending this kind of universality does not commit one to understand principles as having the kind of very specific or narrow scope of application as Scanlon appears to think it does. I will get back to the quote from Scanlon at the start of this footnote later in this §2 because Sheinman raises a different objection to it than the point I have made here.

> (*unrejectable* principles) require A do X (equivalently: just when A's failure to do X violates principles). Additionally/alternatively: A's doing X is right (permissible) just when A's doing X conforms to unrejectable principles. (Sheinman, 2011, pp. 292–293).

According to Sheinman, Scanlon, as all other "modern contractualists", is an example of someone who defends a version of rule-contractualism (Sheinman, 2011, p. 292). Sheinman contrasts rule-contractualism with "act-contractualism":

> Act-Contractualism: A has an obligation to do X (and A's failure to do X is wrong) just when no one could reasonably reject A's doing X as a particular instance of behavior (in short, when it is *unrejectable*). (Sheinman, 2011, p. 295).

What is the common denominator between act- and rule-contractualism? It is what Sheinman calls "foundational contractualism":

> Foundational Contractualism: What matters ultimately in action is unrejectability. (Sheinman, 2011, p. 298).

Sheinman argues that whereas act-contractualism directly reflects foundational contractualism, rule-contractualism reflects foundational contractualism indirectly at best (2011, pp. 299–302). In fact, Sheinman argues that in rule-contractualism, acts are two steps removed from the fundamental moral property of unrejectability: first, because acts are justified by being in accordance with the right rules and, second, because principles in turn rely on "the societal *acceptance/institution* of these very [rules] as public standards of behavior[.]" (2011, p. 301). And so, Sheinman argues, rule-contractualism directly reflects not foundational contractualism, but "foundational [rule]-contractualism":

> Foundational Rule-Contractualism: What matters ultimately in action is conformity to rules whose societal acceptance/institution as public standards of behavior is unrejectable. (Sheinman, 2011, p. 302).

Sheinman's ultimate or main contention is that the rule worship objection to rule-consequentialism – that, come what may, we should stick to following the rules even if that (predictably) leads to

disasters – equally applies to Scanlon's contractualism (p. 292).[28] For that strategy to work, Sheinman needs to argue that Scanlon's contractualism and rule-consequentialism are sufficiently similar, which Sheinman does: "both have an asymmetric two-stage structure that gives rules justificatory priority." (2011, p. 291). How are rule-contractualism and rule-consequentialism different?

> Assuming that 'rules' and 'principles' are coreferential, the difference between modern contractualism [of which Scanlon's contractualism is one example] and rule-consequentialism is that between contractually and consequentially justifiable principles. The former are principles no one could reasonably reject as public standards of behavior even if their general acceptance as such standards would fail to promote the overall good; the latter are principles whose general acceptance as public standards of behavior would promote the overall good even if someone could reasonably reject them as such standards. (Sheinman, 2011, pp. 290–291).

Although I agree with the way Sheinman contrasts rule-contractualism and rule-consequentialism, it is his assumption of 'coreferentiality' between Scanlonian principles and rules that I take issue with.

As I have discussed in §1, Scanlon has a specific understanding of principles. Principles have the narrow scope of application to one kind of act and ones very similar – if not identical – to it. The rules defended by rule-consequentialists, such as Brad Hooker (2009), tend to be much more general than Scanlonian principles. My suspicion is that Sheinman, though understandably so, gets too hung up on Scanlon's requirement of principles that they ought to serve as public standards of behaviour.[29] So, again as discussed in §1, here Scanlonian contractualists run into the choice between endorsing a set of infinite principles and bringing principles to a higher level of generality so that the set of principles can always be evaluated (as well – in addition to the specific principle

[28] For an excellent overview of the strengths, objections to, and replies to those objections regarding rule-consequentialism, see Brad Hooker (2016).

[29] Some further textual support of my point here: "But regardless of the particular way in which principles are accepted as public standards, one cannot reasonably reject them as such standards just by pointing to one person who can reasonably violate them on one abnormal occasion." (Sheinman, 2011, pp. 294–295). This just seems false to me. Given Scanlon's emphasis on considering all the relevant standpoints – which means going well beyond the standpoints of those directly affected – what Sheinman suggests is ruled out on Scanlonian contractualism is precisely what is entailed by it: The Greater Burden Claim and Veto. It is exactly the 'odd' standpoint in some "abnormal" version of a sufficiently like case for the principle to still apply that gets Scanlonian contractualists into trouble with regards to the number problem and the aggregation problem. Both the number problem and the aggregation problem I will discuss in chapter 2.

focused on for reasonable rejection). I have argued that, following Hooker, Scanlonian contractualists should opt for the second alternative: it must be possible to evaluate the set of principles.

So, it is important to note that Sheinman could only have this revised version of Scanlon's contractualism as the appropriate target for his criticism that mirrors the rule worship objection to rule-consequentialism. Although with Scanlon's repeated talk of and emphasis on (the importance of) principles, his (1998) contractualism is much closer to act-contractualism than Sheinman portrays it to be. But since I have argued that Scanlonian contractualists should revise their view in this way, that means Sheinman's rule worship criticism has bite here.[30]

Still, Scanlon considers the question of act vs. rule versions of contractualism and says the following:

> Taking familiar controversies about act and rule utilitarianism as a background, it would be natural to ask why justification of our actions to others should proceed by way of principles at all. Why not consider individual acts instead? Put in this way, the question is *misconceived*. To justify an action to others is to offer reasons supporting it and to claim that they are sufficient to defeat any objections that others may have. To do this, however, is also to defend a principle, namely one claiming that such reasons are sufficient grounds for so acting under the prevailing conditions. (Scanlon, 1998, p. 197) [*my italics*].

Sheinman focuses his attention on the last sentence from the passage from Scanlon above and argues that for it not to amount to a tautology it must be interpreted as follows:

> To offer undefeated reasons for doing some particular act X is also to defend the act-general principle that there are undefeated reasons for doing every act that is G, where G is some non-trivial property of X. (Sheinman, 2011, p. 309).

[30] Of course, there is a middle ground between 'rules' and 'infinite principles'. However, for the sake of argument I am ignoring this complication here, especially because it would come to my aid rather than Sheinman's. And I want to give Sheinman's objection to Scanlon's contractualism the best chance of sticking to it.

Sheinman argues, I think correctly, that Scanlon, as such, fails to engage with the choice between act-contractualism and rule-contractualism as both these versions of contractualism entail his interpretation of Scanlon's claim above:

> If Rule Contractualism is true, whenever you offer undefeated reasons for doing some particular act X you defend the act-general principle that there are undefeated reasons for doing every act the non-performance of which violates unrejectable principles. If Act Contractualism is true, whenever you offer undefeated reasons for doing some particular act X, you defend the act-general principle that there are undefeated reasons for doing every unrejectable act. (Sheinman, 2011, p. 309).

Sheinman summarises his argument as follows:

> The crucial asymmetry between Act and Rule Contractualism comes to this. While both are moral theories of acts, only Act Contractualism justifies acts non-derivatively, by their own unrejectability; Rule Contractualism justifies acts only derivatively, by their conformity to principles whose societal acceptance/institution is unrejectable. Rule Contractualism is vulnerable to the worship charge because it justifies acts derivatively, regardless of whether they are rejectable. Nothing of this sort can be said of Act Contractualism. Whatever Act Contractualism justifies it justifies by its own unrejectability. And Act Contractualism justifies acts only. In particular, it does not justify principles. Now, again, it may or may not make sense to extend Act Contractualism to collective acts. If it does, then Act Contractualism can be said to justify the societal acceptance/institution of principles through their promulgation, teaching, and enforcement. But then such acceptance is just another 'act', to be justified non-derivatively by its own unrejectability. Justification under Act Contractualism proceeds in exactly one stage; the only justification is direct or non-derivative. (Sheinman, 2011, pp. 312–313).

In the end, I struggle to see whether Scanlon and Sheinman disagree. What Sheinman says regarding how the "crucial asymmetry" must be understood and how he goes on to argue why act-contractualism is better I understand as not a complaint against, but a proper description of Scanlon's (1998) view. I fear that Sheinman imposes a much more general view of principles than Scanlon in fact defends.

More important for my purposes is that the distinction does not appear to have bite in the context of aggregation. Act- and rule-contractualists would, I think, agree both that the aggregation problem and the more specific questions within the context of aggregation, which I will discuss in the next chapter, need to be answered. However, that means that on aggregation their positions converge.

Having discussed the distinction between act-contractualism and rule-contractualism, I have argued why I think that the distinction, though important at a more general level of discussion, does not make a difference regarding the aggregation problem for Scanlonian contractualists.

Conclusion

In this first chapter, I have discussed all I consider necessary for the discussion of aggregation and the problem it presents to Scanlon's contractualism without discussing the details of the aggregation problem itself. What I have done in this first chapter in detail are the following two things.

First, I have presented my understanding of Scanlon's contractualism both in terms of what kind of contractualism it is and its distinctive features. Using Quong (2017), I have identified Scanlonian contractualism as a version of what Quong calls "fairness contractualism", and so separated it from both "consent contractualism" and "rationality contractualism". What is distinctive of Scanlon's "substantive" contractualism is that it has as its domain of application the subdomain of morality that is interpersonal morality, or "what we owe to each other".

Within this domain, Scanlon assumes the contractors to be motivated to stand in the positively valuable relationship of "mutual recognition". We stand in this relationship of "mutual recognition" if we abide by principles that no one could "reasonably reject". The 'reasonable rejectablity' test for any principle is determined via a comparative question which weighs the "objections to prohibition" against the "objections to permission". Either kind of objection should be understood in terms of "burdens" which would affect each person as an individual – this is the Individualist Restriction. Parfit has shown that Scanlon's contractualism therefore entails the Greater Burden Claim and I have argued that Veto is the inverse principle of the Greater Burden Claim, and so equally entailed by Scanlon's contractualism.

I have argued that the aggregation problem for Scanlon's contractualism is that Veto is a very strong claim in at least the kind of cases where we can benefit one of either two groups of people, but we cannot benefit everyone. The combination of the fact that we cannot benefit everyone and Veto entails some very implausible results for Scanlon's contractualism. There are four important issues related to this which I will discuss in the next chapter. First, the details of these kinds of cases. Second, why more precisely they put questions of aggregation at the front and centre of the discussion of contractualism. Third, why Scanlon defends his Individualist Restriction in response to what aggregation appears to entail in some of the cases of the kind described. And, fourth, how his Individualist Restriction (and Veto as entailed by it) turns out to be a cure seemingly worse than the 'disease' of always permitting aggregation.

The second main point I have discussed in this chapter, relying on Hanoch Sheinman's (2011) discussion, is that I have addressed the question of whether Scanlon's contractualism is an act-contractualist or rule-contractualist account. I have argued why I believe that Scanlon's contractualism both is and should be a rule-contractualism. However, I have also argued that regarding the aggregation problem, the distinct positions of act-contractualism and rule-contractualism converge on what requires justification.

Having set the stage, I will now move onto the second chapter to start my direct address of the central question of this thesis: what is the aggregation problem, and could Scanlonian contractualists solve the aggregation problem?

University of Reading

The Number Problem and the Aggregation Problem for Scanlonian Contractualism

Abstract This chapter starts the discussion of Scanlon's contractualism and aggregation. I argue that the standard problem discussed in the philosophical literature, the "numbers problem", is an aggregate of two problems that are best kept separate: "the number problem" and "the aggregation problem". Scanlon's proposed solution to the aggregation problem, his Individualist Restriction, has been persuasively objected to by his critics as a cure that is (far) worse than the 'disease' of always permitting aggregation. I argue that insofar Scanlonian contractualists could cherish any hope of finding one overarching solution to both problems, the focus should be on the aggregation problem, rather than the number problem. My point that a solution to the aggregation problem should come first is the reverse of the understandable yet flawed strategy of non-consequentialists in the literature: they standardly first suggest a solution to the number problem and then, if they do, hint at how this solution might be extended to the aggregation problem. Finally, I discuss what I believe is the best account that meets these criteria, Alex Voorhoeve's (2014) "Aggregate Relevant Claims". Voorhoeve's account brings such promise for Scanlonian contractualists that I will discuss whether Scanlonian contractualists could incorporate Voorhoeve's ARC in chapter 3.

Keywords Aggregation • Contractualism • Relevance • Scanlon • Voorhoeve

Introduction

Having set the stage in chapter 1 by introducing Scanlon's contractualism and its key elements, I seek to show why aggregation issues present such a problem for his contractualism in this chapter.

As a reminder, my talk of "aggregation" refers to "interpersonal aggregation" only unless specified otherwise. Interpersonal aggregation I understand as follows:

> Interpersonal Aggregation: The benefits (or burdens) that come to different individuals may be taken together (i.e. they may be 'aggregated') and so considered as one total benefit (or burden).

Aggregation involves 'adding' together benefits or burdens across persons, but aggregation need not be 'linear', like utilitarian aggregation. To illustrate linear aggregation, suppose that we could give either a benefit of size 100 to one person or one hundred benefits of size 1 equally distributed to one hundred other people. According to linear aggregation, we should be indifferent between these two alternatives. Absent other considerations, it is equally good to give an aggregate benefit of 100 by either giving each of a hundred people a benefit of size 1 or to give one other person an individual benefit of 100. The total benefit we bestow in either alternative is 100. Standard utilitarian accounts are linear aggregative accounts.

But there are other options besides linear aggregation. For example, prioritarianism – the view according to which benefits that would come to those people who are or would be worse off matter more – entails, like linear aggregation, that any size of individual benefit can be aggregated. But, unlike linear aggregation, prioritarians give greater weight to those benefits that would come to the worse off people. To use the earlier example, if the one person who would get the 100-benefit is or would be worse off than all of the hundred other people who would each get a 1-benefit, then the one person should get the 100-benefit.

What is important to note is that the added moral weight given to benefits that come to those who are or would be worse off does not change the (numerical) size of either the individual benefit(s) or the aggregate sum of them. On this prioritarian account of aggregation, regardless of who would be worse off and to whatever extent, we would in my example still either give a 100-benefit to one individual or one hundred 1-benefits to one hundred other people. If the one person were worse off, then it would not be true that we would give her the benefit because in virtue of being worse

off her individual benefit would be, say, 110 rather than 100 in size. Thomas Porter (2011) calls this the "standard conception" of prioritarianism, which gives us the following general structure:

> 'Standard' Prioritarianism: [A]ny unit that takes a person from utility level x to level (x+1) will have the same prioritarian value, regardless of the overall utility level of the person involved. But this prioritarian value of any such unit is greater than that of any unit that takes a person from level (x+1) to (x+2). This is because the prioritarian weighting of each individual unit diminishes even as the weighting of other units are being enjoyed by the person remains the same. (Porter, 2011, p. 202).[1]

Two further options, and the final ones I will discuss here for the sake of illustration, are close to but distinct from prioritarianism as I have described it. A third version of aggregation is a view according to which we should discount each person's additional benefit. So, for example, if we could give a same-sized benefit to either Anna or Bella, but we have already given a benefit to Bella, then – despite the 'same-sizedness' of the benefit – we should give the benefit to Anna. This third aggregative view is, therefore, also non-linear.

According to a fourth view, we should discount according to each additional person, such that benefitting one hundred people with a 1-benefit each is not double the value of benefiting fifty other people with a 1-benefit each.

Having illustrated that even if we accept aggregation across the board there are still at least these different options, specifically a choice between linear and non-linear aggregation, let me make a point about the structure of this second chapter. In this chapter, I argue that in the philosophical debate on aggregation two problems have been unhelpfully – and quite ironically – lumped together, usually under the heading of the "numbers problem". The first problem is what I call "the number problem". The second problem is what I call "the aggregation problem".

Both the number problem and the aggregation problem arise in the context of the distribution of a scarce resource and some number of competing claimants such that we, as the distributor, cannot

[1] For examples that illustrate the "standard conception" of prioritarianism in a similar way as I have aimed to do on the preceding paragraph, see Porter (2011, pp. 199, 203).

satisfy all claims.[2] If we were able to satisfy all claims, then we simply should. At least for my purposes, no interesting moral question would arise when the resource isn't scarce. But when the resource is scarce, then there are two variants of this kind of case which mark the difference between the number problem and the aggregation problem.

The number problem is the problem of whether the number of people you could benefit matters morally. Phrased as a question: if everyone has a claim of equal strength (i.e. if all would receive the same-sized benefit if their claim were satisfied), then are we required to satisfy the greater number of claims because they are the greater number?[3]

The aggregation problem is a problem downstream of the number problem. That is, the aggregation problem only arises if, as a first condition, we think that we are required to satisfy the greater number when claims or benefits are equal. To get the aggregation problem, we change one feature of the original number case scenario: the competing claimants have claims of differing individual strength. For the purposes of this chapter, I will assume the groups to be "homogenous", rather than "heterogeneous" – that is, all members of the same group all have the exact same claim.[4]

The aggregation problem is the problem of how to have one's account of aggregation respond to this added complication of having claims of unequal strength between the groups which one could save. For those advocating a "pure aggregation" view, such as utilitarians, this is not much of a problem in the sense of the simplicity of their answer: whether we have claims of equal or unequal strengths is irrelevant to the question of aggregation, we should aggregate in both kinds of cases. We should always aggregate and we should choose that alternative which yields the greatest aggregate total of good. So, one response to the aggregation problem is:

> Pure Aggregation: There is no restriction on aggregating claims, no matter how weak these claims are.

[2] But, to be precise, in chapter 1, I had been speaking only of "benefits", whereas here I speak of people having "claims". Although more commonly you can be benefitted in many more ways than you have a claim to, in my discussion I will not make this distinction because it would distract from the question of aggregation. Therefore, I use "claims" and "benefits" interchangeably – they perfectly and fully overlap in the cases I will discuss.

[3] Again, it isn't a given that the strength of your claim either is or matches the degree to which you could be benefited, but for this discussion I will assume that they overlap/match.

[4] Patrick Tomlin (2017) highlights the importance of this feature of the standard aggregation cases.

But Scanlon presents the following counterexample to purely aggregative views, such as utilitarianism:

> The World Cup Case: Suppose that Jones has suffered an accident in the transmitter room of a television station. Electrical equipment has fallen on his arm, and we cannot rescue him without turning off the transmitter for fifteen minutes. A World Cup match is in progress, watched by many people, and it will not be over for an hour. Jones's injury will not get any worse if we wait, but his hand has been mashed and he is receiving extremely painful electrical shocks. Should we rescue him now or wait until the match is over? Does the right thing to do depend on how many people are watching—whether it is one million or five million or a hundred million? (Scanlon, 1998, p. 235).[5]

Scanlon's pair of questions above are meant to be rhetorical, but I need the spelled-out answers for my discussion. According to Scanlon, we should obviously save Jones now and we should do so regardless of however great the number of people watching might be (1998, p. 235). But, and this brings out the general point Scanlon intends to make with his World Cup Case, advocates of purely aggregative views cannot draw this conclusion.[6] They are committed to the conclusion that if the number of people watching is sufficiently large so that displeasing them by turning off the transmitter would lead to a greater total of (negative) value in the aggregate than the total of (negative) value we would get if we didn't rescue Jones, then we should not rescue Jones. We would in that case be required to keep the transmitter turned on until the match is over and only then save Jones. Hence, Scanlon claims, purely aggregative views are mistaken.[7]

---

[5] Although he in admirable honesty refuses to take credit for having come up with the following humorously astute remark himself, Campbell Brown has suggested to me that there is a way of reading Scanlon's case which makes it a non-trade-off case: if there is an hour left in a football match, then it is currently half-time. And half-time in football usually lasts 15 minutes. So, we can perfectly turn off the transmitter, rescue Jones, and have the broadcast back on exactly as the players come walking back onto the pitch after their half-time cup of tea. No viewer would have to miss a thing. If anything, you will do them favour of saving them from the half-time adverts. So, to have this case work for Scanlon, we cannot interpret it in this way. The actual match must be currently ongoing.

[6] An exception is Dale Dorsey's (2009) view, which relies on a form of lexical superiority between values which avoids the "Lives for Headaches" conclusion. I find Dorsey's view unsatisfactory because it only 'works' by exploiting the inherent vagueness surrounding the nature of value and what makes for a good human life. For a more general discussion of "value superiority", see Gustaf Arrhenius and Wlodek Rabinowicz (2015).

[7] It is important to note here that Scanlon seems to be strictly talking of purely linear aggregative views, as his target is utilitarianism. I will leave the discussion of purely non-linear aggregative views, such as prioritarianism, until chapter 4. For several succinct defences of pure (linear) aggregation, see Alastair Norcross (1997; 1998a; 1998b).

So, how does Scanlon avoid the for him unpalatable result entailed by purely aggregative views? By rejecting pure aggregation. He rejects pure aggregation via what Scanlon himself considers a central commitment of his contractualism, which I have presented in chapter 1, and which Derek Parfit (2011b) has labelled as "the Individualist Restriction". As you will recall from chapter 1, the Individualist Restriction entails the Greater Burden Claim, which in turn implies the inverse principle I have labelled Veto:

> The Individualist Restriction: In rejecting some moral principle, we must appeal to this principle's implications only for ourselves and for other *single* people (Parfit, 2011b, p. 193).[8]

> Greater Burden Claim: 'it would be unreasonable…to reject a principle because it imposed a burden on you when every alternative principle would impose much greater burdens on others.' (Parfit, 2011b, p. 192).

> Veto: it would be reasonable to reject a principle because it imposed a burden on you when at least one alternative principle would impose smaller burdens on everyone else.

In the World Cup Case, Jones can veto any principle which would not entail his immediate rescue simply because he has the most at stake. Importantly, Jones retains his veto regardless of the number of viewers – we could add viewers *ad infinitum*, we would still be required to immediately rescue Jones on Scanlon's contractualism. Now, in the World Cup Case and like cases – e.g. saving one person's life or preventing any number of other people from each having to suffer a mild headache – you might think, like Scanlon, that this just gets things right. More generally, Scanlon's view seems to be the following – which is the second possible response to the aggregation problem:

> Anti-Aggregation: Aggregation of claims is never permissible, no matter how strong these claims are.

[8] In Scanlon's own words: "the justifiability of a moral principle depends only on various *individuals'* reasons for objecting to that principles and alternatives to it." (1998, p. 229). This restriction stems from Thomas Nagel's (2012) idea of "pairwise comparison", which is the idea that in assessing the rightness of a particular way of acting, we should compare everyone's reasons one-on-one (i.e. "pairwise") (Nagel, 2012, p. 125).

However, Scanlon's Individualist Restriction (and Veto as entailed by it) turns out to be a cure seemingly worse than the 'disease' of always permitting aggregation. To illustrate this objection, we should increase the strength of the weaker claims in the kind of case of which the World Cup Case is one version.

Suppose that we have a World Cup Case where we hold what happens to Jones fixed, but what happens to each of the viewers is much worse than in Scanlon's original version of the case. For example, the moment we turn off the transmitter to rescue Jones, each of the viewers receives a painful though – and this is crucial – still weaker electrical shock than the shocks Jones would have suffered – from their television set. The generalisable point is that we should assume that although Jones retains the strongest claim, the relative gap between his claim and the individually weaker claims of each of the many is much smaller.

But, and this is the objection to Scanlon's contractualism, Jones would retain his veto in this revised World Cup Case. No matter how great the number of viewers who will receive electrical shocks, we should save Jones instead. But this is a highly counterintuitive result. In this case, we do want the numbers to matter. And it seems that to make the numbers matter we need to permit aggregation. The fact that we are not even permitted to save the many to-be electrocuted World Cup viewers counts as a *reductio* of Scanlon's contractualism. Scanlon's answer to the aggregation problem is at least as implausible as the answer given by pure aggregationists. More generally, Anti-Aggregation seems at least as implausible as Pure Aggregation.

Scanlon anticipates this fatal-looking objection, and responds as follows:

> [I]t seems that our intuitive moral thinking is best understood in terms of a relation of "relevance" between harms. If one harm, though not as serious as another is nonetheless serious enough to be morally "relevant" to it, then it is appropriate, in deciding whether to prevent more serious harms at the cost of not being able to prevent a greater number of less serious ones, to take into account the number of harms involved on each side. But if one harm is not only less serious than, but not even "relevant to," some greater one, then we do not need to take the number of people who would suffer these two harms into account in deciding which to prevent, but should always prevent the more serious harm. Thus it might be claimed, for example, that missing half an hour of exciting television is not relevant when we are

> deciding whether to save a person in front of us who is in extreme pain, but that total paralysis or blindness is relevant to the even more serious harm of loss of life. So it could be wrong to save one person's life when we could instead have prevented a million people from going blind or becoming paralyzed. (Or, at least, it would be permissible to prevent this harm to the greater number of people even though it would not have been permissible to prevent this lesser harm rather than the greater if the numbers involved had been the same.) (…) I am not certain how such an argument would go, but it does not seem to me to be excluded in advance by the general idea of contractualism. (Scanlon, 1998, pp. 239–240, 241).

Scanlon, however, does not say any more than this. My main aim in this thesis is to pick up where Scanlon left off: to consider (a) whether an argument in support of the 'limited' aggregative view he is hinting at here could be given and (b) even if such an argument could be given, whether this argument would be consistent with (the rest of) Scanlon's contractualism.

So, my motivation is to establish the following view on behalf of Scanlon which consists of two general and related claims:

> Limited Aggregation: if the relative gap between the competing individual benefits become too large, then there is no number of people for the group of the many weaker claimants that could make it the case that we, as the distributor, should benefit them instead of the one (or each of a few) competing claimant(s) who has a much stronger claim. If the relative gap between the competing individual benefits is insufficiently large, then we aggregate the claims and do what yields the greater aggregate good.

But before I focus on the main task at hand, I should note that not just in this chapter, but for the remainder of this thesis, I will focus exclusively on harms of various degrees of seriousness, such as death, paralysis, and mild headaches.[9] Naturally, there are many interesting (philosophical) questions regarding the nature of pain and how, if at all, harms are commensurate. I cannot go into these questions in this thesis, given that my focus is on the moral question of interpersonal aggregation. Therefore, I will assume throughout that different kinds of harm are commensurate.

[9] This focus avoids Scanlon's powerful objections to well-being, specifically his objection that well-being cannot serve as a "master value", see Scanlon (Scanlon, 2003b) and his third chapter of *What We Owe to Each Other* (Scanlon, 1998, pp. 108–143). In (Scanlon, 2003b), Scanlon, by his own choice, finds himself in the swamp or sinkhole of the "objective" vs. "subjective" distinction. However, I should say that Scanlon does as well as possible in navigating this distinction.

However, to emphasise, this is not a thesis in bioethics. The focus on harms is there because of its illustrative potency. But what I go on to say about aggregation is meant to apply to any good which, if received, benefits and, if denied, burdens any individual to some extent.

So, I will assume that (a) in all my cases, on any plausible view of well-being, well-being is affected (where I will stick to physical harm/health); and (b) on any plausible view of morality, well-being matters.[10]

I now turn to the number problem to discuss it in greater detail, how Scanlon suggests we deal with it in a non-aggregative manner (his "tie-breaker argument"), why Scanlon's argument fails, and how the plausible alternative solution to the number problem via permitting some form of aggregation sets up the aggregation problem for his contractualism and for any like-minded philosopher on aggregation.

§1. The Number Problem

The number problem arises in cases in which we face the choice of aiding either one of two groups of people where the individual burdens faced by each member of either group is equivalent.[11] I assume that all people in either group are complete strangers to us and hold all else equal. Furthermore, I will assume that death is a harm. So, for example, we can imagine that without our help, all people in both groups will die, and furthermore that all these people (a) are currently the same age; (b) will, if saved, live to the same age; and (c) will, if saved, have equally good lives.[12]

[10] So, to be more precise, I understand good (physical) health as immensely valuable because it is a means to pursuing other things for oneself or others. But good (physical) health seems only valuable as a means. If it is only valuable as a means, then I cannot believe it is more valuable than the things that are valuable as ends. Note how often people risk or even sacrifice their good health for the sake of other things, such as friendship, achievement, autonomy, etc. And people who are 'health nuts' are criticised precisely for putting too high a priority on their own health. Well-being, by contrast, is (immensely) intrinsically valuable. A further contrast with (physical) health is that well-being is constituted by several elements, such as friendship, achievement, and autonomy. So, although closely related, (physical) health and well-being are distinct things. Scanlon agrees with this view on health and well-being (Scanlon, 2003b, pp. 74–75). Furthermore, my exclusive focus on this physical-harm/health-aspect of well-being avoids Scanlon's persuasive criticism of "welfarist contractualism", which focuses on well-being at large (1998, pp. 217–218, 242-243). For more on the elements of well-being and a "sympathy test" for arguably determining whether something is an element of well-being, see Brad Hooker (2015). For an excellent discussion of Scanlon's contractualism and well-being, see Jonathan Wolff (2004). For an excellent discussion which breaks "health" down into a plurality of values to respect its complexity and shows the (philosophical and/or theoretical) challenges this presents us with, see Daniel M. Hausman (2015). Given my purposes, I had to simplify the value of health as a monism, rather than a pluralism.

[11] I say "equivalent", rather than "(exactly) equal" in this first instance precisely to pre-empt the discussion of 'rough equality' or 'parity'. For an excellent overview collection of the debate on incomparability and incommensurability see Ruth Chang (1997a). Chang herself has presented influential arguments in favour of 'parity' (1997b; 2002) For persuasive criticisms of Chang's arguments, see John Broome (1997) and Luke Elson (2014). For an alternative defence of the vagueness account, see Cristian Constantinescu (2016). I thank Luke Elson for having shared his expertise on the topics of incommensurability, incomparability, and vagueness with me.

[12] Using Quong's method that I have presented in chapter 1, I can say that the outcome of all people dying would constitute the non-agreement point to be avoided here.

Finally, for these cases to illustrate the number problem, I need to assume that there is a difference in the number of people we might save, depending on the particular group we choose to save. For my purposes, it is enough to imagine that the choice involves only three people, divided one vs. two, so that the central number problem scenario is as follows:

> Rescue Case You are faced with a choice between saving either the life of one person, A, or the lives of two others, B and C. All three individuals are complete strangers to you and rescuing them would be at negligible cost to you. Furthermore, there are no morally relevant differences between the two groups (e.g. none of them was responsible for ending up in this situation). Finally, you are the only one that can save some of them. What ought you to do?

One thing that contractualists and utilitarians would agree on is that you certainly are required to save someone.[13] For contractualism, the principle that permits you to not save anyone in this situation can be reasonably rejected by any of the three individuals whose life is at stake. From the reasonable rejection of the principle permitting you not to save anyone, we can deduce that the principle requiring you to save someone cannot be reasonably rejected. This follows from what Scanlon calls "the comparative nature of the question of reasonable rejection" (1998, p. 195), which I have discussed in detail in chapter 1.[14]

The justification for this principle requiring you to save someone is provided by pointing to the features of the scenario, which provide you with reasons.[15] For instance, that without your help people will die and that you can help them at virtually no cost to yourself. It is hard to imagine a weightier reason to the contrary. For example, you are not risking life and limb if you were to try to rescue either A or B and C. Therefore, contractualists conclude that the principle that you ought

[13] As I will discuss in chapter 6, John Broome's (1984; 1990; 1998; 2003a; 2004) theory of fairness – and Broomean accounts that have followed in his wake, e.g. Iwao Hirose (2015) and Christian Piller (2017) – would state that in this situation the fairest thing to do would be to save no one. As Broome would agree, that is a *reductio* of the view that fairness would be all that matters (morally).

[14] Scanlon, though, notes that there are some exceptions to this 'comparative enabling' of a principle through reasonable rejectability of its counterpart. One such exceptional case is the case of two people swimming from a sinking ship and one of them finds a life jacket floating in the water, where both the principle permitting the other to take the life jacket by force and the principle prohibiting her to take it by force can be reasonably rejected (1998, pp. 195–196).

[15] In this respect, Scanlon is a constructivist (2012). But Scanlon is adamant to emphasize that he is not a Kantian (2011).

to save someone in cases like Rescue Case cannot be reasonably rejected.[16] Having established that you ought to save someone in situations like Rescue Case, the following important question arises: given that you cannot save everyone, whom ought you to save here? We should ask the following question:

> Equal Harms: If all else is equal, are we required to save the greater number of lives because they are the greater number? In other words, should we save more people (e.g. B and C) rather than fewer (e.g. A) when we cannot save everyone?

"Yes, we are/should", does sound like a very plausible answer. Scanlon believes the answer to Equal Harms is "yes" (1998, p. 234). Scanlon's argument for saving the greater number has been helpfully formalised by Michael Otsuka (2000):

> Scanlon's Tie-Breaking Argument:
> (P1) The claims of A, B and C should be accorded equal and positive weight (Scanlon, 1998, pp. 232–233).
> (P2) We accord them equal and positive weight if and only if we add C's claim to B's (Scanlon, 1998, pp. 232–235, 397 fn. 35).
> (C1) We should add C's claim to B's (From P1 and P2).
> (P3) Adding C's claim to B's tips the balance in favour of saving B and C (Scanlon, 1998, p. 232).
> (P4) This tipping of the balance in favour of saving B and C justifies saving B and C (Scanlon, 1998, p. 232).
> (C2) Saving B and C is justified (From C1, P3 and P4). (Otsuka, 2000, p. 290).

Unfortunately for Scanlon, not only is this argument for saving the greater number in Equal Harms invalid, but any such argument seems unavailable to his contractualism. First, this particular argument fails because, as Otsuka (2000) points out, P2 is incorrect: it is not true that only by adding C's claim to B's, we "accord them equal and positive weight", since we would also do this when we would hold either (i) a weighted lottery, giving an 1/3 chance of being saved to A and a

[16] There is a question whether it is every single act that needs to pass the contractualist test of reasonable rejection or whether it is 'merely' the principles that need to do so. I have given my reasons for why I think the 'reasonable rejectability' test should be applied to principles when I discussed the act- vs. rule-contractualism distinction in §2 of chapter 1.

2/3 chance to B and C, or (ii) an unweighted lottery, giving all three individuals a chance of 1/2.[17] Second, and more important, it appears that Scanlonian contractualists cannot accommodate any argument for saving the greater number in Rescue Case because of the Individualist Restriction.

This is because, in Rescue Case, the Individualist Restriction requires of us that we compare A's claim to be saved with B's and A's claim with C's separately. As a result, we see that all claims are of equal weight: each of the three individuals faces death. Since the Individualist Restriction prevents both B and C from appealing to the fact that if either one of them is saved so is another person, the result for contractualism appears to be that we are morally required to either hold a lottery or simply choose who to save. More specifically, we are required to hold an unweighted lottery. A weighted lottery aggregates individual chances which is ruled out by the Individualist Restriction.

This so-called "numbers scepticism" is what contractualists are committed to due to the Individualist Restriction.[18] This is the opposite of the conclusion that Scanlon wishes to reach in the kind of equal-claims-different-number cases that gives rise to the number problem, namely saving the greater number. The main point, however, is that the Individualist Restriction prevents contractualists from justifying an affirmative answer to Equal Harms – that we are required to save the greater number in Rescue Case. The failure to provide the justification for the requirement of saving the greater number in Rescue Case and, more generally, the failure to justify its affirmative answer to Equal Harms is the number problem for Scanlon. Although one could argue that for Scanlonian contractualism we might disagree with Scanlon and aim to justify a different answer to Equal Harms, I think Scanlon's answer to Equal Harms is most plausible. Consider for example an equal harms case that were not 1 vs 2, but 1 vs 1,000. Therefore, I will assume that on behalf of Scanlonian contractualists the aim is to justify an affirmative answer to Equal Harms. In other

[17] Iwao Hirose (2015) agrees that Scanlon's argument fails to argue for saving the greater number and that it is consistent with the weighted lottery (Hirose, 2015a, p. 155). However, Hirose also argues that Scanlon's argument is able to reject the unweighted lottery (p. 155). As I will argue below, I disagree with Hirose. Scanlon's argument is unable to reject the unweighted lottery. For the most (in)famous defence of the unweighted lottery in Equal Harms, see John Taurek (1977). For defences of the weighted lottery, see F. M. Kamm (1998) and Ben Saunders (2009). For a rather unique "Individualist" lottery, according to which we don't pool individual chances but should save the other group members as well if the lottery winner is a member of a larger group, see Jens Timmermann (2004). Because Scanlon does not aim to argue for a lottery, whether weighted or unweighted (or anything seemingly in between, like Timmermann's lottery), as the solution to the number problem I will not look to defend a lottery-solution on behalf of Scanlonian contractualists. However, in chapter 6, I will, as mentioned earlier in footnote 13, discuss John Broome's theory of fairness in detail where lotteries will receive ample attention. Lotteries are considered not as an alternative to be argued for, but rather as a possibly unavoidable result if fairness is taken seriously.

[18] For examples of (intended) number sceptical accounts, see G. E. M. Anscombe (1967), John M. Taurek (1977), Véronique Munoz-Dardé (2005), Alan Thomas (2012), and Tyler Doggett (2013).

words, on behalf of Scanlonian contractualists, I will aim to justify saving the greater number in equal-claim-different-number cases.

By contrast, utilitarians have an easy time in justifying saving the greater number as a response to the number problem. This is because utilitarians are concerned with maximising utility.[19] In the cases discussed here, well-being is the utility at stake, and it would be maximised by preventing the greatest total of harms. More importantly and unlike contractualists, utilitarians have the means to justify the claim that you ought to save the greater number in Rescue Case. This is because utilitarians accept pure aggregation and pure aggregation, in turn, provides the utilitarian with the justification of why in cases like Rescue Case you ought to save the greater number. If you can either have a utility of one life saved or a utility that is five times greater by saving five different lives, and if utilitarianism requires that we maximise utility (understood here as the value of well-being), then we ought to save the greater number in Rescue Case. Importantly, this result generalises. Utilitarians have a justification for an affirmative answer to Equal Harms. Therefore, utilitarians have support for the plausible reply to the number problem that we should save the greater number in these cases.

In summary, the number problem is the problem for contractualists that in all cases that have a structure like Rescue Case, they fail to argue for the principle requiring us to save the greater number. At most, contractualists can justify the principle that in Rescue Case, you ought to save someone. However, given the serious nature of Rescue Case, it does not seem unreasonable that we demand some further guidance on what we are permitted and/or required to do from a moral theory. The range of possible responses to Equal Harms seems to be as follows, where each response might be permitted or required dependent upon one's view of aggregation:

(1) Save the greater number;
(2) Weighted lottery;
(3) Unweighted lottery;
(4) Choose who to save (i.e. 'plumping' or 'picking' – no lottery)

At this point, the Individualist Restriction commits contractualists to require option (3), the unweighted lottery and so shows that they are, unintentionally, number sceptics. Even worse for

[19] I am ignoring John Broome's (forthcoming) otherwise persuasive arguments that "utility" should not be used to denote a person's good or well-being (forthcoming, pp. 1–5). I ignore Broome's advice because it has no effect on the discussion of aggregation and, I fear, would distract from it.

Scanlonian contractualists is the fact that there is an alternative moral theory, namely utilitarianism, which entails that we are required to do what Scanlon believes we should do in cases like Rescue Case – option (1) of saving the greater number – while also having a justification for it. Therefore, contractualists should provide a different argument to justify its affirmative answer to Equal Harms – that option (1) is required – or alter their contractualism on this point. Before I discuss such alternative justifications for the requirement of saving the greater number, I will first, in the next section, discuss the second, related problem for contractualism, the aggregation problem. For it is in the context of the aggregation problem that the Individualist Restriction has drawn its appeal to begin with, as I will discuss.

§2. The Aggregation Problem

Like the number problem, the aggregation problem is best explained through an example. Suppose that you are faced with a situation in which people require your help. Suppose, furthermore, that the same assumptions hold in this case as they did in Rescue Case, such as that you have a general duty to help those in need, that helping would involve only negligible cost to you, and so on.

The crucial difference between this case and Rescue Case, however, is that the harms (and so the strength of the individual claims of people to not have to suffer them) are not equivalent. We imagine that the one individual on the one side, A, will suffer a worse harm than each of the members of the larger group on the other side. I will refer to the larger group as "the many". A further assumption we need to make in order to mark the contrast between contractualism and utilitarianism is that aiding the many would result in bringing about a greater total of utility compared to the outcome where we save A from a worse harm. This is important because, according to utilitarianism, the greater total utility of that outcome would determine that we are morally required to aid the many. This is a consequence of utilitarianism's acceptance of pure aggregation and its commitment that rightness is determined by what's best. Bearing all this in mind, we should first ask the following important question in this context for contractualism, which utilitarians answer with an unqualified "yes":

> Unequal Harms: If all else is equal, are we (at least) permitted to help the many if and because helping the many would amount to a greater total of utility if we cannot help everyone?

In this unqualified sense, Scanlon answers this question with a "no". This is because, for Scanlon, a difference in the size of the individual benefits for each member of the many plays a crucial role in whether we should or shouldn't help them (1998, pp. 235–241). This point that the individual benefit matters, I believe, shows the appeal of the Individualist Restriction.

For the Individualist Restriction imposes a constraint upon Unequal Harms which prevents utilitarian conclusions such as that in Unequal Harms we ought to aid the many rather than the one individual, even if as a result the one individual will die and each of the many is only spared from a mild headache. For this reason, Scanlon states that "this feature [the Individualist Restriction] is central to the guiding idea of contractualism, and is also what enables it to provide a clear alternative to utilitarianism and other forms of consequentialism." (Scanlon, 1998, p. 229). But as I have argued in the previous §1, one problem of the Individualist Restriction is that it leads to the number problem. An additional problem that the Individualist Restriction leads to in the context of Unequal Harms, and to which I can now point, is that it does not allow us to prevent an impersonally worse outcome in terms of total utility in which the harms faced by the many are, though less serious than the harms faced by the one individual, still very bad.

Scanlon argues that in cases where the lesser harm is still very bad, we should at least be permitted to save the many, if they are a significantly large number (1998, p. 239). Scanlon explains this, as I have presented in my introduction, by claiming that "our intuitive moral thinking is best understood in terms of a relation of 'relevance' between harms" and that when unequal harms are considered relevant to each other, they belong to the same "broad category of moral seriousness" (1998, pp. 238–240). So, for example, it would mean that, whereas paralysis is relevant to death and so a certain number of it would make it at least permissible to prevent this outcome rather than someone's death, there is no number of mild headaches such that it would be even permitted to prevent these rather than to save someone's life.

However, Scanlon should explain these "categories", rather than to presume them and use them as a justification for his intuitive judgments. So, as good as the ideas of "relevance" and "broad categories of moral seriousness" may sound, these notions need further explanation and justification.

Having described in greater detail the pieces of the puzzle on aggregation and Scanlon's contractualism, let me put the puzzle into full view. There are two kinds of central cases, the Rescue

Case and the World Cup Case. The Rescue Case motivates the view which I have called Pure Aggregation. Scanlon presents the World Cup Case as a counterexample to Pure Aggregation and proposes his Individualist Restriction to limit aggregation. But the Individualist Restriction limits aggregation absolutely and so entails the view which I have called Anti-Aggregation. Anti-Aggregation is not the view of aggregation that Scanlon intends to advocate, nor should Scanlonian contractualists advocate Anti-Aggregation, as I have argued. Anti-Aggregation plausibly rules out saving the many in the World Cup Case, but most implausibly also rules out saving the greater number in the Rescue Case. I have called the view of aggregation that (a) Scanlon intends to advocate and (b) which requires saving the greater number in the Rescue Case but prohibits saving the many in the World Cup Case, Limited Aggregation. The view of Limited Aggregation provides two challenges: (1) is there a principled way to get Limited Aggregation to work? And (2) even if such a principled justification for Limited Aggregation could be given, then can that justification be endorsed from within Scanlon's contractualism?

Let me now present the options for Scanlonian contractualists regarding aggregation. The first choice contractualists face is regarding the three broad views of aggregation: Pure Aggregation, Anti-Aggregation, and Limited Aggregation. If contractualists opt for Pure Aggregation, then their view provides no alternative to utilitarianism. Therefore, I will treat opting for Pure Aggregation by Scanlonian contractualists as a non-starter. Contractualists could opt for Anti-Aggregation. Opting for Anti-Aggregation means that Scanlonian contractualism is a distinct alternative to utilitarianism, which is a version of Pure Aggregation. However, given how implausible the number-sceptical view of Anti-Aggregation is, I think this is not the view that Scanlonian contractualists should opt for.[20] That leaves Limited Aggregation. Limited Aggregation ticks both boxes for contractualists in terms of being a genuine alternative to utilitarianism and being a plausible view of aggregation.

The second choice contractualists face, if they want to hold on to contractualism and the intuitions underlying Limited Aggregation, are as follows. The first option is to reject aggregation altogether and find a different way of justifying the requirement of saving the greater number in the Rescue Case and the prohibition of saving the many in the World Cup Case. The second option is to rely on aggregation as a means of justification but to try and limit it in such a way that it justifies the

[20] Although we might imagine someone saying: "I am a contractualist and I am properly anti-aggregation and anti-numbers. Scanlon should have the courage of his convictions and follow the Individualist Restriction where it leads." Scanlon might think, like me, that this contractualist is wrong, but she is not inconsistent. This is one instance of how Scanlon's own view and what is entailed by his contractualism might come apart.

'right' answers in both the Rescue Case and the World Cup Case. If this were done successfully, contractualists could justify the now simply presumed notions of "relevance" between certain harms, which places these harms in the same "broad category of moral seriousness". I fail to see how the first option could be successful. Therefore, in the remainder of this thesis, I will explore the contractualist's options within the parameters of the second option. More specifically, I consider two alternative Limited Aggregation strategies.

The first strategy is to develop and justify the idea of "relevance", as suggested by Scanlon, to limit aggregation to apply to what I call "close harm cases" (e.g. death vs. paralysis), but not to what I call "distant harm cases" (e.g. the World Cup Case and death vs. mild headaches). The second strategy is to instead limit aggregation by pointing to values other than utility, specifically fairness, that might make it the case that in close harm cases we may aid the many, but that in distant harm cases we may aid the one or few. So, the most important difference between the two strategies, as I see it, is that the first is a monist strategy – it operates on the assumption that there is only one value at play, utility, which is in distant harm cases limited due to the lack of relevance between individuals harms and so prohibits aggregation – whereas the second strategy is a pluralist strategy – it is not so much aggregation but the impact of utility that is limited because of the counterweight that at least fairness provides.[21]

There will likely be other options and other strategies within these two (and possibly further) options that Scanlonian contractualists or any Limited Aggregation advocate could consider. However, to repeat, for the remainder of this thesis, I will concern myself with solely the two strategies under the second option as canvassed above and consider the two challenges I have mentioned before: (1) is there a principled way to get Limited Aggregation to work? And (2) even if such a principled justification for Limited Aggregation could be given, then can that justification be endorsed from within Scanlon's contractualism?

However, before I explore these two strategies, I first want to briefly discuss the relation between the number problem and the aggregation problem in the next section.

[21] For my purposes, I cannot go further into the distinction between "monism" and "pluralism" here. For an excellent discussion on monism and pluralism about value, which is instructive of the monism vs. pluralism (structural) distinction as such, see Chris Heathwood (2015).

§3. Why a Solution to the Aggregation Problem has Priority

Part of the trouble for contractualism in answering both the number problem and the aggregation problem stems from the way in which Scanlon discusses the cases of pain and death outside of cases of conflict as presented above.

For example, in his more general discussion of pain, Scanlon argues that "the idea that value and disvalue are a matter of 'to be promoted' and 'to be prevented' has its greatest plausibility with respect to experiential states of this kind." (1998, p. 82). Furthermore, he argues that the case of pain is the exception to the rule that reasons for our actions do not always flow from the value of the state of affairs they bring about: "A state of affairs in which I am in pain is, for that reason, a worse state of affairs for me, and this fact gives rise to reasons to do what is necessary to prevent it." (1998, pp. 92–93). Now, Scanlon's remarks here do not support the view that in cases of conflict, it is the greatest sum of pain that should be prevented. But I do find Scanlon's remarks at least hint at that view.

Thus, these remarks seem helpful as a starting point for contractualism in answering the number problem, since the harms there are equal for all individuals and contractualists want to prevent as many of these harms as possible. At the same time, however, these same remarks appear to be a non-starter for contractualism regarding cases where the harms are unequal. For it does not say anything about the need to qualify the affirmative answer contractualists want to give to Unequal Harms given contractualists' different take on close harm cases vs. distant harm cases. And without such qualification, rejecting the idea that there is, for example, a number of mild headaches that morally outweighs the loss of another individual's life is unjustified. What do Scanlon's remarks teach us about how contractualists should tackle the two problems they face?

I argue that these remarks by Scanlon show that contractualism should not attempt to derive a solution to the aggregation problem from a solution to the number problem. To illustrate why, consider Tom Dougherty's (2013) "Rational Ends" account. Dougherty presents his account as a non-utilitarian solution to the number problem, and so contractualism might consider adopting it.[22] Put rather crudely, Dougherty's account consists of two steps in a scenario such as Rescue Case:

[22] For similar arguments, see Adam Hosein's (2009) chapter 1, "Numbers, Fairness and Beneficence" (pp. 5-31) and Joseph Raz's (2004) 'Degrees-of-Compliance-with-Reason' account. Like Dougherty, Hosein's and Raz's solutions solely target the number problem but, unlike Dougherty, Hosein and Raz do not consider whether and how their

Rational Ends:

(1) You are morally required to want the survival of each stranger for its own sake.

(2) You are rationally required to achieve as many of these ends as possible, if you have these ends.

When we try to apply this solution to Unequal Harms, the conclusion that there will be some number of people suffering mild headaches such that it requires us to prevent these headaches rather than to save someone's life, seems unavoidable.

To see this, consider first a case of non-conflict in which you can prevent a certain number of strangers from suffering mild headaches. You can choose to (i) do nothing, (ii) help some strangers but not all, or (iii) help all strangers. Rational Ends would require you to do (iii). Next, consider a case of conflict, again with all strangers facing mild headaches where you can choose to (i) do nothing, (ii) help the one stranger on the one side, or (iii) the many strangers on the other side. Rational Ends would again require you to do (iii). Finally, consider a case of conflict where the lone stranger faces death and the many others face mild headaches. Now, surely you should not want the prevention of a mild headache as much as the survival of a stranger, but it seems that, according to Rational Ends, you should want it to some extent. If we were not required to care about the prevention of mild headaches, then in the first two cases, we would perform supererogatory acts when we would do (ii) or (iii) and we would be equally permitted to do (i). But that seems quite counterintuitive, so Rational Ends is best understood as requiring us to want the prevention of a mild headache to some extent.

However, since the second step in Dougherty's argument clearly is one of maximisation, then there will be a number of mild headaches such that it is morally better to prevent these headaches rather than to save another person's life. And, as mentioned above, when the idea of maximisation is in place, attempts to then limit it seem arbitrary and lacking support. In fact, Dougherty argues that

solutions might apply to the aggregation problem. The objection I here raise to Dougherty's solution, applies equally to Hosein and Raz their accounts. Iwao Hirose (2015b) argues that F. M. Kamm's (2007) "Argument for Best Outcomes" (2007, p. 32) and Thomas C. Shelling's (2006) "*probabilistic argument*" are two successful "nonaggregative" solutions to the number problem (Hirose, 2015b, pp. 312–313). However, Hirose provides indirect support for my claim that we should start with attempting to find a solution to the aggregation problem, rather than the number problem, by admission of the following: "These arguments work only in the cases where the level of harm is the same. When the level of harm is different from one person to another, these arguments do not say anything. Therefore, the scope of these arguments is very limited." (Hirose, 2015b, p. 313). I fail to see how either Kamm's or Schelling's argument could be plausibly extended to the aggregation problem without also (at least) running into the same issue I address in this thesis. I will go on to argue in this section why it is (expectable more) imprudent to start at the number problem (rather than the aggregation problem) in an attempt to find an overarching solution to both the number problem and the aggregation problem.

his account avoids the utilitarian conclusion of mild headaches outweighing a life by stating that "[*Rational Ends*] applies only to cases where an agent pursues *relevantly* similar ends." (2013, p. 425) [*my italics*]. So, again, it is the notion of "relevance" that prevents the commitment to utilitarian conclusions, but it is precisely this notion that requires further explanation if non-utilitarians, such as contractualists, want to answer the aggregation problem.

In response to this, contractualists might argue that this simply shows that we have good reason to sever the ties between the number problem and the aggregation problem. The argument would presumably be that, since solving the number problem will require maximising whatever notion is put forth as acceptable to the contractualist (e.g. claims, reasons, or rational ends) and it is the idea of maximisation that brings contractualism into trouble with regard to the aggregation problem, separate solutions for the two problems are required. However, I argue that there are at least two reasons for first looking to find one broad solution that respects the ties between the two problems.[23]

First, the numbers play a significant role in both contexts. Regarding the number problem, this is straightforward. But also with regard to the aggregation problem, contractualists want to argue that in certain cases of Unequal Harms it is because of the sheer number of people that might be burdened by the lesser harms that makes it the case that we are at least morally permitted to prevent this outcome, rather than to prevent the outcome in which a greater harm is imposed upon another individual.

Second, the only difference between the cases that bring out each problem, is whether the burdens for all individuals are equivalent or not. But this also means that if contractualists can solve the aggregation problem via a justification of Limited Aggregation, then they have also solved the number problem. After all, if contractualism can argue why we should rather prevent a certain number of lesser harms being suffered by some than preventing another individual from suffering a greater harm, then surely it should be possible to also justify from this argument why we should save a greater number of people from suffering an equivalent harm, rather than just one person.

Can there be a solution to the aggregation problem that does not solve the number problem? As far as I can see, it would require a view that permits aggregation of unequal-close harms but

[23] On the interpretation of Scanlon's contractualism as discussed and rejected by me in chapter 1 that there is an infinite number of principles, opting for separate solutions to the number problem and the aggregation problem would be more tempting and/or seem a more natural fit.

prohibits aggregation of equal harms. Such a view both seems bizarre to me and I cannot see what the justification for such a view would be. I conclude, therefore, that contractualists should first try to respect the relation between the two problems and find one broad solution that covers both contexts.

In summary, I have argued that it seems difficult to derive a solution for the aggregation problem from a solution to the number problem because, if treated separately, the number problem requires some form of maximisation and when this is then applied to cases of Unequal Harms, any attempt to restrict it seems arbitrary and lacking support. Consequently, contractualists might consider providing one solution to the number problem and another, separate solution to the aggregation problem. I have argued that contractualism should first attempt to provide a solution to the aggregation problem. Why? A solution to the aggregation problem will be a solution to the number problem because the number problem is simply a subset of the aggregation problem. In other words, contractualists ought to treat Equal Harms as a special, easier instance of Unequal Harms.

I will now discuss what I take to be the most promising account for Scanlonian contractualists to justify the idea of "relevance", Alex Voorhoeve's (2014) "Aggregate Relevant Claims", or ARC for short. Voorhoeve's ARC is a proposal for a principled way to distinguish between pairs of unequal harms in which the harms are relevant to each other and pairs for which such a relation does not hold.

§4. Aggregate Relevant Claims

Alex Voorhoeve's account of "Aggregate Relevant Claims", or ARC for short, is spelled out as follows:

> ARC:
>
> 1. Each individual whose well-being is at stake has a claim on you to be helped. (An individual for whom nothing is at stake does not have a claim.)
> 2. Individuals' claims *compete* just in the case they cannot be jointly satisfied.
> 3. An individual's claim is *stronger*:
>   a. the more her well-being would be increased by being aided; and
>   b. the lower the level of well-being from which this increase would take place.
> 4. A claim is *relevant* if and only if it is sufficiently strong relative to the strongest competing claim.

5. You should choose an alternative that satisfies the greatest sum of strength-weighted, relevant claims. (Voorhoeve, 2014, p. 66).

Voorhoeve claims that ARC can provide a qualified answer to Unequal Harms: we should aggregate in some cases, close harm cases, but not in others, distant harm cases. If Voorhoeve's claim is correct and if Voorhoeve's ARC can be endorsed from within Scanlonian contractualism, then contractualists would have their justification for their preferred answers in both the Rescue Case and the World Cup Case, and so would have a solution to the aggregation problem.

There are three important issues that need to be kept distinct here: (1) Does Voorhoeve's ARC work on its own terms? (2) Can Voorhoeve's ARC be accommodated within Scanlon's contractualism? (3) If Voorhoeve's ARC cannot be accommodated within contractualism as is, how fundamental would any alterations to contractualism need to be?

Unsurprisingly, in Voorhoeve's ARC, it is the fourth step, which brings in the notion of "relevance", that especially requires further argument. Fortunately, Voorhoeve does not shy away from this daunting responsibility, as he provides the following argument in support of this idea of relevance.

The first thing that Voorhoeve points out regarding relevance is that he assumes the truth of 'the permissibility of greater self-concern' (2014, p. 71). For instance, when I am faced with the choice of either saving myself or saving some stranger from death, I am morally permitted to save myself, if no other moral considerations stand in the way. Importantly, this permissibility extends beyond cases in which I and the stranger face the same harm. So, for example, when I am faced with the choice of either saving myself from (the less bad harm of) becoming bedridden for the rest of my life or saving some stranger from death, I am still morally permitted to save myself. Given that it also applies to these cases where I have the weaker claim, the permissibility of greater self-concern is a non-consequentialist idea, since a consequentialist would demand that I satisfy another's stronger claim over my own weaker claim.[24] However, Voorhoeve stresses that this self-favouring is not unrestricted (2014, p. 71). If, for instance, I am faced with the choice of either sparing myself from being bedridden for a day or saving some stranger from death, I should save the stranger.

[24] See William Godwin (2000, pp. 81–82) for defence of such extreme impartiality as required by morality.

From this discussion, Voorhoeve takes what he calls "the permissible personal perspective", which is "the point of view that results from taking on a person's maximally permissible degree of self-concern and minimally required other concern" (2014, pp. 71-72). Voorhoeve's next move in his argument is then to say that, based on this perspective, a criterion can be formulated for when a claim is relevant:

> Criterion for Relevance: A's claim is relevant to B's competing claim just in case, in a one-to-one comparison of their claims from A's permissible personal perspective, A's claim is at least as important as B's. (Voorhoeve, 2014, p. 72).

The criterion transfers the permissibility of greater self-concern in cases where the choice is between either saving yourself or saving a stranger, to cases where you must choose which of either group of strangers you will save. In the first self-regarding choice we are considering here, it is permissible to save yourself from complete paralysis rather than to save a stranger's life, when you consider this question from "the permissible personal perspective". The idea that it is permissible to save yourself from the less bad harm means that your weaker claim to be saved from complete paralysis is relevant to the stranger's stronger claim of having her life saved. Following ARC, you should, in an other-regarding choice where you can save either one stranger's life or save many other strangers from complete paralysis, aggregate the set of weaker claims – given that they are relevant to the stronger competing claim – and do what satisfies the greatest sum of strength-weighted claims.

In the second self-regarding choice, on the other hand, the idea is that is not permissible to spare yourself from a mild headache rather than to save a stranger's life, when "the permissible personal perspective" is taken up. And so, given that you are not permitted to save yourself from this less bad harm, this means that your weaker claim to have your mild headache prevented is not relevant to the stranger's stronger claim of having her life saved.

In the other-regarding choice where you can either save one stranger's life or prevent many other strangers from suffering a mild headache, ARC requires you to save the one stranger's life. Importantly, given that there is no relevance between the claim to have one's life saved and the claim to be spared from a mild headache, ARC will require you to save a stranger's life irrespective of any number of other strangers you might have spared from each suffering a mild headache instead.

Voorhoeve's way of illustrating how the first-person choices should impact on the third-person choices is to ask whether there is "unanimous agreement" from taking up the permissible personal perspective of all individuals involved (2014, p. 74). He argues that in a death vs. mild headaches case, from the perspectives of all individuals facing a mild headache, all parties will agree that satisfaction of the competing claim by the lone individual of having her life saved takes priority over the satisfaction over each claim to be spared from a mild headache.

By contrast, in death vs. complete paralysis cases "there is no such unanimity" (Voorhoeve, 2014, p. 74). Each person facing complete paralysis is permitted to insist on her claim to be saved from this harm, while the single person facing death does so too. Voorhoeve claims that, after finishing our "round" of identifying with the perspective of each, "one can see that there is no solution that does full justice to the personal perspective" (Voorhoeve, 2014, p. 74).

What does ARC require considering this disagreement from each person's perspective? It tells us to switch to the aggregative approach, so that we satisfy the greatest sum of claims (Voorhoeve, 2014, p. 75). In summary, ARC seems able to justify exactly the responses that contractualists have to death vs. paralysis and death vs. mild headaches by using a principled form of Limited Aggregation, and so the aggregation problem seems to have been solved successfully.

Before I discuss the question of whether ARC is compatible with Scanlonian contractualism and present several objections that others have raised to ARC and Voorhoeve's "permissible personal perspective" specifically in the next chapter, I would like to stress the ingenuity of Voorhoeve's solution. I want to highlight how he uses the non-consequentialist idea of 'the permissibility of greater self-concern' as a basis for a solution in an area where non-consequentialists (including contractualists) struggle, namely the area of numbers and aggregation. Going through the stages of first pointing to the permissibility of greater self-concern, to then formalising this into the "permissible personal perspective", which in turn generates the criterion of relevance, shows an incredible advancement in the quality of the justification that a non-consequentialist (again, e.g. a contractualist) can now give in response to the aggregation problem.

However, some powerful objections have been raised to Voorhoeve's ARC. Furthermore, even if Voorhoeve's ARC stands up to scrutiny and I could answer these objections on Voorhoeve's behalf, that still leaves the question whether Scanlonian contractualists could implement Voorhoeve's ARC. These questions I address in the next chapter

Conclusion

In this second chapter, I have presented two related but distinct problems which plague Scanlon's contractualism, the number problem and the aggregation problem.

Regarding the number problem, which arises in cases of Equal (read as "equivalent", not "exactly equal") Harms, I have argued that Scanlon's argument fails to justify the principle which requires saving the greater number in cases with the structure of Equal Harms. Even stronger, Scanlon's argument seems to make a stronger case for requiring an unweighted lottery.

Regarding the aggregation problem, which arises in the contrast between the two subsets of cases under the heading of Unequal Harms – close harm cases (contractualists: "aggregate!") and distant harm cases (contractualists: "always save the one/few!") – I have argued that the Individualist Restriction has some appeal here. The Individualist Restriction has some appeal precisely because it successfully prevents utilitarian conclusions in distant harm cases: that, when their membership is large enough, we should save the many rather than the one/few.

Unfortunately for contractualism, the Individualist Restriction is too successful in limiting aggregation. This is because in close harm cases, the Individualist Restriction commits contractualism to the principle that requires us to save the one individual from, e.g. death, no matter how many other people might become, e.g. completely paralysed as a result of that choice. Considering this, I have argued that if contractualists want to both hold on to contractualism and Scanlon's intuitions in the Rescue Case and the World Cup Case have two options: to either reject aggregation and find an alternative way of answering the two problems, or to accept aggregation but successfully limit it. I have made clear that I will for the remainder of this thesis explore two alternatives under the second heading of Limited Aggregation: The first strategy is to develop and justify the idea of "relevance", as suggested by Scanlon, to limit aggregation to close harm cases, not to distant harm cases. The second strategy is to instead limit aggregation by pointing to values other than utility, specifically fairness, that might make it the case that in close harm cases we aid the many, but in distant harm cases we aid the one or few.

I have argued that, regardless of what strategy for finding a solution to these two problems is chosen, the aggregation problem has priority over the number problem. Why? The number problem identifies a subset of (easier) cases that fall under the broad heading of aggregation. One might object that we should work from the easier to the more difficult cases, and so the number problem has priority. However, taking Dougherty's Rational Ends account as an illustration, I

suggest that the most intuitive solution is a purely aggregative solution or to maximise whatever non-utilitarian unit (e.g. rational ends) is inserted. But once pure aggregation or maximisation has a foothold, constraining it seems both most difficult to justify and *ad hoc*. Therefore, I am sticking to the strategy of finding a solution to the aggregation problem first.

Finally, I have presented what, I think, is the most promising attempted Limited Aggregation solution of the first kind to the aggregation problem, Alex Voorhoeve's (2014) "Aggregate Relevant Claims". The obvious next task is to assess the objections that have been raised to Voorhoeve's account. I will do so in my next chapter.

All in all, my discussion here has shown that Scanlonian contractualists are still in need of an account that solves both the number problem and the aggregation problem. If contractualists fail to solve both these problems, then contractualism should either concede defeat to (at least either act- or rule-) utilitarianism or instruct us to make sure that we always carry a coin with us in case we might find ourselves in a situation where equal claims conflict between any number of people divided between two groups. Contractualists could bring dice to be covered for multiple-groups situations with equal claims all around.

University of Reading

Sequential Claims-Matching: A (More) Precise Version of the (Local) Relevance View

Abstract One of the key features of Alex Voorhoeve's "Aggregate Relevant Claims" account is the "permissible personal perspective", which I abbreviate as "PPP". First, I will present Voorhoeve's PPP and how it is supposed to work. Voorhoeve's PPP raises two broader questions. First, how (morally) demanding is Scanlonian contractualism? Second, could Scanlonian contractualists justify Voorhoeve's PPP within their moral theory? Having discussed these two broader questions, I present objections that have been raised to Voorhoeve's account, and his PPP specifically, by Julius Schönherr (2018) and Johanna Privitera (2018). I present answers to both objections on Voorhoeve's behalf. Next, I turn to a distinction drawn by Victor Tadros (2019) between "Global Relevance" and "Local Relevance". Patrick Tomlin (2017) spells out why Global Relevance seems a dead end and so Relevance View advocates should develop their account along the lines of Local Relevance. However, Tadros has left the alternative of Local Relevance underspecified. In the second half of this chapter, I present a more precise version of Local Relevance, which I call "Sequential Claims-Matching". Although Sequential Claims-Matching vindicates Scanlon's intuition in the World Cup Case and other one vs. many cases, there are at least two ambiguities within Sequential Claims-Matching that an advocate of SC-M would need to answer. I present some options a SC-M advocate has regarding either ambiguity without taking a stance on them myself. I end with presenting an alternative account of Local Relevance, which I call "Strongest Decides".

Key words Aggregation • Contractualism • Relevance • Sequential Claims-Matching

Introduction

The recent developments regarding interpersonal aggregation focus on the defensibility of the view that interpersonal aggregation is only sometimes permissible or only permissible in certain cases – i.e. Limited Aggregation. At the centre of recently developed and defended accounts of the limited (im)permissibility of interpersonal aggregation is the idea of "relevance", as I have discussed in my previous chapter.

Alex Voorhoeve (2014), through his "Aggregate Relevant Claims" account, has not only propelled forward the discussion on interpersonal aggregation as constrained by relevance, but also suggested an ingenious way of determining whether there is a relation of relevance between two unequal claims: the "permissible personal perspective" (which I abbreviate to "PPP" in this chapter). As a reminder, in §1, I will present Voorhoeve's PPP and how it is supposed to work. In §2, I investigate whether Voorhoeve's ARC is a solution to the aggregation problem that can be endorsed from within Scanlonian contractualism. In §3 and §4, I will present some recent objections that have been raised to Voorhoeve's account by Julius Schönherr (2018) and Johanna Privitera (2018) respectively. The discussion of these objections will help to provide context to Voorhoeve's account, specify the account's scope, and identify its limitations whilst retaining enough potential applicability to make it worth discussing further.

In §5, I present Patrick Tomlin's (2017) account, which presents important objections to Voorhoeve's ARC account that serve as the basis for a new variant of the Relevance view, which I call "Sequential Claims-Matching" and which I introduce in §6.[1]

In §6.1.-6.2., I present two ambiguities in Sequential Claims-Matching. I then present some options a SC-M advocate has regarding either ambiguity without taking a stance on them myself. I end my chapter, in §7, with presenting an alternative account of Local Relevance, "Strongest Decides".

§1. Voorhoeve's PPP – Non-Arbitrarily Establishing Relevance

An important objection the Relevance View faces is an arbitrariness objection: why draw the line between these two competing unequal claims with the weaker being 'irrelevant' to the stronger, but have these ever so slightly different competing unequal claims be relevant to each other? Alex

[1] I first developed and explored this view in van Gils and Tomlin (forthcoming). I am most grateful to Patrick for his confidence in me and his generous offer to take on this collaborative project, for his patience and understanding in the writing process from start to finish, and for a most fun 'relay' presentation at the 6th Annual Workshop for Oxford Studies in Political Philosophy at the University of Pavia, June 13-15, 2018. My thanks also to the OSPP editors, David Sobel, Peter Vallentyne, and Steven Wall, and all other participants of the workshop at Pavia.

Voorhoeve's "permissible personal perspective" (PPP) is an ingenious non-consequentialist (friendly) reply to this objection, as I have discussed in my previous chapter. Because the PPP is non-consequentialist and because Scanlonian contractualism is a non-consequentialist moral theory, the PPP should at least in principle be compatible with Scanlon's contractualism.

Whenever from a third-personal perspective we are confronted with the question of whether this pair of competing unequal claims are relevant to each other, Voorhoeve suggests that we place ourselves in the position of the weaker claimant and ask the following question: "would I be permitted to satisfy my own weaker claim, even though that will leave the stronger claim of this stranger unsatisfied?" If the answer is "yes", then the claims are relevant and we should aggregate these claims as a third-party arbitrator between these competing claims. If the answer is "no", then the claims are not relevant (i.e. irrelevant) to each other and we should not aggregate them. So, if the competing unequal claims are irrelevant via Voorhoeve's PPP method, then there could be any number of weaker claims but still we would never be permitted to aggregate them. And, given that we are disallowed from aggregating these weaker claims, that entails we will satisfy the stronger competing claim.

Before I can move on to discuss objections that have been raised to Voorhoeve's PPP, we first need to determine whether Voorhoeve's ARC can be endorsed from within Scanlonian contractualism, and, if not, what modifications would need to be made to Scanlonian contractualism. There are two important hurdles that Scanlonian contractualism would need to clear in order to be able to endorse ARC. The first is that ARC makes use of, and assumes, PPP. But it is not clear that Scanlonian contractualism can endorse PPP, or a very generous PPP, due to the Veto. In other words, Scanlonian contractualism might be very demanding, and therefore provide little or no room for PPP, which, even when coupled with ARC, would make it co-extensive, or near co-extensive, with utilitarianism in aggregation cases. Second, ARC, once unanimity between parties is not achieved, makes use of aggregative reasoning. This is in conflict with the Individualist Restriction, which would need to be amended.

§2. (Moral) Demandingness and Justifying PPP for Scanlon's Contractualism

Let me note in advance that not only will my treatment of both these issues be unsatisfactorily coarse-grained, but it has to be so. Both the question of demandingness and the question of justifying Voorhoeve's PPP within Scanlonian contractualism give rise to some of the most fundamental and fascinating questions in both moral and political theory. I cannot even attempt

to properly present these big questions here, let alone attempt to answer them. Therefore, I will discuss only those bigger questions that I consider of direct importance to the question of aggregation within Scanlon's contractualism in a manner that presents them as the relevant background without, to my regret, discuss any of these fundamental questions in satisfactory detail. First, (moral) demandingness.

§2.1. Scanlonian Contractualism and (Moral) Demandingness

First, I should explain why the discussion of demandingness is relevant to Voorhoeve's PPP. The demandingness discussion is relevant here due to the combination of the following reasons. The main purpose of Voorhoeve's PPP is to establish whether there is a relation of "relevance" between any given pair of unequal claims/harms. The answer is provided by considering this relevance question as a question whether satisfaction of this weaker claim would be self-favouring that is nevertheless morally permissible, if the weaker claim were one's own claim. This case/context-sensitive question of morally permissible greater self-concern is an instance of (the broader topic of) "partiality". So, therefore, I need to discuss what Scanlon says on the topic of partiality. And, like several other non-consequentialists, Scanlon discusses (the moral permissibility of) partiality in the context of demandingness. And so, I need to discuss the topic of demandingness in relation to Scanlonian contractualism.

The issue of Scanlonian contractualism being arguably at least as demanding as (strictly impartial maximising) utilitarianism arises because of Scanlon's own doing. It is because of the following principle which Scanlon himself introduces as part of his contractualism:

> The Rescue Principle: [I]f you are presented with a situation in which you can prevent something very bad from happening, or alleviate someone's dire plight, by making only a slight or even moderate sacrifice, then it would be wrong not to do so. (Scanlon, 1998, p. 224).

But with his Rescue Principle, rather than having put to bed (further) questions regarding demandingness and how demanding his contractualism is, Scanlon has opened up the floodgates. The most prominent manner in which Scanlon specifies the clause "slight/moderate sacrifice" is by appeal to (the intuitive moral permissibility of) partiality. So, partiality is meant to set a limit to the otherwise high level of demandingness of the impartial duty to aid. The reason why Scanlon appeals to partiality as constraining moral demandingness is because Scanlon aims to have his

contractualism not be at least as demanding as (strictly impartial maximising) utilitarianism. In other words, the aim is to have Scanlonian contractualism be a distinct alternative to utilitarianism. Scanlon dismisses the suggestion that his contractualism is so demanding by appealing to the intuitive plausibility of being partial toward ourselves, our loved ones, and our life projects (compared to strangers and their loved ones and life projects) as follows:

> I think that it would be reasonable to reject principles that required us always to give the same weight to the lives and interests of strangers as to those of our loved ones, or for that matter, our own. (Scanlon, 1998, p. 397, fn. 36).

Thomas Nagel (1999) provides a more precise response on behalf of Scanlonian contractualists:

> While no one could reasonably reject some requirement of aid from the affluent to the destitute, the cumulative effect on an individual life of an essentially unlimited requirement to give those who are very much worse off than yourself, whatever other affluent people are doing, would simply rule out the pursuit of a wide range of individualistic values – aesthetic, hedonistic, intellectual, cultural, romantic, athletic and so forth. Would the certain abandonment of all these things provide reasonable ground for rejection of a principle that required it – even in the face of starving millions? The question for Scanlon's model would be whether it could be offered as a justification to each one of those millions, and my sense is that it perhaps, that one could say: 'I cannot be condemned as unreasonable if I reject a principle that would require me to abandon most of my life to save yours.' (Nagel, 1999, pp. 3–6).[2]

The first thing that matters is that I think that Nagel spells out the full answer as implied by Scanlon. The second, and more important, thing is the issue that Voorhoeve's PPP relies on partiality and there is a further question whether Scanlon can justify partiality. The reason why it seems Scanlon cannot justify partiality is because his view implies that the strongest claimant always get to veto, which then makes for a very demanding moral theory.[3] Fully spelled out, the idea is

[2] This kind of reply to the threat of morality being (implausibly) demanding by appeal to partiality can be traced to Bernard Williams's (1973) excellent critical response to J. J. C. Smart's (1973) utilitarianism as presented in the same book (1973, pp. 3–74), see especially Williams's section 5 "Integrity" (1973, pp. 108–118).

[3] For further development of this critique of Scanlon's contractualism, see Elizabeth Ashford (2003). Ashford's critique is, in my view, both one of the most fundamental and underappreciated criticisms of Scanlonian contractualism, especially because of the wider implication of her criticism: whichever moral theory we opt for, morality looks to be (very) demanding. I thank Ashford for discussing her article and Scanlon's contractualism with me.

that Veto, as implied by the Greater Burden Claim and the Individualist Restriction, suggests that any principle which permits partiality could reasonably be rejected by the worst-off.

Even if Scanlonian contractualists could revise Veto to be a weaker or more concessive claim such that contractualists could grant both that people should be able to pursue one or several life-projects of their choosing (rather than providing equal support for strangers' projects) and that they may favour themselves and their loved ones (rather than strangers), then I think this would be insufficient. This is because on any plausible account of partiality there is a limit to this kind of favouring – akin to the limits within Voorhoeve's PPP.[4]

But whereas Voorhoeve permissibly assumes those limits set by the intuitive judgment through his PPP, Scanlonian contractualists will have to provide those limits through the contractualist apparatus and answer the further question where contractualism would place those limits.

Here is where I will leave the question whether Scanlon could, via another route, justify a degree of partiality within his contractualism, especially to ward off the 'threat' of being committed to a (very) demanding understanding of morality and moral requirements. As I have indicated, whether Scanlonian contractualists could justify (a degree of) partiality is beside the point anyway on my presupposed understanding of partiality as *sui generis* 'intensifiers' of impartial considerations: given a large enough number of strangers who are in need, whether it is your own on your loved-ones interests at stake, does not make enough of a difference to not be required to help the strangers.

[4] This is how I understand Thomas Hurka's (n.d.) account of partiality: reasons of partiality are *sui generis* 'intensifiers' of impartial considerations (e.g. vulnerability or need). I thank Hurka for having discussed his account with me during Ralf Bader and Ben Lange's excellent workshop on "Partiality" held at the University of Oxford on 6th-7th July 2019. Errol Lord (2016) proposes a similar view to Hurka. So, both Hurka and Lord, although Lord is more explicit about doing this, adopt Jonathan Dancy's (2004; 2018) more sophisticated view of the weight of reasons, which includes "modifiers" such as "intensifiers". Like Lord's and Hurka's view, Simon Keller (2013) relies on Dancy's sophisticated view of reasons and his idea of modifiers specifically to propose his account of partiality. But Keller goes for a different modifier, namely an "enabler" to do the work for his account. However, Lord provides the following compelling objection why Keller's proposal should be rejected: "[Keller] simply *denies* one of the claims that generate the puzzle [about partiality]. This is the claim that the value of strangers provides me with some good reason to act well towards them. According to Keller, this is not necessarily true. I think this is deeply problematic. The impartialists are right to insist that the value of other people, whether they are strangers or not, has moral significance for all of us – that value provides us all with reasons. To think otherwise is to be partial in a pejorative way." (Lord, 2016, pp. 582–583). I am persuaded by Lord's objection to Keller. Niko Kolodny's (2011) 'Resonance View' provides an important supplementary account for Hurka's and Lord's kind of account of partiality. Kolodny's idea of "resonance" provides a way of filtering the 'right' kind of relationships deserving of partiality from the 'bad' kind that do not. For Kolodny's discussion on familial relations specifically, see Kolodny (2010). For an excellent critical discussion of Scanlon's handling of partiality within his contractualism, which differs from my own, see Susan Mendus (2003). For an excellent discussion of the impartiality vs. partiality distinction within morality and/or moral theory, see Brian Feltham and John Cottingham (2010).

At this stage, I should note that "relevance" does not necessarily have to be grounded in "partiality", whether along the lines of the account suggested by Voorhoeve or an alternative account. Given the apparent problems Scanlonian contractualists have in justifying partiality, it therefore seems that Scanlonian contractualists could not make use of Voorhoeve's PPP to justify "relevance". And so, the question arises whether, rather than "partiality" more generally, Voorhoeve's PPP specifically – which is supposed to ground this idea of "relevance" – could be justified within Scanlon's contractualism.[5]

§2.2. Justifying PPP within Scanlonian Contractualism

The first thing to note is that Scanlon creates the perfect opening for something like Voorhoeve's PPP to be slotted into Scanlonian contractualism:

> Contractualism, as a theory, does not tell us exactly what level of sacrifice is required by [the Rescue Principle]. I would not say, for example, that we would be required to sacrifice an arm in order to save the life of a stranger. But here a judgment is required, and I do not think that any plausible theory could eliminate the need for judgments of this kind. (Scanlon, 1998, p. 225).

Voorhoeve presupposes that his PPP offers a principled way of judging whether as a third-party we are permitted or even required to impose sacrifices of the kind Scanlon talks about – losing an arm – on some people for the sake of saving the life of another person. I will not contest Scanlon's

[5] Another option, which I will not explore here, is that rather than holding the (most minimal) partial compliance with our impartial moral duties to aid the global destitute fixed – such as Nagel does – the demandingness of these individual duties should be cashed out in terms of requiring significant effort on each of us to increase overall compliance with these impartial duties to aid. So, for example, by setting up fund-raisers, voting for MPs and political parties standing on manifestos that will increase foreign aid budgets, etc. In general, to change the political landscape and institutional arrangements of one's society such that more of the flows of capital will find their way to the global destitute. Thus, my suggestion here operates within the framework according to which it is much better to understand a crisis such as global poverty as a collective action problem and that, therefore, the solution should be a collective solution as well. So, even if morality is very demanding, it is just not true that the only way for any individual to meet such an obligation would be to sell (nearly) all her possessions and abandon all life-projects that make her life recognisably hers. And for my suggestion here I am ignoring that the appeal to limited compliance, such as done by Nagel, is a red herring anyway. Just because other people do not donate to foreign aid or charities, does not mean that your donation would not be money well spent or make a difference (cf. Parfit, 2011b, p. 205). Liam B. Murphy (1998), in his criticism of Rawls (1999; 2005), agrees with my suggestion here that "people's responsibility for justice [is] in an important sense collective: we together, through our institutions, secure justice." (Murphy, 1998, pp. 256–257). However, as part of his attack on "moral dualism" of which he takes Rawls to be the most influential representative, he adds the following: "But we can accept that justice is fundamentally a matter of collective obligation without accepting that it is fundamentally a matter of institutional design." Although Murphy's caveat is obviously true, I believe there are several moral and prudential reasons why we should not sever the link between collective action and institutional design. Alas, I cannot go further into that discussion here.

claim above that it is a necessary condition for a theory's plausibility to make room for judgment. Rather, granting Scanlon that room for judgment needs to be made, this brings forth what I think is a fundamental distinction within Scanlon's contractualism that is insufficiently highlighted by Scanlon: what I call "institutional contexts" vs. "non-institutional contexts" within the subdomain of morality that is "what we owe to each other".[6]

The strongest textual support for this "institutional" vs. "non-institutional" distinction within Scanlonian contractualism is due to Scanlon's description of the area of morality his contractualism is concerned with as being "broader" than the area with which Rawls's 'political' contractualism is concerned: setting up a just 'basic structure' of society and the just distribution of 'primary social goods' (Scanlon, 1998, p. 242). Scanlon elaborates on this as follows:

> The contrast I am drawing here [between my contractualism and Rawls's contractualism] is not between moral philosophy on the one hand and political philosophy on the other but between very general claims about morality, or about what we owe to each other, and claims about what is right and wrong in specific kinds of circumstances (claims about the justice of social institutions being one example). (Scanlon, 1998, p. 245).

Let me now explain how I think which context we presume to be in, institutional or non-institutional, matters crucially for the question of whether and how Voorhoeve's PPP could be justified within Scanlon's contractualism.

A defining characteristic of "(moral) judgment" is a degree of "liberty" or "discretion" for the individual to decide what she may or should do. Now, as I have discussed, Scanlonian contractualism has its domain of application the subdomain of morality that is interpersonal morality or what Scanlon calls "what we owe to each other". A crucial question presented by Scanlon's suggested scope for his contractualism is as follows: are there any areas within the domain of interpersonal morality (i.e. "what we owe to each other") that are best understood as (at least also) "political"? The answer, I think, is "yes". More importantly, if there were any general area or topic that I would have to suggest for which this is true, then it would be aggregation. For

[6] At best, Scanlon addresses the question of "institutional" vs. "non-institutional" contexts indirectly at a few stages of defending his contractualism. The closest to a direct address of this question is in his response to the objection that any contractualist theory should avoid "upward" arguments, which are arguments relying on casuistry and appeal to judgments about what is "reasonable" (in respective cases) (Scanlon, 1998, pp. 242–247).

our account of aggregation most fundamentally shapes our view of distributive justice, and it is hard to think of a more well-known broad area of political theory than distributive justice.

But, and this is the first of two crucial steps, even if we suppose that, for the sake of argument, Voorhoeve's PPP can be justified to have a place within Scanlonian contractualism, then this raises an issue which needs to be addressed for the purpose of defending Limited Aggregation within Scanlonian contractualism. I believe it is plausible to presume that within the parameters of Voorhoeve's PPP, different people can reasonably differ or disagree about at least the following two points in exercising their individual moral judgment:

> (1) Having greater, lesser, or equal concern for themselves and their own interests as compared to their concern for strangers and strangers' interests.
> (2) Whether the pair of unequal harms or claims under consideration is relevant or irrelevant to each other.[7]

Importantly, these two points of reasonable disagreement are separate: for any set of individuals with an identical attitude as presented under (1) could still reasonably disagree about any pair of unequal harms or claims whether they are relevant or irrelevant to each other.

Here is the second crucial step: it matters whether we are addressing any question of aggregation (e.g. any trade-off case) using Voorhoeve's PPP as a "political" question – or, as I call it, within an "institutional" context – or as an "personal" question - which would make it a "non-institutional" context.

To illustrate what I mean by this distinction, imagine two agents, Y and Z, each facing an identical trade-off situation. For example, whether to save one person's life or preventing one million other people from becoming blind. Suppose, furthermore, that the number of weaker claims is beyond reasonable doubt: if the weaker claims are judged relevant, then we should satisfy them rather than the single life-claim. So, all that remains is the question whether a claim to not become blind is relevant to a claim of having one's life saved or not. This brings us back to our two agents, Y and

[7] Another important question – whether, if relevant, the number of weaker claims under consideration is sufficiently great to balance or even outweigh the stronger competing claim – does not allow for reasonable disagreement. At least, not on Scanlon's contractualism. Why? At rock-bottom of Scanlon's account are "reasons" to which we need to respond to. Each weaker claim provides us with a reason of a corresponding weight. The fact that Scanlon is a constructivist, rather than a realist, about reasons does not make his view more nuanced in a way where we could reasonably disagree whether a number of claims does or does not balance/outweigh a set of competing claims.

Z. The key point is that, before each of Y and Z takes up Voorhoeve's PPP to consider this relevance question, we either (1) ask each of these agent to approach this question of relevance as a "political" question or as "personal" question, or (2) it is just that both Y and Z find themselves in an "institutional" or "non-institutional" context.

Let's start with the relevance question as a personal question in a non-institutional context. Suppose it is Y who is faced with the relevance question in this context. What we are asking her is the following rather straightforward question: "Do you think that a claim to not become blind is relevant to a claim of having one's life saved or not?" Whether Y's answer is "yes" or "no" does not matter. What matters is that either way we might find at least either ourselves or another person reasonably disagreeing with Y's answer. But in response to supposed criticism of the answer she has given, Y can justifiably say: "Look, you asked me to judge whether I think these harms/claims are relevant or not. I have given you my answer, which was up to me (i.e. it fell within the morally permissible limits set by the degree of individual discretion or liberty that is just part of moral judgment). The fact that you disagree with me is both irrelevant and just tough luck for you." To this, finally, I think any of us who disagreed with Y's initial "yes" or "no" answer to the relevance question of this trade-off case could only reply: "Fair enough."

However, now imagine the exact same scenario to unfold with Z who faced an identical trade-off case, but we now change the one feature that for Z the relevance question is "political" or that she finds herself in an "institutional" context. If Z offers the exact same reply as Y did when faced with our criticism of her answer to the relevance question, then, unlike in the case with Y, I think we would be too kind to reply to Z with: "Fair enough." Most importantly, this is because the stakes in an institutional/political context are much higher because it either lays down or reinforces the rule. A requirement of any moral code, certainly concerning political questions or institutional contexts, is that it provides "assurance".[8]

Importantly, Scanlon appears to agree with me on the importance of "assurance" given his acknowledgment of "the (…) class of cases (…) in which there is a need for some principle to govern a particular kind of activity[.]" (1998, p. 339). Questions of relevance between unequal claims/harms in the context of aggregation seem to me certainly to be within this class of cases.

[8] For arguments in support of and a proposal on how to cash out the idea of "assurance", see Brad Hooker (2009), especially his chapter 5, "Predictability and Convention" (2009, pp. 112–125).

Therefore, reasonable disagreement in (at least) the institutional context on any relevance question for any trade-off case just will not do. The matter needs to be settled and settled appropriately. Exactly which collective and or procedural decision-making mechanisms are permitted or (even) required within Scanlonian contractualism (and a potential ranking between them should there be more than one permissible method) I cannot go into here.

However, should there be several different non-reasonably-rejectable principles that could govern the practice in question, then, in settling the matter, Scanlon appears to implicitly endorse a "practice-dependent" view because he defends the following principle:[9]

> The Principle of Established Practices: in situations of this kind [i.e. where there are several different non-reasonably-rejectable principles to choose from], if one of these (nonrejectable) principles is generally (it need not be unanimously) accepted in a given community, then it is wrong to violate it simply because this suits one's convenience. (Scanlon, 1998, p. 339).

By way of ending this subsection, let me note one important complication due to two closely related notions from Scanlon's contractualism which would make my distinction between institutional and non-institutional contexts (or, alternatively, between political and apolitical questions) a distinction without a difference. To be clear, these two notions should be added to Scanlon's comment that the focus of his contractualism is "broader" than Rawls's, which already suggests that Scanlon's contractualism is a version of what Liam B. Murphy (1998) calls "moral monism" – "all fundamental normative principles that apply to the design of institutions apply also to the conduct of people." (Murphy, 1998, pp. 251, 254). The first part that gives rise to the complication is the "universality of reason judgments", as advocated by Scanlon. This is the following idea:

[9] Andrea Sangiovanni (2008) formulates his "practice-dependence thesis" as follows: "The content, scope, and justification of a conception of justice depends on the structure and form of the practices that the conception is intended to govern." (2008, p. 138). In later work, Sangiovanni (2016) tells us that it was (obviously unintentionally) "misleading" of him to present the "practice-dependence" vs. "practice-independence" distinction as a contrast between two different "*approaches* to political theory" (2016, p. 15, fn. 32). He now thinks it is better to understands this distinction as "two different *modes of applying* higher-level principles and values[.]" (2016, p. 15, fn. 32). I thank Jamie Draper, Rob Jubb, and Alex McLaughlin for sharing their expert opinions on this topic of practice-dependence with me.

> Universality of Reason Judgments: any set of factors, G, which give me reason to X in circumstances C entails the universal claim that G gives anyone else a reason to X in circumstances C. (Scanlon, 1998, p. 73).

On my understanding of this principle, then – if we allow for the distinction within interpersonal morality or "what we owe to each other" between what I have called "institutional" vs. "non-institutional" contexts – there is never any difference between what we ought to do in like situations regardless of whether the context is institutional (i.e. political) or non-institutional (i.e. apolitical). If there still were the possibility to have a difference in what we ought to do due to a difference in context, then the Universality of Reason Judgments seems misnamed – it would not be "universal" in what I take to be the most natural understanding of that term.

But perhaps I am being too strict and the term "universal" allows for some flexibility in terms of being somewhat context sensitive. Scanlon certainly thinks so, as is clear from his discussion of (moral) relativism (1998, pp. 328–361). The following can be taken as Scanlon's reply to my (perhaps 'too') strict reading of "universalism":

> According to relativism, moral appraisals of actions, insofar as they are to make sense and be defensible, must be understood not as judgments about what is right or wrong absolutely, but about what is right or wrong relative to particular standards that are made relevant by the context in question, or by the context of the judgment itself. It is important that the standards in question here are *ultimate* standards. Any plausible moral view would allow for the fact that actions are right in one place can be wrong in another place, where people have different expectations, or where different conditions obtain. Failing to help a person whose car has broken down, for example, would be a serious wrong in a place where someone who is stranded overnight is likely to freeze to death, but not in a safe country with a mild climate. A view that allows for such variations in what is right, by applying a fixed set of substantive moral principles to varying circumstances, is not relativism but rather what I will call "parametric universalism." (Scanlon, 1998, p. 329).

But even if we suppose that my concern regarding Scanlon's Universality of Reason Judgments is dispelled because Scanlonian "universalism" should be understood as "parametric universalism",

the complication nevertheless arises due to Scanlon's rich understanding of "principles", as I have discussed in §1 of my first chapter.[10]

As I noted in that section, Scanlon argues that for his contractualist principles we should not just be concerned with those directly affected by the principles in question, but both those who might be affected on a future occasion and even those of whom we would somehow know that they would never be directly affected by this particular principle (1998, pp. 202–203). And the justification for this (very) wide scope of his contractualist principles appeals to the potential huge impact that "widespread performance of acts of a given kind" could have, both directly and indirectly (Scanlon, 1998, p. 203). If all of this determines whether a given principle can reasonably be rejected or not within the entire domain of interpersonal morality, then this significantly impacts the extent of any individual's "liberty" or "discretion" as part of exercising her moral judgment, e.g. as part of Voorhoeve's PPP.

So, the combination of the Universality of Reason Judgments and Scanlon's understanding of principles, I take to imply that we can reasonably disagree about the answer to any relevance question for any trade-off case but that we should continue our discussion until, via one (legitimate/non-coercive) way or another, we will have settled on a single answer for any relevance question that we both can think of in advance and as they present themselves through practice (e.g. we might discover some new injury or illness that we have not encountered before and so cannot have settled its relevance to different, familiar injuries or illnesses).[11]

In conclusion, and putting it in contractualist terms, leaving relevance questions open-ended – even within parameters set by "reasonable disagreement" – can reasonably be rejected. We would still be too close to be living in the "non-agreement point" on Quong's model of contractualism, as discussed in my chapter 1, and we would be too far away from having made Scanlonian contractualism sufficiently action-guiding on the crucial question of aggregation and when the numbers count. One non-rejectable way of settling relevance questions is Voorhoeve's PPP. And so, coming full circle, use of Voorhoeve's PPP within Scanlonian contractualism cannot reasonably

[10] Therefore, I will not elaborate as much on this point here as I have in §1 of my first chapter.
[11] Given both this universality claim and both institutional contexts as well non-institutional contexts plausibly being part of interpersonal morality – the subdomain of morality to which contractualism applies – it seems to me that Scanlonian contractualists have to reject Christian List and Laura Valentini's (forthcoming) intriguing proposal that the areas of moral theory and political theory are each concerned with different normative facts. List's and Valentini's proposal is a version of what Liam Murphy (1998) labels "moral dualism" and which Murphy rejects in favour of "moral monism".

be rejected if there is no alternative account on offer which appears at least as plausible as Voorhoeve's PPP. Only if there were at least one such alternative would we be required to at least consider this alternative account. I am unaware of any such alternative account.[12] And even if such an alternative account were on offer at a later stage, then given Scanlon's Principle of Established Practices proponents of this alternative account would face an uphill battle to replace Voorhoeve's PPP which we will have been relying on (at least) until then.

Having discussed the important broad issues of demandingness and whether and how Voorhoeve's PPP could be justified within Scanlonian contractualism, I now turn to objections that have been made to Voorhoeve's PPP as such.

§3. Schönherr's Objection: PPP is Question-Begging

According to Voorhoeve, his solution

> allows for a powerful explanation to a person whose claim is judged irrelevant. For one cannot complain that one's claim is not satisfied by an impartial third party when it would not even fall within one's personal prerogative to satisfy it oneself if no moral considerations apart from the minimally required concern for the stranger's well-being stood in one's way. (Voorhoeve, 2014, p. 74).

Julius Schönherr (2018) presents the following powerful objection to these exact words of Voorhoeve's. Schönherr argues as follows:

> This argument, it seems to me, begs the question. If my own claims aren't strong enough to justify satisfying my own interests straight away, then this could possibly explain why an impartial third party should not satisfy my claim *straight away* or give both claims *equal* weight. However, it does not entail that an impartial third party should not give it *proportional* weight. Analogously, why would it not be permissible for me to give my weak claim *proportional* weight? Why would I not be permitted to flip an *extremely* biased coin in such a conflict situation? The answer, presumably, is that this would be impermissible because my claim is irrelevant. This shows that Voorhoeve's

[12] I can think of plenty worse kinds of alternative accounts than Voorhoeve's PPP. These would be accounts which ground the relevance-distinction on unhelpfully vague and contestable notions such as "respect", "flourishing", "a decent/good human life", "human nature", etc.

> argument *relies* on the idea of irrelevance of certain claims, rather than explaining it. (Schönherr, 2018, p. 217).

Schönherr's objection, if not answered or avoided, appears fatal to Voorhoeve's account. I agree with Schönherr that there is an important 'non-transferral' principle that is relevant from the moral perspective: even if we accept and allow for a personal prerogative, then it is still left as an open question whether or not whatever is permitted or required from my personal perspective is also permitted or required from the perspective of an impartial third party arbitrator.[13] Furthermore, I agree with Schönherr's presumption that holding a weighted lottery in response to the established irrelevance between the competing claims would likely be ruled out on Voorhoeve's account precisely because we have already established these competing claims to be irrelevant to each other. If Voorhoeve were to take this presumed route, then he and his account would of course run into Schönherr's question-begging objection.

But, on Voorhoeve's behalf, I would note that he does not have to take this route of again appealing to the irrelevance between the claims. Voorhoeve could simply acknowledge that his account is incomplete. An independent argument needs to be given to justify the view that we turn to the straightforward satisfaction of the stronger claim after we have established, via the PPP, that the competing weaker claims are irrelevant to this stronger claim. But that's fine. To have the PPP provide the justification for what to do after (ir)relevance between competing unequal claims has been determined would be to hyperextend the PPP. The PPP is and only should be in the business of establishing whether the relationship of relevance obtains between competing unequal claims.

Before the principled use of Voorhoeve's PPP has been established by any community of moral agents as to determine the relation of (ir)relevance between competing unequal claims, we could easily imagine them discussing separately the question of what to do next whenever (ir)relevance between competing unequal claims has been determined. Here, and this would be the alternative route I would suggest for Voorhoeve's account, the argument would be that we need a way to

[13] Schönherr appears to raise this point of 'non-transferral' of moral judgments between the first-personal and third-personal perspective as a separate objection to Voorhoeve's account from the one I have discussed here, see (Schönherr, 2018, pp. 216–217). However, and I expect Schönherr to agree with this, it is hard to strictly separate these two objections – the objection of non-transferral and the question-begging objection. In my view, as I have tried to argue, the question-begging objection is the more concerning objection, but it can be addressed on Voorhoeve's behalf precisely by acknowledging that Voorhoeve's original argument is harmlessly incomplete. "Harmlessly" because it seems quite plausible to me to suggest that an independent argument could be given that we need to have a satisfactory way of resolving these kinds of conflicting claims cases without having to then further consider and solve questions of whether we should directly ignore the irrelevant weaker claims or whether we should still give them some (disproportionate) chance in a weighted lottery.

resolve these difficult competing claims cases. For the sake of efficiency, effectiveness, the possibility of teaching and internalising the moral rule, and, perhaps most importantly, to simply have the conflict resolved, the moral agents would agree that once irrelevance between competing unequal claims has been established, the stronger claim should directly be satisfied. Now, of course, this argument for direct satisfaction of the stronger claim may be objected to, for example, by arguing that efficiency should not have such a decisive moral role to play. However, for my purposes, such a reply simply underscores that Voorhoeve's PPP is not question-begging. As such, I believe I have addressed Schönherr's objection.

§4. Privitera's Objection: PPP is Implausible in Three(+)-Groups-Cases

Johanna Privitera (2018) persuasively argues as follows how Voorhoeve's own rationale for his account presents a problem for cases where there are (at least) three groups:

> The rough rationale says that if competing claims considerably differ in strength we can expect from the person holding the weaker claim that she withdraws her claim, if she is aware of the situation. The idea, here, seems to be that the person holding her weaker claim should understand that instead of helping *her*, the agent should rather help the person holding the stronger claim. She should withdraw her claim, *in order that* the person holding the strongest claim be helped. But this reasoning does not apply to the *C*-persons [the weakest claimants] in three-option-cases. For, they have no reason to think that, if they withdraw their claims, the agent will help person *A* [the strongest claimant] (whose claim takes priority over theirs). Rather, if the *C*-persons withdraw their claims, the agent will help the *B*-persons [the 'mid-strength' claimants], whose claims do not take priority over the *C*-persons' claims. So, in three-option-cases it cannot be expected from the *C*-persons that they withdraw their claims and consequently it cannot be inferred that [helping the *B*-persons] is the required course of action. (Privitera, 2018, pp. 235–236).

I cannot think of a satisfactory rebuttal to this excellent objection. It is puzzling, at least to me, why the shift from two-group-cases to cases where there are at least three groups present such tremendous problems for any version of the Relevance View that I know of. My suspicion is that the workings of something like pairwise comparison on which the Relevance View seems unable to avoid relying on structurally and/or systematically inhibit the Relevance View from adequately dealing with three-(plus)-groups cases. The best response, on Voorhoeve's behalf, is to

acknowledge Privitera's point and concede that the Relevance View only applies to two-group-cases.[14]

Let's take stock of what I have discussed so far. Voorhoeve's PPP offers a way of non-arbitrarily determining whether a relation of relevance holds between competing unequal claims. Schönherr's important concern regarding the question what to do once (ir)relevance has been established serves as a reminder of the point and purpose of Voorhoeve's PPP. The PPP is not and should not be in the business of justifying what to do in response to irrelevant claims. A further argument needs to be provided for completing Voorhoeve's account as an account of which claims to satisfy and why.

Privitera's objection regarding three-(plus)-group cases stands, not just for Voorhoeve's account but I believe for any version of the Relevance View. On Voorhoeve's behalf, I have argued that the point should be conceded: his account with the PPP to determine 'relevance' applies only to two-group-cases. However, although that means the scope of the account is limited, it does not mean that the PPP has been shown to be false.

§5. Tomlin's Objection: Ambiguously 'Anchoring' Claims

In this section, I will briefly outline the ambiguity which Patrick Tomlin (2017) identifies in the Relevance View, and the principles which he shows different versions of the Relevance View to violate. These principles will be used to test refinements to the Relevance View as we go through. In this section, I will also introduce Victor Tadros's (2019) idea of "Local Relevance", which is a response to the problems Tomlin introduces.

As a reminder, for the Relevance View to get off the ground, it needs to distinguish between those claims that are relevant within a choice situation and those that are not. For a claim to be relevant, it needs to be sufficiently strong in comparison with some other claim. Let us call this latter claim the 'anchoring claim'. In standard cases in the literature, it is always clear what the 'anchoring claim' is. For example, in Scanlon's World Cup Case, the 'anchoring claim' is clearly Jones's claim against suffering agonizing shocks, whilst in the Rescue Case the 'anchoring claim' is the claim against death.

[14] For this reason and others, (e.g. cf. Horton, 2018), I believe that so-called "mixed" or 'pluralist' solutions to the aggregation problem for Scanlonian contractualists should be explored. I do this in chapter 5.

Patrick Tomlin (2017) exposes an ambiguity in the Relevance View. He shows that when we consider complex cases involving groups of claims of diverse strength, the Relevance View suffers from an important ambiguity and seemingly fatal counter-intuitive entailments. There are two different anchoring rules we might endorse. In simple cases, like the World Cup Case and Rescue Case, these two anchoring rules identify the same anchoring claim. But when we consider more complex cases, these two rules suggest different answers to the question of which claim is the anchoring claim.[15] These two anchoring rules are:

> Anchor by Strength: in order to be relevant, a claim must be sufficiently strong relative to the *strongest overall* claim in the competition. (Tomlin, 2017, p. 239).

> Anchor by Competition: in order to be relevant, a claim must be sufficiently strong relative to the strongest claim *with which it competes*. (Tomlin, 2017, p. 239).

Tomlin shows that these anchoring rules violate the following compelling principles. Anchor by Competition violates:

> Equal Consideration for Equal Claims: all claims of equal strength ought to be given equal weight in determining which group to save. (Tomlin, 2017, p. 241).

Anchor by Strength violates:

> The Principle of Addition: merely adding a claim to a group of claims cannot *lessen* that group's choice-worthiness, compared with a fixed alternative. (Tomlin, 2017, p. 245).

In addition, Anchor by Strength violates a similar though importantly different principle that Tomlin does not discuss:

> The Principle of Strengthening: merely strengthening a claim within a group of claims cannot lessen that group's choice-worthiness, compared with a fixed alternative.

[15] For Tomlin's discussion of these more complex cases which show that either anchoring rule violates either one of the two compelling principles presented below, see Tomlin (2017, pp. 232–260)

In his (2019) 'Localized Restricted Aggregation', Victor Tadros introduces an important distinction between "Global Relevance" and "Local Relevance". Under Global Relevance, claims that are judged to be irrelevant to an anchoring claim are irrelevant to the overall decision concerning which group to save. By contrast, under Local Relevance, if a weaker claim is judged irrelevant in comparison with an anchoring claim, while the weaker claim cannot counter-balance or outweigh the anchoring claim, it can counter-balance or outweigh other, weaker, claims with which it competes and thus remain part of the overall decision about which group to save. In sum, according to Global Relevance claims are either relevant or irrelevant *simpliciter* to a decision. According to Local Relevance, claims can be relevant to some competing claims but not others. Both Anchor by Competition and Anchor by Strength are Global Relevance views.

The core idea of Local Relevance is interesting, but it admits of many potential interpretations.[16] We know that claims are not relevant or irrelevant *simpliciter*, but this does not tell us how to decide between two groups of claims in messy cases where the groups are made up of claims of diverse strength. It is clear from Tadros' paper that we must 'match up' claims, allowing them to counter-balance, or neutralize, claims to which they are relevant. However, there are many ways that we might go about doing this 'matching'. In the next section, I will introduce a more precise version of Local Relevance, "Sequential Claims-Matching", which specifies how to match claims, before testing it against Tomlin's (2017) principles.[17]

I will show how Sequential Claims-Matching is able to meet the challenges presented by Tomlin (2017). However, Sequential Claims-Matching faces difficulties and ambiguities of its own. In particular, I show that whilst it can deliver intuitive results in Tomlin's cases and Scanlon's World Cup Case and other one vs. many cases, there are two important ambiguities within the view that any would-be advocates will need to address. Finally, I briefly present an alternative version of Local Relevance, "Strongest Decides".

§6. Sequential Claims-Matching

The central insight of Local Relevance is that claims are not to be judged relevant or irrelevant *simpliciter*. This insight is captured by this more precise view, Sequential Claims-Matching.[18]

[16] In response to Joe Horton (2018), Tadros (2019, Section IV) refined his own version somewhat. I will discuss Horton's dilemma for "partially aggregative views" (i.e. Limited Aggregation) in my next chapter.
[17] For an alternative Tomlin-inspired response to Tadros, see Horton (2018).
[18] This was initially suggested by Garrett Cullity. Tomlin (2017, pp. 252–257) suggested a similar 'matching' procedure for 'quasi-competition' cases (cases where some subset of one large group must be saved) but did not consider this

Sequential Claims-Matching provides a procedure through which we can decide which of two competing groups to save. Furthermore, this way of proceeding, in which we start with the strongest claim in the competition, as if it has a pro tanto claim to be saved which must be matched or defeated by claims relevant to it, seems like a natural extension of the Limited Aggregation approach. Limited Aggregation advocates seek to walk the middle ground between Anti-Aggregation and Pure Aggregation, and to combine personal and impersonal perspectives.[19] According the strongest claim a pro tanto claim to be saved seems to do justice to the personal, Anti-Aggregative, perspective.

Sequential Claims-Matching:

I. Identify the strongest claim-type T1. Does one group contain more individuals with claims of type T1 than the other?

   If not, remove all T1-claims from consideration, and proceed to Step V.

   If so, match each T1-claim from the group with fewer T1-claims to a T1-claim from the group with more T-1 claims, and remove the matched claims from consideration.

II. Now consider the remaining T1-claims. Does the other group contain claims of types that are relevant (i.e. sufficiently strong relative) to claim-type T1?

   If not, you should decide in favour of the group with the remaining T1-claims.

   If so, proceed to Step III

III. Do the relevant competing claims outweigh the T1-claims?

   If not, you should decide in favour of the group with the remaining T1-claims.

   If so, proceed to Step IV.

IV. Match the set of remaining T1-claims to a set of relevant competing claims with comparable weight, and remove the matched claims from consideration.

V. Now consider the remaining unmatched claims. Of these, identify the strongest claim-type T2. Repeat the above procedure.

VI. Continue until either:

---

kind of procedure for 'competition cases' (where we must save one group or the other). Tomlin argues that quasi-competition cases are actually more analogous to many policy decisions.

[19] For similar approaches in the literature that look to combine and/or balance personal and impersonal perspectives, see Kamm's (five different versions of) "Sobjectivity" in her chapters 8-10 (1998, pp. 144–195; 2007, pp. 34–36, 39, 50) and Thomas Nagel's (1989; 1995) "View from Nowhere" vs. "View from Somewhere." Alan Thomas (2009), in his essential guide to Nagel's entire philosophical body of work, identifies the dichotomy between a subjective and an objective point of view as "the fundamental theme of [Nagel's] philosophy." (Thomas, 2009, p. 2).

a. one group contains unmatched claims, in which case you should decide in favour of that group; or
b. neither group contains unmatched claims. Then it is not the case that you should decide in favour of one group over the other (though you must save one).[20]

Sequential Claims-Matching is, as you can see, complicated. Nevertheless, the key idea is fairly simple: can the anchoring claim be matched by claims relevant to it? If not, we should meet the anchoring claim. If so, the anchoring claim is 'matched' to claims relevant to it, and all those claims are then set aside, since they counter-balance one another.[21] We now identify the strongest remaining claim, that becomes the new anchoring claim, and the process begins again.

Before testing Sequential Claims-Matching against Tomlin's principles identified above, it may be helpful to show how it works in a simple case. This will also allow me to introduce the way I will present cases in this chapter. In the following table, on the left-hand side is the strength level of the claim, Level 1 being the strongest. Under headings Group A and Group B are the numbers of claims at each level within each group, each claim being held by a distinct individual. For shorthand, I will refer to claims at Level 3 in Group B as B3 claims.

Suppose that relevance extends two levels up.[22] That is, Level 3 claims are relevant to Level 1 claims, but Level 4 or Level 5 claims are not relevant to Level 1 claims. Furthermore, imagine that two claims at Level *X* precisely match one claim at the level above. So, for example, one claim at Level 1 will be precisely matched by two claims at Level 2, and those will be precisely matched by four claims at Level 3. Unless I state otherwise, these stipulations also apply to all cases hereafter. Here is the initial case:

[20] I leave it open here whether in such a situation an agent has a free choice, or whether they must choose via a fair procedure such as a coin toss.
[21] 'Set aside' here does not mean 'cancelled'. If the claims were cancelled, then two equally strong claims would cancel each other and nobody would have a claim to be saved. I am grateful to Mike Otsuka for useful comments here.
[22] Although I don't have to rely on it for my argument here, we could suppose that these stipulated relevance-relations have been non-arbitrarily justified via Voorhoeve's "permissible personal perspective".

Case 1.

Level	Group A	Group B
1	1	
3		5
5	3	

Here is how Sequential Claims-Matching would handle this case.

First, it would identify A1 as the initial anchoring claim. There are no B1 claims to match A1 with, so we see whether there are weaker but relevant claims in Group B. There are – the five B3 claims. Four of these B3 claims will precisely match the A1 claim, so we take these four B3 claims and the A1 claim which are then set aside, as they counter-balance one another. The one remaining B3 claim then becomes the new anchoring claim. There are no A3 claims with which to match the remaining B3 claim, and so we look to see if there are weaker but relevant claims in Group A. There are – the three A5 claims. However, three A5 claims do not match one B3 claim (since four Level 5 claims precisely match one Level 3 claim), and so we would save Group B in this case.

Sequential Claims-Matching has attractive implications in key cases that Tomlin uses to undermine Anchor by Competition and Anchor by Strength. Against Anchor by Competition, he shows that adding equal numbers of equally strong claims to two groups will force us to switch groups (2017, pp. 240–241). Even stranger, when one additional claim is added to one group (Group A), and a billion equally strong claims are added to another (Group B), Anchor by Competition may require us to switch from saving Group B to saving Group A (Tomlin, 2017, p. 242). Both of these counter-intuitive entailments result because Anchor by Competition violates Equal Consideration for Equal Claims: it allows that claims in one group will be relevant, while claims of equal strength in another will not, as they compete with different 'anchoring claims.'

Sequential Claims-Matching does not have these odd implications in the cases that Tomlin considers. This is because it allows that as soon as claims of strength *X* in one group are relevant, all such claims are relevant.

Tomlin shows that Anchor by Strength has even odder implications: adding a very strong claim (such as a claim against death) to a group of weaker claims would force us to switch away from saving that group, even if the competing group is not altered at all (2017, pp. 244–247). This

violates the Principle of Addition. Again, Sequential Claims-Matching would avoid this implication. Since relevance is determined by whether or not a claim is relevant to claims it competes with, adding a single claim could not suddenly rule as irrelevant claims within its own group.

Thus far we see that Sequential Claims-Matching is a better-specified version of Tadros' Local Relevance view, and that it is preferable to both Anchor by Competition and Anchor by Strength in that it provides more intuitively attractive judgments in the key cases presented by Tomlin (2017). However, Sequential Claims-Matching faces difficulties and ambiguities of its own.

§6.1. Ambiguity 1: How Should We Match Claims in Sequential Claims-Matching?

I noted above that Tadros's explanation of the Local Relevance view seems to require us to 'match up' claims with one another, so that they can counter-balance other claims to which they are relevant. Sequential Claims-Matching is a clear advance on the loosely specified Local Relevance View because it clearly outlines a process by which we should match up claims. Nevertheless, it still contains an ambiguity concerning how to 'match up' claims. The relevant step within the Sequential Claims-Matching view is as follows:

IV. Match the set of remaining T1-claims to a set of relevant competing claims with comparable weight, and remove the matched claims from consideration.

This step shows how anchoring claims must be compared with claims that are relevant to them. However, more than one kind of claim can be relevant to a single anchor, and in such a case we need a procedure to decide which of the weaker claims to compare with the anchor. Consider this case:

Case 2.

Level	Group A	Group B
1	1	
2		3
3	2	
4	4	
5		
6		20

> Case 2 Upshot: A B2 claim can be matched with either A3 or A4 claims. Whether the B2 claim is matched with A3 or A4 claims determines which group is saved. Sequential Claims-Matching provides no guidance as to which is preferable.

The A1 claim would be the initial anchoring claim. This claim is matched by two B2 claims, leaving the remaining B2 claim as the new anchoring claim. Here is where the ambiguity concerning how to match up claims comes in. Both the A3 and A4 claims are relevant to the B2 claim. Either the two A3 claims or the four A4 claims would exactly counterbalance the one remaining B2 claim. However, whether we match the B2 claim either with the two A2 claims or the four A4 claims will decide which group we will save.

Suppose, first, that we match the B2 claim with the two A3 claims. These are all the A3 claims there are, and so we move one level down in order to identify the next anchoring claim. The A4 claims become the new anchoring claims. What is crucial here is that the twenty B6 claims are relevant to the A4 claims and outweigh them. So, matching in this way means that we would save Group B.

But now let's suppose that when we have the remaining B2 claim as the anchoring claim that we match it with the four A4 claims. That would mean that the only claims that are left in the competition are two A3 claims and twenty B6 claims. But B6 claims are not relevant to A3 claims. There could be any number of B6 claims but they could never outweigh a single A3 claim. So, the two A3 claims are unmatched. Therefore, were we to match claims in this way, we would save Group A.

This case shows that there is an ambiguity in how we match claims on Sequential Claims-Matching that could be decisive in deciding which group to save. What Sequential Claims-Matching requires, therefore, is a principled way in which to determine how we match claims in this kind of case. Here are three rival possibilities, all of which have some plausibility, but which would give conflicting advice in Case 2.

The first possibility is the following:

> Match to the Strongest Competing Claim: Anchoring claims should be matched with the strongest unmatched claims with which they compete.

This rule has some plausibility. Sequential Claims-Matching in general has a structure which encourages us to look at the 'next level down' – that is, after all, the next claim in the sequence. It also seems a non-arbitrary and clear way to decide how to match up the claims. In Case 2, Match to the Strongest Competing Claim would have us match the B2 claim to the A3 claims, leaving the A4 claims to be outweighed by the B6 claims. Therefore, we would save Group B.

The second possibility is the following:

> Match in the Interest of the Overall Strongest Claimant: Anchoring Claims should be matched in whatever way is in the interest of strongest overall claimant.

Limited Aggregation generally has, of course, much in common with Anti-Aggregation. Anti-Aggregation is concerned, above all else, with the person who possesses the strongest claim. Limited Aggregation doesn't hold this concern above all else – it would allow the weaker claims to outweigh the strongest claim in death vs. paralysis – but it is plausible to think that it would inherit this general priority to the worst off. Match in the Interest of the Overall Strongest Claimant says that when there is a plurality of ways of matching up claims within the confines of Sequential Claims-Matching, and where different ways of matching would produce different outcomes (as in Case 2) we should, in essence, allow the person possessing the overall strongest claim in the competition to choose how we should match up the claims. In Case 2, A1 is the strongest claim. If we match the B2 claim with the A4 claims, this leaves the A3 claims unmatched, since they are 'out of reach' of the B6 claims. And so, Group A would win, which is, of course, what A1 would want.

The third possibility is as follows:

> One or the Other: when there is an ambiguity over how to match up claims, and different ways of doing so would require saving different groups, then it is not the case that you should decide in favour of one group over the other (though you must save one).

This view takes the ambiguity we have noted in this section not as something to be rectified, but rather as to reveal that we are not required to save one group over the other. Since the Sequential Claims-Matching process can be used to justify saving either group, we cannot say that we are

required to save one group over the other. In such circumstances, it may be that we are permitted to choose, or are required to toss a coin.

If Local Relevance is to become a fully fleshed out answer to the aggregation problem, we need an account of how to match up claims with one another. Sequential Claims-Matching provides such an account. However, there remains an ambiguity in that even within the strictures of that more precise theory, there can be more than one way to match up claims. So, advocates of Sequential Claims-Matching would need to further specify the view.

§6.2. Ambiguity 2: Partial Leftovers

Thus far I have been considering cases in which two claims at one level precisely match one claim at the level above. However, it is unlikely that things will always be quite so neat, especially given how small the differences between claims can be. For example, even if one accepts that two claims at Level 2 will precisely match one claim at Level 1, imagine if one of the Level 2 claims becomes ever so slightly stronger (for example, if it is a claim against partial paralysis, the loss of use of one more finger). Does this remain a Level 2 claim with no additional weight? Surely not. But equally implausible is that it would become a Level 1 claim, and double in weight.

To put things more concretely, many health systems use Quality Adjusted Life Years (QALYs) or Disability Adjusted Life Years (DALYs) to measure the strength of claims. Imagine a case in which a single person facing the loss of 50 QALYs faces three people facing the loss of 20 QALYs each. The single strongest claim outweighs two of the weaker claims and is defeated by three of the weaker claims. 2.5 of the weaker claims precisely match the stronger claim.[23]

If whole numbers of weaker claims do not precisely match stronger claims, we need to decide how to handle the 'partial claims' that are left over once the matching has taken place. To make things clearer, consider this case. In this case 2.5 claims at Level *X* precisely match one claim at the level above.

[23] For related discussion see Kamm (2005), section VII (15–18). Kamm's discussion clearly indicates that lesser claims must compete with left-overs but is ambiguous between my Full Claim Relevance and Partial Relevance as spelled out on the next page.

Case 3.

Level	Group A	Group B
1	1	
2		3
3		
4		
5	20	

> Case 3 Upshot: Following Sequential Claims-Matching, 0.5 of a B2 claim is left over from the first round of 'matching' and becomes the new anchor. But it is ambiguous in Sequential Claims-Matching whether the twenty A5 claims need to be relevant to that 0.5 of a B2 claim or to a full B2 claim. Which route we take, determines whether the twenty A5 claims are relevant and so which group we save.

A1 is the initial anchoring claim. It is matched by 2.5 B2 claims, seemingly leaving half a B2 claim as the new anchoring claim. There is an ambiguity about what kinds of claims should be considered relevant to this partial B2 claim. The key question is whether the twenty A5 claims are relevant to the half B2 claim. If it were a full B2 claim, then there would be no ambiguity, since Level 5 claims are not relevant to Level 2 claims. However, in this case, there is half a Level 2 claim remaining. In order to be relevant, should we demand that the A5 claims are relevant to a full B2 claim or a half B2 claim? This suggests two possible principles.

> Full Claim Relevance: To be relevant, a claim must be relevant to the claim with which it competes. (For example, the full B2 claim.)

> Partial Claim Relevance: To be relevant, a claim must be relevant to the proportion of the claim with which it is in competition. (For example, the half B2 claim.)

On the one hand, we could imagine one of the B2 claimants objecting to small A5 claims being considered relevant to her far weightier claim. On the other hand, it seems unfair to expect A5 claims to have to be relevant to a full B2 claim when they are only in competition with half a B2 claim.

Partial Claim Relevance seems, to me, the right view. However, this raises the issue of very weak claims (e.g. claims to be spared from a sore throat) breaking a tie.[24] Imagine a case in which there is a partial leftover of 1/1,000 of a very serious claim. Very weak claims, which would only need to be relevant to (and not necessarily equal to) the leftover, could then tip the balance. This supports the case for advocates of Partial Aggregation letting go of the tie-break intuition.

§7. Strongest Decides: An Alternative Version of Local Relevance

I began by highlighting how Tadros's Local Relevance was promising. I have sought to provide a more concrete version through Sequential Claims-Matching and have discussed at length the merits and problems of this view. One ambiguity I have highlighted (in §6.1) concerns how to match up claims with one another. One suggestion offered was to Match in the Interest of the Overall Strongest Claimant. However, I have only considered this within the confines of Sequential Claims-Matching. But why not match in the interest of the overall strongest claimant more generally? So, whenever we face cases with groups of claims of diverse strength, if there is any way to match the claims such that the strongest claimant would have her claim met, then we should allow this. The basic idea here is that we take from Anti-Aggregation the priority for the worst off and we allow that to dictate the way in which claims are matched. Call this view "Strongest Decides". Here is an example where this view would deliver a different verdict from Sequential Claims-Matching:[25]

Case 4.

Level	Group A	Group B
1	1	
3	3	3
5		10

Case 4 Upshot: Sequential Claims-Matching would tell us to save Group B. Strongest Decides, by contrast, would tell us to save Group A.

[24] For a detailed discussion of tie-breaker cases, such as Kamm's Sore Throat Case, and how Kamm's intuition that a sore throat could never break a tie cannot be defended on Sequential Claims-Matching, see van Gils and Tomlin (forthcoming).

[25] I draw here on cases presented in Horton (2018), though we put them to different ends here. I will discuss Horton's (2018) account in detail in my next chapter.

Sequential Claims-Matching would first match two B3 claims with the A1 claim, leaving one B3 claim as the new anchoring claim. That would be matched by one A3 claim, leaving two A3 claims. Finally, these two A3 claims would be outweighed by the ten B5 claims. So, we would save Group B.

Strongest Decides, by contrast, would have the A1 claimant decide how we match claims in this case. It is in A1's interest to have the three A3 claims be matched with the three B3 claims, as such taking all these claims out of consideration. This would leave A1's own claim and ten B5 claims. Level 5 claims are not relevant to a Level 1 claim, and so we would save Group A.

Here is a potential problem with Strongest Decides: it may violate Equal Consideration for Equal Claims.[26] It is clear that, in some cases, there will be more than one way of 'matching up' claims, and the way that would be favoured by the person with the strongest claim may be one in which some, say, Level 5 claims are matched with claims to which they are relevant, while other Level 5 claims are matched with claims to which they are not relevant. Therefore, in the 'matching scheme' the strongest claim would select, some Level 5 claims will be counted (as they are relevant to the claims they are matched with) and others will not (as they are not relevant to the claims they are matched with).

While that much is clear, it is not clear whether this in fact violates Equal Consideration for Equal Claims. Tomlin makes it clear that any plausible view will not advocate Equal Treatment for Equal Claims, for when we have two groups with equal claims on both sides, we will end up saving one group and not the other (2017, p. 241). What Tomlin objected to, and Equal Consideration for Equal Claims attempted to articulate, was that at the outset some views ruled that some claims of a certain strength mattered, and other claims, of the same strength, did not. Strongest Decides does not do this.

Strongest Decides only really comes into play when there is a plurality of ways in which to 'match up' claims with one another. Since equal claims can always be matched to each other, claims of equal strength will never be ruled as irrelevant at the outset: there is always a potential 'match up' in claims in which all claims of the same strength are relevant. The strongest claim can dictate which 'match up' is chosen as a tie breaker. This may involve choosing a 'match up' in which some

[26] I am grateful to Bastian Steuwer for pressing this worry.

Level 5 claims are relevant and others are not. Is this like accepting that we will save some Level 5 claims and not others, which is surely acceptable? Or is it like ruling out some Level 5 claims but not others at the outset? Or, is it like neither? Equal Consideration for Equal Claims may need to be further refined. Or perhaps it must be abandoned. Perhaps it didn't properly articulate what Anchor by Competition was getting wrong. At any rate, this is an issue that Strongest Decides must confront.

Conclusion

In this chapter, I have further explored how to avoid biting the bullets of the extreme Anti-Aggregation and Pure Aggregation positions. I began by discussing Alex Voorhoeve's ARC. I investigated whether Scanlonian contractualists can endorse ARC and the PPP specifically. I have argued that Scanlonian contractualists face a tough challenge in justifying Voorhoeve's PPP within their contractualism because the PPP is a version of "partiality" and it is not at all clear that Scanlonian contractualists could justify partiality within their account, certainly not with the ease that Scanlon himself suggests. So, Scanlonian contractualists need to revise their account such that either some plausible version of partiality could not reasonably be rejected which would automatically make room for Voorhoeve's PPP within contractualism or justify an alternative account that grounds "relevance" which does not rely on partiality like Voorhoeve's PPP does. For the purposes of further exploring the Relevance View and aggregation, I have presupposed there is a way Voorhoeve's PPP could be justified within Scanlonian contractualism. Obviously, more work needs to be done here by or on behalf of Scanlonian contractualists, but I could not further pursue this here because that would make "partiality" rather than "aggregation" the main topic of my thesis.

I then defended Voorhoeve's ARC against two recent criticisms. However, ARC is a version of what Tadros calls Global Relevance, as Tomlin has shown Global Relevance is open to potentially devastating objections. Tadros's version of Limited Aggregation, "Local Relevance", is attractive, but can be fleshed out in a variety of ways. It is this idea that I have sought to precisify and explore. I have done this by, first, providing a clearer and more tightly specified version of the idea, Sequential Claims-Matching, and, second, by subjecting that view to scrutiny. I have found that Sequential Claims-Matching stands up well against alternative versions of Limited Aggregation in terms of Tomlin's (2017) cases.

However, while Sequential Claims-Matching is clearer than Tadros's view on how to match claims, there are two important ambiguities within Sequential Claims-Matching: an ambiguity about how to match up claims, and a further ambiguity about how to handle what I have called partial left-overs. Thus, I have shown that Sequential Claims-Matching, at best, needs to be further specified. In addition, I have articulated an alternative version of Local Relevance: Strongest Decides. This view tracks more closely to the Anti-Aggregation view, giving the strongest claim under consideration a defeasible right to be saved. However, this view potentially conflicts with the Equal Consideration for Equal Claims principle. I recommend further scrutiny of this view, and the principle it seems to violate.

In conclusion, I offer an interim report on Local Relevance. The general idea admits of several possible interpretations. We need a clear sense of which versions appear to be the most plausible, both in terms of how they handle key cases, and the sort of deeper justifications which can be offered on their behalf. Even within these versions, there are ambiguities, for example, on how to handle partial leftovers.

But should we even attempt to answer these further questions that arise for Local Relevance and the specific version of Sequential Claims-Matching I have presented? Thanks to Joe Horton's (2018) work we might think that there is no need to bother with these further questions. This is because Horton presents a dilemma for any version of Limited Aggregation, which I will discuss in my next chapter.

University of Reading

A (Non-) Fatal Dilemma for Relevance and a Prioritarian 'Escape'

Abstract This chapter presents Joe Horton's (2018) seemingly fatal dilemma for any partially aggregative view (i.e. Limited Aggregation). I argue that Horton's dilemma fails. But, for the sake of argument, I continue by granting Horton that neither version of Local Relevance that I have considered in my previous chapter, Strongest Decides and Sequential Claims-Matching, could escape his dilemma. If we presume Horton's dilemma to hold, could Scanlonian contractualists nevertheless maintain an alternative plausible position on aggregation? I argue they can. Some versions of prioritarianism, although they are purely aggregative views, fit nicely with the Scanlonian contractualist rationale. In the second half of this chapter, I make the case for prioritarianism by contrasting it with its main non-utilitarian rival: egalitarianism. For the general discussion of egalitarianism vs. prioritarianism, I rely on the egalitarian accounts as defended by Michael Otsuka and Alex Voorhoeve (2009; 2018) and Voorhoeve and Marc Fleurbaey (2012), and the prioritarian accounts as defended by Derek Parfit (1997; 2000; 2012) and Thomas Porter (2011; 2012). Having set the parameters of the general debate on egalitarianism vs. prioritarianism, my next step is to compare the two most sophisticated versions of either view: Andrew Williams's (2012) "Restricted Prioritarianism" vs. "the Competing Claims View" as defended by Fleurbaey, Otsuka, and Voorhoeve. Against Benjamin Lange (2017), I argue that the debate between these two views has not been settled – specifically, we lack sufficient justification for the Competing Claims View being better than Restricted Prioritarianism. I contend that Restricted Prioritarianism is more plausible than the Competing Claims View because it is more practically feasible. However, that also means that the proof is in the pudding: only real-world application and evidence can settle the matter, bar an as-of-yet undiscovered/devised philosophical counterexample to either view.

Keywords Egalitarianism • Horton • Prioritarianism • Relevance • Voorhoeve

Introduction

This chapter has two aims. First, I discuss Joe Horton's (2018) dilemma for any partially aggregative view (i.e. Limited Aggregation).[1] The dilemma applies to all versions of the Relevance View as I have discussed them so far in my thesis. I explain why I think that Horton's dilemma fails: it is guilty of a fallacy of composition and presumes the Relevance advocate makes an assumption which she in fact rejects. Although it is hard to pinpoint what the exact assumption is, it is an assumption closely related to the Relevance View advocate's rejection that "ought to be saved" is transitive. Despite lacking an exact label for this assumption, I think that discussion of it using exemplary cases below will make the point sufficiently clear. However, for the sake of argument, I will continue by presuming that neither Sequential Claims-Matching nor Strongest Decides, as discussed in the previous chapter, escape Horton's dilemma.

Given the presumed fatal blow Horton's dilemma delivers to any partially aggregative view, the second aim of this chapter is to consider what defeat would look like for Scanlonian contractualists if they had to concede that they should have a purely aggregative view in the context of distributive justice. My main point here is to argue that having to endorse a purely aggregative view might not be as bad for Scanlonian contractualists as they might have initially feared. Utilitarianism is the most famous version of a purely aggregative view but accepting Pure Aggregation does not entail accepting utilitarianism. Prioritarianism is another version of a purely aggregative view and I will argue that if no partially aggregative view can be reasonably defended, then Scanlonian contractualists should embrace prioritarianism.

For the general discussion of egalitarianism vs. prioritarianism, I rely on the egalitarian accounts as defended by Michael Otsuka and Alex Voorhoeve (2009; 2018) and Voorhoeve and Marc Fleurbaey (2012), and the prioritarian accounts as defended by Derek Parfit (1997; 2000; 2012) and Thomas Porter (2011; 2012). Having set the parameters of the general debate on egalitarianism vs. prioritarianism, my next step is to draw on Benjamin Lange's (2017) work and compare the two most sophisticated versions of either view: Andrew Williams's (2012) "Restricted Prioritarianism" vs. "the Competing Claims View" as defended by Fleurbaey, Otsuka, and Voorhoeve. Against Lange, I argue that the debate between these two views has not been settled – specifically, we lack sufficient justification for the Competing Claims View being better than Restricted Prioritarianism. I contend that Restricted Prioritarianism is more plausible than the

[1] So, what I call "Limited Aggregation", Horton labels as "Partial Aggregation" or "partially aggregative views". In the following I use these interchangeably. Nothing hangs on the label here.

Competing Claims View because it is more practically feasible. However, that also means that the proof is in the pudding: only real-world application and evidence can settle the matter, bar an as-of-yet undiscovered/devised philosophical counterexample to either view.

§1. Horton's 'Fatal' Dilemma for Partially Aggregative Views

Joe Horton (2018) presents a series of four cases in which advocates of any partially aggregative account cannot escape the choice between two seemingly implausible options. For all cases, we deal with the following harms: death, lost arm, lost finger. The further assumption, to get a partially aggregative view going, is that lost finger is relevant to lost arm, lost arm is relevant to death, but lost finger is not relevant to death. Finally, the ratios between the three different harms are key: 20 complaints against a lost arm outweigh 1 death-complaint, and 4,000 complaints against a lost finger outweigh 20 complaints against a lost arm (Horton, 2018, p. 169). The first case is Case 6, which has two stages (Horton, 2018, p. 168):[2]

> *Stage 1*: You can save either group A, which contains 1 person facing death, or group B, which contains 4,000 people facing a lost finger.
>
> *Stage 2*: 20 people facing a lost arm are added to each group.

In this case, Horton rightly suggests that "partially aggregative views imply that you should save A at Stage 1." (2018, p. 169). Should we expect this verdict to change at Stage 2? A change in verdict would seem strange given that it is the exact same number of people, 20, who all face the same prospective harm, a lost arm, that are added to either group. Horton suggests that "if you should save A at Stage 1, [then] it must be at least permissible to save A at Stage 2." (2018, p. 169). The problem is how Local Relevance could capture this judgment. According to Local Relevance, if a weaker claim is judged irrelevant in comparison with an anchoring claim, while the weaker claim cannot counter-balance or outweigh the anchoring claim, it can counter-balance or outweigh other, weaker, claims with which it competes and thus remain part of the overall decision about which group to save. Local Relevance could be applied in at least two of the following ways (Horton, 2018, p. 169):

[2] Meaning that Horton obviously discusses five cases preceding this case. However, to avoid unnecessary confusion and make my reference to Horton as clear as possible, I have stuck with his case numbers. So, I will discuss Horton's Cases 6-9 (Horton, 2018, pp. 168–171).

Option 1:

Group A	Group B
1 person facing death	4,000 people facing a lost finger
~~20 people facing a lost arm~~	~~20 people facing a lost arm~~

Option 2:

Group A	**Group B**
1 person facing death	~~2,000 people facing a lost finger~~
~~20 people facing a lost arm~~	2,000 people facing a lost finger
	20 people facing a lost arm

The strikethroughs in both options show which sets of competing claims are the first to be balanced and set aside. Under Option 1, we first balance and set aside the same number of the same claims: 20 complaints against a lost arm in Group A vs. 20 complaints against a lost arm in Group B. Under Option 2, we pick the same set of claims in Group A: 20 complaints against a lost arm. But for the set of competing claims we now instead pick 2,000 out of the 4,000 lost-finger claims. So, under Option 2, the first sets of competing claims we balance and set aside are 20 complaints against a lost arm from Group A vs. 2,000 complaints against a lost finger from Group B. These are at least two options of matching claims that we have according to Local Relevance.

Horton notes that if Local Relevance is applied per Option 2, then it has the implausible implication that we should switch from saving group A at Stage 1 to saving group B at Stage 2 because the 20 lost-arm claims in B outweigh the one death-claim in A (2018, p. 169). So, to avoid this result, Local Relevance must provide us with a reason to match claims as per Option 1. Horton merely supposes for the sake of argument that an advocate of Local Relevance could provide us with a justificatory reason as to why we should follow Option 1 and not Option 2. I have presented the view of Strongest Decides in the last chapter as providing a rationale which would justify Option 1: the strongest claimant should take moral priority and so she decides how claims in the competition should be matched. In Case 6, she would decide to match according to Option 1 and not Option 2 because that would have the result that her group A would be saved.

But whether Strongest Decides is considered plausible is not essential here. Horton's supposition that his Local Relevance opponent could somehow provide a reason to follow Option 1 is enough to kickstart his following three cases, Cases 7-9 (Horton, 2018, p. 169):

> *Case 7*: You can save either group C, which contains 1 person facing death, or group D, which contains 20 people facing a lost arm.
>
> *Case 8*: You can save either group E, which contains 20 people facing a lost arm, or group F, which contains 4,000 people facing a lost finger.
>
> *Case 9*: You can save either group C + E or group D + F.

Horton points out that Local Relevance implies saving group D in Case 7 and group F in Case 8 (2018, pp. 169–170). Therefore, surely Local Relevance implies saving group D + F in Case 9, right? Not so, as Horton shows when presenting the choice we face in Case 9 (2018, p. 170):

Group C + E	Group D + F
1 person facing death	4,000 people facing a lost finger
~~20 people facing a lost arm~~	~~20 people facing a lost arm~~
Group C + E	**Group D + F**
1 person facing death	~~2,000 people facing a lost finger~~
~~20 people facing a lost arm~~	2,000 people facing a lost finger
	20 people facing a lost arm

The problem is that group C + E is equivalent to group A in Case 6 and group D + F is equivalent to group B in Case 6 (Horton, 2018, p. 170). The point of Case 6 was to stick with group A (= group C + E) at Stage 2, or to have it at least remain permissible to save group A at that second stage in that Case 6. But now, here in Case 9, on behalf of Local Relevance advocates, we should be desperate to avoid the implication of even being permitted to save group C + E (= group A). Why? It would turn these two 'loser-groups' – C lost in its competition with D and E lost in its competition with F – into 'winner-groups'. But surely D and F should remain the winners in Case 9? After all, all that has been done is that these groups have been grouped together. As Horton argues, "it would be inconsistent to apply Local Relevance [according to Option 1] in Case 6 and [according to Option 2] in Case 9. So Local Relevance must have an implausible implication in

one of these cases." (Horton, 2018, p. 170). As a result, not just Local Relevance, but, Horton argues, "all" partially aggregative views face the following dilemma (2018, p. 170):

> To avoid an implausible implication in Case 6 [switching to save Group B at Stage 2] (this is the first horn of the dilemma), a partially aggregative view must balance claims [according to Option 1]. To avoid an implausible implication in Case 9 [saving group C + E] (this is the second horn of the dilemma), a partially aggregative view must balance claims [according to Option 2]. But it is inconsistent to balance claims [according to Option 1] in Case 6 and [according to Option 2] in Case 9. So all partially aggregative views must have an implausible implication in one of these cases. (Horton, 2018, p. 170).

What about Sequential Claims-Matching? How does this (more) precise version of Local Relevance deal with Horton's dilemma? Advocates of this view can neither avoid Horton's dilemma, nor can they choose what horn they could fall on – insofar as they wanted a choice in that matter. Sequential Claims-Matching must fall on the first horn of Horton's dilemma. Why?

Sequential Claims-Matching would require the switch from saving group A at Stage 1 to saving group B at Stage 2 in Case 6. This is because, at Stage 2, the death-claim in A is the initial anchoring claim and is outweighed by the 20 lost-arm claims. Suppose that it takes 19 lost-arm claims to exactly balance a single death-claim. The one remaining lost-arm claim becomes the new anchoring claim. This claim is obviously outweighed by 20 claims of its own kind in A – lost-arms claims balance each other 1:1 – so that leaves 19 lost-arm claims in Group A as the new anchoring claims. These 19 lost-arm claims are then outweighed by the 4,000 lost-finger claims in B. This is how Sequential Claims-Matching implies a switch from A to B in Case 6.

Regarding Case 9, Sequential Claims-Matching, therefore, also implies saving group D + F:

Group C + E	**Group D + F**
~~1 person facing death~~	~~20 people facing a lost arm~~
~~20 people facing a lost arm~~	~~some people facing a lost finger~~
	some people facing a lost finger

I have not made the number of lost-finger claims that remain after sequential matching precise, but that does not matter for the main point. The main point is that there will be some number of lost-finger claims left that will carry the day and imply that group D + F should be saved.

In the choice between Strongest Decides and Sequential Claims-Matching, it seems to me that Local Relevance advocates would be wise to choose Sequential Claims-Matching over Strongest Decides. In discussing Horton's Case 6, I indicated that out of the two options presented, an advocate of Strongest Decides would (have to) take Option 1 because picking that option ensures that the person with the strongest claim is saved. But, to avoid inconsistency, this means that Strongest Decides entails saving Group C + E over Group D + F in Case 9. And this second horn of Horton's dilemma looks terrible.[3]

However, despite its terrible look, I am not convinced that advocates of the Relevance View cannot escape this second horn, and so escape Horton's dilemma. This is because Horton implicitly appeals to what seems very closely related to the transitivity of "ought to be saved". However, an advocate of Local Relevance has already rejected the transitivity of "ought to be saved over". Patrick Tomlin illustrates the point in terms of "ought to be chosen over":[4]

> "[O]ught to be chosen over is not even the sort of thing we would *expect* to be transitive, and therefore much less is at stake in rejecting transitivity in this context. Indeed, even utilitarians should reject transitivity here. Imagine I must choose whom to give some good to, candidate A, B, or C. Considering only the happiness of noncandidates, we would rank them as follows: A > B > C. But it will make C very, very unhappy to lose out to A in particular, so much so that this outweighs the noncandidates' preference for A over C. Therefore, in pairwise competitions, according to utilitarianism, I should choose C over A, A over B, and B over C. The utilitarian may deny that "Choosing A and so upsetting C" is *the same option* as "Choosing A over B." But this is to invoke a moralized way of individuating options, and the proponent of the Relevance View can do the same, thereby preserving transitivity. For example, this is Voorhoeve's response [(2014, pp. 78–79)]. The

[3] Horton vividly illustrates this by asking us to imagine the choice between C + E and D + F in Case 9 as a button-pressing case where what you ought to do is determined by whether your arm span is sufficiently long to simultaneously press buttons C and E, see (Horton, 2018, p. 173).
[4] To be clear, nothing hangs on "saved" vs. "chosen" here. What matters is the structure of the relation "ought to be [blank] over".

> disagreement therefore comes down to what is morally relevant, not whether one theory violates transitivity whereas another does not. (Tomlin, 2017, pp. 236–237, fn. 11).[5]

Although not quite the transitivity of "ought to be chosen/saved over", Horton relies on a very similar transitive relation: "You ought to choose D over C, and F over E, therefore we ought to choose D and F over C and E." But because Horton implicitly relies on the transitivity of this slightly more complex "ought to choose/save over", we, in turn, can adjust Tomlin's counterexample to fit the structure of Horton's dilemma and so respond to it. Here is such a counterexample:

> As the coach of a football team, you ought to choose A over B to be your centre forward, and C over D. However, D is A's younger brother, and loses out to his elder brother a lot. He is young and has a bright football career ahead of him. But it will crush D's football hopes and dreams to pick his brother over him. So, you should not choose A and C over B and D.

To emphasize, what matters here is the structure, not the content of the example. However we flesh out the example, to reject the transitivity of "ought to be chosen/saved over" certainly is not as implausible as it initially appears in Horton's dilemma. Furthermore, it is a defining feature of the Relevance View that it denies that this kind of "ought to be saved/chosen over" relation is transitive. Finally, to presume the transitivity of this relation is essential for Horton's dilemma to work. So, with all these pieces into view, we can say that advocates of the Relevance View can avoid the second horn of Horton's dilemma, and so Horton's dilemma for the Relevance View fails.

There might be one or several other reasons why, as Horton puts it, we should "always aggregate". However, neither his dilemma nor the denial of the transitivity of "ought to be chosen/saved

[5] Earlier in the same footnote, Tomlin contrasts the version of the Relevance View which I have discussed and which rejects the transitivity of "ought to be chosen/saved over" with an "evaluative version". For the "evaluative" version of the Relevance View, as Tomlin notes, we must deny the transitivity of "all things considered better than". Larry Temkin (2012) at the least considers biting this significant bullet. I don't think Temkin or anyone else should bite this bullet. In his introduction, Temkin, in his characteristic self-deprecating sense of humour, shares an anecdote that when he first presented his idea of denying the transitivity of "all things considered better than" to his then teacher Thomas Nagel, Nagel responded as follows: "Larry, I wouldn't understand what someone *meant* who claimed that *all things considered*, *A* is better than *B*, and *all things considered*, *B* is better than *C*, but *all things considered*, *A* is *not* better than *C*." (Temkin, 2012, pp. vii–viii). Though neither with Nagel's "pity" nor "scorn", I second Nagel's response to Temkin in terms of its content (Temkin, 2012, p. vii).

over" by advocates of the Relevance View can be the reason why we should always aggregate. However, for the sake of argument, I will now continue my discussion by presuming that the second horn of Horton's dilemma is nevertheless a no-go for advocates of the Relevance View. Although the advocates of the Relevance View only appear to appeal to a dodgy principle to defend their different choices in Horton's Case 6 vs. his Case 9, the switch in stance regarding these two cases can still be considered counter-intuitive. This brings us to the following question: how bad is the first horn of the dilemma?

Before discussing the first horn on behalf of Local Relevance, I should present Horton's case that "sharpens" that horn, which is a modified version of his Case 6 (2018, p. 171):

> *Stage 1*: You can save either group A, which contains 1 person facing death, or group B, which contains a ginormous number of people facing a lost finger.
> *Stage 2*: 1,000,000 people facing a lost arm are added to group A, and 20 people facing a lost arm are added to group B.[6]

How does Sequential Claims-Matching handle this modified Case 6?[7] At Stage 1, it requires saving group A. But, at Stage 2, it requires a switch to saving group B. Again, this is just due to the sequential matching and the "ginormous" number of lost finger claims in group B. When that respective stage unavoidably arrives, these lost-finger claims become relevant to the new anchoring claim, which will be any single one of the 999,999 lost-arm claims in group A left. Why is it that 999,999 lost-arm claims will be left? This is because 19 out of 20 lost-arm claims added to group B will have been balanced by the single death-claim in group A in the initial matching stage. The one remaining person facing a lost arm in group B becomes the new anchoring claim. Now we have lost arms on either side – one million in A and one in B. Equal claims match each other 1:1. Therefore, 999,999 people facing a lost arm remain in A. And, although of course a huge number,

[6] Although Horton does not try to articulate the following principle, the point of his modified Case 6 here is to show that it appears to violate this principle, which stems from Tomlin's (2017) Principle of Addition:

> The Principle of Net Addition: Adding claims of equal strength but differential numbers cannot make the group to which more claims are added less choice-worthy compared with the group to which fewer claims are added.

[7] Horton discusses only how Tadros's (2019) account would (fail to) deal with this modified Case 6 because that was the – as far as I am aware – only Local Relevance account available to Horton at the time of writing his (2018) criticism. I agree with Horton both that Tadros's account fails and that a concocted 'respect-based' response would do no good here for the advocate of Local Relevance. For discussion and dismissal of that response, see Horton (2018, pp. 171–172). For an excellent but more general criticism of Relevance accounts, including a criticism of a 'respect-based' strategy, see John Halstead (2016).

999,999 isn't a "ginormous" number in the sense Horton means "a ginormous number of people facing a lost finger" for the sake of his example. So, some huge number out of the ginormous number of lost-finger claims will outweigh the 999,999 lost-arms claims. This will leave the remaining number of lost-finger claims as the new and unchallenged anchoring claims. In other words, the lost-finger claims will carry the day and so, at Stage 2, Sequential Claims-Matching requires (the switch to) saving group B.

Is this first horn of Horton's dilemma, even when sharpened, so bad to fall on? Unsurprisingly, I like the rationale for Sequential Claims-Matching, and so I am tempted to bite the bullet here. The proposed justification for the switch to saving group B at Stage 2 would roughly have two parts. First, the point would again be to start with the strongest claim overall and make sure that claim is 'appropriately' matched – which is achieved by matching it with a sufficient number of relevant claims – and then to sequentially move down the chain. Second, the suggestion would be that the addition of lost-arms claims, although quantitatively hugely in favour of group A, changes the overall (morally) 'qualitative' shape of both group A as a whole and, more importantly, the competition of the two groups of claims such that the we need to redo the matching from scratch. In other words, to keep the result from the sequential matching of claims of Stage 1 going forward to Stage 2 would be morally inappropriate. So, my response here is similar to my earlier resistance regarding the second horn of Horton's dilemma.

I am sure it would not be hard to find more than a few people who would argue that Horton's modified Case 6 serves as a fatal counterexample to Sequential Claims-Matching, including those initially persuaded by Sequential Claims-Matching. I will assume it is a fatal counterexample for the sake of argument.

Faced with the prospect of being unable to justify Limited Aggregation (i.e. a partially aggregative view), the Scanlonian contractualist should contemplate what embracing a purely aggregative view would entail. Most importantly, just like consequentialism is broader than utilitarianism, so Pure Aggregation is broader than utilitarianism. It is open to Scanlonian contractualists to accept a non-utilitarian (yet) purely aggregative view. Arguably the most important example of a non-utilitarian purely aggregative view is prioritarianism. In the next section, I make a case for Scanlonian contractualists to adopt a version of prioritarianism by contrasting prioritarianism with its main non-utilitarian rival: egalitarianism.

§2. Distinguishing Prioritarianism, Luck Egalitarianism, and Relational Egalitarianism

The debate on egalitarianism versus prioritarianism is both fascinating and complicated.[8] First, we need a general definition of both egalitarianism and prioritarianism. For such general definitions, I rely on Derek Parfit's (2000) formulations of "the [Telic] Principle of Equality" and "the Priority View":

> The Principle of Equality: It is in itself bad if some people are worse off than others. (Parfit, 2000, p. 84).

> The Priority View: Benefitting people matters more the worse off these people are. (Parfit, 2000, p. 101).

Parfit favours the Priority View over the Principle of Equality because the latter principle fails to avoid the following "Levelling Down Objection", whereas the Priority View does avoid this objection:

> The Levelling Down Objection: [The Principle of Equality] implausibly implies that, if some people are better off than others, it would be in one way better if everyone became much worse off, but the better off people suffered greater misfortunes, so that everyone became badly off. (Parfit, 2012, p. 399).

But regardless of the Levelling Down Objection, the following complication is especially potent in terms of creating confusion. When Parfit (2000) originally raised the question of "Equality or Priority?", he elaborated on the definitions of these two views as follows:

> On the definition with which I began, the Priority View is not Egalitarian. On this view, though we ought to give priority to the worse off, that is not because we shall be reducing inequality. We do not believe that inequality is, in itself, either bad or

[8] For this section, I am indebted to Alex McLaughlin for many invaluable discussions and extensive written comments. For an excellent and (much) more formal overview discussion of egalitarianism vs. prioritarianism vs. utilitarianism, specifically focused on the aggregation of values, see Nils Holtug (2015). Holtug succinctly presents the aggregation problem in his section 14.5. "Worries about Aggregation" (2015, pp. 282–283). In my next chapter, I pick up the gauntlet that Holtug throws down with the following clearly formulated challenge that aggregation presents through the discussion of these different distributive views and which I too think should be met: "It may therefore seem as if a principle less aggregative than utilitarianism, prioritarianism, and (some versions of) Paretian egalitarianism, but more aggregative than leximin is called for." (2015, p. 282). My proposed version of a so-called "mixed solution", I contend, walks the middle ground between these rival views of distributive justice.

> unjust. But, since this view has a built-in bias towards equality, it could be called Egalitarian in a second, looser sense. We might say that, if we take this view, we are *Non-Relational Egalitarians*. (Parfit, 2000, p. 106).

Although what Parfit says here is certainly correct, to label prioritarianism as "non-relational egalitarianism" is unhelpful. I think Parfit himself considered this label a mistake because he later (1997) changed the distinction to "comparative justice" vs. "non-comparative justice" (1997, p. 207).[9] Although all that Parfit has changed is the word "relational" to "comparative", it much more clearly shows that prioritarianism is not a form of egalitarianism. To add "relational" as an adjective to "egalitarianism" is much more helpful in a debate regarding the two main versions of egalitarianism: relational egalitarianism vs. luck egalitarianism.

In the following subsections, I discuss this distinction within egalitarianism for two reasons. First, luck egalitarianism and prioritarianism might at first glance not seem to be very different, and so it is important for me to show how they are different. Second, as I will show, Scanlon's contractualism and relational egalitarianism seem to be a natural fit. To see how they are a natural fit, I need to present what relational egalitarianism is.

Before I discuss these two points, I should say why prioritarianism delivers the softest blow if we assume that the justification for Limited Aggregation via the Relevance View fails, and so why Scanlon might be interested in prioritarianism for his contractualism. The relationship between contractualism and prioritarianism is best brought out by the following further assumption that Parfit makes for his favoured version of prioritarianism.

Parfit assumes that prioritarians should be "Double Prioritarians", according to whom "we have stronger reasons to benefit people (...) the worse off these people are both [(1)] in their lives as a whole and [(2)] when they receive these benefits." (Parfit, 2012, p. 420). Double prioritarianism so combines "Whole Life Prioritarianism", according to which we should focus only on (1) and "Temporal Prioritarianism", according to which we should focus only on (2) (2012, p. 420).

Insofar as Scanlon permits prioritarian reasons within his contractualism – which is only to a very limited extent – he is a "temporal prioritarian". Consider:

[9] Parfit (2000) is a reprint of the Lindley Lecture with the same title that Parfit gave at the University of Kansas on November 21, 1991, which was first published by the University of Kansas in 1995. Hence, Parfit's (1997) "Equality and Priority" is Parfit's "later" view when compared to Parfit's (2000) "Equality or Priority?".

> [T]here is a significant part of contractualist moral argument which is not marked by priority [for the worst off]. (….) The idea of priority for the worst off has greater plausibility in cases in which the aspect of well-being in which some people are worse off is the same as the way in which they can be helped. (…) I suggest, then, that in order for differences in level to affect the relative strength of people's moral claims to help, these differences have to be in an aspect of welfare that the help in question will contribute to. So, if the claims of the worse off sometimes get priority in the way Nagel claims, what is relevant is their level in this particular aspect of well-being, not, as he says, the "quality of their lives as a whole, from birth to death." (Scanlon, 1998, pp. 223, 227).[10]

But given the setup of the trade-off cases as discussed in this thesis, these cases meet the requirements set by Scanlon on permitting prioritarian considerations. So, there is enough room within Scanlon's original formulation of his contractualism to fit it with a plausible version of prioritarianism for my purposes of considering questions surrounding aggregation.

As I have mentioned, prioritarianism cannot deliver a justification for Limited Aggregation. Given that it cannot, one might wonder why Scanlon should be interested in prioritarianism. The reason why Scanlon should be interested in prioritarianism for his contractualism is because not all versions of Pure Aggregation are created equally.

Time and again, Scanlon makes clear that his main aim is to develop a moral theory which is an alternative to utilitarianism. Utilitarianism is a purely aggregative view. So, once we have established that a partially aggregative view cannot be justified and an anti-aggregative view has been ruled out because of its sheer implausibility, then Scanlon's main aim could only be accomplished if there is a purely aggregative view which is not utilitarianism. Prioritarianism is the clearest alternative.

[10] For the full discussion of the question of priority for the worst off in Scanlon's contractualism, see Scanlon (1998, pp. 223–229). Scanlon's limited advocacy of priority for the worst off marks one of the starkest contrasts between his contractualism and John Rawls's (1999) contractualism – which is shaped by either leximin or maximin: justice requires maximising the position of the worst-off person compared to alternative arrangements of society. For a masterclass-like, one-page explanation of Rawls's argument for his Difference Principle, see Scanlon (2018, p. 137). In his discussion of unequal income (2018, pp. 133–151), Scanlon takes a step back from the particulars of that discussion on unequal income to point to his disjunctive necessary condition for features of a basic structure that generate significant inequalities: "[I]t must be either true that these inequalities could not be eliminated without infringing important personal liberties, or that they are required in order for the economic system to function in a way that benefits that all." (p. 141). Scanlon notes that his disjunctive condition is weaker than Rawls's Difference Principle, but still very demanding (p. 142).

Importantly, it would not be an alternative for the sake of keeping Scanlonian contractualism distinct from utilitarianism. Prioritarianism fits nicely with the overall rationale of Scanlonian contractualism in terms of finding principles that "no one could reasonably reject" to so stand in a relation of "mutual recognition" and so have "justifiability to each person". It seems particularly important to be able to justify one's conduct to whomever will be worst off. Concern for the worst off is, of course, the driving force of prioritarianism by definition.

More specifically, in the trade-off cases which are central to the aggregation problem, if we cannot justify a relevance vs. irrelevance distinction to any strongest claimant, no matter how much stronger their claim might be than those claims with which they compete, then it certainly seems more in line with the general idea of justifying our action-guiding principles to each person that prioritarian considerations would be taken into account. Compared to utilitarianism, on prioritarianism a stronger claim is not outweighed by a slightly greater aggregate good of much weaker claims that have been summed. Because the added moral weight of the stronger claimant needs to be taken into account as well, many more weaker claims are required to compensate for this added moral weight given to the stronger claimant.

It is important to note that the strongest claimant is not always the worst-off person. Suppose that in a modified version of Scanlon's World Cup Case Jones is a very high level of well-being and that, even if we don't save him from the painful electrical shocks for the next hour, he will still be much better off than the viewers watching the match. Jones is still the strongest claimant, and quite clearly so. Jones is the strongest claimant because the benefit to him of shutting down the transmission until he is safe from electrical shocks would benefit him far more than continuing with the transmission will benefit any other single individual. But Jones does not receive prioritarian added weight to his claim. In fact, if there are viewers who are much worse off than Jones even if they get to watch the World Cup final, then Jones's case for being saved weakens. His case weakens because each of the weaker claimants receives an added moral weight to her claim because she is much worse off than Jones, and so fewer weaker claims are needed to outweigh Jones's claim.

§2.1. Luck Egalitarianism vs. Prioritarianism

Although luck egalitarianism is closely related to prioritarianism, the two views do not fully overlap. We can see that luck egalitarianism and prioritarianism do not fully overlap by considering the

main claim of luck egalitarianism and comparing that with Parfit's definition of the Priority View above:

> Core Luck Egalitarian Claim: It is unjust if some people are worse off than others through their bad luck. (Lippert-Rasmussen, 2019, p. 245).

The Core Luck Egalitarian Claim above shows that luck egalitarianism is a narrower view than prioritarianism. Any benefit that would come to a worse-off person would have greater moral weight regardless of whether the worst-off person is due to "bad luck". Presumably, prioritarians would recognise a difference in degree of the added moral weight to a benefit if the worse off person is worse off due to bad luck or because of their own choices (allowing either one to be to a certain degree). Luck egalitarians, at least traditionally, would disagree with prioritarians here. Following Ronald Dworkin (2002), luck egalitarians make a strict, binary distinction between "option luck" and "(bad) brute luck":

> Option luck is a matter of how deliberate and calculated gambles turn out—whether someone gains or loses through accepting an isolated risk he or she should have anticipated and might have declined. Brute luck is a matter of how risks fall out that are not in that sense deliberate gambles. If I buy a stock on the exchange that rises, then my option luck is good. If I am hit by a falling meteorite whose course could not have been predicted, then my bad luck is brute (even though I could have moved just before it struck if I had had any reason to know where it would strike). Obviously the difference between these two forms of luck can be represented as a matter of degree, and we may be uncertain how to describe a particular piece of bad luck. If someone develops cancer in the course of a normal life, and there is no particular decision to which we can point as a gamble risking the disease, then we will say that he has suffered bad brute luck. But if he smoked cigarettes heavily then we may prefer to say that he took an unsuccessful gamble. (Dworkin, 2002, pp. 73–74).

Importantly, what Dworkin says regarding these two forms of luck being "a matter of degree" refers to the difficulty we might, plausibly quite often in real life, face in categorizing a piece of bad luck as either "brute" or "optional". However, this difficulty or vagueness in determining which kind it is does not change the fact that the distinction between bad brute luck and option luck remains strictly binary. Once something has been categorized as "option luck", then for a luck

egalitarian it is not true that we should give any moral priority or any greater moral weight to the benefit this worse off person could receive. Again, by contrast, for prioritarians it is an open question whether someone who is worse off solely through bad option luck should still receive some moral priority or some greater moral weight to their potential benefit.

As I understand prioritarianism, the answer could at least be affirmative: worse off people who are worse off strictly due to option luck should nevertheless receive some added moral weight to their claim to a potential benefit. Luck egalitarians could not agree with this. This marks one distinction, as I understand it, between luck egalitarianism and prioritarianism. My view on the distinction between luck egalitarianism and prioritarianism seems in agreement with Richard J. Arneson's (2000) characterisation of the distinction, which he captures in an illuminating way with an imaginative example:

> Since prioritarianism welcomes chance events that increase people's well-being, it is not, strictly speaking, a member of the luck egalitarian family of views, just a close cousin. Prioritarianism prefers the outcome in which a random meteor shower confers benefits costlessly on some already advantaged people to the status quo ante in which well-being is less for some people and better for none, but more equally divided. (Arneson, 2000, p. 341, fn. 6).

However, the most important distinction between luck egalitarianism and prioritarianism is that luck egalitarianism is an egalitarian/comparative/levelling down view, and prioritarianism is not. For the purposes of my discussion, from here on out, I focus on prioritarianism and, unfortunately for the luck egalitarians, I will leave luck egalitarianism behind.[11] But I had to discuss luck egalitarianism to make clear that it is prioritarianism, and not this seemingly closely related alternative of distributive justice, that I suggest Scanlonian contractualist should adopt.

[11] Scanlon juxtaposes his relational view of equality with luck egalitarianism because Scanlon rejects the luck egalitarian idea that "departures from equality are justifiable if they result from choices that individuals actually make." (2018, p. 63, fn. 21). However, I think Scanlon's justification for what matters regarding any person's choices on his own view fails to separate his view from luck egalitarianism as he claims. Commenting on his own account, Scanlon says the following: "What matters on this account is a person's *having* a choice under sufficiently good conditions, rather than her consciously *making* a choice. It can be enough that a person was placed in (good enough) conditions under which she could have gotten a certain outcome by choosing appropriately even if, because she failed to pay attention to the fact that she had this choice, she passed up the option without choosing to do so." (2018, p. 63). I fail to see how what Scanlon says here about "having choices" vs. "making choices" and the relevance of "sufficiently good" background conditions could not be accepted by luck egalitarians. In fact, I would think luck egalitarians would put great emphasis on setting up a basic structure of society or providing background conditions which, if not wholly eliminating any effects of bad brute luck, at least minimize the (potential) impact bad brute luck has.

§2.2. Relational Egalitarianism

What about "relational egalitarianism"? The *locus classicus* for this view comes from Elisabeth Anderson (1999).[12] Strictly speaking, Anderson calls her view "democratic egalitarianism", but we can safely ignore that label here. What matters for my purposes is the way in which Anderson characterises her view of equality:

> [D]emocratic equality is what I shall call a relational theory of equality: it views equality as a social relationship. Equality of fortune [i.e. luck egalitarianism] is a distributive theory of equality: it conceives of equality as a pattern of distribution. Thus, equality of fortune regards two people as equal so long as they enjoy equal amounts of some distributable good—income, resources, opportunities for welfare, and so forth. Social relationships are largely seen as instrumental to generating such patterns of distribution. By contrast, democratic equality regards two people as equal when each accepts the obligation to justify their actions by principles acceptable to the other, and in which they take mutual consultation, reciprocation, and recognition for granted. Certain patterns in the distribution of goods may be instrumental to securing such relationships, follow from them, or even be constitutive of them. But democratic egalitarians are fundamentally concerned with the relationships within which goods are distributed, not only with the distribution of goods themselves. This implies (…) that democratic equality is sensitive to the need to integrate the demands of equal recognition with those of equal distribution. Goods must be distributed according to the principles and processes that express respect for all. People must not be required to grovel or demean themselves before others as a condition of laying claim to their share of goods. The basis for people's claims to distributed goods is that they are equals, not inferiors, to others. (Anderson, 1999, pp. 313–314).[13]

[12] Arneson (2000) provides an excellent critical review of Anderson's account.

[13] Although Anderson hints at the following, she does not explicitly emphasize what I would explain in terms of Andrea Sangiovanni's (2017) work – although Sangiovanni is concerned with moral equality as such, rather than relational egalitarianism: the relational egalitarian commitment to moral equality is best explained by a rejection of cruelty. To have "avoid cruelty for all" as one's aim seems much more precise than the much more vague "express respect for all". Intuitively, there seems to be much greater room for reasonable disagreement for what constitutes as expressing "respect", rather than what constitutes "cruelty". Therefore, to aim for the avoidance of cruelty, rather than the promotion and/or maintenance of respect makes relational egalitarianism a more practically feasible and, therefore, more plausible view. Frank Lovett (2010) presents a similar and equally attractive proposal of interpreting the "arbitrariness" in the arbitrary power concept that is central to conceptions of domination as "procedural", rather than "substantive" (2010, pp. 85–124). I thank Jamie Draper for kindly discussing different conceptions of domination with me and for pointing me to Lovett (2010). A telic version of egalitarianism which is similar to Sangiovanni's proposal in focusing on the avoidance and eradication of negative relationships is Ingmar Persson's (2001) "Anti-inegalitarianism", according to which "the intrinsic value of just equality will be neutral, consisting in the *mere absence* of something intrinsically bad, namely unjust inequality." (2001, p. 30).

Although phrased in terms of 'acceptance' rather than 'reasonable rejection', it is evident from the passage above that Anderson's relational view of egalitarianism is a natural fit for Scanlon's contractualism. How much of a natural fit Andersonian relational egalitarianism is with Scanlonian contractualism is clear from considering the progression throughout his work Scanlon has made towards Anderson's view of equality to now explicitly defending this kind of relational egalitarianism.

As I have mentioned earlier in the thesis, in his earlier (1982) writing, Scanlon advocated a prioritarian view. By the time of (1998) *What We Owe to Each Other*, although not rejecting prioritarianism, Scanlon has significantly minimized the importance of prioritarian considerations within his contractualism. As I have shown earlier, in *What We Owe to Each Other*, Scanlon claims that prioritarian considerations only have a place in the distribution of benefits where the respective benefit improves the potential beneficiary in only exactly that aspect of her well-being. Furthermore, Scanlon rejects the Nagelian view that we should consider how well the potential beneficiaries' lives "as a whole" will (have) go(ne) to determine whom we should benefit. Recently, Scanlon (2018) has completely crossed over from prioritarianism to egalitarianism by explicitly advocating a version of relational egalitarianism in his discussion of why inequality is bad. And Scanlon explicitly states he is in Anderson's camp when it comes to different versions of egalitarianism (Scanlon, 2018, p. 9, fn. 10).[14]

However, at least in the specific context of aggregation, Scanlonian contractualists should not opt for relational egalitarianism. As a positive suggestion, Parfit (2011b) in his criticism of Scanlon's contractualism, especially regarding aggregation, advocates that Scanlon should return to a version of the "Contractualist Priority View". First, we should recognize the understandable bias of Parfit towards prioritarianism, as he himself (2012) defends prioritarianism. Second, despite Scanlon's now clearly being an egalitarian, his (1998) contractualism is compatible with prioritarianism. It is because of Scanlon's recent advocacy of relational egalitarianism that to revise his contractualism to embrace a version of prioritarianism would give rise to the question whether we could still legitimately call this revised version of contractualism "Scanlonian". However, that is different

[14] Scanlon only explicitly addresses the question of how his preferred version of egalitarianism relates to prioritarianism in his overall conclusion of his discussion of *Why Does Inequality Matter?*: "The cumulative effect of these overlapping reasons for objecting to inequality is a form of priority for the worst off. The forms of inequality that there is strongest reason to object to, and to eliminate if possible, are the forms affecting the poor, especially those who are both poor and subject to discrimination. But this does not make what I have offered a prioritarian *as opposed to egalitarian* view. This is because most of the reasons supporting this priority, including in particular objections to unequal status, violations of equal concern, and lack of fairness in political and economic institutions, are themselves egalitarian in character." (2018, p. 154).

from, for example, the suggested revision by Parfit (2011c), that Scanlon should give up the Impersonalist Restriction, apart from the Individualist Restriction. The revision to give up the Impersonalist Restriction would mean giving up on the label "Scanlonian" to describe this revised version of contractualism.

Having discussed (1) the general versions of both relational egalitarianism and prioritarianism, and (2) the relations and differences between prioritarianism, luck egalitarianism, and relational egalitarianism, I can now turn to the specific versions of prioritarianism and a distributive, rather than a relational egalitarianism that I will focus on.

§2.3. Restricted Prioritarianism vs. The (Modified) Competing Claims View

Benjamin Lange (2017) neatly presents the specific versions of prioritarianism and egalitarianism which I will focus on as well. The specific version of prioritarianism is Andrew Williams's (2012) "Restricted Prioritarianism", which Lange concisely presents as follows:

> Restricted Prioritarianism: an individual has a claim to receive a benefit if and only if her well-being is at stake; the strength of her claim is determined by
>
> (i) her potential gain in well-being; and
> (ii) the absolute level from which this gain takes place (with a claim to a given increment in well-being being stronger, the lower the level from which this increment takes place); but
> (iii) her claim is unaffected by how her level of well-being compares to others' level of well-being. (Lange, 2017, p. 139).

The specific version of egalitarianism is the Competing Claims View, which is defended by Michael Otsuka and Alex Voorhoeve (2009; 2018) and Voorhoeve and Marc Fleurbaey (2012):

> Competing Claims View: an individual has a claim to receive a benefit if and only if her well-being is at stake; the strength of her claim is determined by
>
> (i) her potential gain in well-being; and
> (ii) her level of well-being relative to others with whom her interests conflict. (Lange, 2017, p. 140).

So, even though both Restricted Prioritarianism and the Competing Claims View are much more specified than their general ancestors of prioritarianism and egalitarianism, the main difference between the two views remains the same: Restricted Prioritarianism is a non-comparative view, whereas the Competing Claims View is a comparative view.

First, I should note that Williams's proposed version of prioritarianism looks more plausible than both Parfit's (2012) and Thomas Porter's (2012) versions because it is "restricted". Williams's version of prioritarianism is restricted because, unlike Parfit's and Porter's versions, it rejects the second of the following two assumptions:

> (a) Diminishing Value of Benefits: The weight of our reasons to benefit others decreases as the level of advantage of the recipient increases[.]
>
> (b) Invariance across Conflicts: The same weightings apply in both intrapersonal and interpersonal conflicts. (Williams, 2012, p. 318).[15]

Importantly, Williams's rejection of Invariance across Conflicts and his restriction of prioritarianism to interpersonal conflicts only is not *ad hoc*. Williams persuasively argues that there is both a Rawlsian and a Nagelian rationale for restricting the scope of prioritarianism so that it applies only to interpersonal conflict, and not to intrapersonal conflict (Williams, 2012, pp. 324–327). Williams overreaches slightly when he argues that Nagel's view provides no similar backing for Otsuka and Voorhoeve's (2009) view, as the two of them claim. I agree with Lange that Nagel's view is ambiguous so that there is never a clear distinction between non-relational (i.e. prioritarian) and relational (i.e. egalitarian) ways of giving priority to the worst off (Lange, 2017, p. 141). What matters for my purposes is that there is a rationale for both Restricted Prioritarianism and the Competing Claims View and that, importantly, Williams's restrictive modification of the Priority View is not *ad hoc*.

[15] Another option would be to follow Martin O'Neill's (2012) suggestion to revise the Priority View to be 'pluralist', so that it can account for the difference between intrapersonal and interpersonal trade-offs by giving weight to egalitarian considerations. Although that might give us a plausible version of the Priority View as such, in the dialectic of settling whether egalitarianism or prioritarianism is more plausible, going pluralist especially in terms of including egalitarian considerations strikes me as the prioritarian conceding defeat to the egalitarian. Hence, I will not further consider O'Neill's pluralist version of prioritarianism however plausible it might be in its own right. I thank Johann Frick for discussing the merits and drawbacks of pluralist prioritarianism with me.

However, the above still leaves the question unanswered why, with Williams and on behalf of prioritarians, we should reject Parfit's and Porter's insistence that prioritarianism applies to both interpersonal and intrapersonal conflicts of interests. The reason why we should think the failure to restrict the scope of prioritarianism as to apply only to interpersonal conflict is a costly mistake is because of Otsuka and Voorhoeve's (2009; 2018) persuasive objection to such 'unrestricted' prioritarianism. Otsuka and Voorhoeve illustrate their objection with a pair of cases and argue that Parfit's brand of prioritarianism, also defended by Porter, violates the following morally significant fact:

> Unity of the Individual: An individual's life has a unity that renders it appropriate to balance benefits and burdens which accrue to her *for her sake* (in a way that maximizes her expected well-being) but inappropriate to balance benefits and burdens which accrue to her as if they accrued to different individuals' lives. (Lange, 2017, p. 137).

I will not go into the details of Otsuka and Voorhoeve's cases that support their claim that 'unrestricted' prioritarianism, as defended by Parfit and Porter, violates the Unity of the Individual.[16] What matters is that because Williams's restriction legitimately restricts prioritarianism to apply only to interpersonal cases, it respects the Unity of the Individual and he is right to claim that "the Restricted View [of prioritarianism] sails past Otsuka's and Voorhoeve's objections, and refutes their vindication of relational egalitarianism." (Williams, 2012, p. 324).

Here I can jump to the current state of the debate. Lange argues that Restricted Prioritarianism fails to capture a further morally significant fact which a "suitably modified" version of the Competing Claims View can, and so the Competing Claims View is more plausible than Restricted Prioritarianism (2017, p. 138). Lange presents the following case to support his claim:

> Twin Trouble: Imagine you are a soon-to-be parent who will have either of two sets of twins. Either Anne and Betty or Chloe and Daisy will be born. Suppose that the existence of the two pairs is equally likely. You have just received bad news from your doctor: your children are certain to develop different impairments in the case that they are born. If Anne and Betty are born, Anne will develop a moderate respiratory

[16] For discussion of their cases, see Otsuka and Voorhoeve (2009; 2018). Ingmar Persson (2001) criticizes prioritarianism in a way much similar to Otsuka and Voorhoeve in his section on "the inter-personal non-relativity of the (Absolute) Priority View" (Persson, 2001, pp. 36–38). However, Persson fails to separate the intrapersonal context from the interpersonal context to make this criticism of prioritarianism as persuasive as Otsuka and Voorhoeve manage to do.

impairment and Betty, being quite unlucky, will develop a very severe impairment. If Chloe and Daisy are born, Chloe will develop a moderate mobility impairment, comparable in severity to Anne's, but Daisy will be fine. You need to change jobs and must therefore now decide whether you will move to the coast or the countryside. If you move to the coast, Anne could receive treatment in a special medical facility which would somewhat alleviate the effects of her disability. However, there would be no such special facilities to treat Betty who suffers from a very rare and incurable impairment. Neither Cloe nor Daisy would be able to benefit from the move to the coast. Alternatively, you could also move to the mountain region where Chloe would have access to a unique medical facility which would somewhat alleviate the effects of her mobility impairment. But again, neither Daisy nor Anne and Betty would benefit from the move to the countryside. Suppose that the benefit to Anne by moving to the coast would be as great as the benefit to Chloe by moving to the countryside. How should you choose? (Lange, 2017, pp. 142–143).[17]

Lange helpfully summarizes the choice of his Twin Trouble case with the following table:

Utilities for Twin Trouble

	State of the world (equiprobable)			
	S_1 (50%)		*S_2 (50%)*	
Individuals	Anne & Betty	Chloe & Daisy	Anne & Betty	Chloe & Daisy
Move to Coast	0.7, 0.2			0.5, 1
Move to Countryside	0.5, 0.2			0.7, 1

[17] Importantly, and this is recognized by all authors in the debate discussed here, everything changes once we move to non-identity cases. The dialectic in the debate between egalitarianism and prioritarianism as presented here – narrowing our options down to a choice between Restricted Prioritarianism and the Competing Claims View – has progressed in this way because of objections that have been raised and accepted by both camps to telic versions of both egalitarianism and prioritarianism. But it is only telic versions that stand a chance in being applicable to non-identity cases precisely because they are non-person affecting views or views concerned with outcomes/states of affairs. For a specific proposed version of both egalitarianism, "Telic (Anti-in)egalitarianism", and prioritarianism, "Relational Prioritarianism", which arguably deals with non-identity cases, see Persson (2001). In later work, Persson (2008) provides a succinct restatement of both his proposals, see Persson (2008, pp. 297–298). In the context of non-identity cases, Thomas Porter (2011) notes that neither egalitarianism nor prioritarianism can avoid Parfit's (1992) "Repugnant Conclusion" (Porter, 2011, p. 203, fn. 10). For Parfit's discussion of this conclusion, see Parfit's chapter 17, "The Repugnant Conclusion" (1992, pp. 381–390).

Now, to clarify Lange's Twin Trouble table: S_1 and S_2 are equiprobable states of the world. An empty cell represents non-existence. Furthermore, a utility level of 0 corresponds to a quality of life barely worth living, while a utility level of 1.0 corresponds to full health (Lange, 2017, p. 143).

With the Twin Trouble case laid out, Lange presents the one, and so crucial, difference between the equiprobable states of S_1 and S_2:

> [I]n S_2, Chloe and Daisy will exist and Daisy will have a higher level of well-being than Chloe; and in S_1, Anne and Betty will exist and Betty will have a lower level of well-being than Anne. Given that S_1 and S_2, are equally likely to occur and Chloe would fare relatively worse than Daisy while Anne would fare relatively better than Betty, it seems correct that we should move to the countryside in order to help Chloe. (Lange, 2017, p. 143).

Lange supports the claim that in Twin Trouble we should move to the countryside because if S_1 comes about and so Anne would have lost out due to our move to the countryside, we could still provide her with the following justification for our decision to move to the countryside:

> We had stronger reason to move to the countryside, because, if Chloe had existed, she would have been worse off than Daisy *even after having received the treatment for her disability*. In your case, you are still better off than Betty. Since there was a 50:50 chance that you or Chloe would exist, have the same level of well-being and face the same increment in well-being improvement, Chloe's claim was strengthened by the fact that she would fare less well than Daisy. We therefore chose to move to the countryside. (Lange, 2017, p. 144).

Lange rightly argues that advocates of Restricted Prioritarianism could not give the above justification to Anne, should S_1 come about and we have decided to move to the countryside (Lange, 2017, p. 145). The reason they could not provide this justification is because of the combination of the following three reasons: (1) there is an equal probability of improving either Anne's or Chloe's well-being to the exact same degree (+0.2); (2) the baselines or starting points of well-being before improvement are exactly the same (0.5); and, crucially, (3) by definition, (restricted) prioritarians cannot appeal to relative or relational factors. Therefore, advocates of

Restricted Prioritarianism require us to be indifferent between moving to either the coast or countryside.

However, Lange astutely notes that, quite surprisingly, the current version of the Competing Claims View equally fails to provide the justification for moving to the countryside because Competing Claims View advocates must also conclude that the strength of the only competing claims – Anne's and Chloe's – is equal (Lange, 2017, p. 146). Therefore, Lange suggests the following modified version of the Competing Claims View:[18]

> Modified Competing Claims View: The strength of an individual's claim depends in part on her level of well-being relative to *everyone else who exists alongside her*. (Lange, 2017, p. 150).

Now, viewed in isolation, I think that Lange is right that Twin Trouble illustrates how the Modified Competing Claims View can capture a morally significant fact which Restricted Prioritarianism cannot. However, I think it would be a mistake to conclude that, on the basis of Twin Trouble and the Modified Competing Claims View's ability to capture the significant fact illustrated by that case, the Modified Competing Claims View is better than Restricted Prioritarianism. Despite the reason that Twin Trouble gives us to favour the Modified Competing Claims View, I argue in the next subsection that we have countervailing reasons to favour Restricted Prioritarianism instead. The general reason is the greater practical feasibility or action-guiding potential of Restricted Prioritarianism over the Modified Competing Claims View. The more specific reason is that determining different claims' strengths is a lot harder on a relational view (e.g. the Modified Competing Claims View) than on a non-relational view (e.g. Restricted Prioritarianism).

§2.4. The Pressure of Practical Application as a reason to favour Restricted Prioritarianism?

In this subsection, I aim to provide the "additional reason" that Lange claims Restricted Prioritarians owe us to endorse the view apart from its supposed invulnerability to the Levelling Down Objection (2017, p. 151).[19]

[18] Lange suggests another alternative, the "Claim Egalitarian View", but he dismisses it for what I believe are good reasons. (2017, pp. 145–149). Therefore, I will not consider the Claim Egalitarian View.

[19] I say "supposed" because John Broome (forthcoming) claims that the Levelling Down Objection is "a bad argument for prioritarianism". And it is a bad argument, according to Broome, because some versions of prioritarianism are equally guilty of levelling down. Since I lack the skills in formal theory to check Broome's 'axiomatic' argument, I can only mention his critique but not rely on it and I can only phrase it in conditional terms: if Broome's is correct, then the Levelling Down Objection provides no reason to support prioritarianism. Before Broome, Persson (2008) had

Lange claims that the different forms of the Competing Claims View that he considers are compatible with different forms of aggregation (2017, p. 139, fn. 6). Lange's claim is correct, I believe. However, from the perspective of aggregation, adopting any relational view, such as egalitarianism and the more precise version of the (Modified) Competing Claims View, comes at the cost of quite significantly complicating our account of aggregation as compared to prioritarianism, whether as Restricted Prioritarianism or any other version of it.

The complication is due to the most characteristic difference between the two views. On Restricted Prioritarianism, it is enough to know who the worst-off is. On the Modified Competing Claims View, we need to know all the relative gaps between claimants. I contend that as the number of claimants or potential beneficiaries becomes greater, so the gap between Restricted Prioritarianism and the Modified Competing Claims View widens in terms of practical feasibility. This is because it becomes increasingly more difficult on the Modified Competing Claims View to keep track of all the relative gaps between different claimants, which is necessary if we are to figure out which gap is the largest for which individual. It is certainly more difficult than figuring out who is the worst off, full stop. My worry is thus not that Restricted Prioritarianism better accommodates aggregation than the Modified Competing Claims View. My worry is that the essential building blocks for any account of aggregation, individual claim strengths, are much harder to determine on the Modified Competing Claims View than on Restricted Prioritarianism.[20]

Of course, even if I am correct and it is indeed harder to determine the strengths of different claims on the Modified Competing Claims View than Restricted Prioritarianism, then I have not shown that Restricted Prioritarianism is better than the Modified Competing Claims View. But the concern about practical application persists insofar as we would want to insert our theory of aggregation as the underlying justificatory account of our public policies. I believe we should want this for any theory of aggregation.

I claim that if it is more difficult to determine the strengths of different claims on the Modified Competing Claims View, then insisting that we nevertheless should aim to do so might lead to

---

also argued for the claim that Parfit's kind of (absolute) prioritarianism is equally hit by the Levelling Down Objection, or even worse objections. However, Thomas Porter (2011) persuasively disarms Persson's "trilemma" by pointing out that Persson's argument depends upon the claim that indifference to changes in the average prioritarian value of benefits implies indifference to changes in the overall prioritarian value of a state of affairs. But, as Porter argues, "sensible" conceptions of prioritarianism have no such implication and so prioritarianism is not hit by the Levelling Down Objection.

[20] Persson (2008), being a relational egalitarian, nobly admits this point (2008, p. 297, fn. 4).

much worse results (1) as compared with following the guidelines of Restricted Prioritarianism and (2) even if the Modified Competing Claims View better captures what is morally 'true'. The choice between these two views would be an instance of the more general issue between 'best' and 'second-best' options or alternatives, where aiming for what's 'best' might lead to far worse results than aiming for what's 'second-best'.[21]

I have only given a conjecture of how there might be practical reasons for favouring Restricted Prioritarianism over the Modified Competing Claims View. Only practice itself (i.e. real-world evidence from real-world application or implementation) of both these views and comparing results could show whether my conjecture is true at all, and, even if it is true, to what degree it is true. It might be that Restricted Prioritarianism only provides a minimal or negligible practical advantage of the Modified Competing Claims View. If this were so, then the weak reason this would give us to favour Restricted Prioritarianism over the Modified Competing Claims View would probably be outweighed by the reason we have to favour the Modified Competing Claims View over Restricted Prioritarianism instead as stemming from cases like Twin Trouble.

Conclusion

Some time ago, Parfit (2011b) suggested that Scanlon should return to his earlier (1982) prioritarian version of his contractualism. In this chapter, I have provided additional arguments to underscore why Parfit seems to have been correct. However, given Scanlon's own recent defence of (relational) egalitarianism, I accept that the suggested revision of Scanlon's contractualism to incorporate prioritarianism, specifically Williams's Restricted Prioritarianism, might give rise to the question whether such a revised version of contractualism would still be recognisably "Scanlonian".

[21] For one, I believe plausible, proposal of a practically applicable version of prioritarianism, see Daniel Sharp and Joseph Millum (2018). Liam B. Murphy (1998) can be interpreted as arguing that my phrase of "a practically applicable version of prioritarianism" in my previous sentence is misleading because "if we think of distributive justice in terms of weighted beneficence [i.e. prioritarianism] we see no distinction between the principles that govern the design of institutions and those that govern personal conduct." (1998, p. 263). Crucially, Murphy indicates that this does not mean that the "shadow of morality" is obtrusively hanging over our everyday lives and everyday decisions as individuals: "[W]hen designing institutions, it will make eminent sense from within this perspective to remove as much as possible the burdens of promoting weighted well-being from people's day-to-day lives. It will make eminent sense to set up background structures of taxation and transfer such that people can for the most part pursue their own interests, 'secure in the knowledge that elsewhere in the social system the necessary corrections to preserve background justice are being made.' [(Rawls, 2005, p. 269).] If the background institutions are doing their job properly, people will not have to think too much about promoting general well-being, and this liberation is, from the point of view of beneficence, all to the good." (Murphy, 1998, pp. 263–264).

But in any event, unlike Lange, I think the jury is still out on the question whether we should accept a version of the Competing Claims View and reject Restricted Prioritarianism, or vice versa. I also think that this question will in part (have to) be settled by the success or failure of the practical application or implementation of both these views. I have argued why I think any version of prioritarianism will probably outperform any version of egalitarianism in practice. Perhaps unsurprisingly so in a way, my reasons rely on the most characteristic difference between these two views, that prioritarianism is a non-relational view whereas egalitarianism is a relational view.

Where does that leave the Scanlonian contractualist, who has been determined to find a way to limit aggregation? Not in high spirits, I would imagine. However, in my second chapter, I argued that there are two alternative 'limited-aggregation' strategies, not just one. The first, "relevance"-strategy I have considered in detail in my third chapter. I have pushed this "relevance"-strategy as far as presenting the promising specific version of Local Relevance, Sequential Claims-Matching, and have identified the challenges that remain for that view. However, I don't think that Horton's dilemma presents as much of a problem as it initially might seem to do. I have given my reasons why I think his dilemma fails. For the sake of argument though, I have presumed that Horton's dilemma holds. If we presume Horton's dilemma to stick, then I think it is true that if, in turn, Scanlonian contractualists remain determined to 'limit' aggregation, rather than to accept Pure Aggregation even if it were to come in a prioritarian mould, then it is only the second strategy of opting for a pluralism that remains for them.

The proposed pluralist strategy would limit aggregation indirectly, rather than directly via, e.g. a notion of 'relevance'. Aggregation would be limited indirectly on this pluralist strategy because it works on the assumption of a purely aggregative view – all benefits or individual claims can be summed or aggregated – but the argument is that in some cases other values outweigh the greater aggregate good of satisfying the claims of one of the two competing groups, such that all-things-considered we should satisfy the claims of the other group.

To repeat what I had said in my second chapter, the point of the second strategy is to limit aggregation by pointing to values other than utility, specifically fairness, that might make the case that in close harm cases we aid the many, but in distant harm cases aid the one or few. I now turn to this second strategy starting in my next chapter where I will discuss so-called "mixed solutions".

University of Reading

A Mixed Bag of Mixed Solutions

Abstract Mixed solutions (Gertken, 2016; Hirose, 2015; Peterson, 2009; 2010) are proposed solutions to interpersonal trade-off cases of competing claims to an indivisible good. So far, advocates of mixed solutions have focused solely on the subset of these competing-claims-to-an-indivisible-good cases where (a) the claims are equal and (b) the number of claimants between the two groups differs. I will indicate how these mixed solutions could be extended to the more interesting cases where not only do the number of people but also the strengths of their claims differ. However, we first need a good understanding of how mixed solutions work, and so in my discussion of three different versions of a mixed solution proposed by the three authors cited above – the first half of my chapter – I will maintain exclusive focus on competing-claims-to-an-indivisible-good cases where both (a) and (b) are satisfied. In these cases, what these solutions 'mix' are the traditional solutions of (1) saving the greater number and (2) holding an equal-odds lottery. In a nutshell, a mixed solution entails that, if the difference in the number of people we could save is (or becomes) sufficiently great, we should save them. By contrast, if the difference is insufficiently great, we should hold an equal-odds lottery. But what is crucial to understand about mixed solutions is that pointing to the mixing of the two traditional solutions is just a helpful shorthand for what mixed solutions 'truly' mix. Mixed solutions are a (value) pluralism made precise. What they mix are different morally relevant considerations – in the context of the cases discussed the focus is on two competing values: goodness and fairness – to argue for picking one of the two traditional solutions. Mixed solutions have not received the amount of attention they deserve. Especially given Joe Horton's (2018) seemingly fatal dilemma for partially aggregative views, a mixed solution seems the only option left to justify the intuitions that have given rise to the discussion of partially aggregative views. In this chapter, I give mixed solutions the credit that they are due.

Key words Fairness • Gertken • Hirose • Mixed Solutions • Peterson

Introduction

In my second chapter, I had indicated that I would pursue the following two strategies in this thesis:

(1) To develop and justify the idea of "relevance", as suggested by Scanlon, to limit aggregation to apply to what I call "close-harm cases" (e.g. death vs. paralysis), but not to what I call "distant-harm cases" (e.g. the World Cup Case and death vs. mild headaches).
(2) To limit aggregation by pointing to values other than utility, specifically fairness, that might make it the case that in close-harm cases we may aid the many, but in distant-harm cases we may aid the one or few.

The previous chapter was the last to focus on the first of the two strategies above. In this chapter, I turn to the second strategy for the remainder of this thesis. And, as I see it, so-called "mixed solutions" are the main articulation of this second strategy, and so I will focus my discussion on mixed solutions in this chapter.

To introduce mixed solutions, consider the following likely all-too-familiar case, of which I will take Derek Parfit's (2003) formulation:[1]

> Life Boat: White is stranded on one rock, and five people are stranded on another. Before the rising tide covers both rocks, we could use a life boat to save either White or the five. (Parfit, 2003, p. 376).

There are the following two traditional solutions to this case: (1) save the greater number and (2) hold an equal-odds lottery. I am aware of other proposed solutions, some of which we might rightfully call traditional solutions as well, such as the weighted lottery (the main version of which makes chances proportionate according to group-size – so in Life Boat, White would get a 1/6 chance and the five a 5/6 chance of being saved). I am also aware of the many arguments for and

[1] Commonly, John Taurek (1977) is credited with having come up with this case. This is true insofar as regarding the morally irrelevant details of being in a life boat. However, a structurally similarly example in the context of a hospital, too many patients, and a scarce drug supply had already been discussed by Philippa Foot (1967, p. 4). Foot defends the save the greater number view. G. E. M. Anscombe (1967) points to the view that no one member of the many is wronged if we save the one. So, I think it is Anscombe, rather than Taurek, who should be credited for raising doubts about the save the greater number view and paving the way for lottery views in cases like Life Boat in modern moral philosophy (pun intended).

against both (1) and (2). I will not get into that literature here. What suffices for my discussion is that there is some plausible support for both (1) and (2) in Life Boat. (1) is plausible because saving more lives rather than fewer at no further costs just is plausible. (2) is plausible because of two ideas. The first is that no single one of the individuals suffers a worse fate if not saved. The second is that each individual has a claim to be saved that shouldn't be overridden.[2]

Mixed solutions aim to strike a balance between (1) and (2) depending on the difference in number of people in each group who could be saved. If the difference in number is sufficiently great, then we should save the greater number. If the difference is insufficiently great, then we should hold an equal-odds lottery. The first half of my discussion, §1, will focus on discussing and rejecting three different versions of a mixed solution, those by Iwao Hirose (2015), Martin Peterson (2009; 2010), and Jan Gertken (2016).[3] Having done the negative work, in the second half of my discussion, §2, I propose my own version of a mixed solution.

If there is a single lesson to be drawn from the discussion of mixed solutions, then it is how mixed solutions seem of great instrumental value in improving our understanding of how to navigate and/or balance competing morally relevant considerations, such as goodness and fairness. I know there are plenty of other morally relevant considerations, but given the assumptions of cases like Life Boat, it is only the respective weighting of goodness and fairness that makes a difference. This also provides a helpful explanation of both traditional solutions that a mixed solution mixes. Saving the greater-number advocates (implicitly or explicitly) rely on the view that (aggregate) goodness trumps fairness, whereas, by contrast, equal-odds-lottery advocates (again, implicitly or explicitly) rely on the view that fairness trumps (aggregate) goodness.

---

[2] The importance of not overriding claims for the sake of fairness is a central tenet of John Broome's (1990) theory of fairness, which I discuss in my next and final chapter.

[3] For earlier discussions of relevant parts of Hirose's "formal aggregation" account, see (Hirose, 2001; 2004; 2007; 2013). I will restrict myself to using (Hirose, 2015a) as it is both the most recent and comprehensive discussion of his "formal aggregation" and aggregation as such. For an excellent review of Hirose's most recent discussion, see Susanne Burri (2016). There are two other versions of the traditional mixed solution. First, there is Rob Lawlor's (2006) version. The main reason for my not discussing Lawlor's mixed solution is that he impressively presents the relevant considerations and potential objections that could be raised to the weighted lottery as correctly balancing the values of saving the greater number (i.e. goodness) and giving each person an equal chance (i.e. fairness) but does not provide a positive proposal regarding what we should do instead of the traditional weighted lottery. Second, there is Gerald Lang and Rob Lawlor's (Lang & Lawlor, 2015) mixed solution. Lang and Lawlor's proposal relies on the same assumption made by Hirose which I discuss later – that the disvalue of unfairness can be interpersonally aggregated – but is not as neatly presented as Hirose, so I have picked Hirose's version instead as the representative version of that kind of mixed solution. Furthermore, Lang and Lawlor's discussion of Broome's view and Broomean-like views is guilty of straw-manning: I disagree with them that Broomeans claim that "principles [of fairness] must *always* be sidelined in conflict cases." (Lang & Lawlor, 2015, p. 310). Furthermore, their view is guilty of arbitrarily applying fairness considerations to equal-claims cases but not to unequal-claims cases, such as Scanlon's World Cup Case, see (Lang & Lawlor, 2015, p. 310).

My phrase of "trumping" is ambiguous between (i) views that value both (aggregate) goodness and fairness but where one of the two values outweighs the other and (ii) views which value only one of the two, (aggregate) goodness or fairness. This ambiguity will be important starting in the next section when I present Hirose's view because it explains different understandings of (different versions of) a lottery-solution. So, I will get back to it then.

In conclusion, I argue that mixed solutions are a most promising route remaining for those who would like to justify the intuitions that underlie a partially aggregative view without defending a partially aggregative view.[4] Mixed solution advocates could assume a purely aggregative view, but argue that in some cases the aggregate good that comes from satisfying the greater number of claims is outweighed by the other morally relevant consideration of the unfairness we would do to those whose claim(s) we wouldn't satisfy. In other words, the unfairness done would be too big a (moral) price to pay for securing a greater amount of (aggregate) good.[5]

§1. Mixed Feelings about Mixed Solutions

In this first section, I discuss and reject Iwao Hirose's, Martin Peterson's, and Jan Gertken's versions of a mixed solution. In §1.1., I present Hirose's "formal aggregation" and his discussion of how it, according to him, vindicates F. M. Kamm's (1998) intuition of holding an equal-odds lottery in the saving 1,000 vs. 1,001 lives despite the fact that we could saving an additional life outright if we save the 1,001.[6] In §1.2., I present Peterson's criticism of Hirose's account which focuses on the same example. In §1.3., I briefly digress to discuss Peterson's, in my view, misinterpretation of Broome's theory of fairness. I then present Peterson's own mixed solution

[4] In my previous chapter, I have expressed my arguments against Joe Horton's (2018) dilemma. However, for the sake of argument, I have presumed Horton's dilemma to be fatal to any partially aggregative view, or at least for the most sophisticated versions present in the philosophical literature as of yet. This, in part, thus motivates my exploration here of mixed solutions.

[5] For textual support that Scanlon is at least open to the 'weighing-up' or 'trade-off' structure of different values, which is part and parcel of a mixed solution, consider the following: "Perhaps the various forms of fairness and equality can be brought together under one all-encompassing notion of distributive justice which is always to be increased, but is not obvious that this is so. In any event, it would remain the case that attempts to increase fairness and equality can have costs in other terms; they may interfere with processes whose efficiency is important to us, or involve unwelcome intrusions into individuals' lives. In such cases of conflict it does not seem that considerations of fairness and equality, as such, are always dominant. An increase in equality may in some cases not be worth its cost; whether it is will depend in part on what it is equality *of*. Economists often speak of 'trade-offs' between equality and other concerns (usually efficiency). I have in the past been inclined, perhaps intolerantly, to regard this as crassness, but I am no longer certain that it is in principle mistaken." (Scanlon, 2003c, p. 32).

[6] Scanlon, in conversation, has confirmed his consistent application of his tie-breaker argument: if the saving of an additional life tips the balance in a 1 vs. 2 case, then it should also tip the balance in a 1,000 vs. 1,001 case. So, this is one important area of disagreement between Kamm and Scanlon. In this respect, lumping their view together as "the Kamm-Scanlon argument" is misleading because Kamm and Scanlon disagree about tie-breaker argument's scope of application.

and my reasons for rejecting it. Finally, in §1.4., I discuss and reject Jan Gertken's 'aggregative' mixed solution, which Gertken develops as a response to Peterson's mixed solution.

§1.1. Hirose's "Formal Aggregation" and the Large-Scale Rescue Case

The Large-Scale Rescue case gives us the alternatives of either saving 1,000 lives or saving 1,001 different lives. Without naming it as such, F. M. Kamm (1998; 2007) introduces the case to support her idea of "irrelevant utilities."[7] Kamm's point about this case is that, unlike in a 2 vs. 1 saved lives case, the possibility of saving an additional life in the Large-Scale Rescue Case should not push us to saving the greater number (1,001) outright. Instead, for reasons of fairness, Kamm believes we should toss a fair coin (1998, p. 103).[8] Hirose argues that his "formal aggregation" can "justify the case for the coin-toss without claiming that the extra person's life is not a relevant utility." (2015, p. 198).[9] His full argument is as follows:

---

[7] In full, Kamm calls it "the Principle of Irrelevant Utilities." (1998, pp. 146, 150–160). I fail to see how it is a principle. It is just a name for an intuition Kamm has, even if it is a plausible intuition: that in certain cases she does not want to consider the small extra utility to be gained as a factor in the decision-making procedure. It is fair enough to list an intuition, but it is not a principle. Recently, Scanlon seems to agree with the idea of "irrelevant utilities" (or "irrelevant goods", as Kamm later calls them (2007, pp. 34, 38, 290. 409), see Scanlon (2018, pp. 50–51, fn. 9). In my view, Scanlon should not adopt Kamm's Principle of Irrelevant Utilities. I will not further discuss Kamm's "irrelevant utilities" but simply mention it as an option which I will not make use of and move on. For a similar criticism of Kamm's Principle of Irrelevant Utilities, see Hirose (2015, pp. 196–197). For some persuasive arguments as to why we should stop talking about "utility" when talking about a "person's good", see Broome (forthcoming).

[8] Insofar as I understand Kamm's (1998) view, I think her "Ideal Procedure" commits her to defending a weighted lottery solution here, in which she would allow for the 'pooling' of individual chances. I will leave it to others who better understand Kamm's account to correct me on this.

[9] Regarding the contrast between "formal aggregation" and (the more familiar or standard) "substantive aggregation", see Hirose's chapter 3 (Hirose, 2015a, pp. 42–63), and especially section 3.3 (Hirose, 2015a, pp. 51–58). Hirose defines substantive aggregation as follows: "In substantive aggregation, the morally relevant factors of individuals are predetermined, and then combined into the overall value. By 'predetermined'. I mean that morally relevant factors are identified from outside of the aggregative process, and some argument is given concerning what the morally relevant factors are, independently of the aggregative process. That is, the morally relevant factors are determined independently and in advance of the aggregative process. Substantive aggregation then aggregates the given morally relevant factors to identify the goal pursued. An example of substantive aggregation is classical utilitarianism." (Hirose, 2015a, p. 52). Hirose's definition of formal aggregation, as you might expect, is contrasted with his definition of substantive aggregation: "[F]ormal aggregation leaves the morally relevant factors undetermined prior to aggregation and allows for the possibility of including some factors that emerge through the aggregative process. There may be several important factors that emerge from the relations between different people. For example, inequality and unfairness are usually thought to be factors that are not part of any person's well-being or pleasure. Inequality and unfairness are essentially relational, and their badness is not accounted for if we simply calculate each person's well-being or pleasure. Formal aggregation enables us to consider the badness of these interpersonal factors and disperse it to particular individuals. That is, formal aggregation can divide up and distribute the badness of relational factors such as inequality and unfairness across individuals. Formal aggregation merely represents how the morally relevant factors of different individuals are structured. It gives a structure to morally relevant factors. An example of formal aggregation is found in [John Broome's] version of egalitarianism [according to which inequality is bad because it is unfair to the worse off]." (Hirose, 2015a, p. 53). I think Hirose presents us with a false dichotomy in the choice between substantive aggregation and formal aggregation. Substantive and formal aggregation are not mutually exclusive. I contend even that these two kinds of aggregation should inform and shape each other akin to the process of "reflective equilibrium". But even if I grant Hirose that his formal aggregation is only meant to give an accurate description of the different morally relevant factors at play in these cases, rather than substantively prescribing us what to do or what morally relevant factor(s) should rule the day, it nevertheless takes a stance on what is accurate and what isn't as a description of these cases. As I disagree with his mapping of the structure of the morally relevant factors, I

According to the kind of formal aggregation that takes the badness of unfairness seriously, all 2,001 people [in the Large-Scale Rescue Case] have an equally strong claim to being saved. Fairness requires us to satisfy their claims equally. The only way to satisfy their claims equally is to save no one. But it is better to save someone than no one. One way to mitigating the unfairness while still saving someone is to toss a fair coin.[10]

Now the question is whether the expected good from tossing a coin is greater than the goodness of saving 1,001 people for certain. If we toss a coin, the expected good is 1,000 and a half lives saved. On the other hand, if we save the greater number, we can be sure that we will save 1,001 lives. What about the badness of unfairness? The badness of unfairness, done to *each* person, is constant regardless of the size of the group of people concerned. It is a matter of whether each person receives proper respect or not, and this has nothing to do with the number of people concerned. The badness of the unfairness for each person is the same in the original Rescue Case [i.e. Life Boat] and in the Large-Scale Rescue Case. But the *overall* badness of the unfairness is different in the two cases. Unfairness done to 1,000 people is greater than the unfairness done to one person. For simplicity, suppose that we can represent the overall badness of unfairness by the sum of the unfairness done to each person. (…) [L]et us denote the unfairness done to each person by *u*. In the original Rescue Case [i.e. Life Boat], if we directly save five lives, this is unfair to one person: so the badness of unfairness is *u*. In the Large-Scale Rescue Case, if we directly save 1,001 lives, this is unfair to 1,000 people; so the overall badness of unfairness is *u* x 1,000. Thus, an unfairness done to 1,000 people is greater than the same unfairness done to one person.

We are thus led to compare {1,000 and a half lives saved} and {1,001 lives saved – *u* x 1,000}. Is the goodness of half a life saved greater than the badness of the unfairness done to each of 1,000 people? We should toss a coin if *u* x 1,000 is greater than the goodness of half a life saved. Alternatively, we should save 1,001 lives if the goodness of half a life saved is greater than *u* x 1,000. Formal aggregation can, and likely will, judge that the badness of the unfairness done to 1,000 people is greater than

therefore reject his formal aggregation account. For an excellent account of how to understand and use the method of "reflective equilibrium" in moral and political theory, see Carl Knight (2017).

[10] So, here is the first instance where the ambiguity that I have described earlier as either (aggregate) goodness or fairness "trumping" the other is apparent. On Hirose's account, tossing a coin in this 1,000 vs. 1,001 saved lives case constitutes a compromise between (aggregate) goodness and fairness. I think this is mistaken. Fairness *simpliciter* (i.e. fairness by itself) requires the coin-toss here.

> the goodness of half a life saved, and hence that it is right to toss a coin. At a minimum, formal aggregation can claim that it is right to toss a coin when the size of the two groups of individuals is sufficiently large and the relative difference in the size of the two groups is sufficiently small. Although formal aggregation does not claim that an extra life saved is an irrelevant utility, not surprisingly, it can reach the same conclusion that Kamm's principle of irrelevant utilities does in the Large-Scale Rescue Case, although it is not committed to Kamm's conclusion. The fairness-based argument employing only formal aggregation can support the case for saving 1,001 people, provided that the badness of the unfairness to each person is assumed to be very small.
>
> Clearly this reasoning includes [interpersonal] aggregation. It aggregates the unfairness done to each of 1,000 people to give the overall badness of unfairness. It also aggregates the badness of unfairness, done to 1,000 people, and the goodness of saving an extra person. Thus, my reasoning extensively uses aggregation, but can reach the same conclusion as Kamm, who wants to rule out aggregation. Moreover, it seems to me that the reasoning based on formal aggregation in the Large-Scale Rescue Case is more plausible than Kamm's. Thus, I believe that formal aggregation is a useful and appealing tool for accommodating various morally relevant factors within a coherent framework, and that there is no compelling reason to rule out aggregation. (Hirose, 2015a, pp. 199–200).

Central to Hirose's argument is that the (negative) value of unfairness, *u*, can be interpersonally aggregated. Unfairness done to one person not saved is *u* (or, fully spelled out, *u* x1), whereas unfairness done to 1,000 people, as in the Large-Scale Rescue, is *u* x1,000. Hirose's main claim is that his formal aggregation could defend the intuition that we should hold an equal-odds lottery rather than save the greater number in the 1,000 vs. 1,001 case because the expected good of half a life saved is likely outweighed by *u* x1,000.[11] To make a supposed problem with interpersonally aggregating unfairness clear, I will now turn to Martin Peterson's (2010) criticism of Hirose's account on this point.[12]

[11] Remember, emphasis should be placed on "could" in this sentence because Hirose's account is "formal", rather than "substantive". Hirose's hedging or the non-prescriptive nature of his formal aggregation does not affect the (central) descriptive terms of his account: how his account (descriptively) maps the morally relevant considerations in these cases. And, again, how he proposes we map the cases is still a substantive proposal – one which I will reject.

[12] Given the year of Peterson's publication, Peterson's criticism is of Hirose's earlier work – (Hirose, 2001; 2004; 2007) – not Hirose's most recent and comprehensive statement of his view – (2015). However, Hirose's view is the same on this point.

§1.2. Peterson's Criticism of Hirose: Why Unfairness does not Interpersonally Aggregate

Peterson represents Hirose's formal aggregation argument as I have presented it in full above as follows for the two Rescue Cases we are now focused on (Peterson, 2010, pp. 445, 446):

Save 1,000 or 1,001?

	Expected number of people saved	Aggregated unfairness	Total value
Fair Lottery	1,000.5	0	A
Save the greater number	1001	-*u* x1,000	B

Save 1 or 2?

	Expected number of people saved	Aggregated unfairness	Total value
Fair Lottery	1.5	0	C
Save the greater number	2	-*u* x1	D

According to Peterson, "The problem is that (…) [Hirose's] Mixed View relies heavily on the assumption that there is a larger *total amount* of unfairness in the first [table] as compared to the second [table]." (2010, p. 445). Peterson questions whether this is the case by presenting the following counterexample: imagine we face the choice between saving 1,000 lives and 2,000 lives (2010, p. 445).[13] Crucially, the proportion of claims in this case is exactly the same as in the save 1 or 2 lives case: "For each person who is left behind, there are two others who are saved. Therefore, (…), it need not hold true that the unfairness *done to each person* is constant, regardless of the number of people concerned." (Peterson, 2010, p. 445).

The point of Peterson's example is that it is far from clear that the total amount of unfairness is different in the 1 vs. 2 case as compared to the 1,000 vs. 2,000 case despite the numbers in these cases being very different. Why would the total amount of unfairness be the same in both cases despite, and so contra Hirose's account, the numbers in these being different? In both cases we

[13] In §2, I will come back to Peterson's 1,000 vs. 2,000 case to test my own version of the mixed solution, which differs from both Peterson's version and Hirose's version of the mixed solution.

have a 1:2 ratio and so it is true in both cases that "for each person who is left behind, there are two others who are saved." (Peterson, 2010, p. 445).

In response to his own 1,000 vs. 2,000 saved lives case, Peterson suggests that we should accept the following principle:

> Individualism: All moral values, including fairness, are individual values which can at most be aggregated *intra*personally, but not *inter*personally across people. (Peterson, 2010, p. 445).

Peterson helpfully applies this "individualist version of the Mixed Version" to the two focus cases (2010, pp. 446, 447):

Save 1,000 or 1,001 (under Individualism)?

	A1		A1000	B1		B1001
Save 1000	$1 - u$	…	$1 - u$	$0 - u$	…	$0 - u$
Save 1001	$0 - u$	…	$0 - u$	$1 - u$	…	$1 - u$
Fair Lottery	0.5	…	0.5	0.5	…	0.5

Save 1 or 2 (under Individualism)?

	A1	B1	B2
Save 1	$1 - u'$	$0 - u'$	$0 - u'$
Save 2	$0 - u'$	$1 - u'$	$1 - u'$
Fair Lottery	0.5	0.5	0.5

Regarding the two cases under his Individualism, Peterson claims that, in the save 1 or 2 example, "each person treated unfairly faces less unfairness than in [the first example, saving 1,000 or 1,001]." (Peterson, 2010, p. 446). However, Peterson fails to spell out why the size of $u' < u$. He has already rejected the idea that unfairness could be interpersonally aggregated. He has left open the possibility with Individualism that unfairness could be intrapersonally aggregated, but further discussion is warranted. After all, on Broome's theory of fairness that is operating in the background here, we cannot interpersonally aggregate claims either. How exactly u and u' could be different in size is, therefore, left unanswered. But although Peterson leaves this question about (un)fairness unanswered, he still provides his own version of the mixed solution. In other words,

Peterson's tells us how to weigh (un)fairness against (aggregate) goodness even if he has not told us how the size of (un)fairness (e.g. *u* vs. *u'*) is determined.

Before I present my argument why, although a possibility, I reject Peterson's mixed solution on behalf of Scanlonian contractualists, I need to discuss Peterson's interpretation of Broome's theory of fairness and how I think it is mistaken. Presenting what I take to be Peterson's misunderstanding of Broome will be important for the next section where I will present my own account.

§1.3. Peterson's Mixed Solution: A Plausible but Undesirable Option

Considering an example where we could save either a million or ten people, Peterson argues that Broome proposes that we set up a weighted lottery "in which the probability that the first group (of *m* [e.g. one million] people) is saved is $m/m + n$, and the probability that the second group (of *n* [e.g. ten] people) is saved is $n/m + n$." (Peterson, 2009, p. 167).[14] This is false. It is not Broome's view, neither regarding fairness nor regarding morality overall. However, I have an idea why Peterson has understandably misinterpreted Broome here.

Peterson refers to Broome's (1998) discussion of Kamm's account of fairness (in equal-claims cases). But Peterson fails to note that Broome is there reconstructing Kamm's account of fairness and as part of his reconstruction Broome says the following: "for the sake of argument, I am willing to concede [Kamm's] step [in her argument] and permit pooling [of individual chances]." (Broome, 1998, p. 960). But, for his own theory of fairness, Broome certainly does not advocate pooling of chances, nor should he.[15] In fact, Broome's rejection of the weighing of claims because of the separateness of persons and his considering the harm of unfairness to be an individual harm entail that Broome's theory of fairness rules out the 'pooling' of chances.

In general, on Broome's theory of fairness, weightings (1) are proportionate to the strength of the claims involved but (2) ignore the number of those claims. So, for example, say that the claim of

[14] In later work, and independent of the discussion of Broome's theory of fairness, Peterson still only considers the kind of weighted lottery "in which each person is saved with a probability that is proportional to the size of the group to which he or she belongs." (Peterson, 2010, p. 441). Lawlor is equally guilty of presenting proportionalizing chances to the size of the respective groups as the sole interpretation of the weighted lottery, see (Lawlor, 2006, pp. 161–162).

[15] I think the key sentence from Broome's article on Kamm that, very understandably, trips up Peterson is the following: "I see no good grounds for the procedure of proportional chances." (Broome, 1998, p. 961). However, and it would have been helpful if Broome had added this himself, Broome is here talking about Kamm's procedure of proportional chances, in which chances are made proportionate by the number of people in the respective groups (i.e. 'pooling'). That way of 'proportionalizing' Broome, both rightly and consistently with his theory of fairness, dismisses. Scanlon also mistakenly identifies the 'pooling' of individual chances in a weighted lottery as Broome's own view (1998, p. 397, fn. 37).

having your life saved should receive a chance of 90% vs. a claim to be spared from the loss of your hand as 10%. These weightings remain constant whether we can save one life vs. one hand as opposed to one life vs. a million different hands of a million different people. In either case, the weighted chances in the lottery should be according to a 9:1 ratio.[16] Like Broome, I think there are much stronger reasons not to run the lottery in the one life vs. one million hands case than there are in the one life vs. one hand case. But, and this is key, these reasons do not stem from fairness but from another moral consideration, goodness, which likely makes it the case that fairness is outweighed. In either version of the example, perfect fairness is achieved by not helping anyone – letting people die and (plenty of) hands be lost – but 'surrogate satisfaction' of fairness can be achieved by having a weighted lottery where chances are proportionate to the strength of individual and non-aggregated claims.

My disagreement with Peterson on how to understand Broome's theory of fairness thus relates to my earlier-discussed distinction between the views according to which either (i) any lottery is a compromise between the conflicting demands of fairness and goodness or (ii) some lotteries, depending on the particular view defended, are required 'purely' by fairness (e.g. an unweighted lottery) whereas other lotteries compromise between the conflicting demands of fairness and goodness (e.g. weighted lottery).

Having clarified my understanding of these relevant aspects of Broomean fairness, let's turn to Peterson's proposed version of the mixed solution. Peterson starts by saying that "doing the right thing ultimately amounts to balancing different values against each other" and relies on the plausible intuition that in extreme situations such as saving either a million or ten people "the positive moral value gained by arranging a lottery is entirely *outweighed* by the number of people saved when saving the greater number." (2009, p. 167). Broome agrees: in these extreme situations fairness is most certainly outweighed by (aggregate) goodness. This is exactly his view as I have presented it and so, *pace* Peterson, his mixed solution is not novel in this respect.

However, Peterson's next key move is to bring in moral indeterminacy through value incommensurability: "It is, I propose, not always possible to make a trade-off between these conflicting values [e.g. fairness and the number of people saved]. Some conflicting values are

[16] In the next chapter, I argue that there is a good case to be made that if Broomeans incorporate Christian Piller's (2017) proposed distinction between 'outcome' fairness and 'procedural' fairness, then they cannot be indifferent between different representations of the fair ratio. They should, for example, prefer a 9% vs. 1% to a 90% vs. 10% lottery if the first lottery means we have a 90% chance of not giving the indivisible good to anyone because that is the fairest outcome, according to Broome.

incommensurable." (2010, p. 447). Peterson then uses this idea of value-incommensurability to claim that "moral obligations come in degrees", which leads to moral indeterminacy:

> Sometimes one is obliged to a certain degree to save the greater number *and* obliged to another degree to toss a coin. In such cases, the agent's obligations are morally indeterminate. Hence, you are obliged to save two rather than one because that will save the greater number, but you are also obliged to make the decision by tossing a coin because that would be fair. These two obligations, to save the greater number and to respect fairness, cannot be aggregated into an all-things-considered obligation since they are morally incommensurable. Therefore, it is neither right nor wrong to save the greater number, and it is neither right nor wrong to toss a coin. (Peterson, 2010, p. 447).[17]

In further support of the proposal that moral obligations come in degrees, Peterson presents a moral obligation's strength as the counterpart of this distinction regarding moral obligation. (2010, p. 448).[18] So, we have "degree" vs. "strength". Roughly, the strength of an obligation increases the more people you can save – and for unequal-claims cases, I guess, the more aggregate good you can do. Furthermore, and clearly closely related to strength's first feature, "there is no upper limit to the maximum strength of an obligation." By contrast, the degree of a moral obligation is determined by the nature of the circumstances, such as potential cost to you and the profile or characterization of the would-be beneficiaries. To give one of Peterson's examples on the last point: your obligation to save 1,000 people from a burning house would hold to a lesser degree if, instead of being innocently trapped, they "have caused the fire themselves by deliberately igniting a barrel of gasoline and threaten to murder anyone who attempts to rescue them." (Peterson, 2010, p. 448) Peterson claims that the degree of an obligation, in contrast to its strength, has an upper limit. To support this, Peterson gives the following example: "If one is fully obliged to help an innocent refugee from drowning aboard a sinking ship, then no other obligation could possibly obtain to a higher degree." (Peterson, 2010, p. 448).

Two closely related objections arise here. First, granted that any obligation valid to the maximum degree cannot be outweighed by any other obligation, it leaves the problem that there could be

[17] In later work, this is at the heart of Peterson's (2013) more expansive theory of "multi-dimensional consequentialism". For an excellent overview of his view and replies to several commentators, see Peterson (2016).
[18] The imaginative and illuminating examples that Peterson uses in clarifying the distinction between degree and strength, which I will not use here, can be found on the same page.

any other number of other maximally degree-valid obligations on the table. For example, what if we can save one innocent refugee from drowning or five? In such situations, we cannot but act wrongly, on one interpretation of Peterson's view. Here I should note that Peterson precisely (and so coherently) appeals to moral dilemmas as showing that rightness comes in degrees (Peterson, 2010, p. 447).[19] I find this an unappealing result. Second, a theoretical dilemma for Peterson which will be eerily familiar to Scanlonian contractualists: either an obligation valid to the maximum degree cannot be outweighed by any competing obligation regardless of its strength – i.e. the familiar absolutist/Individualist Restriction/number scepticism problems – or, precisely because degree has an upper limit, whereas strength has not, any maximally degree-valid obligation could (eventually) be outweighed by a competing obligation of a certain (increasing) strength – i.e. we just get maximising consequentialism, which is unpalatable to a Scanlonian contractualist.

But ignoring my two objections for the moment, or more charitably assuming Peterson could answer them, I want to turn finally to Peterson's appeal to W. D. Ross's (2002) theory of *prima facie* duties, highlighting where and how Peterson believes his account overlaps and where and how it differs from Ross's (Peterson, 2010, p. 449).[20] Peterson, I believe correctly, notes that both he and Ross hold the view that conflicting values (or, *prima facie* duties) are incommensurable (Peterson, 2010, p. 449).[21] Ruth Chang (2015) provides the following definitions of both "incommensurability" and "incomparability" and she in most helpful simple terms explains the relation between these two tricky concepts:

> Two items are *incommensurable* just in case they cannot be put on the same scale of units of value, that is, there is no cardinal unit of measure that can represent the value of both items. Two items are *incomparable* just in case they fail to stand in an evaluative comparative relation, such as being better than or worse than or equally as good as the other. (…) How do incommensurability—the failure to be measurable by a shared

[19] Discussing the possibility of moral dilemmas and related issues is, of course, another topic in its own right entirely. I have no space to do so here. My own tentative view is that there is no such thing as a genuine moral dilemma. True, there are (very) hard moral choices – and plenty of them, in fact. But one could always act rightly.

[20] For his full account, see Ross (2002). Given his view emerging as providing the common ground for discussing, comparing, and contrasting some of the most well-developed and influential (contemporary) moral theories, W. D. Ross's work's importance is hard to overstate. I will have to leave the discussion of Ross for another occasion. Parfit, in his Sections 57-58, uses Ross's pluralist deontology to represent Common Sense Morality in his discussion on whether the Triple Theory could be expanded to include more theories, most importantly Act Consequentialism and Virtue Ethics (Parfit, 2017, pp. 395–437). For what it's worth, I think that not only Parfit's 'Quintuple' Theory fails but that his Triple Theory reduces to the 'Single' Theory of rule-consequentialism. Though my 'reductionist' view is not defended by himself, I have drawn support from Brad Hooker's (n.d.) excellent impartial discussion of Parfit's ethics.

[21] On the (in)commensurability of goods on Ross's account, see Ross (2002, pp. 142–154).

> cardinal unit of value—and incomparability—the failure to be comparable—relate? (…) While incommensurability does not entail incomparability, incomparability entails incommensurability. If there is no comparative relation that holds between two items, a fortiori, there is no cardinal unit of measurement by which the two might be compared. Being commensurable is simply *one way* in which items might be comparable, and so if items are incomparable, they are incommensurable. Thus while incommensurability does not entail incomparability, commensurability entails comparability, both for value bearers and for abstract values themselves. (Chang, 2015, pp. 205, 206, 207).[22]

As Chang indicates, incommensurability does not entail incomparability, and neither Peterson nor Ross signs up to the latter idea.

The key difference between Peterson and Ross is that Peterson rejects Ross's idea that "even very disparate moral obligations [i.e. one's *prima facie* duties] can be merged into a single 'duty proper', i.e. an all-things-considered obligation, although there is of course no simple algorithm for doing this [according to Ross]." (Peterson, 2010, p. 449). As will become clear in the next section, I side with Ross on this point. But, before discussing my account, let me as a final point of this section present what, according to Peterson, we should do in light of the moral indeterminacy we face in the cases considered. What should we do given that it is indeterminate what we ought to do?[23]

Peterson's answer: randomize! Peterson borrows Ross's idea of thinking of different moral obligations (so, we must be talking about the *prima facie* duties here, and not the duty proper) as being different "forces" (Peterson, 2010, p. 449). On Peterson's view, "the force of an obligation is determined by the degree to which an obligation is valid as well as by its strength." (2010, p. 449). Most helpfully, Peterson explains his view by discussing a familiar example:

[22] Of course, the devil is in the details and so more needs to be said. However, I cannot do so here, given my purposes. The rest of Chang's (2015) discussion provides more details for those interested.

[23] This is one way of putting the central question of the booming literature on moral uncertainty. For an excellent overview of the main idea and positions regarding moral uncertainty, see Krister Bykvist (2017). I thank Brian Weatherson for having taught me on the topic of moral uncertainty during his wonderful classes that were part of the Hebrew University's Centre for Moral and Political Philosophy's Annual Project 2018 "Ethics and Uncertainty". I am grateful to the organiser of that year's project, Ittay Nissan-Rozen, and to the CMPP co-directors, David Attas and David Enoch, for inviting me. Furthermore, I am grateful to all participants of the "Ethics and Uncertainty" project and all other members of the CMPP for being so welcoming and for many great (philosophy) conversations.

> Suppose, for instance, that you face a choice between saving one or two people in the Number Problem. The obligation to toss a coin, which is the fairest option, corresponds to the force F_1, where F_1 is a real number. Moreover, the obligation to save two, i.e. to save the greater number, prompts another force, F_2. Force F_1 points in one direction, whereas F_2 points in a different direction. It is not possible to merge these forces into a single, aggregated force, since the moral values underlying these obligations are incommensurable. Instead, a decision about what to do can be reached by *applying the weighted force rule*. According to this rule, you ought to choose each alternative with a probability that is directly proportional to its force. Hence, in the choice described above you should act such that the probability is F_1/F_1+F_2 that the decision is taken by a coin toss. Moreover, the probability that you choose to save the greatest number should be F_2/F_1+F_2. [Therefore,] as long as all forces are non-zero, the weighted force rule will always recommend a probabilistic act. (Peterson, 2010, pp. 449–450).

According to Peterson, the weighted force rule is plausible insofar as it respects the 'roughness' and 'complexity' of the 'moral landscape' (Peterson, 2010, p. 450). But although I agree with Peterson's attitude to respect the complexity of the moral landscape, I disagree with his non-committal attitude regarding his own proposed weighted force rule. For it is not just that Peterson distances himself from the greater certainty that we get with Ross's 'duty proper' compared with his proposed randomization through the weighted force rule – which basically means that we have a weighted lottery to determine whether to hold an equal-odds lottery or save the greater number. No, Peterson is not even committed to follow (through with) whatever result the weighted force lottery generates. We can still ignore it and do what any single one of the moral values/considerations directs us to do![24] So, for example, if the weighted force rule provides the result to hold an equal-odds lottery, we can, according to Peterson, ignore this result and save the greater number regardless.

The above complaint is on top of the following critical questions to Peterson's account: if we can establish the strengths of F_1 and F_2, then why are we not doing what has greater force? Furthermore, does Peterson's response to incommensurability not presuppose commensurability?

[24] Letting Peterson speak for himself: "Hence, if it turns out that one ought to randomise among a set of pure acts it need not be the case that one is obliged to the highest degree to perform the randomised act prescribed by the weighted force rule. One's ultimate obligation will still come in degrees. All alternatives are neither entirely right nor entirely wrong." (Peterson, 2010, p. 450).

How else could we put a numerical value on any force F which determine the probabilities? And even if Peterson manages to resist the pressure that for his account he needs to put any numerical value on any F or that assigned numbers for different forces stem from different incommensurable cardinal scales, then commensurability still seems to be (illicitly) presumed by him because he adds different, supposedly incommensurable forces together to determine the probability of saving the greater number: "F_2/F_1+F_2".

Now, I will be the first to concur with Peterson that we should account for moral complexity, but this is advocating non-commitment to such an extent as to ground an important complaint against Peterson's view: it is insufficiently prescriptive or action-guiding. At the very least, Peterson should require us to commit to do what the weighted force rule tells us to do in any case. If we don't like the result it has yielded in this case, then that's just too bad. If we ignore the result and go with the directives issued by one of the moral values/considerations instead, then we are saying either one of the two following things: (1) "I now see those other considerations I had previously identified as morally relevant are not morally relevant or they are overruled by this moral value (e.g. fairness) whose directives I am now following." Alternatively, (2) "Well, what the weighted force rule required me to do wasn't entirely right anyway, and what I choose to do now instead is not entirely wrong either, so it's fine."

Regardless of whether following the result of the weighted force rule remains optional or whether we would revise Peterson's account such that we are required to abide by the result it yields, (either version of) Peterson's solution gives up on what I take to be one of, if not the, chief requirement(s) for a moral code: its ability to provide assurance.

Now let's reconsider Ross' pluralist deontology as a closely related but rival position to Peterson's. Philip Stratton-Lake makes the following insightful comment on Ross's alleged qualms with his self-created idea of a "*prima facie* duty":

> Because the term '*prima facie* duties' is so misleading Ross was keen to find some other term. (…) At the timing of writing *The Right and the Good*, however, Ross could not think of a better term. (…) In *The Foundations of Ethics*, however, he accepts that the term 'responsibility' is a better way of expressing what he was getting at with the term '*prima facie* duty'. *Prima facie* duties are, he there states, responsibilities to ourselves and

> to others, and what we should do (our duty proper) is determined by the balance of these responsibilities. (Ross, 2002, p. xxxiv).[25]

I would add to the above that a Rossian duty proper is the ultimate (moral) responsibility. And responsibility is to be taken. Comparing Ross's account with Peterson's account, I find Ross' view a better balance between descriptive and prescriptive: it tells us to a sufficient extent what we should do, rather than just describe how rough and complex the moral landscape is. On Peterson's account, there seems an imbalance between descriptive and prescriptive where there is so little prescription as to be problematic.

In conclusion, Peterson's version of the mixed solution is, I believe, an improvement upon Hirose's. Most importantly, I agree with Peterson's criticism of Hirose's claim how the value of (un)fairness interpersonally aggregates. However, I disagree with the extreme conclusion Peterson draws from his criticism of Hirose: that we should reject the interpersonal aggregation of the value of (un)fairness. Furthermore, Peterson's mixed solution faces objections in terms of its theoretical incoherence with Broomean fairness on which it relies and its practical implausibility: a morally impermissible amount of work is left to chance and there is an implausibly large tolerance for the built-in lack of its ability to guide our actions: (almost?) anything goes because nothing is completely right or wrong.[26]

§1.4. Gertken's 'Aggregative' Mixed Solution

In this final subsection of the first half of this chapter, I discuss Jan Gertken's (2016) 'aggregation-friendly' version of the mixed solution. Gertken develops his version of the mixed solution in response to Peterson's.

Gertken's main objection to Peterson's mixed solution focuses on Peterson's advocacy of Individualism. As a reminder:

> Individualism: All moral values, including fairness, are individual values which can at most be aggregated *intra*personally, but not *inter*personally across people. (Peterson, 2010, p. 445).

[25] The quote is from Stratton-Lake in his Introduction to Ross (2002). So, this is Stratton-Lake commenting on Ross, but for the reference I had to use "Ross".

[26] This might sound like an overstatement, but I don't think it is an overstatement because I have argued earlier that, on Peterson's account we neither have to do that which "force" has the greatest strength nor do we have to follow through with whatever result his "weighted force rule" generates.

And now reconsider the 1,000 vs. 1,001 case:

Save 1,000 or 1,001 (under Individualism)?

	A1		A1000	B1		B1001
Save 1000	$1 - u$	…	$1 - u$	$0 - u$	…	$0 - u$
Save 1001	$0 - u$	…	$0 - u$	$1 - u$	…	$1 - u$
Fair Lottery	0.5	…	0.5	0.5	…	0.5

Gertken points out that on Peterson's account "if not giving everyone a chance is unfair to a very high degree, there will be a point at which holding a lottery in order to settle the question of whom to save is *better for everyone* involved than all relevant alternatives." (2016, p. 62). And, Gertken continues, that this condition is satisfied in the 1,000 vs. 1,001 case:

> Peterson suggests that $u > 0.5$ if group *B* is directly saved in 1000 vs. 1001, and that this implies that the outcome represented by $\{1, u\}$ is worse for everyone in group *B* than the outcome described by $\{0.5, 0\}$. It would follow that the lottery option is better for everyone in group *B* than the option of being directly saved. By the same reasoning, one can argue that a lottery is better for everyone in group *A* than the option of being directly saved (since $u^* > u$). Thus, accepting Peterson's description and evaluation of the relevant outcomes allows one to conclude that the lottery option is Pareto better than all relevant alternatives (…). Given that there are no further relevant facts, we can plausibly infer from this that all relevant reasons, which are based on the impact that each action has on what is good for the respective individuals, count in favour of the very same action, i.e. holding a lottery. (Gertken, 2016, p. 62).

But as Gertken objects, it is a most counterintuitive implication of Peterson's account that holding the lottery is Pareto better than being directly saved for the members of the group who we would directly save: "the suggestion that it is at least as good for you to get e.g. a 50 % chance of survival as it is to be unfairly saved for sure is very implausible, and Peterson's account implies similar results for cases where the chances of winning the lottery are *very* low." (Gertken, 2016, p. 64).[27]

[27] Gertken's own view is that it cannot be unfair to "get more than one can rightfully claim", but he grants Peterson this assumption for the sake of argument (Gertken, 2016, p. 64). In my next chapter, I argue why I think it is unfair to receive a surplus chance to or portion of a good to which one has no claim. Getting one's fair share cannot (also) mean getting more than you deserve, or so I think at least.

Furthermore, on his account, Peterson cannot accommodate the plausible intuition that a lottery is ruled out on moral grounds in cases where the difference in number is huge, e.g. a 1 vs. 1 million case (Gertken, 2016, p. 63).

For his own positive proposal for a mixed solution, which he formulates in terms of moral reasons, Gertken explicitly makes the following two assumptions:

- Each person involved in [an equal-claims-indivisible-good] case has a positive claim to be saved, which provides the agent with a corresponding reason to save her/him.
- Each person involved in [an equal-claims-indivisible-good] case has a claim not to be treated unfairly, which corresponds to a reason against giving this person a lesser chance than is given to anyone else. (Gertken, 2016, p. 65).

Gertken then applies his mixed solution to two cases, a 1 vs. 1,000 case and the 1,000 vs. 1,001 case, which are represented in the following two tables (Gertken, 2016, pp. 65, 66):

1 vs. 1000

Group A (n = 1)	Group B (m = 1000)
a_1: gives the agent a reason to save A	b_1: gives the agent a reason to save B, … … b_{1000}: gives the agent a reason to save B
a_1: gives the agent a reason against unfair treatment of a_1	b_1: gives the agent a reason against unfair treatment of b_1, … … b_{1000}: gives the agent a reason against unfair treatment of b_{1000}

1000 vs. 1001

Group A (n = 1000)	Group B (m = 1001)
a_1: gives the agent a reason to save A, … … a_{1000}: gives the agent a reason to save A	b_1: gives the agent a reason to save B, … … b_{1000}: gives the agent a reason to save B b_{1001}: gives the agent a reason to save B

a_1: gives the agent a reason against unfair treatment of a_1, … … a_{1000}: gives the agent a reason against unfair treatment of a_{1000}	b_1: gives the agent a reason against unfair treatment of b_1, … … b_{1000}: gives the agent a reason against unfair treatment of b_{1000} b_{1001}: gives the agent a reason against unfair treatment of b_{1001}

Gertken's final move is to explain these two tables and how his proposal outperforms Peterson's at least on the point that in cases where the difference in number is great, e.g. 1 vs. 1000, so that a lottery is ruled out because "it does not assume that members of the larger group give the agent reasons that count in favour of holding a lottery in the cases under consideration." (Gertken, 2016, p. 65, fn. 10):

> The lowercase letters with indices stand for different members of the respective groups. For each person involved, the tables mention the reasons that this person provides the agent with. I have highlighted the descriptions of those reasons that need to be compared and balanced if the agent wants to decide whether or not he should save the greater number in each case. Whereas all reasons corresponding to individual claims to be saved also count in favour of saving the respective groups, the fairness-based reasons corresponding to the claims of the members of the smaller group count against saving the greater number.
>
> In the 1000 vs. 1001 case, there are more reasons counting against saving the greater number than there are reasons in favour of this action, given that both the reasons in favour of saving members of group *A* and the reasons against the unfair treatment of the members of *A* also count against directly saving group *B*. In the 1 vs. 1000 case, this is not so. This leads to a mixed solution if we assume that, therefore, the relations count against saving the greater number are strong enough to tip the balance in the 1000 vs. 1001 case, but are outbalanced by the reasons in favour of saving the larger group in the 1 vs. 1000 case. In other words, a mixed solution results if we assume that the balance of reasons is sensitive to the number of reasons that count in favour of and against the various actions. (Gertken, 2016, p. 65).

Gertken's mixed solution seems most plausible.

However, as he himself anticipates, Gertken's 'aggregative' proposal runs into what he calls "the *problem of small losses*" (2016, p. 67). Gertken argues this problem comes in two forms. First, as a tie-breaker, such as in F. M. Kamm's "Sore Throat" case (1998, p. 101). And, second, in a one vs. many cases, such as T. M. Scanlon's World Cup Case, where the many would suffer a greater aggregate total of very small losses compared to the one's great loss (1998, p. 235).[28] Gertken, understandably, given the purposes of his paper, does not further pursue these issues apart from pointing to Kamm's principle of irrelevant utilities (1998, p. 101) and Alex Voorhoeve's (2014) account. That is, Gertken makes an appeal to "relevance" to solve his "problem of small losses" but does not discuss how this actually solves his problem. As I have argued in §2 of the second chapter of my thesis, this is not good enough. "Relevance" needs to be justified, not simply appealed to or presumed to be correct.

However, given the main topic of my thesis, I will, in the next §2, apply my own version of the mixed solution to both equal-claims cases as well as unequal-claims cases. Unlike Gertken, I will not make any appeal to "relevance" for my mixed solution.

But as I understand his mixed solution, there seems to be a problem with Gertken's proposal already in equal-claims cases. I fear that the combination of his understanding of fairness-based reasons and the aggregative structure of his account will make it difficult for him to justify the lottery in the 1,000 vs. 1,001 case. To explain this, remember that Gertken claims that the different kinds of reasons can be aggregated and then reconsider the following tables that represent the same two cases that Gertken is concerned with. All I have done is made explicit the total number of reasons we have of (1) each kind – so saving lives and unfairness reasons, and (2) the two kinds combined to give us the total number of reasons we have for directly saving (a) group *A* and (b) group *B*:

Group *A* (n = 1000)	Group *B* (m = 1001)
a_1: gives the agent a reason to save *A*, …	b_1: gives the agent a reason to save *B*, …
…	…
a_{1000}: gives the agent a reason to save *A*	b_{1000}: gives the agent a reason to save *B*
	b_{1001}: gives the agent a reason to save *B*

[28] Patrick Tomlin and I (forthcoming) argue why it is better to keep Kamm's tie-breaking cases and Scanlon's one vs. many cases separate.

(1) Subtotal of 1000 reasons in favour of saving *A* (because 1000 lives saved)	**(1) Subtotal of 1001 reasons in favour of saving *B* (because 1001 lives saved)**
a_1: gives the agent a reason against unfair treatment of a_1, … … a_{1000}: gives the agent a reason against unfair treatment of a_{1000} **(2) Subtotal of 1000 reasons in favour of saving *A* (because unfair treatment of each member of *A* should be avoided)**	b_1: gives the agent a reason against unfair treatment of b_1, … … b_{1000}: gives the agent a reason against unfair treatment of b_{1000} b_{1001}: gives the agent a reason against unfair treatment of b_{1001} **(2) Subtotal of 1001 reasons in favour of saving *B* (because unfair treatment of each member of *B* should be avoided)**
Total number of reasons in favour of saving *A*: 2000	Total number of reasons in favour of saving *B*: 2002

So, rather than advocating a 'traditional' mixed solution which would entail saving the greater number if the difference in the number of people between the two groups is sufficiently great and holding an equal-odds lottery if the difference in number is insufficiently great, Gertken's proposal stacks the deck in favour of saving the greater number in all cases. The reason for this is that his proposed fairness-based reasons should be added to the number of life-saving reasons we already have for saving that group. By the same reasoning, in the 1 vs. 1,000 case, we have two reasons to save the one and two thousand reasons to save the 1,000. So, although in the 1 vs. 1,000 case Gertken's argument enshrines the result that we intuitively would like to have – saving the greater number – for cases where despite the difference in numbers we, for the sake of argument, seek to justify holding a lottery – e.g. 1,000 vs. 1,001 case – Gertken's proposal actually makes it harder for us to justify the lottery and that is, ironically, because he has introduced fairness-based reasons which need to be added to the goodness reasons of saving lives we have for either group.

Despite my criticisms of all three versions of a mixed solution discussed here, I think that the work done by each of Hirose, Peterson, and Gertken has been invaluable. Each's proposal has contributed to making more precise the two morally relevant considerations of goodness (as in saving lives) and (un)fairness and the relation between them in these trade-off cases. More

generally, mixed solutions help us to consider specific proposals of how to understand the relations between different morally relevant considerations or (value) pluralism as such.[29] In the next section, relying on different parts from each of the three proposed versions of a mixed solution in this first section, I will argue for my own version of the mixed solution.

§2. My Mixed Solution

My proposed version of the mixed solution relies on the following three assumptions.

First, I argue that unfairness both intrapersonally as well as interpersonally aggregates. Hence, I reject Peterson's Individualism.

Second, and this is what separates my version of the mixed solution from the three versions I have discussed and rejected in §1, I argue that the intrapersonally aggregated total of unfairness, which I will represent with "*U*", is determined by two factors: one 'number-dependent' factor, and one 'number-independent' factor. These two factors that determine the intrapersonally aggregated size of *U* are:

(1) The number of "pairwise fairness violations" (number-dependent); and
(2) The proportionate strength of the unsatisfied claim if we ran a (fairly) weighted lottery (number-independent).

Only once the intrapersonally aggregated total of unfairness, or *U*, has been determined for each claimant whose claim wouldn't be satisfied do we interpersonally aggregate *U* to get us the total amount of unfairness.[30] Let me now elaborate on both the originality of my proposal as compared to the three mixed solutions discussed in §1 and the details of my proposal.

All three of Hirose, Peterson, and Gertken take the (negative) value of (un)fairness into account in each of their versions of the mixed solution, although of course in different ways. Though there

[29] In this way, I take the project of mixed solutions to be different from the project of 'consequentialising' (first-order) moral theories. In mixed solutions, and this is I believe for the better, two open questions are (1) whether we can weigh competing/different moral considerations or values; and (2) whether, even if we answer (1) affirmatively, we should do what is of greatest total value. On the 'consequentialising' project, we are showing our moral fault by not questioning the affirmative answers to both these questions. I thank Ralf Bader for his patient explanation to me of how exactly to understand the 'consequentialising' project and why he thinks it ultimately fails.

[30] So, to be clear, *U* in my account refers to the "intrapersonally aggregated total of unfairness", whereas *u* in the three accounts I have discussed in §1 refer to "unfairness", full stop. For clarity's sake, I have therefore taken *U*, rather than *u*, in my account.

is both explicit discussion between the three authors whether unfairness should or shouldn't be interpersonally aggregated – with Hirose and Gertken arguing unfairness should be interpersonally aggregated and Peterson rejecting such interpersonal aggregation of unfairness via his Individualism – and all permitting the intrapersonal aggregation of unfairness, none of the authors specify how unfairness is to be intrapersonally aggregated. The specification I propose for measuring the intrapersonally aggregated total of unfairness, which I represent with *U*, is determined by two factors that centre around the relative gap between the competing individual claims. As a next step, this intrapersonally aggregated total of unfairness, or *U*, is then interpersonally aggregated by the number of claimants whose claims we wouldn't satisfy. Let me explain.

The first factor determining the size of *U* is in terms of the number of what I shall call "pairwise fairness violations" that might occur.[31] So, for instance, in this case say that A is the stronger claimant and B_1, B_2,…$B_{1,000,000}$ are all individual members of the group of the many weaker claimants. If we satisfy the many weaker claims, rather than A's claim, then A suffers a pairwise fairness violation with each of B_1, B_2,…$B_{1,000,000.}$ So, the more of these pairwise fairness violations there are, the larger the (intrapersonal) aggregate sum of unfairness, *U*, suffered by A. By contrast, if we satisfy A's claim instead, then each of the Bs only suffers a single pairwise fairness violation because it is only A who has her claim satisfied. So, the number of pairwise fairness violations is the first part which determines the size of *U*. And *U*, in turn, accompanies each alternative that is part of the distributive choice we face.

The second factor which determines the size of *U* is a 'number-independent' factor: the proportionate strength of the unsatisfied claim if we ran a (fairly) weighted lottery.

To illustrate the proportionate strength of the unsatisfied claim if we ran a weighted lottery as relevant to determining the size of *U*, take first the simpler one vs. one equal claims case. In this equal-claims case, which of the two claims we do not satisfy does not make a difference in

[31] This idea is based on the following persuasive proposal by Julia Mosquera (2017): "While eliminating deprivations leads to the elimination of inequalities, reducing the incidence of deprivation leads to an uneven distribution of the pairwise relations of inequality of a population, which leads to the concentration of pairwise relations of inequality in the worse off." (Mosquera, 2017). More specifically, I look to apply her proposed principle of "*distribution-sensitive pairwise comparison of inequality* (DSP)" in the context of unfairness (Mosquera, 2017, p. 965). Given that, like Broome, I think that the badness of inequality is explained by the harm of unfairness – something Mosquera appears to suggest as well (Mosquera, 2017, p. 963) – I believe Mosquera's account of badness of this dimension of inequality reduces to my account of the disvalue of unfairness. The idea of 'pairwise comparison' stems from Thomas Nagel (2012), see especially (2012, pp. 125–127). F. M. Kamm (2005) discusses two different versions of pairwise comparison, "the Context-Aware View" and "the Blinker View", see (2005, pp. 8–12).

determining the size of *U* in this respect. *U* would be ½ or 0.5 because that would be the chance that the claimant should have received if we had held a lottery. I need to say more about this number-independent factor which determines the size of *U* in unequal-claims cases, but I will leave that for when I get to those unequal-claims cases later in this §2.

The main positive implication of this way of measuring unfairness is that any value must be 'personal'. Only values that are related to people, things that are good or bad for them, should make a difference. No impersonal value should impact our decision in these cases, at least in terms of "what we owe to each other". As hinted at in the previous sentence, this personal/impersonal distinction makes for a nice fit with Scanlon's contractualism, as Scanlon makes the same distinction and similarly rules out impersonal value from (at least directly) determining "what we owe to each other" (Scanlon, 1998, pp. 218–223).[32] To be clear, this is not an implication that separates my version of the mixed solution from any of the three versions I have discussed in §1. I believe all three authors got this right. But therefore, it is important that my proposed version of the mixed solution gets this right as well. Again, what is distinctive of my mixed solution is how *U* is measured.

My third assumption is that all claims, so including unequal claims, are commensurable. The upshot of this is that the satisfaction of any claim can be represented with one and the same value, I will take *G* for "goodness". Like unfairness, goodness both intrapersonally as well as interpersonally aggregates. Next, more generally, between unequal claims we can then stipulate a commensurable goodness ratio. Of course, I understand that a lot of work needs to be done to justify the ratios that I am stipulating, but I hope you will grant me the point for the sake of argument.

Finally, and to repeat, I will include cases where we have claims of unequal strength. To be clear, this is not a necessary assumption for my mixed solution but rather to see how it performs in unequal-claims cases. In other words, it is to immediately test the plausibility of my mixed solution as applied to cases which I think matter more than equal-claims cases because either there are more of these cases or they are more likely to occur than equal claims cases. However, I will first discuss the three equal-claims cases that Hirose, Peterson, and Gertken have relied on for both their accounts and criticisms of each other.

[32] I understand my view as an extension of John Broome's (2004) "Principle of Personal Good", which I take to rule out impersonal values (2004, p. 120).

§2.1. Case 1: Equal Claims – 1 vs. 5 (e.g. Life Boat)

For this Case 1 – and the remaining two cases hereafter, Case 2 and Case 3 – we are dealing with equal claims. One important consequence in dealing with equal claims is that there is no relative-strength gap between the competing claims – the claims are equal after all (e.g. 50/50 in a weighted lottery). And, therefore, a related second important consequence in equal-claims cases is that the size of *U* will only vary between the alternatives of our choice in terms the difference in the number of pairwise fairness violations. Now, consider Table 1 below.

Table 1.

	Total Goodness	Intrapersonally Aggregated Unfairness suffered by (each member of the group of) the unsatisfied claimant(s)	Total Unfairness (after interpersonally aggregating *U*)	Total Moral Value of Action
Save the One	*G*	*U**	*U**x5	*G*-(*U**x5)
Save the Five	*G*x5	*U*	*U*	(*G*x5)-*U*

Regarding total goodness, if we satisfy the many claims, then total goodness is five times greater compared to when we would satisfy the one claim instead. Regarding *U*, let's consider the two factors again each in turn.

First, if we satisfy the one claim, then there is a single pairwise fairness violation suffered by each of the five. If we satisfy the five claims instead, then the one claimant suffers five pairwise fairness violations – one with each of the five.

Second, as mentioned, *U* and *U** are equal in terms of the 'number-independent' factor which in part determines either their size.

Combining the two points on determining the size of *U* and *U** for either option in Case 1 we get the following: if we satisfy the one claim, then *U**=(1/2). If we satisfy the five claims, then *U*=5x(1/2). *U*=5x(1/2)=2.5. So, if we save the five, then *U* is 2.5. *U**=(1/2)=0.5. So, *U* is five times greater than *U**.

Now for the interpersonal aggregation of *U* and *U** to get the total amount of unfairness for either alternative. (*U**=0.5)x5=2.5. So, if we save the one, then total unfairness is 2.5. (*U*=2.5)x1=2.5. So, if we save the five, then total unfairness is also 2.5. So, the total unfairness of saving the one rather than the five is exactly equal to the total unfairness of saving the five rather than the one. This result provides original support for G. E. M. Anscombe's (1967) and John Taurek's (1977) idea that in one important sense – namely that in terms of unfairness – it is no worse to save the one rather than the five.

But fairness isn't everything. In terms of goodness, we get more to the ratio of 5:1 if we save the five rather than the one. Since unfairness gives us no reason to save either the one or the five – we should be indifferent between these alternatives, and so this explains why fairness requires an equal-odds lottery in this case – but goodness does give us reason to save the five rather than the one, and assuming that fairness *simpliciter* and goodness *simpliciter* are equally important, then we have greater reason to save the five directly rather than to hold an equal-odds lottery.

But why should we not hold a weighted lottery where we proportionalize chances according to the relative number of people in each group – i.e. 'pooling of chances', such as advocated by Kamm (1998, pp. 130–131) – as a compromise between goodness and fairness? The answer is that we have already accounted for all relevant aspects of fairness. Insisting on this kind of weighted lottery here, I think entails a commitment to valuing fairness *simpliciter* to some greater extent than goodness *simpliciter*. However, I cannot think of a good reason why I would accept this implied assumption. I find it much more plausible to assume that all morally relevant values matter equally.

The previous two paragraphs helps explain why for Taurek and the numbers sceptic's view more generally it is essential to claim that it makes no sense to say that in terms of goodness it is better to save more people rather than fewer.[33] If the number sceptic were to grant the point that goodness matters (at least) as much as fairness, then she would seem committed to the view that goodness determines what we ought to do in this case. Again, this is because goodness favours saving the five over the one, whereas fairness can only recommend indifference between the outcomes of saving the one and saving the five. So, fairness does not require the equal-odds lottery directly. The total unfairness of each alternative is the same size on my account. And if it is the same size, we should be indifferent between saving the one and saving the five from the perspective of fairness at least in terms of the outcome. But rather than plumping for either

[33] For an excellent discussion of this "No Worse Claim", arguably defended by Taurek, see Weyma Lübbe (2008).

alternative, we hold an equal-odds lottery to decide what to do because fairness also requires not to override claims. This is the justification for the equal-odds lottery.

§2.2. Case 2: Equal Claims – 1,000 vs. 1,001 (Kamm's Large-Scale Rescue Case)

As I have mentioned before, Kamm presents her Large-Scale Rescue Case of 1,000 vs. 1,001 equal-claims to argue for her idea that satisfying the additional claim does not outweigh the unfairness that would result from not having given the 1,000 their proportionate chance of having their claims satisfied (1998, p. 103). Kamm's idea is that satisfying an additional life-claim in a 1 vs. 2 case would be sufficient to outweigh the unfairness done to the one claimant's proportionate chance she would have had in the lottery. Let's map the Large-Scale Rescue Case using my account, as shown in Table 2 below.

Table 2.

	Total Goodness	Intrapersonally Aggregated Unfairness suffered by (each member of the group of) the unsatisfied claimant(s)	Total Unfairness (after interpersonally aggregating *U*)	Total Moral Value of Action
Save the 1,000	*G*x1,000	*U**	*U**x1,001	*G*x1,000-*U**x1,001
Save the 1,001	*G*x1,001	*U*	*U*x1,000	*G*x1,001-*U*x1,000

Regarding *G*, if we satisfy the 1,001 claims, then total goodness is slightly greater compared to when we would satisfy the 1,000 claims instead. Regarding *U*, let's consider the two factors again each in turn.

First, if we satisfy the 1,000 claims, then there are 1,000 pairwise fairness violations suffered by each of the 1,001. If we satisfy the 1,001 claims instead, then each of the 1,000 claimants suffer 1,001 pairwise fairness violations.

Second, again, *U* and *U** are equal in terms of the 'number-independent' factor which in part determines either their size because both each of the 1,000 and the 1,001 have a claim to have her life saved.

Combining the two points on determining the size of *U* and *U** for either option in Case 2 we get the following: if we satisfy the 1,000 claims, then *U**=1,000x(1/2). If we satisfy the 1,001 claims, then *U*=1,001x(1/2).

*U**=1,000x(1/2)=1,000x0.5=500. So, if we save the 1,000, then *U** is 500. *U*=1,001x(1/2)=1,001x0.5=500.5. So, if we save the 1,001, then *U* is 500.5. So, *U* is ever so slightly greater than *U**.

Now for the interpersonal aggregation of unfairness: *U**(=500)x1,001=500,500. So, if we save the 1,000, then total unfairness is 500,500. *U*(=500.5)x1,000=500,500. So, if we save the 1,001, then total unfairness is 500,000. So, total unfairness of saving the 1,000 rather than saving the 1,001 is equal to the total unfairness of saving the 1,001 rather than the 1,000.

This could be an interesting result. Kamm's intuition, which Hirose also seeks to justify, is that the additional good of one more life saved if we save the 1,001 should not override the claims the 1,000 have to a lottery. But here matters already get slippery. The 1,000 have claims to have their lives saved. They don't have claims to a lottery (as well). Now it is true that we avoid any procedural unfairness if we hold an equal-odds lottery in deciding whether to save the 1,000 or the 1,001 because all have 2,001 people have equally strong claims.[34] However, again outcome unfairness is unavoidable because we either will save the 1,000 or the 1,001 despite all having equally strong claims.

But in terms of both these outcomes, my mixed solution suggests that in terms of preventing unfairness we should be indifferent between saving the 1,000 over the 1,001 and saving the 1,001 over the 1,000. With total goodness slightly in favour of saving the 1,001 over the 1,000 – this point both Kamm and Hirose would agree with – there is thus additional merit to the claim that we should have the 1,001 over the 1,000. Our fairness-based reasons tell us to be indifferent between saving 1,000 and 1,001. But our goodness-based reasons tell us to slightly favour saving

[34] Here, I am relying on the distinction between outcome (un)fairness and procedural (un)fairness as suggested by Christian Piller (2017), which I will discuss further in my next chapter.

the 1,001 rather than the 1,000. Therefore, having paid due diligence to both fairness and goodness, we should save the 1,001 over the 1,000, directly.

My account subtly but importantly makes precise the question raised by Kamm and Hirose for this kind of Large-Scale Rescue Case. The question should be made more precise such that we are asked whether (1) it is more important that all 2,001 claimants should have a chance, proportionate to the strength of each of their claims, to have their claims satisfied – and so we hold an equal-odds lottery; or (2) having considered (a) goodness-based reasons slightly favouring the outcome of saving the 1,001 rather than the 1,000, and (b) fairness-based reasons favouring neither the outcome of saving the 1,000 over the 1,001 nor the outcome of saving the 1,001 over the 1,000, all-things-considered we should save the greater number. It seems like we could justify an equal-odds lottery in this case only if we either disproportionately overvalue (procedural) fairness or disproportionally undervalue goodness. Since I find neither of these options attractive, I think we should save the 1,001 directly rather than hold an equal-odds lottery in this Large-Scale Rescue Case. At the very least, Kamm and Hirose owe us an argument on why we should attach either greater importance to (procedural) fairness or lesser importance to goodness.

§2.3. Case 3: Equal Claims – 1,000 vs. 2,000 (Peterson's Counterexample to Hirose)

For my third case, I focus on Peterson's criticism of Hirose's account, when comparing a 1 vs. 2 equal-claims case and a 1,000 vs. 2,000 equal-claims case. Peterson's criticism was as follows: "For each person who is left behind, there are two others who are saved. Therefore, (…), it need not hold true that the unfairness *done to each person* is constant, regardless of the number of people concerned." Consider how my account maps this case as shown in Table 3 below.

Table 3.

	Total Goodness	Intrapersonally Aggregated Unfairness suffered by (each member of the group of) the unsatisfied claimant(s)	Total Unfairness (after interpersonally aggregating *U*)	Total Moral Value of Action
Save the 1,000	*G*x1,000	*U**	*U**x2,000	*G*x1,000-*U**x2,000
Save the 2,000	*G*x2,000	*U*	*U*x1,000	*G*x2,000-*U*x1,000

Leaving the story regarding total goodness aside, as it would be explained akin to Cases 1 and 2 above and because it is non-pertinent to Peterson's objection, let's focus exclusively on unfairness. As hopefully familiar by now, let's consider the two *U*-factors each in turn.

First, if we satisfy the 1,000's claims, then there are 1,000 pairwise fairness violations suffered by each of the 2,000. If we satisfy the 2,000 claims instead, then each of the 1,000 claimants suffers 2,000 pairwise fairness violations.

Second, again, *U* and *U** are equal in terms of the 'number-independent' factor which partly determines either their size because each of the 1,000 and the 2,000 in either group have a claim to have her life saved.

Combining the two points on determining the size of *U* and *U** for either option in Case 3 we get the following: if we satisfy the 1,000 claims, then *U**=1,000x(1/2). If we satisfy the 2,000 claims, then *U*=2,000x(1/2). *U**=1,000x0.5=500. So, if we save the 1,000, then *U** is 500. *U*=2,000x(1/2)=2,000x0.5=1,000. So, if we save the 2,000, then *U* is 1,000. So, *U* is twice as great as *U**.

Now for the interpersonal aggregation of *U* and *U** to get the total amount of unfairness for either alternative. (*U**=500)x2,000=1,000,000. So, if we save the 1,000, then total unfairness is 1,000,000. (*U*=1,000)x1,000=1,000,000. So, if we save the 2,000, then total unfairness is also 1,000,000. So,

the total unfairness in saving the 1,000 rather than the 2,000 is exactly equal to the total unfairness of saving the 2,000 rather than the 1,000.

So, just as in Case 1 and Case 2, here in Case 3 we have equally strong reason in terms of wanting to avoid (total) unfairness to save the 1,000 over the 2,000 as to save the 2,000 over the 1,000. Surprisingly, my account goes against Peterson's intuitive suggestion that there is no difference between the 1 vs. 2 and the 1,000 vs. 2,000 case because the ratio of claims is 1:2 in both cases. According to my account, these cases are the same because in terms of total (un)fairness there is no reason for favouring either alternative over the other in either case. On my account, there is no reason to favour one alternative over the other in either case because of the interplay between the intrapersonal and interpersonal aspects of unfairness.

But, like Case 1 and unlike Case 2, here in Case 3 total goodness strongly favours saving the 2,000 over the 1,000. To repeat, my mixed solution only presents the options of our choice in any trade-off case considered, rather than necessarily recommending going for either option. But on the plausible assumption that goodness and fairness are equally important, I think my mixed solution helps to show why it makes good sense to save the 2,000 directly rather than holding an equal-odds lottery.

§2.4. Case 4: Unequal Claims – One (Stronger) vs. Many (Weaker)

For my fourth case, suppose that we have one stronger claim competing with one million weaker claims. Suppose, furthermore, that the ratio of unequal claims is 1:1,000. So, the one strong claim needs a thousand of these weaker claims to be balanced in terms of total goodness. 1,001 weaker claims will mean that total goodness is slightly greater for satisfying these 1,001 weaker claims than total goodness would be if we satisfied the one stronger claim instead.

So, applied to this one stronger claim vs. one million weaker claims, there are two separate points that we need to take into account for goodness (*G*). First, *G* is multiplied by the number of claims that are satisfied. Second, to account for the ratio of unequal claims, which I suppose here to be 1:1,000, we need to divide the number of weaker claims satisfied by that ratio. So, to illustrate, if we satisfy the million weaker claims, then we have *G*x1,000, rather than *G*x1,000,000. If we had "*G*(x1)" for satisfying the one stronger claim and "*G*x1,000,000" for satisfying the million weaker claims, then I would have failed to account for the difference in strength between these competing claims at the individual level. "*G*(x1)" vs. "*G*x1,000,000" would represent the alternatives of

satisfying one claim vs. satisfying one million claims where all 1,000,001 claims in the case would be equal. But the claims are not equal in this case. Therefore, with a million weaker claims competing with a single stronger claim at this 1:1,000 ratio, we get *G*x1,000 and not *G*x1,000,000 if we satisfy the one million claims. The alternatives in terms of goodness in this one stronger claim vs. one million weaker claims case are "*G*(x1)" vs. "*G*x1,000".

Next, the number-independent factor that determines the size of *U* in a one vs. one case where the one claim, A's claim, is stronger than the other claim, B's claim, to the ratio of 1:1,000. It would have taken 1,000 of these weaker claims to exactly balance this stronger claim (held by A). However, there is only one such weaker claim. Although perhaps counterintuitive in terms of what is fair, for determining the size of *U* in this respect, we need to consider the proportionate chances each claimant would receive if we held a weighted lottery. After all, it is Broomean fairness upon which other advocates of mixed solutions have relied and a key component of Broomean fairness, which I will discuss in greater detail in the next chapter, and I have followed these other mixed solution advocates in relying on Broomean fairness.[35]

Proportionately fair chances in a weighted lottery for any unequal-claims case would, I contend, track the goodness ratio that captures the difference in strength between the competing claims. So, in this case, if we held a weighted lottery, the stronger claimant should receive a 999/1,000 chance and the weaker claimant a 1/1,000 chance of winning the lottery.

Finally, our one stronger claim (A) vs. one million weaker claims (B_1, B_2,…$B_{1,000,000}$) case, which is the case I want to focus on in this subsection. In this case, I presume that the goodness ratio is exactly the same 1:1,000 ratio as it was in the one stronger vs. one weaker claim case discussed in the previous two paragraphs. So, the starting point for determining the size of *U* for either alternative of our choice in this case is that again we take the fair individual chances these claimants should receive if we held a weighted lottery. So, if we save all the Bs, then *U* is 999/1,000 because that would have been the chance that A should have received in the weighted lottery. If we save

[35] It seems to me that the account of fairness relied on in any mixed solution does not have to be Broomean. If correct, then this would be crucial for the evaluation of the prospects of mixed solutions. It might be that the mixed solutions I have discussed in this chapter, including my own, are in (large) part deemed implausible because these solutions build on Broome's account of fairness. It would be both a shame and a mistake to draw the dismissive conclusion that, therefore, further exploration of mixed solutions is not worth our while. It might be that, for example, a mixed solution which relied on a Hookerian account of fairness would be more plausible. I will not explore this alternative Hookerian mixed solution here in this chapter. I thank Stephanie Collins both for presenting me with her objections to Broomean fairness and helping me to see that the account of fairness upon which a mixed solution relies does not necessarily have to be Broomean.

A instead, then the size of *U* will be matching the size in strength of any single one of the Bs' claims.

One crucial question that arises in this case because there is more than one weaker claimant: should we take into account that there are a million weaker claimants for *U* in the respect under discussion, for example, by aggregating the individual chances of the Bs (i.e. permit 'pooling' of individual chances amongst the Bs)? I contend that we should not. I consider pooling of individual chances to be a violation of the separateness of persons. None of the Bs has a claim stronger than the claim she actually has. Why do I state explicitly this obvious fact? I do so because one explanation of what goes on when we permit pooling of individual chances is precisely that each of the Bs receives a disproportionately greater chance to have her claim satisfied because each of B_1, B_2, and so on were permitted to pool their individual chances. The main reason why I think this is unfair is because A, who has a much stronger claim, would suffer a disproportionately smaller chance in the lottery as a result of the Bs pooling their chances. Why the strongest claim rather than the interpersonally aggregated sum of all claim-strengths of those whose claims we wouldn't satisfy, e.g. all the Bs? No single member of the group whose claims we wouldn't satisfy, e.g. B_2, has herself a claim that is of the greater size that would result if we interpersonally aggregated her claims-strength with that of at least one other member's claim, e.g. B_3.

What if the group of claimants whose claims we wouldn't satisfy was 'heterogenous' – variance of claim-strengths within the group – rather than 'homogenous' – all claims within a group being of equal strength? I contend that within the heterogeneous group of weaker claimants, it should be the strongest claimant who is taken as the 'representative' of the group of weaker claimants.

But this gives rise to the following critical question in response: would it not be unfair to the member(s) of the group whose claim(s) we consider satisfying to let weaker claims within the group of unsatisfied claims 'piggyback' on the strength of strongest claimant within their group of losers? After all, this will partly determine the size of *U*, and then in a domino effect the size of total unfairness of that alternative and ultimately the total moral value of that alternative.

My answer here requires further justification which I cannot give here, but I would appeal to the same separateness-of-persons line of reasoning to justify why it is the strongest member of the group of weaker claimants, and she alone, who determines the size of *U*. Suppose, for the sake of illustration, that it is B_1 who has the strongest claim out of the group of weaker claimants, the Bs.

If we were to, for example, instead take the average strength of the Bs' claims to determine the size of *U* in this respect, then this would be unfair to (at least) B_1. It would be unfair to B_1 because she will necessarily receive a disproportionately smaller chance because the average strength of the Bs' claims will be smaller than the proportionate chance she would receive given the strength of her own claim. Why? The group of Bs is heterogenous and B_1 herself has the strongest claim. This entails that the other Bs will have (varying degrees of) weaker claims, which will necessarily So, the complaint B_1 would have against this taking-the-average-strength-claim-of-the-group alternative is structurally identical to the previously considered complaint A would have against the pooling of individual chances by the Bs: it would be unfair to B_1 in her competition with A's claim.

Apart from what I have discussed, heterogenous groups bring further complexities that I cannot consider here in the initial discussion of whether mixed solutions are at all plausible. Only if the initial case for mixed solutions is sufficiently persuasive will it be worth our while to consider the added complexity of heterogenous groups and whether the above-described 'piggybacking' should be accounted for within the overall measure of (un)fairness.

Returning to our homogenous one vs. one million case, since the one million claims in group B are all of equal strength, we can randomly pick anyone of the members' claims of group B, e.g. B_{11}'s claim, to represent *U*. If we save group B instead, then the size of *U* will be matching the size of A as she has the strongest (because only) claim in group A. If we save A, then the fact that there are another 999,999 claimants with the same claim-strength as B_{11} competing with A's stronger claim is irrelevant for determining the size of this second factor that determines the intrapersonally aggregated total of *U*.

The important point is that *U* is much smaller if we satisfy A's claim and so, consequently, leave the claims of the Bs unsatisfied than if we save group B and so leave A's claim unsatisfied. That seems highly plausible to me. Any single one of the claims of the many, the Bs, is, after all, much weaker. Not satisfying, say, B_{11}'s claim is still unfair to some extent if we do satisfy A's claim – the extent being proportionate to the strength of B_{11}'s claim.[36] So, on that note, I have explained both the 'number-dependent' and the 'number-independent' factors that determine the (varying) size of *U*. To keep track of the varying size of intrapersonally aggregated unfairness as in part dependent

[36] This is because, as I will discuss in greater detail in my next chapter, I accept the Broomean view that overriding any claim, no matter how weak, is unfair.

on individual claim-strength, I will represent this as either *U* or *U** depending on the alternative of our choice considered (i.e. which of either group's claims wouldn't be satisfied).

In summary, the size of *U* is determined by two factors. One factor is 'number-dependent', whereas the other factor is 'number-independent'. These two factors that determine the intrapersonally aggregated size of *U* are:

(1) The number of pairwise fairness violations (number-dependent); and
(2) The proportionate strength of the unsatisfied claim if we ran a (fairly) weighted lottery (number-independent).

Once we have determined the intrapersonally aggregated total size of *U* by combining (1) and (2) above, then we can determine the total size of unfairness. We determine the total size of unfairness by taking *U* and multiplying that by the number of claimants whose claims we wouldn't satisfy.[37] By multiplying *U* by the number of claims we would not satisfy, we have interpersonally aggregated unfairness as well. The interpersonal aggregation of *U* is necessary to account for the fact that from the perspective of fairness it is worse for more rather than fewer people to suffer unfairness, regardless of the individual strength(s) of these claims, if we know we will satisfy some claims but not others.

Now, for Case 4. Remember: one vs. one million claims and the ratio for *G* is 1:1,000, so that the single stronger claim is balanced by a thousand of these weaker claims.[38] Comparing the options of either saving the one or saving the many, the results are represented in Table 4 below. As mentioned earlier, to help keep track of the fact that the size of *U* will vary in unequal-claims cases, whereas the size of *G* will not (inequality in claims' strength in terms of *G* is, after all, accounted

[37] It is true that the number of claimants whose claims we would not satisfy might be much smaller than the number of claims we would not satisfy, since some claimants might have multiple claims (in turn, possibly with different grounds). For the sake of clarity, I will presume that in my cases each claimant has only one claim, so that the number of claims and the number of claimants is equal.

[38] Patrick Tomlin (2017) has ingeniously spotted the hidden ambiguity in the Relevance View due to its (understandable) focus on homogenous group cases (2017, p. 238). Making the ambiguity explicit presents what Tomlin calls the "anchoring problem" in the Relevance View (2017, p. 238). Either interpretation of two possible interpretations, "Anchor by Strength" and "Anchor by Competition", is shown by Tomlin to violate highly plausible-looking principles and lead to highly implausible verdicts in cases considered (Tomlin, 2017, pp. 239–250). Joe Horton (2018), building on Tomlin's principles, has devised a seemingly fatal dilemma for any partially aggregative view. Although my account is looking to justify the same key intuitions that have propelled the Relevance View as an instance of a partially aggregative view, my account is a version of the mixed solution, not the Relevance View or a partially aggregative view. Whether my account avoids or answers both Horton's and Tomlin's objections is to be seen and tested, but at least it will not run afoul of them in its initial setup.

for by a ratio which has already been incorporated into *G*), I will use *U* and *U** in both my Case 4 here and my final Case 5, which I discuss in the next subsection.

Table 4.

	Total Goodness	Total Intrapersonally Aggregated Unfairness suffered by (each member of the group of) the unsatisfied claimant(s)	Total Unfairness (after interpersonally aggregating *U* or *U**)	Total Moral Value of Action
Save the One (Stronger)	*G*	*U**	*U**x1,000,000	*G*-(*U**x1,000,000)
Aid the Many (Weaker)	*G*x1,000	*U*	*U*	(*G*x1,000)-*U*

Regarding *G*, if we satisfy the many weaker claims, then *G* is a thousand times greater compared to when we would satisfy the one stronger claim instead. Regarding *U*, given that it is determined by two different factors on my account, two points need to be made.

First, if we satisfy the stronger claim, then there is a single pairwise fairness violation suffered by each of the million. If we satisfy the weaker claims instead, then the stronger claimant suffers one million pairwise fairness violations – one with each of the many.

Second, and independent from the number of pairwise fairness violations suffered, the size or seriousness of the pairwise fairness violation suffered by either A or each of the Bs is different. Why is it different? The size or seriousness of the pairwise fairness violation depends on the proportionate strength of each individual's claim. If we were to hold a weighted lottery between aiding group A or group B with chances proportionate to the strength of the individual claims in

either group, then group A should get a 999/1,000 chance and group B a 1/1,000 chance.[39] I propose that $U > U^*$ by this exact ratio of the 999/1,000:1/1,000 proportionate chances that we would assign if we were to hold a weighted lottery. Therefore, *U* is much greater than *U**.

Combining the two points on determining the size of *U* for either option in Case 4 we get the following: if we satisfy the one stronger claim, then $U^*=1/1{,}000$. If we satisfy the many weaker claims, then $U=1{,}000{,}000\text{x}(999/1{,}000)$. So, *U* is much larger than *U**.

But remember that both *U* and *U** represent only the intrapersonally aggregated total of unfairness for either alternative of our choice. To get the total amount of unfairness, we need to multiply each of *U* and *U** by the number of claimants who would each suffer this intrapersonally aggregated total of unfairness. Since *U* would only be suffered by A, we get *U*x1, or simply (still) *U*. By contrast, *U** is suffered by each of the million Bs if we don't satisfy their claims. So, *U**x1,000,000.

A couple of key question that now arise: (1) once multiplied by a million, what is the gap that remains between the two total sums of unfairness? And (2) once multiplied by a million, is it now even the case that the total unfairness is greater if we save A than if we aid all the Bs because *U**x1,000,000 is greater than *U*(x1), despite *U** being much smaller than *U*?[40]

To answer questions (1) and (2) in combination, consider the following: $U^*=1/1{,}000=0.001$. $0.001\text{x}1{,}000{,}000=1{,}000$. So, if we save the one, then total unfairness is 1,000. $U=1{,}000{,}000\text{x}(999/1{,}000)$. $999/1{,}000=0.999$. $0.999\text{x}1{,}000{,}000=999{,}000$. So, if we aid the many, then total unfairness is 999,000. So, total unfairness in aiding the many, rather than saving the one, is greater to the ratio of 999:1.

[39] As I will argue in my next chapter, on the Broomean account of fairness on which I am relying, it is impermissible to have 'pooling' (i.e. interpersonal aggregation) of individual chances. For this aspect of fairness, determining proportionate chances in a weighted lottery, it is irrelevant whether group B contains a million people or has just one member. In a 1 vs. 1 case with the same difference in individual claim strengths, the Broomean weighted lottery would also require a 999/1,000 chance for individual A and a 1/1,000 chance for individual B. In short, I think that F. M. Kamm (1998; 2007) is correct that we should permit interpersonal aggregation in the realm of fairness, but she jumps the gun. We shouldn't be interpersonally aggregating the individual proportionate chances. Weighted lotteries on her account I understand as a compromise between the values of goodness and fairness, not as a compromise within – or balance struck between different aspects of – fairness.

[40] I say "saving" with regards to A because she has a much stronger claim than any of the Bs. It sounds somewhat perverse to me to talk of "saving" the Bs when their claims are so much weaker. "Aiding" the Bs seems more appropriate. For my argument, nothing hangs on this bit of language though.

In terms of goodness, or *G*, the difference between our alternatives is as follows: total goodness in aiding the many, rather than saving the one, is greater to the ratio 1,000:1.

I am only minimally committed to the claim that my mixed solution presents us with the question whether we have greater moral value if we save the one or if we aid the many, without necessarily taking a (substantive) stance on that question. In this way, my proposal is akin to Hirose's formal aggregation. Like him, I am drawing a map of the (moral) terrain, but I am not necessarily telling you where to go. However, I think my proposal, more precisely than Hirose's, captures the two different considerations of goodness and fairness that pull in opposite directions in a one (much) stronger claim vs. many (much) weaker claims case. Everything hangs on how we should weigh goodness and (un)fairness against each other. But without either favouring or dismissing either consideration in advance, my proposal shows that in this Case 4 goodness and fairness pull us in opposite directions to almost the same degree (of strength): a 1,000:1 ratio vs. a 999:1 ratio.

I will now (re)consider a further case to further show how my mixed solution works in structuring the relations between the morally relevant considerations of goodness and fairness in these kinds of trade-offs between claims.

§2.5. Case 5: Unequal Claims – Many (Stronger) vs. One (Weaker)

Finally, Case 5, in which we have a million stronger claims competing with one weaker claim. Therefore, Case 5 is Case 4 but in reverse. You might wonder why on earth we would have to reflect on this case? Just save the many! Please note that I agree with that judgment, but for the purposes of providing an account of the interplay of the relevant values, goodness and (un)fairness, we need to take a step back and analyse this case and see what my account suggests.

In Case 5, we have one million claims vs. one claim and the ratio for *G* is 1:0,001, so that the single weaker claim is balanced by one thousandth of (the strength of) a single stronger claim (or, conversely and like in Case 4, one stronger claim would be balanced by one thousand of these weaker claims).[41] Comparing the options of either aiding the one or saving the many, the results are represented in Table 5 below.

[41] As a clarificatory note on the overall structure of this thesis, this would be a point where the discussion of "partial leftovers" as discussed in §6.2. of chapter 3 would be relevant. Alas, I cannot go into that discussion (again) here.

Table 5.

	Total Goodness	Total Intrapersonally Aggregated Unfairness suffered by (each member of the group of) the unsatisfied claimant(s)	Total Unfairness (after interpersonally aggregating *U* or *U**)	Total Moral Value of Action
Aid the One (Weaker)	*G*	*U**	*U**	*G*-*U**
Save the Many (Stronger)	*G*x1,000,000,000	*U*	*U*x1,000,000	*G*x1,000,000,000-*U*x1,000,000

Regarding total goodness, if we satisfy the many stronger claims, then total goodness is a billion times greater compared to when we would satisfy the one weaker claim instead. Regarding *U*, let's consider the two factors again each in turn.

First, if we satisfy the weaker claim, then there is a single pairwise fairness violation suffered by each of the million. If we satisfy the stronger claims instead, then the weaker claimant suffers one million pairwise fairness violations – one with each of the many. So, in terms of the number of pairwise fairness violations, Case 5 is exactly the same as Case 4.

Second, $U^* < U$ by the ratio of the 1/1,000:999/1,000 proportionate chances that we would assign if we were to hold a weighted lottery. Simply put, *U** is smaller than *U* to the same extent as *U* was greater than *U** in Case 4. It is the ratio of Case 4 in reverse.

Combining the two points on determining the size of *U* for either option in Case 5 we get the following: if we satisfy the one weaker claim, then $U^*=1x(999/1,000)$. If we satisfy the many stronger claims, then $U=1,000,000x(1/1,000)$. So, *U* is larger than *U**. However, *U* is not larger than *U** by even remotely the same proportion as it was in Case 4. The reason why? The million pairwise fairness violations are now suffered by the one (much) weaker claimant, not the one

(much) stronger claimant. But remember that we still need to interpersonally aggregate unfairness as well.

*U**=999/1,000=0.999. 0.999x1,000,000=999,000. So, if we aid the one, then total unfairness is 999,000. *U*=1,000,000x(1/1,000). 1/1,000=0.001. 0.001x1,000,000=1,000. So, if we save the many, then total unfairness 1,000. So, total unfairness in aiding the one, rather than saving the many, is greater to the ratio of 999:1. The numbers match how descriptively Case 5 is Case 4 tilted on its head, at least in terms of unfairness.

In terms of total goodness, the difference between our alternatives is as follows: total goodness in saving the many, rather than aiding the one, is greater to the ratio 1,000,000,000:1.

These numbers seem to make good sense. Compared with Case 4, first, there is much more at stake in terms of total goodness between the alternatives of our choice in Case 5. Second, in terms of wanting to avoid greater unfairness, we have further very strong reason in Case 5 that we should satisfy the many stronger claims rather than the one weaker claim. Fairness has a role to play and is duly considered. What it does is strengthen the case for saving the many, which was already overwhelmingly plausible in terms of considering only goodness.

§2.6. Summary

In summary, in both equal-claims cases and unequal-claims cases, two values are at play which exhaust the moral terrain considered: (aggregate) goodness, which is generated through the satisfaction of claims, and unfairness, which is generated through the failure to satisfy (some) claims whilst other claims are satisfied.[42] What is distinctive of my mixed solution is how the intrapersonally aggregated total of unfairness, or *U*, is measured. One factor in determining the size of *U* is 'number-dependent', whereas the other factor is 'number-independent'. These two factors that determine the size of *U* are:

[42] I think it is true that if people do not have legitimate claims on benefits ('goodness gains') beyond the threshold of need unless these people deserve the benefits or have been promised the benefits, then these unneeded and undeserved goodness-gains are outside the scope of 'satisfaction of claims'. These goodness-gains are part of impersonal good. However, precisely because I have drawn up my proposed mixed solution within the parameters of "what we owe to each other", I have ruled out impersonal values, such as impersonal good, from playing a role in this sub-domain of morality. Intuitively, I don't want to reject the thought that when there is a lot of stake in terms of impersonal good and/or morality which is not "what we owe to each other" these considerations outweigh what we are required to do according to 'personal' good or "what we owe to each other". Hence, I don't think I'm a Scanlonian contractualist.

(1) The number of pairwise fairness violations (number-dependent); and
(2) The proportionate strength of the unsatisfied claim if we ran a weighted lottery (number-independent).

Once we have determined *U* for either alternative of our choice in the case under consideration, we need to multiply *U* by the number of claimants whose claims we wouldn't satisfy (i.e. we interpersonally aggregate unfairness) to get the total amount of unfairness of choosing either alternative in this case.

One intriguing general result my mixed solution seems to generate is the following: if competing claims are equal, then, even if the number of people we could save per alternative differs, in terms of preventing the greatest sum of total unfairness we have no reason to favour the saving of the many over the one/few nor for saving the one/few over the many.[43]

In determining what we ought to do we need to compare total goodness-total unfairness of each alternative of our choice and do what gives us the largest sum of moral value. Can we simply read off what the largest (numerical) sum of moral value is? No. It is an open question whether goodness and fairness are commensurable such that they could be put on the same cardinal scale. However, even if we consider goodness and fairness to be incommensurable, my account helps systematize the dynamic interaction between these two values which might pull us in different directions in different cases. At the end of the day, we need to make a judgment about what is right. My account helps us make the decision in an informed way by pointing us to the overall structure and structurally relevant factors that we need to consider.

For the bigger picture of my thesis, I want to highlight why I think the prospects of my mixed solution look much better than any version of the Relevance View. This is because my mixed solution does not have its plausibility plummet in the way this happens to any version of the Relevance View when we switch from two-groups cases to three(+)-groups cases.[44] The

[43] This result could be important insofar as Johann Frick (2015), in light of his criticism of Scanlon's contractualism, has suggested that Scanlon's contractualism is better understood as governing only a sub-domain of interpersonal morality or "what we owe to each other" having to do with "fairness" only, rather than the whole domain of "what we owe to each other". If there is anything to my proposed mixed solution here, then Scanlonian contractualists would be wise to reject Frick's suggestion. On my account of fairness discussed here in these cases, it would tie Scanlonian contractualists to the number scepticism that they are (sensibly) looking to avoid.
[44] For illustration, consider the version of this objection to the Relevance View as presented by Privitera (2018), which I have discussed in §4 of my third chapter.

explanation why is as straightforward as it could be: apart from Gertken's version, mixed solutions don't rely on "relevance".

As I had indicated when I set out to explore mixed solutions as my second strategy to justify Scanlon's intuitions which underpin Limited Aggregation, mixed solutions do not limit aggregation (whether via "relevance" or any other notion). Mixed solutions operate on the (background) presumption of Pure Aggregation. With mixed solutions, the idea is that rather than limiting aggregation in any way, it is rather by incorporating at least one more value – (un)fairness – which seems part of the moral landscape that it does not lead to the (for Scanlon at least) counterintuitive results that Pure Aggregation leads to in some cases. Even though my mixed solution cannot justify that there is, e.g. no number of viewers watching the World Cup that makes it the case that we should save Jones.

But, I contend, the following two considerations make it worth Scanlonian contractualists' while to give up on Scanlon's intuition in the World Cup Case. First, it puts saving the many in an entirely different perspective than on a 'traditional' linear aggregation account (e.g. classical utilitarianism). Second, the (much more) promising outlook of my mixed solution because of its (greater) practical feasibility (than the Relevance View) – it could be applied to any case with any number of groups of competing claimants (instead of being limited to two-groups cases). I urge Scanlonian contractualists to take this deal on my mixed solution.

On a final note, my mixed solution could be combined with prioritarianism (e.g. Restricted Prioritarianism, which I have defended in my fourth chapter). However, my mixed solution could equally be combined with linear aggregation, which may or may not, in turn, be complicated by adopting, e.g. the Competing Claims View on another version of (relational) egalitarianism instead of prioritarianism. The important general point here is that my mixed solution can be combined with any version of a purely aggregative view.

Conclusion

In this chapter, I have discussed three different versions of the so-called "mixed solution", Hirose's "formal aggregation", Peterson's 'higher-order' or 'incommensurable' mixed solution, and Gertken's 'aggregative' mixed solution. All mixed solutions consider balancing goodness against fairness to determine what we ought to do in trade-off cases with the aim of justifying the intuition

that when the difference in number is sufficiently great, we save the greater number, but when the difference in number is insufficiently great, then we hold an equal-odds lottery.

My 'hidden' agenda in discussing mixed solutions has been to see if any of these three proposals could justify the intuitions underlying Scanlon's contractualism as a partially aggregative view. Although I have rejected each of Hirose's, Peterson's, and Gertken's proposals for different reasons, I do think both that their work has been invaluable and that exploring different versions of the mixed solution is the way forward, especially considering the absolute mountain of counterexamples and plausible principles violated by different versions of the Relevance View. It is better to develop a principled(!) value-pluralism, rather than attempting to constrain a value-monism via a relevance/irrelevance-distinction.

For my own proposal of a mixed solution, I have relied on what I take to be a key insight from Mosquera (2017) in the debate on egalitarianism: to develop the idea of what I called "pairwise fairness violations". The number of these pairwise fairness violations are the first determinant out of two which tell us what the size of intrapersonally aggregated unfairness, or *U*, is in trade-off cases. My account of unfairness as applying to these cases is as follows:

> The size of *U* is determined by two factors. One factor is 'number-dependent', whereas the other factor is 'number-independent'. These two factors that determine the size of *U* are:
>
> (1) The number of pairwise fairness violations (number-dependent); and
> (2) The proportionate strength of the unsatisfied claim (number-independent).
>
> Once we have determined *U* for any alternative of the choice faced in the case under consideration, we interpersonally aggregate unfairness by multiplying *U* by the number of claimants whose claims we wouldn't satisfy. This gives us total unfairness for that alternative of our choice.

But fairness isn't everything. Total unfairness needs to be balanced with total goodness. Total goodness of any alternative in these cases is the sum of satisfied claims of the people we would help.

If anything, I consider my account as an improved version of Hirose's "formal aggregation". Rather than immediately taking a stand, my account better maps the relevant normative terrain.

Reconsidering the bigger picture of my thesis, my mixed solution requires Scanlonian contractualists to give up on the Individualist Restriction, but not on the rationale that I understand as underlying it: at least within the moral domain of "what we owe to each other", what matters most that is each of us is considered as an individual – justifiability to each person.[45] We should aggregate individual claims in terms of goodness, and so need to reject the Individualist Restriction.[46] If we assume they have claims, however much weaker they might be than the claims of their competitors, we should interpersonally aggregate claims to have mild headaches cured and watching the World Cup Final without interruption. Such interpersonal aggregation will give us some aggregate total sum of goodness for any alternative in the trade-off choice we face.

But goodness isn't everything. Even if give up on the Individualist Restriction, in terms of considering the unfairness done to any individual who will see their claims go unsatisfied in trade-off cases we are paying due diligence to the idea of justifiability to each person. My account presents a principled way of both structuring (i) both goodness and fairness separately and (ii) the relation between them.

The general lesson here is that there is nothing wrong with the rationale for the Individualist Restriction, that rationale being justifiability to each person. But the Individualist Restriction, understood as all-encompassing (i.e. as applying to all values that make up our pluralism) simply is not the way in which we should want to cash out that rationale, given the implausible results entailed by the Individualist Restriction. Ironically, we need to aggregate within different values, such as fairness and goodness, to act in accordance with the regulative ideal of justifiability to each person.

[45] I have deliberately used the phrase "justifiability to each person" because Parfit thought that out of the many valuable lessons Scanlon's (1998) contractualism has taught us, this was the most important one. I think Parfit was right about this. Even though I disagree with most substantive parts of Scanlon's contractualism, the original way Scanlon has argued for justifying our conduct to each person as at least one central regulative ideal that should guide our practical reasoning is at the centre of Scanlon's lasting philosophical legacy. I thank Parfit and Scanlon for their conversations on contractualism with me.

[46] In conversation, Scanlon has appeared open to the suggestion that at least claims relating to well-being may be interpersonally aggregated in his contractualism.

In conclusion, we should always aggregate claims, but we should not always maximize the good nor should we always do that which is (perfectly) fair. Instead, we should always do what is right, which in aggregation cases requires striking the right balance between these two different, possibly incommensurable, values.[47]

[47] To further illustrate using a poetic phrase from Henry Sidgwick (1981), although it not being in the broader terms of "self-interest" and "duty" but the narrower terms of "goodness" and "(un)fairness", this is one (more) way in which "practical [or moral] reason [is] being divided against itself[.]" (1981, p. 506). I thank Kent Hurtig for having educated me on Sidgwick's view and how we might 'enrich' Sidgwick's "dualism of practical reason" to be a 'triad' if the requirements of rationality were normative. For an excellent account which raises the question of the normativity of rationality, see John Broome (2013). Furthermore, I therefore reject Parfit's suggestion – not clearly explicitly advocated by him – in his closing remarks of his *On What Matters: Volume Three*: I believe we could never have sufficient impartial reason to act wrongly in any way. (Parfit, 2017, pp. 434–435). If we have sufficient impartial reason to act a certain way, then we must be acting rightly.

University of Reading

Measuring Unfairness

Abstract In this final chapter, I focus on John Broome's (1984; 1990; 1998; 2003a; 2004) theory of fairness to see if it could justify the intuitive judgments underlying Limited Aggregation independently from it being a crucial cog in the machinery of a mixed solution. Discussing Broome's theory of fairness as such is worthwhile because, as I will argue, Scanlon's understanding of what (un)fairness is cannot justify Scanlon's own intuitions. To have a well-focused discussion of Broome's account, I critically evaluate Christian Piller's (2017) take on Broomean fairness. I argue that if Christian Piller's distinction between 'outcome fairness' and 'procedural fairness' is incorporated within Broomean fairness, then this distinction vindicates a familiar Broomean result: fairness requires, as a second-best solution, that we always have a weighted lottery to distribute any indivisible good. Contrary to what Piller claims, Broomean fairness never requires directly giving the indivisible good to the stronger claimant over holding the weighted lottery. The upshot of Piller's work is that the intuitively plausible distinction between outcome fairness and procedural fairness can be integrated in Broome's theory of fairness, which increases the degree of plausibility Broome's theory as an account of fairness already had.

Key words Broome • Claims • Fairness • Lotteries • Piller

Introduction

Having discussed mixed solutions in the previous chapter, the focus of this final chapter is to get a better understanding of the (negative) value of (un)fairness. And to get a better understanding of fairness, it is helpful to first consider Scanlon's understanding of fairness.

Given Scanlon's aim of defending an account which is supposed to be an alternative to utilitarianism, I find it surprising that he has a rather 'minimal' understanding of fairness which is 'utilitarian-friendly'. Let me explain why I describe Scanlon's understanding of fairness as such. Scanlon argues that principles which "arbitrarily favour" one person over others are in this respect unfair and can for that reason be reasonably rejected:

> [E]ach person has reason to prefer partial principles that would favor him or her. If one of these principles is made binding, with no further reason to support it, then one person's reason for wanting to be favored is given precedence over others' similar reasons, without justification. This is what makes such a choice arbitrary, and makes the principle rejectable. This substantive objection applies to principles that make essential use of proper names as well as to those relying on "rigged" descriptions. (Scanlon, 1998, p. 212).[1]

On Scanlon's understanding of arbitrary favouring constituting unfairness, utilitarianism is a perfectly fair moral theory. Given utilitarians' strict adherence to impartiality and having every single person count equally, utilitarians do not arbitrarily favour some people over others. Therefore, Scanlon's understanding of fairness does not help him in separating his contractualism from utilitarianism. Therefore, I find this move by Scanlon unnecessarily self-undermining for his self-assigned purpose of presenting his contractualism as an alternative to utilitarianism.[2]

[1] Consider the following alternative formulation by Scanlon, just to drive the point home that he understands fairness as the 'mere' avoidance of arbitrariness: "[I]t is reasonable to object to principles that favor others arbitrarily. (…) We have reason to object to principles simply because they arbitrarily favor the claims of some over the identical claims of others: that is to say, because they are unfair. In the process of moral reflection that contractualism describes, this provides a perfectly understandable reason for finding partial principles objectionable, a reason that does not depend on a prior idea that such principles, or the practices they would permit, are wrong." (Scanlon, 1998, p. 216). With the remark that no appeal may be made to prior (intuitive) ideas of right and wrong in the rejection of principles, Scanlon shows that he respects what Derek Parfit (2011a) calls "the Deontic Beliefs Restriction"(2011a, pp. 366–370). If contractualists were to violate this restriction, their contractualism would be viciously circular.

[2] My surprise at Scanlon's 'minimal' account of fairness is, I believe, further supported by Scanlon's allusion to a richer account of fairness that he suggests by contrasting "fairness" with "equality" as two aspects of "distributive justice" in other work: "Rather than speaking of 'distributive justice', which can encompass a great variety of considerations, I will speak instead of fairness, as a property of processes (e.g. of competitions), and equality, as a property of resultant distributions." (Scanlon, 2003, p. 30). Of course, as I have discussed, Scanlon would no longer, given his (Scanlon,

Consider again Scanlon's own World Cup Case. One plausible justification for why we should always save Jones regardless of the number of viewers – and so, regardless of the greater aggregate good we would achieve if we were to satisfy the viewers' weaker claims – is that it would be unfair to Jones not to save him for the sake of helping a (very) large number of people each with much weaker claims than his.[3] Given that Scanlon's 'minimal' understanding of (un)fairness blocks an appeal to this kind of justification in one vs. many cases of which the World Cup Case is one example, I argue that Scanlonian contractualists should have a 'thicker' understanding of fairness. 'Thicker' fairness might then be of help to Scanlonian contractualists in justifying saving the one (or few) rather than the many in distant-harm cases.

Crucially, and unlike the first 'relevance' strategy, the pluralist nature of the second strategy – i.e. the appeal to further values or morally relevant considerations apart from aggregate good realised through the satisfaction of claims – does not limit aggregation. The second strategy operates on the presumption that we always aggregate claims (i.e. Pure Aggregation). But what the appeal to other values, specifically the (negative) value of (un)fairness, is supposed to do is to have the overall balance of morally relevant considerations – (aggregate) goodness vs. (un)fairness – in some cases point to saving the many, but to saving the one (or few) in other cases. Specifically, we should have a justification for saving the many in what I have called close harm cases and a justification for saving the one (or few) in distant harm cases. What account of fairness, unlike Scanlon's, could possibly deliver in terms of the above?

In this chapter, I focus on a recent defence of John Broome's (1984; 1990; 1998; 2003a; 2004) theory of fairness by Christian Piller (2017) to have focused discussion of considering an alternative account of fairness to Scanlon's. One reason for considering Broome's theory of fairness as an alternative to Scanlon's is that Scanlon does so himself, albeit ever so briefly (1998, p. 397, fn. 39). More importantly, Scanlon mentions Broome's account in the context of the Rescue Case and aggregation. Quite uninformatively, Scanlon's dismisses Broome's "more narrow" understanding of fairness in favour of his own because "if a decision is based on consideration of the relevant reasons, then it is not unfair in the sense I specified." (1998, p. 397, fn. 39). With this phrase, Scanlon refers to his tie-breaker argument in the Rescue Case and equal-harms cases more

2018), defend this telic egalitarianism. But what matters here is that Scanlon's quoted understanding of "fairness" here seems to be specifically "procedural fairness", but which in turn does not necessitate the narrow understanding of fairness as 'non-arbitrariness'. The narrowing-down of fairness as non-arbitrariness is a move Scanlon made in *What We Owe to Each Other* and it is still unclear to me why Scanlon did so.

[3] I think the language of "unfairness" offers another way of formulating Privitera's earlier-discussed (in my §4 of chapter 3) objection to Voorhoeve's account in three(+)-groups-cases.

generally. However, I have already discussed in my second chapter how Scanlon's tie-breaker argument fails because it is question-begging and makes an implicit and illicit appeal to aggregation. So, the alternative of Broomean fairness is still very much on the table and I will consider it here on behalf of Scanlonian contractualists. Again, for the sake of having a manageable scope of discussion, I focus on Christian Piller's defence and extension of Broome's theory of fairness.

Piller argues that Broome has been radically misinterpreted by his critics. Most importantly, Piller argues that within the Broomean framework we should distinguish between 'outcome fairness' and 'procedural fairness'. I have my doubts about whether this distinction is a fair reflection of Broome's original view and I will discuss these doubts in more detail in §1 below.

But I should admit from the start that this is a question of exegesis in which I am less interested. Independent of the question of whether the distinction between outcome fairness and procedural fairness is part of Broome's view, it is an interesting distinction. Although I find such a distinction intuitively plausible, in §2 and §3 I will discuss some worries I have about the distinction, at least within a Broomean framework of fairness.

My main conclusion is that if Piller's distinction between outcome fairness and procedural fairness is nevertheless incorporated within Broomean fairness, then this distinction vindicates a familiar Broomean result: fairness requires, as a second-best solution, that we always have a weighted lottery to distribute any indivisible good. Contrary to what Piller claims, Broomean fairness never requires directly giving the indivisible good to the stronger claimant over holding the weighted lottery. The upshot of Piller's work is that the intuitively plausible distinction between outcome fairness and procedural fairness can be integrated in Broome's theory of fairness, which increases the degree of plausibility Broome's theory as an account of fairness already had.

§1. Piller versus Broome on Fairness

The core principle of Broome's theory of fairness is 'treat claims proportionately [according to their strength]'. For my discussion, I will focus on cases with an indivisible good and two competing claims of unequal strength. So, these cases have the structure of aggregation cases discussed so far in this thesis, but they have been helpfully simplified to have only a single claimant on either side (i.e. in either 'group').

The main reason why, for the purposes of my thesis topic of aggregation, I do not need my discussion to extend beyond one vs. one cases to one vs. many cases is because, in the language of the previous chapter on mixed solutions, I take the discussion here to be a more detailed discussion of what falls under the heading of the number-independent dimension of unfairness. Therefore, whether there is only a single weaker claimant or as many as a million weaker claimants does not make a difference in this discussion of fairness.

One view about unequal-claims-to-an-indivisible-good cases is that fairness requires satisfying the stronger claim.[4] Broome, however, argues that in such cases, in order to 'treat claims in proportion to their strength', fairness requires that we should hold "a weighted lottery in which chances correspond to claim-strength: the stronger your claim, the bigger your chance ought to be of getting the [indivisible] good." (Piller, 2017, p. 227).[5] Piller suggests that, from within Broome's theory, this is not quite right. Piller claims that this is the right way to understand the implications of Broome's view when the two claims are close in strength but argues that it is not what the view would recommend when the claims are far apart in strength.

Piller is troubled by the idea that, in unequal-claims-to-an-indivisible-good cases where the competing claims are strongly unequal, we would be required by fairness to hold a weighted lottery. Piller is concerned about even giving a chance to the (much) weaker claimant to have her claim satisfied. Piller claims the following: "When claims are strongly unequal, Broome does not recommend holding a weighted lottery. Such a lottery would, on Broome's theory of fairness, be less fair than satisfying the stronger claim." (Piller, 2017, p. 229).

Let me just briefly comment on how I see a structural similarity between Piller's worry regarding Broome's theory of fairness as explained in the previous paragraph and Scanlon's (hopefully by now familiar) worry about permitting aggregation within his contractualism. The worry is that if we hitch our wagon to the respective notion – so, either aggregation or Broomean fairness – then this notion will commit us to some results we intuitively want to resist. Of course, we have to suppose here, for the sake of argument that we share either Piller's and Scanlon's intuition when

[4] This is the view defended by Brad Hooker (2005), and also by Peter Stone (2007), as Piller notes (Piller, 2017, p. 227, fn. 30). I have omitted Piller's discussion of Hooker vs. Broome as that discussion understandably boils down to a disagreement about intuitions on a particular case; see (Piller, 2017, pp. 227–229) for that discussion.

[5] This interpretation is correct because Broome explicitly rules out the permissibility within fairness of overriding weaker claims with stronger claims (Broome, 2003, p. 196). By elimination, that leaves the weighted lottery as the way to go in indivisible-good-unequal-claims cases if we have also ruled out withholding the good and want chances to (proportionately) reflect individuals' claim-strength.

it comes to either Broomean fairness or aggregation. I have already explained how Scanlon, in my view, unsuccessfully attempts to avoid the aggregation conclusions he finds unpalatable by appeal to "relevance" and "broad categories of moral seriousness" that need to be justified rather than simply presumed. An important difference that should be noted before proceeding to discuss Piller's account is that Scanlon is trying to resist all-things-considered entailments, while Broome is only committed to "fairness-requires-that"-judgments, which are not all things considered. The stakes are considerably lower and so the bullets easier to bite.

But having noted the difference in stakes in the two discussions, we can focus on the following question: how does Piller attempt to avoid what he believes are counterintuitive results that Broomean fairness seems to imply? He does so by arguing that there is a distinction within Broomean fairness which needs to be (further) explicated, which Piller does in the context of indivisible-good-unequal-claims cases:

> Let us first consider the outcome-related way of treating people. If a weighted lottery rewards the person with the stronger claim, there will be some unfairness – the weaker claim has been neglected. The weaker the claim the smaller this unfairness will be. If the weighted lottery rewards the person with the weaker claim, there will again be unfairness – the stronger claim has been neglected. Neglecting the stronger claim involves weightier unfairness than neglecting the weaker claim. The bigger the difference between the stronger and the weaker claim, the bigger the unfairness will be if the weaker claim wins. This is what it looks like when we consider ways of treating people's claims that refer to the distribution that will result from our thus treating them. The second relevant aspect we need to consider is how lotteries treat people's claims in terms of being a procedure of distribution. In this sense, the weighted lottery satisfies the proportionality principle. In this sense…weighted lotteries are fair. [Broome] says that the fairness achieved by the weighted lottery might *mitigate* the distributional unfairness. The fairness achieved by the weighted lottery will, however, have to be weighed against the fairness loss a lottery brings, which will vary depending on whether the weaker or the stronger claim wins. (Piller, 2017, pp. 229–230)[6]

[6] In the context of equal-claims cases, Rob Lawlor (2006) voices a similar concern regarding the possibility that the weighted lottery might appoint the single claimant in a one vs. a million lives saved scenario, see (2006, p. 162).

Although I think Piller in the above passage outlines a plausible view on fairness, insofar as it purports to be a Broomean theory of fairness, I believe it is almost entirely incorrect. The only exception is the claim near the end that "the fairness achieved by the weighted lottery might *mitigate* the distributional unfairness." Weighted lotteries do indeed mitigate the distributional unfairness, on Broome's account.

Piller presents a contrastive example involving a "small-difference" case (51 per cent vs. 49 per cent chances in the weighted lottery appropriately matching the respective claim-strengths) and a "large-difference" case (90 per cent vs. 10 per cent) (Piller, 2017, p. 230). In both cases, we achieve (perfect) procedural fairness because the chances given perfectly correspond with the claims' strengths.[7]

If we are going to distribute the indivisible good, then outcome unfairness is unavoidable. Withholding the indivisible good is the only way to prevent any unfairness, whether procedural unfairness or outcome unfairness, in this kind of case.[8] However, if we distribute the indivisible good, then on Piller's interpretation of Broomean fairness the outcome unfairness may be greater or smaller depending upon whether the stronger or the weaker claim wins the lottery. The main fairness-question to be addressed is whether we should run the weighted lottery or whether we should directly satisfy the stronger claim. Broome believes his theory provides us with the following ordering, in fairness terms:

[7] For the remainder of this chapter, I will omit the "(perfect)" when speaking of "procedural fairness" or "outcome fairness". 'Imperfect' outcome or procedural fairness equates to (some degree of) outcome or procedural unfairness.

[8] "The heart of my suggestion is that fairness is concerned only with how well each person's claim is satisfied *compared with* how well other people's are satisfied. It is concerned only with relative satisfaction, not absolute satisfaction. (…) [I]t is not *unfair* if [claims] are not [satisfied], provided everyone is treated proportionally. (…) [Now, consider cases with an indivisible good.] Take a case, first, where all candidates have equal claims. It would be possible to satisfy their claims equally, as fairness requires, by denying the good to all of them. (…) [But if we were to distribute the indivisible good, then] the candidates' claims cannot all be equally satisfied, because some candidates will get the good and others will not. So some unfairness is inevitable. (…) Next, take a case where several people have claims to a good that are roughly, but not exactly equal. Perhaps, for instance, they all need the good, but not exactly equally. And suppose again that there is not enough to go round them all. Fairness requires satisfaction in proportion to their claims. So if the good goes to the people with the strongest claims, the others will not have been fairly treated; their claims will have been overridden. And if it goes to other people, the unfairness will be worse. So unfairness is once again inevitable. But once again it can, if the circumstances are right, be mitigated [note, not alleviated or undone!] by giving everyone a chance of getting the good. Ideally, each person's chance should be in proportion to the strength of her claim: the lottery should be unequally weighted. (At first, it is particularly puzzling how a weighted lottery could be fair. If it is fair for some people to have a greater chance than others, that means they more ought to have the good. So why not let them have it without a lottery? My theory explains why not). (Broome, 1990, pp. 95, 97, 98). Although Broome specifies the two conditions for holding a lottery when (1) "it is important to be fair" and (2) "the candidates' claims are equal or roughly equal", I believe his theory can plausibly be extended to govern cases where the claims are unequal (Broome, 1990, p. 99). It is at least as important to be fair in unequal-claims cause as it is in (roughly-)equal-claims cases.

1. Do not distribute
2. Distribute according to weighted lottery
3. Distribute to satisfy the strongest claim.

Therefore, if distributing the good, Broome states that, so far as fairness is concerned, 2 ought to be preferred to 3. Piller, however, argues that 3 ought to be preferred to 2 in some cases, namely those where the claims differ significantly in strength. Piller argues for this as follows:

> We need an ethical judgment that compares these fairness losses [in either case]. What is worse in terms of fairness: reducing a 49% chance to 0 or having a 49% chance of satisfying a slightly weaker claim? Broome thinks that reducing the 49% to 0 would be more unfair. Thus, in small-difference cases it is fair to use an appropriately weighted lottery. (…) What is more unfair: reducing a small [10%] chance to 0 or risking that a much weaker claim is satisfied? (…) I, for one, think that reducing a very small chance to 0 is a small fairness sacrifice. The fairness loss it brings is considerably smaller than the risk of satisfying the much weaker claim, as satisfying the much weaker claim would be very unfair. (Piller, 2017, pp. 230–231).

Having discussed these cases, Piller draws the following general conclusion:

> Thus, Broome's theory [of fairness] is fully compatible with the normative view Hooker and I have endorsed: it would be unfair to use a weighted lottery when the strength of claims differs significantly. More precisely, given the right judgement about how fairness losses compare, it would be *fairer*, on Broome's proportionality theory, to satisfy the stronger claim than to hold a weighted lottery. (Piller, 2017, p. 231).

Before further discussing Piller's suggested distinction, I note the obvious problem that Piller's account mandates that we do something which Broome has explicitly told us not to do from the perspective of fairness: override weaker claims (Broome, 2003, p. 196).

However, if we accept Piller's view, then within Broome's theory of fairness, for any case involving (i) an indivisible good and (ii) unequal competing claims, we need to ask the following question: should we prioritize outcome fairness (i.e. giving the good to the stronger claimant) over

procedural fairness (i.e. holding an appropriately weighted lottery)?[9] Piller's answer, on behalf of Broome, is that for certain cases, we should. The range of cases in which we should prioritize outcome fairness is where the weaker claim is so much weaker than the stronger claim that the situation does not merit the holding of a weighted lottery for the sake of fairness.

§2. Measuring Outcome Unfairness

Piller's view relies on contrasting outcome unfairness and procedural unfairness, and aiming to reduce unfairness overall. In considering the lottery, and the outcome unfairness, Piller focuses on how unfair (in outcome terms) it would be for the lesser claim to win. However, it is important to note that, when considering a weighted lottery, we must not look only at how unfair the outcome of the lesser claimant receiving the good would be. We also must look at how likely that would be. In other words, in assessing a lottery, we must focus not on outcome unfairness, but expected outcome unfairness.

Let's stick to Piller's contrastive example of a 51 per cent vs. 49 per cent case and a 90 per cent vs. 10 per cent case. Piller is right that it would be worse from the perspective of outcome fairness if the 10 per cent-claimant were to win her lottery (vs. the 90 per cent-claimant) than if the 49 per cent-claimant were to win her lottery (vs. the 51 per cent-claimant). Generally, the weaker the claim of the weaker claimant, the worse it would be if she won the lottery: the outcome unfairness would be greater.

However, we should not forget that we are running lotteries. And that if we are running a lottery, any outcome unfairness will be subject to a probability of occurring. That probability is of course exactly the chance given to either claimant in the lottery. Now, we are guaranteed that we will have some outcome unfairness precisely because we have an indivisible good and competing claimants – that's the tragedy of the situation.

Suppose, for the sake of illustration, that the value of the indivisible good is 100 units. Following F. M. Kamm's (2005) "method of virtual divisibility" for indivisible goods, we can suppose that each claimant has a claim to the respective number of units – even though of course we cannot actually divide the 100-unit indivisible good. So, in the one competition the 90 per cent-claimant

[9] As a reminder, the fairest thing to do, on Broome's theory, in any case involving an indivisible good and competing claims, whether these claims are of equal or unequal strength, is to withhold the good. For the sake of the discussion here, please assume that for non-fairness reasons that is not an option here. We must distribute this indivisible good but now we want to know what the fairest way to do so is. This is why the weighted lottery is a second-best solution at best.

should ideally get 90 out of 100 units and the 10 per cent-claimant should get 10 out of 100 units, whereas in the other competition the 51 per cent-claimant should ideally get 51 out of 100 units and the 49 per cent-claimants should get 49 out of 100 units.

The upshot of using Kamm's "virtual divisibility" is that we can put a number on the 'unfair surplus' that the winning claimant would receive when awarded with the 100-unit indivisible good.

For the sake of argument, let's assume that outcome unfairness is measured by this 'unfair surplus'. To repeat, the unfair surplus is unavoidable if we distribute the indivisible good. In the large-difference case, if the stronger claim wins, the outcome unfairness will be 10 units – the ideal distribution was 90, but the claimant received 100. If the weaker claim wins, the outcome unfairness will be 90. It is this possibility of the weaker claimant winning the lottery that Piller is worried by and thinks that this is so important that we ought to forgo procedural unfairness to protect against it.

However, consider the chances that are part of the weighted lottery. In the large-difference case, these are 90 per cent vs. 10 per cent. The expected outcome unfairness of the stronger claimant winning the lottery is 9 (because there is a 90 per cent chance she will receive 10 surplus units – units to which she has no claim). The expected outcome unfairness of the weaker claimant winning the lottery is also 9 (because there is a 10 per cent chance she will receive 90 surplus units).

Let us now look at the small-difference case (51 per cent vs. 49 per cent). The expected outcome unfairness of the stronger claimant winning the lottery is 24.99 (because 0.51x49=24.99). The expected outcome unfairness of the weaker claimant winning the lottery is also 24.99 (because 0.49x51=24.99).[10]

There are two interesting results here. The first is that in the large-difference case, the expected outcome unfairness of satisfying the strongest claim is the same as the expected outcome unfairness of satisfying the weaker claim. Therefore, even though if the outcome eventuated it would, on this measure, be more unfair to satisfy the weaker complaint, because this is so unlikely, the expected outcome unfairness of both eventualities is the same. Piller is therefore wrong to be

[10] There must be an interesting deeper explanation for the symmetry in expected outcome unfairness in either case (9 vs. 24.99). Alas, I cannot go into exploring that here.

especially concerned by the risk of satisfying the weaker claim. The risk of satisfying either is the same.

The second is that there is much more at stake, in terms of outcome unfairness, in small-difference cases than in large-difference cases. In cases where the competing claims are closely matched in strength (e.g. 51 per cent vs. 49 per cent), then based on outcome fairness-considerations we could much more strongly argue that we should not hold a weighted lottery. This is because the expected outcome unfairness is much greater in a small-difference case than in a large-difference case. This, in turn, is because it is certain that we will either give a claimant nearly double the value of the good she has a claim to (i.e. if the 51 per cent-claimant receives the good) or give another claimant just over double the value of the good she has a claim to (i.e. if the 49 per cent-claimant receives the good). By contrast, because there is such a small chance that a much weaker claimant (e.g. 10 per cent-claimant) will win the lottery and so will be awarded with the good, the expected outcome unfairness is much lower.

Therefore, with respect to expected outcome fairness only, the case for a lottery is in fact much stronger in the large-difference case. Although Piller focuses on the dangers of the large-difference case, the dangers of holding a lottery in the small-difference case are in fact even greater.

Of course, a final determination on which procedure (weighted lottery or giving to the strongest claim) is better will also need to take account of procedural fairness. In order to do that, we will need (a) an account of how to measure procedural unfairness, and (b) a way of trading the two types of unfairness off against one another.

In order to make good on his claim that the lottery is the fair procedure for small-difference cases, then the procedural fairness of the lottery (or the procedural unfairness of not holding the lottery) will have to be vastly greater in the small-difference cases than in the large-difference cases. If, for example, the fairness value of lotteries were constant, then we would have a stronger case for lotteries in the large-difference cases.

In conclusion, Piller is correct that the actual outcome unfairness (or outcome unfairness *simpliciter*) would be greater the weaker the weaker claimant's claim is and that, therefore, there is a stronger fairness-case to not hold the lottery the weaker the weaker claimant's claim is. However, I have argued that Broomeans should not focus on actual outcome unfairness, but on expected outcome

unfairness. After all, we should not ignore the fact that we will run a weighted lottery for the sake of fairness. Once we have rightfully shifted our focus from actual outcome unfairness to expected outcome unfairness, the intuitively odd result following from Piller's distinction between outcome fairness and procedural fairness is that the support for his intuitively plausible judgment that it is fairer to hold a weighted lottery in a small-difference case than in a large-difference case is incomplete. More needs to be said.

However, I have only discussed outcome (un)fairness so far. I will now turn to procedural (un)fairness. I argue that there are broadly four possible ways of measuring procedural unfairness. Piller simply assumes one of them. I will do my best to carefully outline all four and argue that there are more plausible and/or more Broomean ways of measuring procedural unfairness than Piller suggests.

§3. Four Ways of Measuring Procedural Unfairness

In overview, the four possible ways of measuring procedural unfairness are as follows:

1. The Subtraction Measure:
    - Procedural unfairness is the chance-loss in absolute terms to the weaker claimant.
2. The Subtraction and Unfair Addition Measure:
    - Procedural unfairness is the chance-loss in absolute terms to the weaker claimant plus the unfair surplus chance given to the stronger claimant.
3. The Proportional Loss to the Weaker Claimant Measure:
    - Procedural unfairness is the proportional chance-loss compared with the chance the weaker claimant would ideally have had.
4. The Comparative Measure:
    - Procedural unfairness is the chance-loss compared to the ideal distribution of chances (i.e. a distribution of chances that matches the respective claims-strengths' proportionate ratio).

Obviously, it matters which measure of unfairness is used by a Broomean theory. I will explain each in the following subsections. It is not clear to me what measure Piller favours. As I read his account, he either favours the first or the second measure. However, I will not further pursue the

question of how to most plausibly interpret Piller on this point because I am concerned with how to best understand Broome's theory of fairness instead.

§3.1. The Subtraction Measure

First, let's reconsider Piller's words on the procedural aspect of fairness:

> The second relevant aspect we need to consider is how lotteries treat people's claims in terms of being a procedure of distribution. In this sense, the weighted lottery satisfies the proportionality principle. In this sense…weighted lotteries are fair. [Broome] says that the fairness achieved by the weighted lottery might *mitigate* the distributional unfairness. The fairness achieved by the weighted lottery will, however, have to be weighed against the fairness loss a lottery brings, which will vary depending on whether the weaker or the stronger claim wins. (Piller, 2017, p. 230).

Let's see how to best understand procedural (un)fairness on Broome's theory of fairness if our aim is to incorporate the distinction between procedural fairness and outcome fairness in Broome's theory of fairness.

The first possible interpretation is to understand the size of procedural unfairness done to an individual as the absolute difference between (1) the proportional chance she deserves given the strength of her claim and (2) the actual chance she is given. So, taking the large-difference case as an example, we suppose the weaker 10 per cent claimant's chance in the lottery is removed. On this first possible way of measuring procedural unfairness, the Subtraction Measure, the size of procedural unfairness done to the weaker claimant is 10. This first possible way of measuring procedural unfairness is analogous to the way in which Piller explicitly suggests we measure outcome unfairness.

However, the Subtraction Measure is not plausibly Broomean. The Subtraction Measure is not Broomean because it fails to capture the concern for the relative treatment, rather than the absolute treatment, of claims in the example of the large-difference case as I have just given it. The unaddressed issue of relative treatment gives rises to the second possible way of measuring procedural unfairness, the Subtraction and Unfair Addition Measure.

§3.2. The Subtraction and Unfair Addition Measure

The unaddressed issue of relative treatment of claims that the Subtraction Measure gives rise to stems from the following question that is left open on that first measure: what to do with the weaker claimant's chance (10 per cent in the large-difference case) that, for the sake of fairness on Piller's account, we have taken away from her? There are two main options here.

First, we can leave the 10 per cent unassigned. However, leaving the 10 per cent chance unassigned seems to me to be equivalent to having a 10 per cent chance that we will not distribute the good. The revised lottery is then that we give the stronger claimant 90 per cent chance to win the lottery, proportionate to the individual claim's strength, on the one hand. And, on the other hand, we have a 10 per cent chance of not distributing the good. This non-distribution possibility is significant on Broomean fairness. Not distributing the indivisible good when there are at least two competing claimants, regardless of whether the strength of their individual claims is equal or unequal, results in outcome fairness.

However, the point relevant for the present purpose of how to understand and measure procedural unfairness is that on the first option the size of procedural unfairness is determined solely by the chance taken away from the weaker claimant. So, the first option is the first suggested measure, the Subtraction Measure. So, in the large-difference case the procedural unfairness would be 10 (because we take away the 10 per cent-chance from the weaker claimant) and in the small-difference case the procedural unfairness would be 49 (because that is the 49 per cent chance taken from the weaker claimant).

However, there is a second option here, which gets us the second way to measure procedural unfairness, the Subtraction and Addition Measure. Sticking to the large-difference case as our example, we can add the 10 per cent chance we have taken away from the weaker claimant to the stronger claimant's chance. Now, of course, whatever the respective strengths of the individual claims are, whenever we add the weaker claimant's chance to the original proportionate chance of the stronger claimant, then we give the stronger claimant a 100 per cent chance in the weighted lottery. This makes running the lottery unnecessary. We are effectively giving the indivisible good directly to the stronger claimant. Of course, directly giving the indivisible good to the stronger claimant is precisely the result Piller wants (for Broome and Broomeans) in some cases.

What is the most important difference between the Subtraction Measure and the Subtraction and Addition Measure? The Subtraction and Addition Measure tells us that it is more unfair to give the taken-away chance to the stronger claimant than merely taking away the chance from the weaker claimant. The Subtraction Measure tells us that it is equally (unfair) to (1) take away the chance from the weaker claimant or (2) take away the chance from the weaker claimant and give that taken-away chance to the stronger claimant.

Furthermore, on the Subtraction and Addition Measure, the size of procedural fairness is doubled as compared to the Subtraction Measure. This is because there are now two counts of procedural unfairness. First, like the Subtraction Measure, there is the procedural unfairness done to the weaker claimant of having her chance taken away. But, second, and unlike the Subtraction Measure, there is the unfair surplus chance that is given to the stronger claimant. So, for example, in the large-difference case, procedural unfairness is 20 (because we have both taken away the weaker claimant's 10 per cent chance and we have made the stronger claimant's chance disproportionate by an additional 10 per cent).

Piller only explicitly argues for the need of an ethical judgment as to how bad it would be to take away the weaker claimant's chance in the case considered. As I have presented his view, how bad that would be depends on the size of the proportionate chance the weaker claimant has. The weaker the weaker claimant's claim, and so the smaller the proportionate chance she would have in the weighted lottery, the less bad it is, according to Piller, to take her chance away for the sake of fairness.

Now, importantly, Piller argues that it less bad to take this chance away the weaker the weaker claimant's claim is because of the risk there is that such a weak claimant will win the lottery and so receives the good. So, the badness of taking away the weaker claimant's chance is determined on the basis of a comparative question: what is the 'fairness-damage' done if we take away this weaker claimant's chance (procedural unfairness) given the possible 'fairness-damage' that she, as the weaker claimant, might win the weighted lottery (outcome unfairness)?

The key point here is as follows: Piller sets up the comparison between (risk of) satisfying the weaker claim and the unfairness of removing the weaker claimant's proportional chance which is owed to her from the perspective of fairness. Neither of these is quite right: as shown above, the risk of the weaker claim being satisfied is no worse than the risk of the stronger claim being

satisfied; and as shown in the previous paragraph, taking away the weaker claimant's chance is unfair but so is adding the taken-away chance to the stronger claimant's proportional chance.

However, as I have indicated in parentheses in the question above, this reveals the distinction Piller draws between outcome fairness and procedural fairness. Therefore, these two features of the comparative question that are operative in Piller's argument can be isolated from each other. For the purpose of considering different possible ways of measuring procedural unfairness, we then focus on the first part of the comparative question: whether the weaker claimant should be given (or permitted to keep) her proportionate chance.

But, as I have argued in these first two sub-sections, there is a further important fairness-question left unaddressed by Piller: if we decide to take away the weaker claimant's proportionate chance, what should we do with that chance? There are two separate issues here. First, there is the question of what we should do. Second, there is the question of how we should measure fairness. It strikes me that out of the two options considered, the Subtraction and Addition Measure is the right way to measure procedural unfairness. However, on this same view, it would be fairer to only subtract the weaker claimant's chance because we do not compound the procedural unfairness by giving the stronger claimant a chance that is disproportionally larger than is owed to her given the strength of her claim. If we want to minimize procedural unfairness in this way, then the Subtraction and Addition Measure and the Subtraction Measure will converge in terms of their recommendation of what we should do, even though they conflict in terms of how we should measure procedural unfairness.

On the Subtraction Measure, the size of procedural unfairness equates to the proportionate chance the weaker claimant would have had if we had run the weighted lottery. The Subtraction and Addition Measure, as a rule, doubles the size of procedural unfairness under the Subtraction Measure if the chance is redistributed to the other claimant. The size of procedural unfairness is doubled on the second measure because apart from considering the proportionate chance taken away from the weaker claimant, we factor in the disproportionately greater chance the stronger claimant now has. So, in the large-difference case, procedural unfairness is 10 according to the Subtraction Measure but 20 according to the Subtraction and Addition Measure. In the small-difference case, the Subtraction Measure gives us 49 for procedural unfairness, whereas the Subtraction and Addition Measure doubles this to 98.

Now, as I read his account, Piller measures as according to the Subtraction Measure, but he supposes that the chance will in fact be redistributed, which is more in line with the Subtraction and Addition Measure. Therefore, it is unclear to me how to best understand Piller's view. In any case, out of the first two discussed measures, the Subtraction and Addition Measure is the more Broomean measure of procedural unfairness in one important respect. It is the more Broomean measure because, unlike the Subtraction Measure, the Subtraction and Addition Measure is explicitly concerned with the relative treatment of claims.

However, once we bring the essential Broomean tenet of relative treatment into focus, we are also reminded of the importance of proportionality on Broome's theory of fairness. In this manner, a third and fourth measure of procedural unfairness suggest themselves.

§3.3. The Proportional Loss to the Weaker Claimant Measure

The third possible measure of procedural unfairness is the Proportionate Loss to the Weaker Claimant Measure. As always, the starting point is the proportionate chance that the weaker claimant would receive if we were to hold a weighted lottery (i.e. were things procedurally fair). However, according to this third measure, we need a different understanding of any departure from this proportionate chance that the weaker claimant's claim merits. Recall that Broome's theory is concerned with proportionate satisfaction, so it seems plausible that in measuring procedural unfairness, we should be looking at the proportionate, not absolute, loss of chances. The different understanding is the proportionate procedural fairness loss to the weaker claimant which is determined by the difference between (1) the proportionate chance the weaker claimant deserves in a weighted lottery and (2) the actual chance we leave her with.

As an example of how this approach would work, imagine someone's procedurally fair chance would be 10 per cent. If we run a lottery in which they are given a 5 per cent chance, then this would be a 5 per cent reduction in absolute terms, but a 50 per cent reduction in proportionate terms.

In the context of Piller's discussion, this leaves us with an immediate complication. Piller only considers whether we should completely take away the weaker claimant's chance. So, whether we should, for the sake of fairness overall, bring the weaker claimant's proportionate chance (whether it is 49 per cent in the small-difference case or 10 per cent in the large-difference case) all the way down to zero.

On both of the first two suggested measures, the Subtraction Measure and the Subtraction and Addition Measure, we had different sizes of procedural unfairness. This was because of a feature that both accounts shared: the varying strength of the weaker claimant's claim. So, taking the small-difference and large-difference cases again, the procedural unfairness would be 49 or 98 and 10 or 20, even though in both cases the weaker claimant's chance was reduced to zero.

But if we (1) reduce a weaker claimant's chance to zero and, to have our account of procedural fairness be more closely aligned with Broome's theory of fairness, (2) argue that the procedural fairness loss should be understood proportionately, then we cannot (or no longer) differentiate between the small-difference case and the large-difference case (or any other case, for that matter). The procedural fairness loss will always be the same, regardless of the varying strength of the weaker claimant. Let me explain.

In the small-difference case, we have a 51:49 ratio. In the large-difference case, we have a 90:10 ratio (or, more simply, a 9:1 ratio). In the large-difference case, Piller suggests we should in fact reduce the weaker claimant's chance to zero (for the sake of fairness). If we were to do so, then we would move from a 9:1 ratio to a 9:0 ratio. Here is the first crucial point: whereas 9 is nine times greater than 1, 9 is infinitely greater than 0 (precisely because it is zero). The second crucial point is that any positive number will be infinitely greater than zero. So, considering the small-difference case and supposing for the sake of argument that we reduce the weaker claimant's chance from 49 per cent chance to 0 per cent, then we will have moved from a 51:49 ratio to a 51:0 ratio. And, like 9 is infinitely greater than zero, so 51 is infinitely greater than zero.

Comparing the small-difference case shows an important general feature of the Proportional Loss to the Weaker Claimant Measure: we can ignore the question as to whether we should award the taken-away chance from the weaker claimant as a surplus to the stronger claimant, for example, moving her from 51 per cent to 100 per cent in the small-difference case or from 90 per cent to 100 per cent in the large-difference case. 100, like 51 and 9 (or 90), is infinitely greater than zero. Since any positive number is infinitely greater than zero, whether we do or don't award the taken-away chance as a surplus chance, the resulting ratio amounts to the stronger claimant having an infinitely greater claim.

Importantly, it does not immediately follow that an infinite ratio stands for an infinite amount of procedural unfairness. However, without an argument to the contrary – and I won't have time to

consider one here – I find it a natural interpretation to say that with an unjustifiable (because disproportionate) infinite ratio we have infinite procedural unfairness.

But now we might wonder whether my explanation of the third proposed measure also serves as a *reductio* of the measure. Regardless of whether the taken-away chance of the weaker claimant is added as a surplus chance to the stronger claimant, whatever the relative strength of the weaker claimant's claim, doesn't reducing that claim to zero (whether e.g. from 49 per cent or 10 per cent) always amount to an infinite ratio? I would like to make three points in response.

First, the Proportional Loss to the Weaker Claimant Measure captures an important phenomenon that the first two measures fail to capture. The important phenomenon is that not all reductions in chance are equal. On the first two measures I have considered, there is nothing special regarding reducing a 1 per cent chance to 0 per cent as, for example, compared to reducing a 49 per cent chance to 48 per cent. In both cases, we have reduced the chance by 1 per cent. However, there does seem to be a significant difference between receiving a (disproportionately – and so, procedurally unfairly) smaller chance and receiving no chance at all. The third possible measure of procedural unfairness potentially captures this phenomenon by means of the infinite ratio. Furthermore, it underscores Broome's idea that there is something especially bad from the perspective of fairness whenever we would override (weaker) claims.

Second, unlike the first two measures and in line with the core argument of Broome's theory of fairness, the third measure has proportionality as its central feature. So, even if we find the results this third measure leads to unpalatable, that is no good reason to reject the measure given that the aim was to make sense of procedural fairness in a Broomean picture.

Third, Piller treats the question of procedural fairness in the small-difference and large-difference cases that he considers as a binary question: either we give or leave the weaker claimant with her proportionate chance and run the weighted lottery or we reduce the weaker claimant's chance (all the way) to zero. However, this is a false dilemma. We could reduce the weaker claimant's chance in a variety of ways without reducing it to zero. More importantly, if we give or leave the weaker claimant with at least some chance, no matter how small, then we will avoid the infinite ratio.

But despite these three points in response, we might still view the infinite ratio (as plausibly implying infinite procedural unfairness) in cases where and because the weaker claimant's chance

is reduced to zero as a problem that must be addressed. Is there an answer to this infinity problem or does the problem show that we should abandon the idea of understanding the size of procedural unfairness as the proportional loss suffered by the weaker claimant? I do not know an answer to this problem, and so am (further) tempted to reject the proposed distinction between outcome fairness and procedural fairness. So, I turn to a fourth possible measure of procedural unfairness. It is the most plausible and Broomean of all four possible measures of procedural unfairness considered in this chapter.

§3.4. The Comparative Measure

The fourth possible measure has two main parts: (1) procedural fairness means giving the right proportionate/ratio of chances comparatively, and (2) procedural unfairness should be measured by departures from the appropriate ratio. This is the Comparative Measure. I will now unpack both parts of this fourth possible measure.

Consider again the large-difference case where proportionate chances would be 90 per cent for the stronger claimant and 10 per cent for the weaker claimant. The chance-distribution, therefore, mirrors a 9:1 ratio. But now I would like to make a new point regarding the proportionate ratio.

Since Broomean fairness focuses on the relative treatment of claims, then on Broome's theory the chance-distribution 'merely' needs to respect the proportionate ratio.[11] How we express the proportionate ratio in the chance-distribution is an open question for our case at hand. Furthermore, Broomeans should be indifferent between the various possible expressions of the proportionate ratio in the chance-distribution for the sake of procedural fairness.

For example, in the large-difference case, where the ratio is 9:1, not only would a 90 per cent vs. 10 per cent chance-distribution secure procedural fairness, so would giving the stronger claimant a 9 per cent chance and the weaker claimant a 1 per cent chance, while reserving a 90 per cent chance for destroying the good. This second arrangement reduces the chances for both claimants. However, because it does so proportionately, it retains procedural fairness. In fact, this proportionate reduction of chances might be fairer on Broome's theory precisely because we have increased the chance that neither claimant will receive the good, which is the fairest outcome of all.

[11] See footnote 8 for textual support from Broome that, for his theory of fairness, he is solely concerned with "relative", rather than "absolute" satisfaction of claims.

§4. In Summary

I have criticized Piller's interpretation of Broome as implicitly relying on an outcome fairness vs. procedural fairness distinction. Consider the following example of Piller on how we ought to decide who should get the indivisible good of going home early from work:

> If one person has worked nine hours toward completing a shared task and the other person only put in one hour, we'd know exactly which weights to set in a weighted lottery, if, that is, Broome recommended one (which he does not). (Piller, 2017, p. 234).

I agree. Broome does not recommend a weighted lottery all things considered. Broome does not believe fairness to be the only value. His theory of fairness, however, does recommend a weighted lottery if and only if an indivisible good is to be distributed. After all, the fairest solution, according to Broome's theory of fairness, is to not give the indivisible good to any of the competing claimants. Furthermore, we do not know exactly which weights to set. What we know exactly is the ratio of the weights to be set, 9:1. But whether we ought to run a 90 per cent vs. 10 per cent lottery, which guarantees someone receives the indivisible good (in this case, leaving work early), is a separate question. Given that we are closer to achieving perfect fairness by not satisfying either claim, it would be fairer to give the stronger claimant a 9 per cent chance and the weaker claimant a 1 per cent chance and assign a 90 per cent chance to neither person going home early (i.e. withholding the good).[12]

Therefore, insofar as Broome has considered adopting this distinction within his theory, I have presented several reasons why he should not. The distinction between outcome fairness and procedural fairness unnecessarily complicates Broome's otherwise neat theory of fairness. However, I find the distinction between outcome fairness and procedural fairness intuitively attractive, and so have considered Piller's proposal of incorporating this distinction within Broome's theory of fairness.

If we are to adopt this distinction, we require a way of measuring procedural fairness and a way of measuring outcome unfairness. We then also require a way of comparing them against each other to reach the fairest method of distribution overall.

[12] Therefore, let's agree not to discuss this example with either your or my employer in the room.

Regarding outcome fairness, Piller focuses on the risk of the weaker claimant's receiving the good. I have pointed out that Piller seems to have overlooked the fact that Broomeans should be concerned with expected outcome unfairness, rather than actual outcome unfairness. Once we see this, the risk of the stronger claimant receiving the good is just as strong and, furthermore, large-difference lotteries are overall less concerning than small-difference lotteries when it comes to expected outcome fairness.

Regarding procedural fairness, I have considered four possible measures whereas Piller simply assumes one, namely the Subtraction and Addition Measure. Given the Broomean emphasis on proportionality and relative satisfaction of claims, I have given reasons why the fourth measure, the Comparative Measure is both the most plausible and the most Broomean way of understanding and measuring procedural unfairness.

Conclusion

The more intricate the subject discussed, the more complex an account of the matter will have to be. Christian Piller has made a valiant effort to deal with important but difficult cases and apparent counterexamples to Broome's defence of weighted lotteries. In this chapter, I have attempted to show why I believe Piller hyperextends Broome's theory of fairness with the proposed distinction between outcome fairness and procedural fairness.

I heartily appreciate the important knowledge gained from Piller's excellent work. But I have argued that we should not interpret Broome's theory of fairness as implicitly relying on an outcome fairness vs. procedural fairness distinction in the way Piller suggests. Furthermore, however intuitively plausible I find such a distinction within fairness independent of Broome's theory, the distinction between outcome and procedural fairness leads to some unpalatable results within a Broomean account of fairness.

The result buttresses part of Broome's philosophical legacy: if indivisible goods are to be distributed, then fairness always requires they be distributed by weighted lotteries. However, as Broome has repeatedly noted, sometimes what fairness requires is not what is morally right, all things considered.[13]

[13] Broome's theory of fairness still faces significant objections. Patrick Tomlin (2012) presents the most damning-looking objections, in my view. Most importantly, I understand Tomlin's objections as implying that Broome's theory of fairness should have a much more restrictive scope: Broomean fairness only applies in cases of what Tomlin calls

Bibliography

Anderson, E. S. (1999). What Is the Point of Equality? *Ethics*, *109*(2), 287–337. https://doi.org/10.1086/233897

Anscombe, G. E. M. (1967). Who is Wronged? Philippa Foot on Double-Effect: One Point. *Oxford Review*, *5*, 16–17.

Arneson, R. J. (2000). Luck Egalitarianism and Prioritarianism. *Ethics*, *110*(2), 339–349. https://doi.org/10.1086/233272

Arrhenius, G., & Rabinowicz, W. (2015). Value Superiority. In I. Hirose & J. Olson (Eds.), *The Oxford Handbook of Value Theory* (pp. 225–248). Oxford University Press. https://doi.org/10.1093/oxfordhb/9780199959303.013.0013

Ashford, E. (2003). The Demandingness of Scanlon's Contractualism. *Ethics*, *113*(2), 273–302. https://doi.org/10.1086/342853

Ashford, E., & Mulgan, T. (2018). Contractualism. In *Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/archives/sum2018/entries/contractualism/>

Broome, J. (forthcoming). Lessons from Economics. In M. D. White (Ed.), *The Oxford Handbook of Ethics and Economics* (pp. 1–18). Oxford University Press. http://users.ox.ac.uk/~sfop0060/pdf/Lessons%20from%20economics.pdf

Broome, J. (1984). Selecting People Randomly. *Ethics*, *95*(1), 38–55.

Broome, J. (1990). Fairness. *Proceedings of the Aristotelian Society*, *91*, 87–101. JSTOR.

Broome, J. (1997). Is Incommensurability Vagueness? In R. Chang (Ed.), *Incommensurability, Incomparability, and Practical Reason* (pp. 67–89). Harvard University Press.

Broome, J. (1998). Kamm on Fairness. *Philosophy and Phenomenological Research*, *58*(4), 955–961. JSTOR. https://doi.org/10.2307/2653739

"non-voluntary claims or NV-claims". For a recent criticism of Broomean fairness which I believe misses the mark because it relies on a misunderstanding of Broome's view, see Wintein and Heilmann (2018). I say "misunderstanding" because Wintein and Heilmann rely on a 'Pillerian' interpretation of Broome's theory of fairness.

Broome, J. (2003). *Weighing Goods: Equality, Uncertainty and Time* (transferred to digital print). Blackwell Publ.

Broome, J. (2004). *Weighing Lives*. Oxford University Press.

Broome, J. (2013). *Rationality Through Reasoning*. Wiley Blackwell.

Burri, S. (2016). Moral Aggregation, Iwao Hirose. Oxford University Press, 2015, xiii + 234 pages. *Economics and Philosophy*, *32*(01), 149–156. https://doi.org/10.1017/S0266267115000437

Bykvist, K. (2017). Moral Uncertainty. *Philosophy Compass*, *12*(3), e12408. https://doi.org/10.1111/phc3.12408

Chang, R. (Ed.). (1997a). *Incommensurability, Incomparability, and Practical Reason*. Harvard University Press.

Chang, R. (1997b). Introduction. In R. Chang (Ed.), *Incommensurability, Incomparability, and Practical Reason* (pp. 1–34). Harvard University Press.

Chang, R. (2002). The Possibility of Parity. *Ethics*, *112*(4), 659–688. https://doi.org/10.1086/339673

Chang, R. (2015). Value Incomparability and Incommensurability. In I. Hirose & J. Olson (Eds.), *The Oxford Handbook of Value Theory* (pp. 205–224). Oxford University Press. https://doi.org/10.1093/oxfordhb/9780199959303.013.0012

Constantinescu, C. (2016). Vague Comparisons. *Ratio*, *29*(4), 357–377. https://doi.org/10.1111/rati.12146

Dancy, J. (2004). *Ethics Without Principles*. Oxford University Press. https://doi.org/10.1093/0199270023.001.0001

Dancy, J. (2018). *Practical Shape*. Oxford University Press. https://doi.org/10.1093/oso/9780198805441.001.0001

Darwall, S. (Ed.). (2003). *Contractarianism/Contractualism*. Blackwell.

Darwall, S. (2006). *The Second-Person Standpoint: Morality, Respect, and Accountability*. Harvard University Press.

Doggett, T. (2013). Saving the Few. *Noûs*, *47*(2), 302–315.

Dorsey, D. (2009). Headaches, Lives and Value. *Utilitas*, *21*(01), 36–58. https://doi.org/10.1017/S095382080800335X

Dougherty, T. (2013). Rational Numbers: A Non-Consequentialist Explanation Of Why You Should Save The Many And Not The Few. *The Philosophical Quarterly*, *63*(252), 413–427. https://doi.org/10.1111/1467-9213.12047

Dworkin, R. (2002). *Sovereign Virtue: The Theory and Practice of Equality* (4. print). Harvard Univ. Press.

Elson, L. (2014). Heaps and Chains: Is the Chaining Argument for Parity a Sorites? *Ethics*, *124*(3), 557–571. https://doi.org/10.1086/674844

Feltham, B., & Cottingham, J. (2010). *Partiality and Impartiality*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199579952.001.0001

Foot, P. (1967). The Problem of Abortion and the Doctrine of the Double Effect. *Oxford Review*, *5*, 1–7.

Freeman, S. (2007). *Rawls* (Reprinted). Routledge.

Frick, J. (2015). Contractualism and Social Risk. *Philosophy & Public Affairs*, *43*(3), 175–223.

Fried, B. H. (2012). Can Contractualism Save Us from Aggregation? *The Journal of Ethics*, *16*(1), 39–66. https://doi.org/10.1007/s10892-011-9113-3

Gauthier, D. (1986). *Morals by Agreement* (Repr). Clarendon Pr.

Gertken, J. (2016). Mixed Feelings About Mixed Solutions. *Ethical Theory and Moral Practice*, *19*(1), 59–69. https://doi.org/10.1007/s10677-015-9660-y

Godwin, W. (2000). *Enquiry Concerning Political Justice*. Batoche Books. http://ebookcentral.proquest.com/lib/reading/detail.action?docID=3117779

Halstead, J. (2016). The Numbers Always Count. *Ethics*, *126*(3), 789–802. https://doi.org/10.1086/684707

Harsanyi, J. C. (1975). Can the Maximin Principle Serve as a Basis for Morality? A Critique of John Rawls's Theory. *The American Political Science Review*, *69*(2), 594–606. JSTOR. https://doi.org/10.2307/1959090

Hausman, D. M. (2015). The Value of Health. In I. Hirose & J. Olson (Eds.), *The Oxford Handbook of Value Theory* (pp. 338–355). Oxford University Press. https://doi.org/10.1093/oxfordhb/9780199959303.013.0019

Heathwood, C. (2015). Monism and Pluralism about Value. In I. Hirose & J. Olson (Eds.), *The Oxford Handbook of Value Theory* (pp. 136–157). Oxford University Press. https://doi.org/10.1093/oxfordhb/9780199959303.013.0009

Hirose, I. (2001). Saving the Greater Number without Combining Claims. *Analysis*, *61*(4), 341–342. JSTOR.

Hirose, I. (2004). Aggregation and Numbers. *Utilitas*, *16*(01), 62–79. https://doi.org/10.1017/S0953820803001067

Hirose, I. (2007). Weighted Lotteries in Life and Death Cases. *Ratio*, *20*(1), 45–56. https://doi.org/10.1111/j.1467-9329.2007.00345.x

Hirose, I. (2013). Aggregation and the Separateness of Persons. *Utilitas*, *25*(02), 182–205. https://doi.org/10.1017/S0953820812000386

Hirose, I. (2015a). *Moral Aggregation*. Oxford University Press.

Hirose, I. (2015b). Skepticism about Value Aggregation. In I. Hirose & J. Olson (Eds.), *The Oxford Handbook of Value Theory* (pp. 300–313). Oxford University Press. https://doi.org/10.1093/oxfordhb/9780199959303.013.0017

Hobbes, T. (2009). *Leviathan*. Floating Press. http://public.eblib.com/choice/publicfullrecord.aspx?p=441243

Holm, S. (2018). The Luckless and the Doomed. Contractualism on Justified Risk-Imposition. *Ethical Theory and Moral Practice*, *21*(2), 231–244. https://doi.org/10.1007/s10677-018-9882-x

Holtug, N. (2015). Theories of Value Aggregation. In I. Hirose & J. Olson (Eds.), *The Oxford Handbook of Value Theory* (pp. 267–284). Oxford University Press. https://doi.org/10.1093/oxfordhb/9780199959303.013.0015

Hooker, B. (n.d.). *Parfit's Final Ethics*.

Hooker, B. (2003). Contractualism, Spare Wheel, Aggregation. In M. Matravers (Ed.), *Scanlon and Contractualism* (pp. 53–76). Frank Cass.

Hooker, B. (2005). Fairness. *Ethical Theory and Moral Practice*, *8*(4), 329–352. https://doi.org/10.1007/s10677-005-8836-2

Hooker, B. (2009). *Ideal Code, Real World: A Rule-Consequentialist Theory of Morality* (Repr). Clarendon Press.

Hooker, B. (2015). The Elements of Well-Being. *Journal of Practical Ethics*, *3*(1), 15–35.

Hooker, B. (2016). Rule Consequentialism. In *Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/archives/win2016/entries/consequentialism-rule/>

Horton, J. (2017). Aggregation, Complaints, and Risk: Aggregation, Complaints, and Risk. *Philosophy & Public Affairs*, *45*(1), 54–81. https://doi.org/10.1111/papa.12084

Horton, J. (2018). Always Aggregate. *Philosophy & Public Affairs*, *46*(2), 160–174. https://doi.org/10.1111/papa.12116

Hosein, A. (2009). *The Significance of Fairness* [Dissertation/Thesis, Massachusetts Institute of Technology]. https://search.proquest.com/docview/305099302?accountid=13460

Hosein, A. (2013). Contractualism, Politics, and Morality. *Acta Analytica*, *28*(4), 495–508. https://doi.org/10.1007/s12136-013-0198-0

Hurka, T. (n.d.). *Love and Reasons: The Many Relationships*.

James, A. (2012). Contractualism's (Not So) Slippery Slope. *Legal Theory*, *18*(3), 263–292. https://doi.org/10.1017/S135232521200002X

Kamm, F. M. (1998). *Morality, Mortality Volume I: Death and Whom to Save From It*. Oxford University Press. https://doi.org/10.1093/0195119118.001.0001

Kamm, F. M. (2005). Aggregation and Two Moral Methods. *Utilitas*, *17*(01), 1–23. https://doi.org/10.1017/S0953820804001372

Kamm, F. M. (2007). *Intricate ethics: Rights, responsibilities, and permissible harm*. Oxford University Press. http://public.eblib.com/choice/publicfullrecord.aspx?p=415447

Keller, S. (2013). *Partiality*. Princeton University Press.

Knight, C. (2017). Reflective Equilibrium. In A. Blau (Ed.), *Methods in Analytical Political Theory* (pp. 46–64). Cambridge University Press.

Kolodny, N. (2010). Which Relationships Justify Partiality? The Case of Parents and Children. *Philosophy & Public Affairs*, *38*(1), 37–75. https://doi.org/10.1111/j.1088-4963.2009.01173.x

Kolodny, N. (2011). Which Relationships Justify Partiality? General Considerations and Problem Cases. In B. Feltham & J. Cottingham (Eds.), *Partiality and Impartiality: Morality, Special Relationships, and the Wider World* (pp. 169–193). Oxford University Press. http://dx.doi.org/10.1093/acprof:oso/9780199579952.001.0001

Kumar, R. (1999). Defending the Moral Moderate: Contractualism and Common Sense. *Philosophy & Public Affairs*, *28*(4), 275–309. JSTOR.

Kumar, R. (2001). Contractualism on Saving the Many. *Analysis*, *61*(2), 165–170. JSTOR.

Kumar, R. (2003a). Reasonable Reasons in Contractualist Moral Argument. *Ethics*, *114*(1), 6–37. https://doi.org/10.1086/378034

Kumar, R. (2003b). Who Can Be Wronged? *Philosophy & Public Affairs*, *31*(2), 99–118. https://doi.org/10.1111/j.1088-4963.2003.00099.x

Kumar, R. (2009). Wronging Future People: A Contractualist Proposal. In A. Gosseries & L. H. Meyer (Eds.), *Intergenerational Justice* (pp. 251–272). Oxford University Press.

Kumar, R. (2015). Risking and Wronging. *Philosophy & Public Affairs*, *43*(1), 27–51. https://doi.org/10.1111/papa.12042

Kumar, R. (2017). Contractualist reasoning, HIV cure clinical trials, and the moral (ir)relevance of the risk/benefit ratio. *Journal of Medical Ethics*, *43*(2), 124–127. https://doi.org/10.1136/medethics-2015-103127

Lang, G., & Lawlor, R. (2015). Numbers scepticism, equal chances and pluralism: Taurek revisited. *Politics, Philosophy & Economics*, *15*(3), 298–315. https://doi.org/10.1177/1470594X15618967

Lange, B. (2017). Restricted Prioritarianism or Competing Claims? *Utilitas*, *29*(2), 137–152. https://doi.org/10.1017/S0953820816000182

Lawlor, R. (2006). Taurek, Numbers and Probabilities. *Ethical Theory and Moral Practice*, *9*(2), 149–166. https://doi.org/10.1007/s10677-005-9004-4

Lenman, J. (2008). Contractualism and Risk Imposition. *Politics, Philosophy & Economics*, *7*(1), 99–122. https://doi.org/10.1177/1470594X07085153

Lippert-Rasmussen, K. (2019). Precís of Luck Egalitarianism. *Critical Review of International Social and Political Philosophy*, *22*(3), 245–252. https://doi.org/10.1080/13698230.2018.1438783

List, C., & Valentini, L. (forthcoming). What Normative Facts Should Political Theory Be About? *Oxford Studies in Political Philosophy*, *6*.

Lord, E. (2016). Justifying Partiality. *Ethical Theory and Moral Practice*, *19*(3), 569–590. https://doi.org/10.1007/s10677-015-9675-4

Lovett, F. (2010). *A General Theory of Domination and Justice*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199579419.001.0001

Lübbe, W. (2008). Taurek's No Worse Claim. *Philosophy & Public Affairs*, *36*(1), 69–85. JSTOR.

Marshall, S. (2003). Scanlon and Reasons. In M. Matravers (Ed.), *Scanlon and Contractualism* (pp. 13–32). Frank Cass.

Matravers, M. (2003). Introduction. In M. Matravers (Ed.), *Scanlon and Contractualism* (pp. 1–12). Frank Cass.

Mendus, S. (2003). The Magic in the Pronoun 'My'. In M. Matravers (Ed.), *Scanlon and Contractualism* (pp. 33–52). Frank Cass.

Moreau, S. (1998). Contractualism and Aggregation. *Ethics*, *108*(2), 296–311. https://doi.org/10.1086/233806

Mosquera, J. (2017). An Egalitarian Argument against Reducing Deprivation. *Ethical Theory and Moral Practice*, *20*(5), 957–968. https://doi.org/10.1007/s10677-017-9842-x

Munoz-Dardé, V. (2005). The Distribution of Numbers and the Comprehensiveness of Reasons. *Proceedings of the Aristotelian Society*, *105*, 191–217. JSTOR.

Murphy, L. B. (1998). Institutions and the Demands of Justice. *Philosophy & Public Affairs*, *27*(4), 251–291. JSTOR.

Nagel, T. (1989). *The View from Nowhere* (First issued as an Oxford University Press paperback). Oxford Univ. Press.

Nagel, T. (1995). *Equality and Partiality*. Oxford University Press. https://doi.org/10.1093/0195098390.001.0001

Nagel, T. (1999). One-to-One. *London Review of Books*, *21*(3), 3–6.

Nagel, T. (2012). Equality. In *Mortal Questions* (Canto ed, pp. 106–127). Cambridge University Press.

Norcross, A. (1997). Comparing Harms: Headaches and Human Lives. *Philosophy & Public Affairs*, *26*(2), 135–167. JSTOR.

Norcross, A. (1998a). Great Harms from Small Benefits Grow: How Death Can Be Outweighed by Headaches. *Analysis*, *58*(2), 152–158. JSTOR.

Norcross, A. (1998b). Speed Limits, Human Lives, and Convenience: A Reply to Ridge. *Philosophy & Public Affairs*, *27*(1), 59–64. JSTOR.

Oberdiek, J. (2017). *Imposing Risk*. Oxford University Press.

O'Neill, M. (2012). Priority, Preference and Value. *Utilitas*, *24*(3), 332–348. https://doi.org/10.1017/S0953820812000118

Otsuka, M. (2000). Scanlon and the Claims of the Many versus the One. *Analysis*, *60*(3), 288–293. JSTOR.

Otsuka, M., & Voorhoeve, A. (2009). Why It Matters That Some Are Worse off than Others: An Argument against the Priority View. *Philosophy & Public Affairs*, *37*(2), 171–199. JSTOR.

Otsuka, M., & Voorhoeve, A. (2018). Equality Versus Priority. In S. Olsaretti (Ed.), *The Oxford Handbook of Distributive Justice* (Vol. 1, pp. 65–85). Oxford University Press. https://doi.org/10.1093/oxfordhb/9780199645121.013.1

Parfit, D. (1992). *Reasons and Persons* (Repr). Clarendon.

Parfit, D. (1997). Equality and Priority. *Ratio*, *10*(3), 202–221. https://doi.org/10.1111/1467-9329.00041

Parfit, D. (2000). Equality or Priority? In M. Clayton & A. Williams (Eds.), *The Ideal of Equality* (pp. 81–125). Macmillan Press ; St. Martin's Press.

Parfit, D. (2003). Justifiability to Each Person. *Ratio*, *16*(4), 368–390. https://doi.org/10.1046/j.1467-9329.2003.00229.x

Parfit, D. (2011a). *On What Matters: Volume One*. Oxford University Press. https://doi.org/10.1093/acprof:osobl/9780199572809.001.0001

Parfit, D. (2011b). *On What Matters: Volume Two*. Oxford University Press. http://www.oxfordscholarship.com.idpproxy.reading.ac.uk/view/10.1093/acprof:osobl/9780199572816.001.0001/acprof-9780199572816

Parfit, D. (2012). Another Defence of the Priority View. *Utilitas*, *24*(3), 399–440. https://doi.org/10.1017/S095382081200009X

Parfit, D. (2017). *On What Matters: Volume Three*. Oxford University Press.

Persson, I. (2001). Equality, Priority and Person-Affecting Value. *Ethical Theory and Moral Practice*, *4*(1), 23–39. https://doi.org/10.1023/A:1011486120534

Persson, I. (2008). Why Levelling Down could be Worse for Prioritarianism than for Egalitarianism. *Ethical Theory and Moral Practice*, *11*(3), 295–303. https://doi.org/10.1007/s10677-007-9102-6

Peterson, M. (2009). The Mixed Solution to the Number Problem. *Journal of Moral Philosophy*, *6*(2), 166–177. https://doi.org/10.1163/174552409X402331

Peterson, M. (2010). Some Versions of the Number Problem Have No Solution. *Ethical Theory and Moral Practice*, *13*(4), 439–451. JSTOR.

Peterson, M. (2013). *The Dimensions of Consequentialism: Ethics, Equality and Risk*. Cambridge University Press. https://doi.org/10.1017/CBO9781139519243

Peterson, M. (2016). The Dimensions of Consequentialism: Reply to Schmidt, Brown, Howard-Snyder, Crisp, Andric and Tanyi, and Gertken. *Ethical Theory and Moral Practice*, *19*(1), 71–82. https://doi.org/10.1007/s10677-015-9662-9

Piller, C. (2017). Treating Broome Fairly. *Utilitas; Cambridge*, *29*(2), 214–238. http://dx.doi.org/10.1017/S0953820816000303

Plato, & Jowett, B. (1892). *The Republic*. Floating Press. http://public.eblib.com/choice/publicfullrecord.aspx?p=413152

Porter, T. (2011). Prioritarianism and the Levelling Down Objection. *Ethical Theory and Moral Practice*, *14*(2), 197–206. https://doi.org/10.1007/s10677-010-9231-1

Porter, T. (2012). In Defence of the Priority View. *Utilitas*, *24*(3), 349–364. https://doi.org/10.1017/S0953820812000039

Privitera, J. (2018). Aggregate Relevant Claims in Rescue Cases? *Utilitas*, *30*(02), 228–236. https://doi.org/10.1017/S0953820817000279

Quong, J. (2017). Contractualism. In A. Blau (Ed.), *Methods in Analytical Political Theory* (pp. 65–90). Cambridge University Press.

Rawls, J. (1999). *A Theory of Justice* (Rev. ed). Belknap Press of Harvard University Press.

Rawls, J. (2005). *Political Liberalism* (Expanded ed). Columbia University Press.

Raz, J. (2004). Numbers, With and Without Contractualism. In P. Stratton-Lake (Ed.), *On What We Owe to Each Other* (pp. 46–66). Blackwell.

Ricoeur, P., & McLaughlin, K. (2008). *Time and Narrative. Vol. 3: ...* (Repr). Univ. of Chicago Pr.

Ricœur, P., & McLaughlin, K. (2009a). *Time and Narrative. Vol. 1: ...* (Repr). Univ. of Chicago Press.

Ricœur, P., & McLaughlin, K. (2009b). *Time and Narrative. Vol. 2: ...* (Repr). Univ. of Chicago Press.

Ross, D. (2002). *The Right and the Good* (P. Stratton-Lake, Ed.). Oxford University Press. https://doi.org/10.1093/0199252653.001.0001

Rüger, K. (2018). On Ex Ante Contractualism. *Journal of Ethics and Social Philosophy*, *13*(3). https://doi.org/10.26556/jesp.v13i3.323

Sangiovanni, A. (2008). Justice and the Priority of Politics to Morality. *Journal of Political Philosophy*, *16*(2), 137–164. https://doi.org/10.1111/j.1467-9760.2007.00291.x

Sangiovanni, A. (2016). How Practices Matter. *Journal of Political Philosophy*, *24*(1), 3–23. https://doi.org/10.1111/jopp.12056

Sangiovanni, A. (2017). *Humanity without Dignity: Moral Equality, Respect, and Human Rights*.

Saunders, B. (2009). A Defence of Weighted Lotteries in Life Saving Cases. *Ethical Theory and Moral Practice*, *12*(3), 279–290. https://doi.org/10.1007/s10677-009-9157-7

Scanlon, T. M. (1975). Preference and Urgency. *The Journal of Philosophy*, *72*(19), 655–669. JSTOR. https://doi.org/10.2307/2024630

Scanlon, T. M. (1982). Contractualism and Utilitarianism. In A. Sen & B. Williams (Eds.), *Utilitarianism and Beyond* (pp. 103–128). Cambridge University Press.

Scanlon, T. M. (1998). *What We Owe to Each Other*. Belknap Press of Harvard University Press.

Scanlon, T. M. (2003a). Contractualism and Utilitarianism. In *The Difficulty of Tolerance: Essays in Political Philosophy* (pp. 124–150). Cambridge University Press.

Scanlon, T. M. (2003b). Preference and Urgency. In *The Difficulty of Tolerance: Essays in Political Philosophy* (pp. 70–83). Cambridge University Press.

Scanlon, T. M. (2003c). Rights, Goals, and Fairness. In *The Difficulty of Tolerance: Essays in Political Philosophy* (pp. 26–41). Cambridge University Press.

Scanlon, T. M. (2003d). Précis of What We Owe to Each Other. *Philosophy and Phenomenological Research*, *66*(1), 159–161. https://doi.org/10.1111/j.1933-1592.2003.tb00249.x

Scanlon, T. M. (2011). How I Am Not a Kantian. In S. Scheffler (Ed.), *On What Matters: Volume Two* (pp. 116–140). Oxford University Press. https://doi.org/10.1093/acprof:osobl/9780199572816.001.0001

Scanlon, T. M. (2012). The Appeal and Limits of Constructivism. In J. Lenman & Y. Shemmer (Eds.), *Constructivism in Practical Philosophy* (1st ed, pp. 226–242). Oxford University Press.

Scanlon, T. M. (2014). *Being Realistic About Reasons*. Oxford University Press.

Scanlon, T. M. (2018). *Why Does Inequality Matter?* (First edition). Oxford University Press.

Schelling, T. C. (2006). Should Numbers Determine Whom to Save? In *Strategies of Commitment and Other Essays* (pp. 140–146). Harvard Univ. Press.

Schönherr, J. (2018). Still Lives for Headaches: A reply to Dorsey and Voorhoeve. *Utilitas*, *30*(02), 209–218. https://doi.org/10.1017/S0953820817000152

Sharp, D., & Millum, J. (2018). Prioritarianism for Global Health Investments: Identifying the Worst Off: Prioritarianism for Global Health Investments. *Journal of Applied Philosophy*, *35*(1), 112–132. https://doi.org/10.1111/japp.12142

Sheinman, H. (2011). Act and Principle Contractualism. *Utilitas*, *23*(03), 288–315. https://doi.org/10.1017/S095382081100015X

Shelby, T. (2007). Justice, Deviance, and the Dark Ghetto. *Philosophy & Public Affairs*, *35*(2), 126–160. JSTOR.

Shelby, T. (2016). *Dark Ghettos: Injustice, Dissent, and Reform*. The Belknap Press of Harvard University Press.

Sidgwick, H. (1981). *The Methods of Ethics*. Hackett Pub. Co.

Smart, J. J. C., & Williams, B. (1973). *Utilitarianism: For and Against*. University Press.

Smith, M. (1995). Internal Reasons. *Philosophy and Phenomenological Research*, *55*(1), 109–131. JSTOR. https://doi.org/10.2307/2108311

Southwood, N. (2010). *Contractualism and the Foundations of Morality*. Oxford University Press.

Southwood, N. (2018). Contractualism for Us As We Are. *Philosophy and Phenomenological Research*. https://doi.org/10.1111/phpr.12500

Stone, P. (2007). Why Lotteries Are Just. *Journal of Political Philosophy*, *15*(3), 276–295. https://doi.org/10.1111/j.1467-9760.2006.00274.x

Stratton-Lake, P. (2004). Introduction. In P. Stratton-Lake (Ed.), *On What We Owe to Each Other* (pp. 1–18). Blackwell.

Suikkanen, J. (2014). Contractualism and the Conditional Fallacy. In M. Timmons (Ed.), *Oxford Studies in Normative Ethics, Volume 4* (pp. 113–137). Oxford University Press. https://doi.org/10.1093/acprof:oso/9780198722144.003.0006

Tadros, V. (2019). Localized Restricted Aggregation. *Oxford Studies in Political Philosophy*, *5*.

Taurek, J. M. (1977). Should the Numbers Count? *Philosophy & Public Affairs*, *6*(4), 293–316. JSTOR.

Temkin, L. S. (2009). Aggregation Within Lives. *Social Philosophy and Policy*, *26*(1), 1–29. https://doi.org/10.1017/S0265052509090013

Temkin, L. S. (2012). *Rethinking the Good: Moral Ideals and the Nature of Practical Reasoning*. Oxford University Press.

Thomas, A. (2009). *Thomas Nagel*. Acumen.

Thomas, A. (2012). Giving Each Person Her Due: Taurek Cases and Non-Comparative Justice. *Ethical Theory and Moral Practice*, *15*(5), 661–676. https://doi.org/10.1007/s10677-012-9358-3

Timmerman, P. (2015). Contractualism and the Significance of Perspective-Taking. *Ethical Theory and Moral Practice*, *18*(5), 909–925. https://doi.org/10.1007/s10677-014-9543-7

Timmermann, J. (2004). The Individualist Lottery: How People Count, but Not Their Numbers. *Analysis*, *64*(2), 106–112. JSTOR.

Timmons, M. (2004). The Limits of Moral Constructivism. In P. Stratton-Lake (Ed.), *On What We Owe to Each Other* (pp. 90–122). Blackwell.

Tomlin, P. (2012). On Fairness and Claims. *Utilitas*, *24*(02), 200–213. https://doi.org/10.1017/S0953820812000143

Tomlin, P. (2017). On Limited Aggregation. *Philosophy & Public Affairs*, *45*(3), 232–260. https://doi.org/10.1111/papa.12097

van Gils, A., & Tomlin, P. (forthcoming). Relevance Rides Again? Aggregation and Local Relevance. *Oxford Studies in Political Philosophy*, *6*.

Velleman, J. D. (2010). *The Possibility of Practical Reason* (Reprint.). Scholarly Publ. Office, Univ. of Michigan Library.

Voorhoeve, A. (2014). How Should We Aggregate Competing Claims? *Ethics*, *125*(1), 64–87. https://doi.org/10.1086/677022

Voorhoeve, A., & Fleurbaey, M. (2012). Egalitarianism and the Separateness of Persons. *Utilitas*, *24*(3), 381–398. https://doi.org/10.1017/S0953820812000040

Williams, A. (2012). The Priority View Bites the Dust? *Utilitas*, *24*(3), 315–331. https://doi.org/10.1017/S0953820812000106

Wintein, S., & Heilmann, C. (2018). Dividing the indivisible. *Politics, Philosophy & Economics; London*, *17*(1), 51–74. http://dx.doi.org/10.1177/1470594X17715248

Wolff, J. (2003). Contractualism and the Virtues. In M. Matravers (Ed.), *Scanlon and Contractualism* (pp. 120–132). Frank Cass.

Wolff, J. (2004). Scanlon on Well-Being. In P. Stratton-Lake (Ed.), *On What We Owe to Each Other* (pp. 32–45). Blackwell.